Do information sharing platforms facilitate coordination? Evidence from the collapse of a cartel*

Robert Clark
University of Toronto

Jean-François Houde
UW-Madison & NBER

October 24, 2026

Abstract

Platforms aggregating sales data face growing scrutiny, yet their effects on collusion and efficiency remain ambiguous. We study Québec's gasoline industry, where cartels operated until the Competition Bureau's 2006 investigation, while a market-intelligence provider collected stations' sales volumes. We model how volume reports complement communication and derive a revealed-preference test: participation should decline when collusion collapses if reports facilitate coordination. Participation fell 13 percentage points but persisted, suggesting both efficiency and collusive value. Volume shifted toward neighborhoods where stations stopped sharing, consistent with more aggressive pricing. We conclude that even passive platforms can facilitate collusion with granular and timely data.

*We thank Dennis Carlton and Michi Igami for excellent discussions, and participants at the IIOC, Greenline IO Day, MaCCi Summer Institute for Competition Policy, CRESSE Conference, Northwestern's Conference on Antitrust and Competition Policy, including David Byrne, Daniel Chaves, Marco Duarte, Joe Harrington, Kei Kawai, Jeanine Miklos-Thal, Rob Porter, and Otto Toivanen, for helpful suggestions. We used Claude (Anthropic) to assist with the writing, the code, and the proofs; the authors are responsible for all content.

1 Introduction

Third-party platforms that aggregate sales data are under growing antitrust scrutiny, yet there is no direct empirical evidence on whether they facilitate collusion or instead generate legitimate efficiencies for firms. The theoretical case for concern is well-established. Firms that observe each other's volumes can distinguish a rival's price cut from an unfavorable demand realization, allowing a cartel to punish deviations without punishing firms for demand fluctuations (Stigler, 1964; Green and Porter, 1984). The case for efficiency is also well-understood. Platforms that circulate quantity and cost data, such as Agri Stats and NielsenIQ, allow firms to benchmark their own performance and to track market conditions, services a firm would value whether or not it colludes (Kühn, 2001).

Distinguishing the two motives empirically requires direct observation of firms' decisions to supply information and a source of variation in the value of coordination that does not move the non-collusive value of the platform's service in the same direction. Existing empirical work typically has neither. Participation is often inferred from pricing patterns or from web scraping (Assad et al., 2024; Calder-Wang and Kim, 2024), and the collusive and the efficiency motives for sharing move together with demand.

Our setting provides both by combining unique data on firms' voluntary decisions to share information with quasi-natural variation arising from an antitrust investigation of collusion, allowing us to examine whether platform participation declined following the cartel's collapse. We study Québec's retail gasoline industry, where alleged cartels operated in Magog, Sherbrooke, Thetford Mines, and Victoriaville from 2002 until the Competition Bureau of Canada issued search warrants in May 2006.[1,2] The alleged cartels involved extensive communication and monitoring, and coordinated prices through a structured order of play that also reallocated market share across members through inter-temporal transfers. The investigation and subsequent convictions were followed by the collapse of collusive margins in the prosecuted markets. During this period, a market intelligence provider operated an information-sharing platform (the *Platform*) that surveyed stations about their sales volumes and sold compiled reports to subscribing firms. Crucially, we observe participation directly at the station-year level, both in the cartel cities and throughout the province.

The cartel's monitoring problem is the result of cost fluctuations. Members agreed on a uniform price, but wholesale costs moved often, so the price had to be adjusted frequently. A persistent price difference is visible at the pump and easy to punish, but the timing of a price change is costly to observe and requires extensive communication. During an adjustment, a station could delay an increase by a few hours, or cut ahead of the others, and gain volume while posting the agreed price for the rest of the week. Such a deviation leaves two kinds of evidence: complaints from rivals who notice it at the pump, which are suggestive but not conclusive, and sales volumes, which

[1] **Legal Disclaimer:** This paper analyses the alleged Québec retail-gasoline cartel case strictly from an economic point of view. We base our understanding of the facts mostly on documents prepared by the Competition Bureau. The Competition Bureau has not brought charges against the Platform. In 2026 it reached a consent agreement with the company that now operates the Platform, which changes how the information is shared with subscribers. The Bureau did not find that coordination had occurred.

[2] See Clark and Houde (2013, 2014) for analysis of these alleged cartels.

are conclusive but which stations cannot observe at their rivals. The court documents show that participants used the Platform to monitor precisely this dimension of compliance.

To formalize the role of information sharing we develop a model of repeated oligopoly pricing in which firms communicate about price adjustments and can share sales volumes through a market intelligence platform. The model captures the key features of the cartel documented in the court record. Prices are adjusted sequentially, with a designated last mover during price increases, who gains a market share transfer as a result. Firms can deviate either publicly, by cutting price, or secretly, by delaying their price increase. The latter is detected imperfectly, from rivals' complaints and from the volume reports. The cartel monitors compliance with a statistic built from the reports, comparing the last mover's volume to the average volume of the reporting members, and initiates a price war when that statistic falls too low. Because this statistic is a ratio, common demand shocks cancel out, allowing the cartel to distinguish deviations from fluctuations in market demand.

The model yields three key insights. First, platform reports and direct communication are complements. The volume reports provide verifiable evidence that screens out the false positives a complaint alone would trigger, raising the value of collusion. Second, the precision of the monitoring signal is determined endogenously by the number of platform participants, and the arrangement is easier to sustain the more of them report. Third, an investigation that makes communication more costly, or lowers the discount factor, lowers the collusive value of sharing information. Participation in the Platform should therefore fall, most among the followers with the least to gain from collusion. At the same time, the platform may retain value for non-collusive purposes, such as helping stations distinguish local demand shocks from changes in their competitive position. The response of participation to the investigation therefore provides a revealed-preference test of whether the Platform facilitated coordination during the collusive period.

We use our data on information sharing to estimate the decline in participation with an event study that compares stations in the four prosecuted markets to stations in cities elsewhere in the province where no allegations of collusion were made. Our findings suggest that sharing in the cartel cities fell by roughly 13 percentage points relative to the control cities by 2010–2011, from a flat pre-period. The market and station fixed-effects estimates agree closely, so the decline reflects incumbent stations changing behavior rather than entry and exit. The decline is gradual over 2006 to 2008, and participation did not collapse. Three quarters of the cartel-city stations still shared their volumes in 2011, which confirms that the Platform also had a non-collusive value.

The composition of the decline points to the stations with the least to gain from collusion. We use court records from the prosecution to classify stations in the cartel cities by ownership type. The stations that withdraw are those with pricing discretion. Most lessee dealers operate a single site and set their own prices, and branded lessee dealers decline by 8 percentage points and unbranded lessee dealers by 12. The cartel leaders, who organized the arrangement and ran several stations each, do not leave, and neither does the low-price guarantee chain that enforced price wars. This is consistent with the model, in which the organizers have the most to gain from collusion. A large province-wide convenience-store chain also contributes to the drop in participation: its sharing

rate falls by 18 points. The court record shows that it fought the charges and was not found guilty. Stations with large convenience stores also have less to gain from collusion, since cross-selling gives them a lower effective marginal cost of selling gasoline (Houde, 2012). We consider three alternative explanations for the decline in information sharing: stations cutting costs to fund their legal defenses, withdrawing to avoid the appearance of cartel involvement, and a decline in the reports' value for competitive purposes. However, none of these can simultaneously account for the fact that the stations most exposed to prosecution continued sharing, participation remained high in competitive markets, and volume moved toward the neighborhoods where stations stopped sharing after the cartel collapsed.

Our second result examines the competitive effects of the breakdown in information sharing. Stations that stopped sharing may have had the least to gain from collusion and therefore the strongest incentives to price aggressively once coordination broke down. We cannot test this directly, since we do not observe the volumes of stations that stopped reporting, and our price data are not sufficiently high-frequency to capture the resulting price wars. Instead, we test whether volume shifted toward neighborhoods in which a station stops sharing. Because nearby stations tend to match one another's prices rather than concede sales, we expect these neighborhoods to gain volume as a whole.[3] Consistent with this prediction, always-sharing stations within 500 metres of a station that stopped sharing gained 0.247 log points of quarterly volume between 2006Q3 and 2008, relative to always-sharing stations elsewhere in the same market and quarter. This reallocation is consistent with a breakdown in coordination that allowed more aggressive pricing in neighborhoods where stations stopped sharing.

Finally, we ask whether the same pattern appears outside the four prosecuted cities. Clark and Houde (2014) identify a group of Québec cities that priced like the cartel cities before 2006 and whose margins fell after the investigation, and the Bureau's documents show participants discussing price changes in some of them. Participation adds a new piece of evidence. In the eight smaller of these markets, sharing falls abruptly in 2007 and 2008, by about as much as in the cartel cities, and the hazard of ceasing to share rises at the same time. In Québec City, the decline starts in 2003 and continues after 2006, and it is concentrated in two unbranded discount chains while the branded networks continue to share. We interpret these results with caution. These markets were never investigated, so we have no record of how prices were set, and our price and volume data for them are too limited to identify the cause of the decline. The timing and the size of the decline are nevertheless what the model predicts if the investigation also disrupted coordination there.

Related literature. Early work on information sharing among oligopolists asks whether firms gain from sharing private cost or demand information before competing (Vives, 1984; Gal-Or, 1985; Raith, 1996). Our model belongs instead to the literature on collusion under imperfect monitoring, which asks what firms must observe in order to punish deviations without punishing bad luck. Stigler (1964) identifies the problem, and Green and Porter (1984) show that a cartel

[3]Our approach is similar in spirit to Hastings (2004) and Dafny (2009), who measure the effect of a merger from the price response of the merging firms' rivals, nearby gasoline stations in the first case and hospitals in the second.

whose members observe only a noisy statistic sustains collusion by reverting to a price war when the statistic is low, so that wars occur on the equilibrium path. Porter (1983) takes this prediction to the Joint Executive Committee railroad cartel. Abreu et al. (1991) and Kandori (1992) show that the precision and frequency of monitoring determine the range of discount factors over which collusion is sustainable, a comparative static that Proposition 2 applies to the number of members whose volumes are reported.

More closely related to the Platform setting we consider, Harrington and Skrzypacz (2007, 2011) study cartels that monitor sales rather than prices, and Awaya and Krishna (2016, 2020) show that cheap talk about private signals raises the collusive payoff. In both, the arrangement must make the reports truthful. In our setting, a surveyor reads the meter, eliminating concerns about truthful reporting. Moreover, a complaint is not cheap talk, but generates evidence the deviator does not control. We investigate what verified reports add to communication, and how deterrence depends on how many members are measured.

On the empirical side, the papers closest to ours are Byrne et al. (2025a,b), who study a subscription platform that shares station-level prices among competing gasoline retailers.[4] On that platform prices are observed every 15 to 30 minutes, so imperfect monitoring plays little role. The Platform we study shared volumes, which reveal the deviation that posted prices hide. Kreps (2026) studies the end of the Agri Stats exchange among turkey processors in 2019 with a structural model that separates the conduct effect of the exchange from what it teaches firms about demand. Earlier cartels exchanged sales statistics through trade associations, as in the Sugar Institute (Genesove and Mullin, 1998, 2001), and Kawai et al. (2026) and Marshall and Marx (2012) give other examples. Those statistics were self-reported and circulated among the members. To our knowledge, ours is the first empirical study of a commercial third party that sells verified, firm-level sales volumes to colluding and competing firms alike, and in which each firm's decision to share is observed. Benchmarking of this kind also has efficiency value. Grennan and Swanson (2020) show that hospitals with access to a database comparing their purchase prices with those of peer hospitals negotiate lower prices from their suppliers. A broader literature studies mandatory public transparency and finds that its competitive effects depend on who receives the information (Albæk et al., 1997; Luco, 2019; Ater and Rigbi, 2023; Rossi and Chintagunta, 2016). Aryal et al. (2022) study public announcements about capacity in the airline industry (see also Duso et al. (2026)).

The Platform shares features with hub-and-spoke arrangements, in which a central intermediary redistributes competitively sensitive information among competitors (Sahuguet and Walckiers, 2017; Garrod et al., 2021; Clark et al., 2024), and with third-party pricing algorithms, whose designers may have an incentive to facilitate coordination (Harrington, 2022, 2025; Calvano et al., 2025; Aleksenko and Miklós-Thal, 2026; Sugaya and Wolitzky, 2026) and which have been linked to coordinated pricing in retail gasoline and rental housing (Assad et al., 2024; Calder-Wang and Kim, 2024). The difference is that the Platform is passive. It aggregates and sells sales data without

[4]Other studies of collusion in retail gasoline include Wang (2008, 2009) and Byrne and deRoos (2019) for Australia, Chaves and Duarte (2025) for Brazil, Sorgard and Steen (2024) for Sweden, and Martin and Verboven (2026) for Belgium. Houde (2012) and Carranza et al. (2015) study competition in Québec's market.

any intent to coordinate, and it facilitated collusion nonetheless. Participation reflected both a collusive and a non-collusive motive, and the collapse of the cartel lets us bound the collusive one.

2 Data and market description

This section characterizes the industry, the price-fixing case, and the three data sources we use. We first describe the industry and the case. We then describe the price data and the classification of markets, the station-level volume and characteristics data for the cartel cities, and the information-sharing data. A timeline of the case and of the data coverage is in Online Appendix Figure G.1.

2.1 Québec's retail gasoline industry and the price-fixing case

Stations differ in their vertical arrangement with refiners and distributors. Branded stations sell under a refiner's trademark and operate under one of three contractual arrangements, which differ in who owns the fuel. At a company-operated station the refiner owns the fuel and runs the station. At a commission-operated station the refiner owns the fuel and the dealer earns a per-litre commission. A lessee dealer buys the fuel and resells it at a price it sets. The largest refiner also advertised a low-price guarantee program, which committed it to match the lowest posted price in each local market. The company relied on commission contracts and set prices centrally.

Unbranded stations, which carry a distributor's brand or none, are run by lessee dealers who operate on their own account, with the larger ones holding supply contracts with one or more refiners and the smaller ones buying at the posted rack price. At a company-operated or commission-operated site, the retail price is set by corporate policy, while lessee dealers, branded or unbranded, set the price independently. We do not observe the identity of the price setter, except in the markets targeted by the investigation, where the court records name the person who set the price at each station they list.

The largest convenience-store chain in the province also operated a network of stations and had contracts with all upstream brands. Most sites were owned by the chain, and prices were decided by a central office. Big-box retailers whose primary business lies outside gasoline also operated stations. They similarly authorized prices through a central pricing office and tended to post lower prices.

The Competition Bureau opened an investigation into price-fixing under Section 45 of the Competition Act. The investigation was triggered by a Victoriaville newspaper article published on June 6, 2004, which contained an interview with a station owner who said he had been harassed by competitors for refusing to join their price-fixing arrangement. The investigation focused on the retail gasoline markets of four cities in the southern part of the province: Victoriaville, Thetford Mines, Sherbrooke, and Magog.[5]

[5]Thetford Mines and Victoriaville are about 70 kilometres apart and are treated as distinct markets. Magog and Sherbrooke are closer together and have overlapping market boundaries. The Sherbrooke metropolitan area, which includes Magog, has a population of nearly 200,000. Victoriaville and Thetford Mines have populations of roughly 40,000 and 25,000. Partly because of a lack of data on Magog's stations, we sometimes refer to the Sherbrooke market

Online Appendix Figure G.1 plots the timeline of the investigation and its overlap with our main sources of data. On March 3, 2005, the Bureau began recording private conversations among the suspected participants and deployed agents to corroborate the recordings. The wiretaps ran from March to June 2005 in Victoriaville and Thetford Mines, and from March to June 2005 and again from December 2005 to April 2006 in Sherbrooke and Magog. On May 29, 2006, the Bureau began executing search warrants. Charges were announced publicly in June 2008, with further rounds in 2010 and 2012. Fourteen companies and thirty-nine individuals were charged, drawn from approximately 130 stations and more than 60 firms.[6] Clark and Houde (2013, 2014) describe the case and the collusive arrangement in detail.

2.2 Retail margins and market classification

We obtain weekly data on retail prices and acquisition costs from the Régie de l'énergie du Québec, the province's energy board responsible for enforcing the price floor regulation since 1997 (Carranza et al., 2015).[7] For our purposes the floor is the benchmark against which we measure margins (the posted price net of the regulated floor). The Régie runs a weekly survey of posted retail prices to administer the policy. We use public data from the survey to build a weekly panel of city-level retail prices from 2002 to 2012. The price for a city-week is the median posted price for regular unleaded across the stations surveyed in that city that week.[8]

We use these data to construct our control group of competitive cities. Following Clark and Houde (2014), we sort cities into three categories: competitive, potentially collusive, and infringement. In that paper, a city is classified as competitive if its median pre-investigation margin is less than 6 cents per liter or the ratio of median price increases over price decreases is less than or equal to two. We replicate the classification at the city level using a k-median clustering algorithm. The two methods agree on all but three cities, which we drop.[9] We also drop competitive-cluster cities located in the same administrative region as a cartel city, to avoid contamination from geographic spillovers. Throughout the paper, "city" denotes a Régie survey unit and "market" a survey unit of the Platform, the data provider introduced in Section 2.3. Online Appendix Table G.4 maps the classified cities into the Platform markets used in the information-sharing analysis.

as including Magog. The volume subsection below states where Magog is included and where it is not.

[6] The allegations of the Competition Bureau have not been proven in court. Details of the charges (file 03079) are available at http://www.competitionbureau.gc.ca/eic/site/cb-bc.nsf/eng/03079.html.

[7] The price floor is updated every Monday and corresponds to the sum of the rack price, taxes, and a small adjustment for transportation cost that varies across administrative regions.

[8] We exclude cities in the remote regions of the province, namely Côte-Nord, Nord-du-Québec, and Gaspésie. The survey has one gap: it records no prices between March 10 and May 26, 2007. For nine cities we fill the gap with data from a second market research firm that surveyed the same markets over the same period, and for the remaining cities those weeks are missing. The nine cities are Drummondville, Gaspé, Montréal, Québec City, Rimouski, Saguenay, Sherbrooke, Trois-Rivières, and Val-d'Or. The gap coincides with the start of the ex-post period (introduced below), so it does not straddle a period boundary.

[9] The three are Sorel-Tracy, Senneterre and Mont-Laurier. Each is classified as non-cyclical by the threshold rule and as collusive by the clustering, and the three are dropped rather than assigned. The classification differs slightly from Clark and Houde (2014) because of differences in aggregation and sample construction, but the substantive grouping is the same.

Table 1: Retail Gasoline Margins and Price Adjustment Ratios by City Group and Period

	Collusive (2002–2006)			*Collapse* (2006–2007)			*Ex-post* (2007–2012)		
	Margin	P90-P10	Ratio	Margin	P90-P10	Ratio	Margin	P90-P10	Ratio
Panel A: Cartel cities									
Sherbrooke/Magog (n=2)	6.78	3.90	2.99	5.48	9.15	2.27	6.05	7.40	2.59
Thetford Mines (n=1)	7.80	4.81	3.00	4.70	9.20	2.12	6.90	6.60	2.77
Victoriaville (n=1)	6.95	6.20	2.67	2.80	8.60	0.51	2.50	8.40	0.87
Panel B: Potentially collusive cities									
Other coll. cities (n=18)	7.44	5.75	3.35	6.23	7.99	2.35	6.98	7.91	2.65
Panel C: Competitive control									
All comp. cities (n=31)	3.94	6.90	1.60	3.51	8.43	1.76	5.66	9.13	1.08

Notes: All statistics are first computed at the city level and then averaged across cities within each group. Margin is the average of city-level median weekly retail margins (cents/litre); weekly observations with |margin| > 20 are excluded as outliers. P90-P10 is the mean of city-level inter-decile ranges $p90 - p10$ of margins across weeks. Ratio is $|\tilde{\Delta}p^{+}/\tilde{\Delta}p^{-}|/|\tilde{\Delta}c^{+}/\tilde{\Delta}c^{-}|$, where $\tilde{\Delta}$ denotes a station-level median of the corresponding price or floor changes; values above 1 indicate asymmetric pass-through (rockets-and-feathers pattern). City groups are assigned via k-medians clustering on median margin, Ratio, and price stickiness in 2002–2006; cities in cartel regions are excluded from the competitive control group.

Table 1 summarizes the distribution of margins and price adjustments in three periods: the collusive period, from January 2002 to the search warrants of May 29, 2006; the collapse window, from May 30, 2006 to March 9, 2007; and the ex-post period, from March 10, 2007 to December 2012.[10] Panel A reports on the cartel cities, Panel B on the other collusive-cluster cities (18), and Panel C on the competitive controls (31). Online Appendix Figure G.2 shows the location of the three groups. The competitive cities are concentrated around Montréal and in the west of the province, and the potentially collusive cities in the east and around Québec City.

Before the investigation, in the prosecuted cities margins are elevated and stable, and adjustment ratios are well above one. The other collusive-cluster cities look the same. In contrast, margins are lower and more volatile in the competitive cities, and adjustment ratios are lower. The investigation changes this. In the collapse window, margins in Victoriaville fall to 2.80 cents per litre and its adjustment ratio falls to 0.51, below the symmetric benchmark of one, consistent with aggressive undercutting. Margins in Sherbrooke–Magog and Thetford Mines decline by less. The larger change is in stability. The interdecile range in the cartel cities rises to between 8.60 and 9.20 cents per litre, from 3.90 to 6.20 before the investigation, as margins begin to swing between near-zero during price wars and temporary spikes during failed attempts to restore coordination.[11] This instability persists beyond the investigation. Margins in Sherbrooke–Magog and Thetford Mines return to within one cent of their pre-investigation level, although margins in the competitive cities rise over the same years, so the premium over them is smaller than before 2006. The other potentially collusive cities follow the same path with smaller magnitudes. The competitive cities show the same pattern in all three periods: low margins and near-symmetric adjustment.

[10] This breakdown follows Clark and Houde (2014), which focused on the collusive and collapse periods.

[11] Online Appendix F reports a city-level difference-in-differences version for completeness.

2.3 Station characteristics and sales volume in the cartel cities

Our station-level data come from a market intelligence provider that surveyed the Canadian gasoline market during our sample period. We refer to it as the *Platform*. The Platform's surveyors visited every station in the cities they covered, whether or not the station subscribed to the company's data service, and recorded posted prices for each fuel grade and station characteristics. Prices and characteristics are observable from the forecourt and the surveyor collected them without the operator's assistance. Volume is the one variable the surveyor could obtain only with the operator's cooperation. We purchased two data sets. The first is a quarterly panel of volumes and station characteristics for the markets targeted by the investigation, described here. The second, described in Section 2.4, is an annual province-wide panel that records whether each station supplied its volume to the Platform.

The quarterly panel covers the four prosecuted cities, Victoriaville, Thetford Mines, Sherbrooke, and Magog, along with Drummondville. Drummondville was not targeted by the investigation. It lies about 45 kilometres from Victoriaville and falls in the potentially collusive cluster of Section 2.2. The panel runs from 2001 to 2008 for Victoriaville, Thetford Mines, and Drummondville, and to 2011 for Sherbrooke. It spans the cartel period, the investigation, and the June 2008 charges. The volume analysis of Section 5, and the summary statistics below, use the years 2001 to 2008, over which the three prosecuted markets with pre-investigation data are observed together. We do not use the Sherbrooke data after 2008.[12]

At each visit the surveyor reads the meter, and if available the Platform reports the total volume sold since the previous reading (i.e., over the quarter). We divide that quarterly total by the number of days between the two readings, so our volume variable is average daily sales in litres. The interval between readings averages 91 days. We sum the four gasoline grades and exclude diesel, which many stations do not sell.

Volume is missing for 17.5 percent of station-quarter observations. The Platform flags each observation as a direct read or as missing, but does not distinguish between voluntary refusals and exogenous shocks such as a meter reset or a pump failure. As a result, at least when a volume first goes missing, subscribers to the survey cannot distinguish a data error from a decision to withdraw from the Platform. For observations without a direct read the Platform supplies its own estimate, constructed from the station's historical average. We do not use this estimate in the analysis.

The survey also records the station characteristics we use as controls in the volume regressions, namely the number of gasoline pumps, the number of islands, the presence and size of a convenience store, a car wash, and a repair shop, the type of service, and the brand. We use station coordinates to construct local market structure variables. For each station and quarter we count the competing stations within 500 metres and within one kilometre, and decompose those counts by type (same-brand rivals, major-brand rivals, and distributor-brand rivals). Competition in retail gasoline is

[12]The Sherbrooke market as the Platform defines it includes Magog. Magog enters the survey only in the fourth quarter of 2005, so we exclude it from the volume analysis, which requires observations from before the investigation. Magog is included in the information-sharing analysis.

local, and these radii are the ones at which a station's volume responds to what its neighbours do.

Table G.5 in Online Appendix G describes the panel by market. The three markets differ in size in the way their populations imply. The average station in Thetford Mines sells 4,264 litres of gasoline a day, against 6,397 in Victoriaville and 6,685 in Sherbrooke, and Sherbrooke has 63 active stations against 23 in Victoriaville and 12 in Thetford Mines. Stations are also larger in Sherbrooke, with 11.1 gasoline pumps on average against 8.5 and 6.8. The brand composition differs in a way that matters for the analysis of who withdraws from the Platform. Major branded stations are the largest group everywhere, and the low-price guarantee chain holds 24 percent of the stations in Thetford Mines and Sherbrooke but only 9 percent in Victoriaville. Victoriaville has the largest share of unbranded and independent stations, 49 percent against 10 percent in Thetford Mines, which is consistent with the court record's description of that market as the one in which the arrangement was hardest to hold together. Local market structure is dense in all three. The average station has about one rival within 500 metres and two to three within a kilometre.

2.4 Information-sharing variables

The second data set purchased from the Platform is an annual province-wide panel of stations, covering 2002 to 2013. It records each station's location and brand, whether the surveyor obtained a volume reading that year, and, when no reading was obtained, the surveyor's stated reason. Coverage is the universe of retail gasoline outlets in the cities the Platform surveyed, roughly 1,800 stations a year across 45 or more markets, of which 30 match a classified city from Section 2.2 and enter our sample. Each annual observation is a snapshot from the first survey round of the year, taken in the first quarter.[13] The timing matters for the event studies. Since the survey is taken in the first quarter, the earliest observation that can reflect a response to the May 2006 warrants is the 2007 snapshot.

Two things determine whether a station shares. The first is the subscription, which required subscribers to provide their cumulative volume readings. The second is physical access to the meter. A station that did not subscribe could still let the surveyor read the meters, and many did. During our sample, stations were replacing analog pumps with digital systems. With the older equipment the surveyor could take the reading without assistance, and the default was to share. Digital systems require the operator to retrieve the figure, which makes sharing a deliberate act. An operator could therefore withhold its volumes either by not subscribing or by denying access to the meters. Prices and station characteristics were collected either way. The decision to share is made by whoever sets the price. At company-operated and commission-operated stations, and in the three chains described in Section 2.1, it is made by the head office. For lessee dealers, branded or unbranded, it is made at the station by the operator. Our data record whether a station supplied its volumes, not whether it purchased the compiled reports.

[13] In the largest markets the survey runs every two months rather than quarterly, on a February, April, June, August, October and December cycle, so the first round of the year is conducted in February. The read dates in the Québec City bi-monthly panel confirm this. The first read of the year falls in late February in every year from 1995 to 2001, the period for which we hold the dated file.

A volume can be missing for two kinds of reasons. Some are unrelated to the operator's decision, such as the surveyor's first visit to a new site, a power outage, a temporary closure for renovation, or an equipment failure (e.g., a meter reset). Others are a decision by the operator, such as an outright refusal, a partial refusal, or a denial of access to the meters. When no reading was obtained, the survey records the surveyor's explanation, and we use it to separate the two. We build a station-year indicator $Share_{it}$ equal to one when station i supplied its volume to the Platform in year t, and we write $NonShare_{it} = 1 - Share_{it}$. A station-year is coded as non-sharing when the surveyor records a refusal. It is coded as sharing when a reading was obtained or when the explanation describes a non-strategic reason. Equipment failures are ambiguous, since a mechanical excuse can also cover a refusal. We treat an isolated failure as genuine, and a failure that repeats in consecutive years or that is adjacent to a refusal as a refusal. Online Appendix H lists the survey codes and their frequencies, states the coding rules, and reports two alternative codings: one that counts only explicit refusals, and one that extends the post-investigation window.

To analyse which stations left, we also build a cross-sectional indicator $StopShare_i$. It equals one for a station that shared in every year up to 2006 and failed to share in at least one year between 2007 and 2009. It is missing for stations that are not observed on both sides of the investigation, except that stations that exited before 2006 are coded as zero. The indicator identifies 86 of the 2,028 stations for which it is defined, and 13 of the 121 stations in the cartel cities. It marks a station type, a marginal participant in the sharing arrangement, rather than a change at a date.

Participation was high before the investigation. Ninety-one percent of the 1,653 stations observed in 2006 across 29 markets supplied their volume. Appendix Table B.1 decomposes this rate with a linear probability model of $Share_{it}$ on station-type indicators and market fixed effects. The ordering follows the contractual structure described in Section 2.1. The baseline rate for independent and unclassified stations, 9 percent of the sample, is 62 percent. The low-price guarantee chain and the stations branded by a major refiner, together 68 percent of the sample, share at about 96 percent. Retail chains share at 91 percent and unbranded distributors at 80 percent. Vertically integrated chains and branded networks participate as a matter of corporate policy, while the stations with the weakest contractual ties to the Platform participate least.

Three further patterns are informative. First, stations that entered after 2002 (5 percent of the cross-section) are 11.2 percentage points less likely to share, conditional on type and market. This is consistent with the technological transition described above, since newer stations have electronic pumps that require the operator to retrieve the figure. Stations that exit before 2013 share at a 7.6 point higher rate, so exit alone lowers measured participation over time. Second, conditional on station type, sites owned by a refiner share at a 6.0 point higher rate than sites owned by a dealer, a co-operative, or a retail chain, and sites owned by a regional distributor at a 6.2 point lower rate.[14] Third, participation varies less across markets than within them. The standard deviation of the market fixed effect is 9.8 percentage points, about a third of the gap between branded stations

[14]The site owner holds the permit for its underground storage tanks in the registry of the Régie du bâtiment du Québec, matched to the survey by address (see Online Appendix G). The variable measures who owns the site and its tanks, not who sets the price, since a refiner-owned site can be run by a lessee dealer.

and the omitted category. Denser markets share more. The rate is 96 percent in central Montréal against 89 percent in the cartel cities.

3 Information sharing and cartel stability

In this section we consider the role of information for cartel stability. We first describe what the court record shows about how the cartel was organized and how it used the Platform's reports. Next, we develop a model of a cartel that monitors compliance with the help of the Platform, and derive the conditions under which participation makes collusion easier to sustain. We then ask what an investigation does to a cartel of this kind, and to the value of the reports to a station that has stopped colluding. The section closes with the predictions the empirical sections test.

3.1 The role of information sharing and communication in the Québec cartel

Clark and Houde (2013) study the collusive arrangement in the target cities using documents that the Competition Bureau submitted to the Superior Court of Québec, which describe the phone calls made around each price adjustment.[15] The cartel solved the coordination and enforcement problems of a price-posting market with two instruments, telephone communication before each adjustment and a fixed order in which the adjustment was carried out. Byrne and deRoos (2019) document price leadership with the same structure in Perth, without direct communication between the retailers.[16]

The order of play assigned each participant a role. Leaders initiated the increases, placed calls, and instructed others what price to post and when. A group of low-cost stations, including a chain with a low-price guarantee and big-box retailers, moved last on increases and first on decreases. By raising prices last, they sold at the old, lower price while rivals had already raised theirs, gaining a volume transfer that induced them to participate. They also led the retaliation by setting prices at the floor, where their costs allowed them a margin.[17] The remaining stations, the followers, adjusted in the prescribed order. Many were lessee dealers who set their own prices, so their compliance had to be obtained.

Because the order of moves transferred volume, compliance had to be assessed by quantities as well as prices. A station that posted the agreed price but raised several hours late captured volume without ever visibly violating the agreement. Rivals might notice and complain, but a complaint was not proof, and in the larger markets no one could watch every station. The Platform's reports supplied the missing evidence, namely the volumes of the other stations. Our reading of the record

[15] The documents report who called whom, at what time, and a summary or extract of what was said. They contain the Bureau's summaries of the recorded conversations and extracts from them, not the conversations themselves. They are close to a complete record of price changes in Victoriaville and Thetford Mines, and less complete for Sherbrooke and Magog, particularly price decreases. The online appendix of Clark and Houde (2013) explains coding of the calls.

[16] Byrne et al. (2026) study a related arrangement in Melbourne, where the leading retailers and the smaller ones bargained over the timing of each price restoration.

[17] In an April 26, 2005 conversation, the Victoriaville leader explained to a follower that in nearby Drummondville, where some stations had deviated from the agreement, the market "ended up at cost for 6 months."

is that this mattered most for the low-cost chains whose prices were set by a central office. The recorded conversations show that the managers of these stations had to clear a price change with a supervisor or a pricing centre, that the chains moved last as a matter of policy, and that the response to a station that did not follow was left to them, because only they could sustain a war.[18] Since central offices were not present to observe late adjustment, we infer that they relied on evidence provided by the reports to detect lost volume before authorizing retaliation.

The Bureau's documents show that participants used the Platform for these purposes. The wiretap period was one of open instability in the cartel cities and two extracts from the wiretaps illustrate the practice of using the Platform in the manner described.

> *On March 8, 2005, Agent 1 (Firm A) told Agent 2 (Firm B) that Agent 3 (Firm C) had called and said he was "not moving" because "they" (Firm C) had looked in the Platform and that in terms of sales volume, it was not in "their interest" (Firm C's interest). [...] Firm C does not want to go above 86.7 cents because Firm D (the dissident) is capturing too much volume, and Firm C does not want to "let it have it."*

> *On April 6, 2005, Agent 1 replied: "My supervisor told me why you are not following him (Independent Station X). I tried to explain the principle at the time... we didn't even worry about (Independent Station X); he wasn't turning over volume back then; he's now up to 5 million liters. My supervisor knows this because he went to look it up on the Platform."*

Both conversations are about establishing whether the volume outcomes implied by the agreed order of play had materialized, and deciding whether a response was warranted. In the first, a participant refuses to follow an increase because the Platform's figures show the dissident taking volume. In the second, a network supervisor uses the Platform to establish that an independent station once too small to be worth policing had grown large, which motivated the decision to stop following it. In both cases the reports are consulted to size up a particular rival and to calibrate the reaction to it.

The dissident named in the first extract is an independent operator in Victoriaville whose location meant that an increase could not hold without him (Clark and Houde, 2013). His refusals to raise price produced the two Victoriaville price wars before the investigation, in October 2004 and in December 2005 (Online Appendix Figure G.3). The second episode kept margins below the regulated floor for three weeks. He is also the operator whose newspaper interview on June 6, 2004 set the investigation in motion (Clark and Houde, 2013). A refusal to raise price on time might be seen from his sign and generate a complaint. What the participants learned from the Platform was the volume the dissident was taking, and therefore whether it was worth initiating a war. He never

[18] Two independent operators in Sherbrooke described this on April 12, 2006, after a corporate station had been slow to follow an increase: *"Agent 1 said that they cannot start a price war in Sherbrooke. [...] Agent 2 said that the corporate stations can bury them much faster. [...] She said that (Firm E) had said that if they go down to $1.09, they will not come back up."*

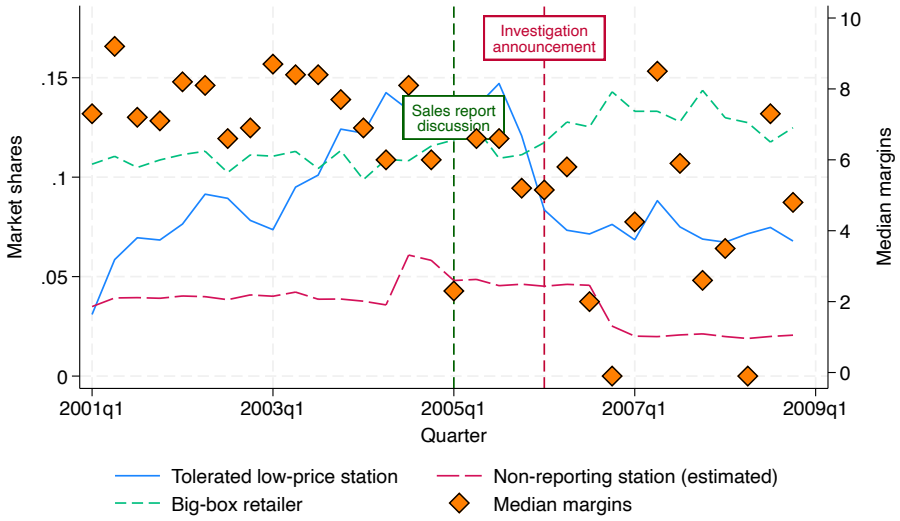


Figure 1: Market shares of non-cooperating stations and median margins in Victoriaville

Notes: Quarterly market shares (left axis) of the larger of the two tolerated low-price stations, of the dissident, and of the big-box retailer, and median margins in cents per litre (right axis), 2001 to 2009. The dissident did not report, so his share is the surveyor's estimate. The first vertical line marks the recorded conversations about the Platform's reports (March and April 2005), the second the investigation announcement. The court record attributes the price wars of October 2004 and December 2005 to the dissident; the wiretaps cited for the second are from the March to June 2005 campaign, and there is no Victoriaville recording from December 2005. In the weekly Régie series, margins fell to between 1.0 and 1.9 cents per litre for three weeks from October 9, 2004, and below the regulated floor for three weeks from December 17, 2005.

supplied his volumes, so the reports showed his gain through the shares the members were losing. The members' case for retaliating against the dissident rested on those market shares, supported by complaints about delays. This is the sequence that the model presented in Section 3.2 will formalize (a deviation, an assessment of its size from the volume reports, and a punishment of about three weeks at or below the floor).

The second extract refers to the two tolerated stations on the periphery. Both sit about 4.5 kilometres from the city center, facing each other on the same street, and each was allowed to post two cents below the collusive price. Both reported their volumes to the Platform. A fixed discount is verified at the pump and needs no data to police, but monitoring the timing of price adjustments is more challenging for more distant stations. The volume reports were where such a gap would show. That is what the supervisor in the second extract found. A station once too small to be worth policing had grown to five million litres, more than the agreed discount alone would be expected to produce.

Figure 1 shows how large these movements were. The larger tolerated station raised its share of the Victoriaville market from 5 to 7 percent in 2001 to nearly 15 percent in 2005. It lost about half of that share in the price war of December 2005, which ended its discount one quarter before the warrants, and did not recover. The big-box retailer that carried out that war moved the other way, from about 11 percent before 2005 to 13 or 14 percent in 2006 and 2007. These are changes of

50 to 70 log points for the tolerated station and about 20 for the big-box retailer, the same order as the reallocation we estimate in Section 5. The dissident's series is the surveyor's estimate. It is flat for long stretches but registers a gain in 2004 and a fall by half after the investigation.

The two extracts, and the price wars of October 2004 and December 2005, show that the arrangement was under strain well before the warrants, and that the participants used the Platform's reports to establish whether a rival had taken volume before deciding to retaliate. The case the Bureau built against the participants rested on wiretap evidence of their communication. Nowhere in the case was participation in the Platform treated as unlawful, and many stations continued to share their volumes after the warrants.

3.2 A theory of information sharing

We now develop a model of the arrangement just described. Its purpose is to illustrate what volume reports do for a cartel whose members can also communicate, and what happens to reporting when the arrangement weakens. Proofs are presented in Appendix A.

Firms, demand, and the arrangement. There are $n \geq 3$ firms selling a homogeneous product. One of them, labeled L, is the low-cost firm, and we normalize its marginal cost to zero. The other $n-1$ firms are the followers, each with constant marginal cost $c \in (0,1)$. The collusive price is normalized to one. A period is a quarter, the interval between two reports of the Platform, and firms discount at δ. Firm j's volume in period t is

$$q_{jt} = s_{jt}\, m_t\, u_{jt}, \tag{1}$$

the product of its share s_{jt}, a market-wide component m_t with mean $\bar{m}$, and an idiosyncratic component u_{jt} that is log-normal with mean one and independent across firms and periods and of m_t. A firm observes its own volume, and cannot distinguish between a weak market and a rival stealing share.

The collusive price has to be adjusted occasionally, for instance because wholesale prices change, and within a quarter it is assumed to be raised a fixed number of times. Adjustments are never simultaneous, so during part of each increase the firm that has not yet raised posts a price Δ_p below the others and serves more of the market. The arrangement assigns that position to L, because letting the low-cost firm move last is how the arrangement transfers share to it (Clark and Houde, 2013).[19] Over the quarter these episodes deliver to L a share $\tau > 0$ above the equal split $s \equiv 1/n$, funded equally by the $n-1$ followers, so shares are $s+\tau$ for L and $s-\tau/(n-1)$ for each follower. Price cuts are matched quickly, so we assume that only increases transfer share. We take τ and Δ_p as given, with $\Delta_p < 1-c$, and we assume that L posts the collusive price, holds the last-mover position, always supplies its volume, chooses the trigger described below and applies it

[19] The cartel leaders actually set the collusive price. For simplicity and to focus on the key mechanism, here we assign this role to L. We assume that the leaders set a price that L is willing to go along with.

mechanically, and does not deviate.[20] If the arrangement breaks down, firms play the static Nash equilibrium of the stage game for W periods: the price is c, L serves the market, and every follower earns zero. The punishment needs no credibility argument because it is a stage Nash equilibrium.

Deviations. A follower can depart from the arrangement in two ways. The first is a public price deviation in which it posts just below the collusive price and serves the market for the period. Prices are posted, so this is seen and attributed with certainty. The second is a secret delay. The follower joins L during an increase, raising late, and takes a share $d \in (0, \bar{d}]$, with $\bar{d} < s + \tau$, of the quarter's volume from L. The other followers are unaffected. It takes that share while still posting the low price, so it sells at a price Δ_p below the collusive one and gains in expectation

$$\bar{m}\, d\,(1 - \Delta_p - c), \tag{2}$$

which is positive when $\Delta_p < 1 - c$. Online Appendix C derives (2).[21]

Monitoring signals. The cartel monitors price adjustment delays in two ways. The first signal is the Platform's report, which publishes the volume of every member that supplied one and flags the members whose volume is missing, supplying its own estimate in its place, without stating a reason. In the model a member that shares supplies its volume with certainty, so a follower's volume is missing from the report only if the follower withheld it. In the data a reading is also missing when the surveyor could not take one, after a pump failure for instance. In the competitive cities, where no station has a collusive reason to hide its sales, this happens for 2 to 3 percent of readings. We do not model accidental missing reports here and return to them at the end of this section. Let r be the number of reporting followers. The cartel compares L's volume with the average of theirs,

$$Y = \log q_{Lt} - \frac{1}{r} \sum_{j \text{ reporting}} \log q_{jt}. \tag{3}$$

Every term in (3) carries $\log m_t$ with the same coefficient, so the market-wide component cancels exactly, whatever its distribution. The Platform's contribution, therefore, is to remove the common component of demand from the monitoring problem. Under the arrangement the mean of Y is $\mu_0 = \log[(s+\tau)/(s-\tau/(n-1))]$, which does not depend on how many followers report, while its variance is $\sigma_u^2(1+1/r)$, which does. Each report is one more observation of the common component, so the benchmark is sharper the more members supply one.

[20] We do not study L's decision to collude or to deviate. Its constraint is not the mirror image of a follower's. During a war L serves the market at price c and earns a positive profit, so its willingness to sustain the arrangement would place a lower bound on the transfer τ. This is not central to the paper and we set it aside. The questions we ask are about the followers.

[21] Both deviations are assumed to last one period, which keeps a single time index. The model could instead be written with a weekly period and a quarterly reporting interval. A public price deviation would then be punished immediately, since posted prices are seen at once, while a secret delay would be detected only after the reporting period. That version carries heavier notation and delivers the same message, so we keep one period throughout. One consequence of the single period is discussed in Online Appendix C: the war needed to deter a public cut is counted in weeks rather than quarters.

The second signal is communication. Because the members monitor prices and talk to one another, the cartel sometimes learns that a follower has been cheating. We model what it learns as a single binary message M, observed by every member at the same time as the report, and taking the value "complaint" with probability $\phi(d)$ after a delay of size d. We assume ϕ is continuously differentiable and increasing on $[0, \bar{d}]$, with $\phi(0) = \phi_0 \in (0,1)$, $\phi(\bar{d}) < 1$ and $\phi'(0) > 0$. The message is informative even at small deviations. That is, a complaint is strictly more likely after a small delay than after none. The message is imperfect in both directions, so it neither convicts nor exonerates on its own. Conditional on d, the message is independent of the volume shocks, so the two signals are different pieces of evidence and not two readings of the same one.

Price war trigger. The cartel's monitoring problem is a one-sided test. Under the arrangement Y has mean μ_0. The cartel tests the hypothesis that every follower complied against the alternative that L's share is lower than the arrangement implies, using Y as the test statistic. It rejects, and initiates a price war, when Y falls too far below μ_0, with a critical value that may depend on the message:

$$\text{war} \iff \frac{\mu_0 - Y}{\text{sd}(Y)} > \kappa_C \ \text{ if } M = \text{complaint}, \qquad \frac{\mu_0 - Y}{\text{sd}(Y)} > \kappa_N \ \text{ if not.} \tag{4}$$

Here $\text{sd}(Y) = \sigma_u \sqrt{1 + 1/r}$ is the standard deviation of the statistic under compliance, and κ_C and κ_N are the critical values in standard-deviation units. We assume that the cartel chooses the two critical values to deter every delay with as few price wars as possible on the equilibrium path.

We restrict the trigger in one respect. The critical values do not depend on the number of missing reports. The number of reports changes the statistic and its standard deviation, and the critical values are written in standard deviations of the statistic actually computed, but a missing report is not itself treated as evidence. We impose this because the Platform does not say why a volume is missing, so a rule that started a war on a missing volume would punish accidents at the same rate as withholdings. The restriction has a consequence that we use in what follows.

Lemma 1 (A deviator withholds)**.** *A follower that takes a secret delay of size $d > 0$ strictly prefers to withhold its volume from that period's report, for every trigger of the form* (4) *with at least one finite cutoff.*

The reason is that its own report is evidence against itself. If the deviator reports, the volume it stole enters the average in (3) and pushes the statistic down twice, once through L's lower volume and once through its own inflated one. If it withholds, only the first effect remains. The statistic is also computed on one fewer report, and because the cartel keeps its critical value in standard-deviation units, the same fall in L's volume is then less likely to cross it. Both effects favor withholding. From here on the secret delay is understood as a pair: delay by d, and withhold that period's volume. Withdrawal from the Platform is not a separate decision in the model. It is part of the deviation.

The two monitoring technologies are complements. The cartel chooses the trigger to make price wars as rare as possible while still deterring both deviations. Write α_0 for the frequency of wars on the equilibrium path and $\alpha(d)$ for the probability of a war after a delay of size d. Deterrence requires the excess $\alpha(d) - \alpha_0$ to be large enough at every d. The two signals enter this excess differently. A delay lowers the mean of Y, so the Platform signal raises the probability that the statistic crosses either cutoff in (4). Communication determines the weight the cartel puts on the two states, complaint and no complaint. Only a rule whose two cutoffs differ can use the message. Compare three cartels that differ only in what they can condition on: (A) both signals, with a cutoff for each value of the message; (B) communication only, so the rule is a pair of war probabilities; and (C) the Platform only, so a single cutoff applies. (C) is available to the cartel of regime (A), and Proposition 1 shows that (A) also improves on (B).

Proposition 1 (Two technologies are superior to one)**.** *Fix parameters at which regimes (B) and (C) are each sustainable. Then the lowest sustainable frequency of price wars satisfies* $\alpha_0^*(A) < \min\{\alpha_0^*(B), \alpha_0^*(C)\}$*, and the lowest discount factor at which collusion can be sustained satisfies* $\delta_{\text{crit}}(A) < \min\{\delta_{\text{crit}}(B), \delta_{\text{crit}}(C)\}$.

When both cutoffs are finite and the constraint against the delay binds at a single delay size, the optimal rule in regime (A) sets $\kappa_C < \kappa_N$ (Lemma C.1 in the Online Appendix). A complaint tells the cartel when a war is least likely to be a mistake. A complaint is not proof, however. Because $\phi_0 > 0$, complaints also arrive when nobody has cheated, and a cartel that acted on them alone would punish the innocent at that rate. The report supplies what a complaint lacks, a measured loss of volume that every member can verify, so it is never ignored in equilibrium. This is how we represent the calls the leader received. The distinction mattered in this cartel because retaliation was delegated to the low-cost chains, whose prices were set by a central office that needed evidence before authorizing a price cut.

Participation makes the arrangement more stable. Suppose k of the followers never report, for reasons outside the model, and let $k \leq n - 3$ so that a statistic in which the market-wide component cancels still exists.

Proposition 2 (Participation and stability)**.** *Fix the message technology and any trigger of the form* (4) *with at least one finite cutoff. Then* α_0 *does not depend on* k*, and for every* $d \in (0, \bar{d}]$ *the excess* $\alpha_k(d) - \alpha_0$ *is strictly decreasing in* k*. The critical discount factor* δ_{crit}*, over all triggers, is therefore strictly increasing in* k*: the set of discount factors that sustain collusion shrinks as more members stop reporting, and it is largest under full participation. Moreover, the smallest on-path war probability* α_0 *at which the cartel can deter every delay is strictly decreasing in the number of reporting followers: a cartel that re-optimizes its trigger starts fewer price wars the more of its members report.*

The statistic compares L's share with the average reporting follower's share. Each additional report averages out more of the demand noise, so the estimate loses precision as the number of

reports falls.[22] A given theft of volume is then less likely to cross the standardized cutoff, and every remaining follower's delay is harder to deter. Reporting is therefore a public good. Each report makes the constraint against the delay slack, and the cartel can take the benefit as fewer price wars, the lever in our model (Proposition 2), shorter wars, or a higher collusive price. Whichever it uses, every member collects its share of the profit the extra report makes sustainable, so a follower's own report has a positive private value, roughly $1/n$ of what the report is worth to the arrangement. This private value is the collusive component of a follower's willingness to share. It is proportional to the collusive profit and to the expected horizon of the arrangement, so a lower discount factor or a less informative message lowers it for every member, and it is zero once the members revert to the static Nash equilibrium. As members stop reporting the cartel moves from regime (A) toward regime (B), where it has communication and nothing else.

Accidental missing reports. Allowing volumes to be missing by accident, at the 2 to 3 percent rate observed in the competitive cities, changes two things. A missing volume becomes evidence, so the trigger may condition on the count of missing volumes. It is not obvious that the cartel wants it to, because every accident then risks a war, and the added wars lower the value of the arrangement and can make it unsustainable. The deviator's choice between reporting and withholding then depends on the size of the delay. A small delay is better hidden inside the report, because the deviator's own volume moves the average by little, and a large delay behind a missing volume. Under the restriction imposed above every delay is withheld, which is Lemma 1. Online Appendix D states the result and gives the argument. The next subsection keeps the restriction and treats a secret delay as the pair of Lemma 1, a delay and a withheld volume.

3.3 Predicted effects of a cartel investigation on information sharing

Our test compares participation in the Platform before and after May 2006. Its logic is that the investigation removed the collusive component of a station's willingness to share its volume and left the non-collusive component intact. This subsection explains how the investigation reduced the reports' collusive value, disrupted the arrangement, and left their non-collusive value intact.

The collusive value falls. The investigation lowers the collusive value of the reports in two ways. First, it makes communication costly. The Bureau's documents show that in the months before the warrants station owners and network managers were already reluctant to agree on prices by telephone, and at least two integrated firms instructed their lessee operators not to discuss prices with competitors. Complaints therefore carry less information about cheating. Second, it shortens the horizon of the arrangement, since the members now expect it to end. In the model the first change flattens the message and the second lowers the discount factor. Both reduce the collusive

[22] On the equilibrium path every follower reports. Under a trigger held fixed in standard-deviation units a complying follower is indifferent about doing so, since α_0 does not depend on the number of reports. Under the optimal trigger it strictly prefers to report, by its share of the price wars its report lets the cartel avoid. We do not model the decision to join because the question we ask is about participation after the arrangement is disrupted.

profit that a report helps sustain, and with it the private value of reporting for every member (Section 3.2). Participation should therefore fall, and it should fall unevenly. Stations for which the reports have value outside the arrangement keep sharing. Withdrawals are concentrated among the followers with the least to gain from collusion, such as a lessee dealer operating a single station or a follower with lower costs than its rivals. For them the reports are the first to stop being worth the risk that a reported volume starts a price war. The organizers, whose stations announce the increases, have the most to lose and should be the last to leave.

Disruption of the arrangement. Two outcomes are possible. One is that the arrangement collapses and the members revert to the static Nash equilibrium, the other is that it survives but in degraded form. Under the second outcome the trigger was set for a more informative message and a longer horizon, and the members cannot retune it, because that would require the kind of meeting the investigation made prosecutable. Delays that the old trigger deterred become profitable, and by Lemma 1 a station that delays also withholds its volume. Each withdrawal lowers the precision of the statistic for the members who remain (Proposition 2), so the next delay is harder to deter. We expect a chain of deviations rather than a single break. The two outcomes differ in what they predict for margins. Under static Nash, margins fall to the competitive level and stay there. Under the degraded arrangement, they move every time a deviation starts a war and every time an attempt to restore the price fails. Margin volatility in the cartel cities rose after 2006, in absolute terms (Appendix Table B.2) and relative to the competitive cities (Online Appendix F), which is consistent with the second outcome. Under that outcome the stations that stop sharing are the ones that should gain market share in the price wars that follow. Section 5 tests this prediction.

The non-collusive value does not fall. The revealed-preference argument needs one further element. We interpret the withdrawals as evidence that the collusive value of the reports fell. That reading is valid only if the non-collusive value did not fall at the same time. This value comes from benchmarking, and it accrues to the stations that subscribe to the reports. We do not observe which stations subscribed (Section 2.4), but a station that supplies its volume without subscribing has no non-collusive reason to do so, so its withdrawal reveals a fall in the collusive value alone. A subscribing retailer observes its own sales but not its demand curve. The volumes of nearby stations separate a weak market from a loss of position, and they are worth more when the environment is volatile, since own history is then a poor forecast of own sales. Online Appendix E gives one illustration, a station that orders before it knows its demand and pays for the error, so that its inventory cost rises with the variance of its forecast of own volume. The investigation made the environment more volatile. Dispersion of station volumes within a market and quarter rose in all three prosecuted markets we observe over the full sample, and the within-quarter standard deviation of the weekly retail margin rose by 1.1 to 1.8 cents per litre (Appendix Table B.2). The non-collusive value of the reports was therefore weakly larger after 2006, and the decline in participation reveals a fall in the collusive value and is a lower bound on that fall.

4 The effect of the cartel collapse on information sharing

Section 3 models a cartel that is more stable the more of its members report to the Platform, and a station whose willingness to share its volume has a collusive and a non-collusive component. The investigation reduced the first and weakly raised the second, so the change in participation after May 2006 is a lower bound on the fall in the collusive value. The model also predicts that participation falls first among the followers with the least to gain from collusion. We test these predictions by measuring the effect of the investigation on information sharing in the four prosecuted markets and by comparing the decline across types of stations. We then discuss alternative interpretations and quantify the possible spillovers of the investigation on other potentially collusive markets the Bureau never investigated.

4.1 Information sharing in the cartel cities

Event study. To estimate the dynamic effect of the cartel collapse on information sharing in the prosecuted markets, we estimate the following event-study specification:

$$Share_{it} = \alpha_i + \eta_t + \sum_{y \neq 2002} \beta_y \left(Cartel_i \times \mathbf{1}\{t = y\} \right) + \varepsilon_{it}, \tag{5}$$

where $Share_{it}$ is an indicator equal to one if station i shares volume data with the Platform in year t, α_i is a station (or market) fixed effect (depending on the specification), and η_t is a year fixed effect. The estimation sample includes stations in the four cartel cities and in the competitive control group. The reference year is the alleged starting date of the cartel (2002). Standard errors are clustered at the station level.

Figure 2 plots estimated coefficients from two specifications, one with market fixed effects and one with station fixed effects. The former identifies the effect from variation across stations within markets, while the latter identifies the effect from within-station variation over time, absorbing any compositional differences due to entry or exit.

The cartel cities exhibit no pre-trend. The coefficients for 2003–2005 are close to zero (0.01 to −0.01) under both specifications. After the investigation, sharing declines progressively, reaching approximately −0.13 by 2010–2011.[23] Because the province-wide data consist of annual snapshots taken during the first-quarter survey round, the earliest observation that can reflect a response to the May 2006 investigation is the 2007 data point.

The magnitude of this effect is substantial given that many stations in the sample are vertically integrated or belong to large retail networks whose participation is set by head-office policy or a corporate contract with the Platform (Section 2.4). A 13 percentage point decline in the overall sharing rate implies a larger behavioral shift among the subset of stations with genuine discretion over sharing. The close agreement between the market- and station-fixed-effects estimates confirms

[23] The six Montreal markets hold 840 of the control group's 1,041 stations in 2006. Excluding them leaves the estimates unchanged, for instance −0.120 (0.050) against −0.126 (0.043) in 2011. Online Appendix Figure G.8 reports the full series.

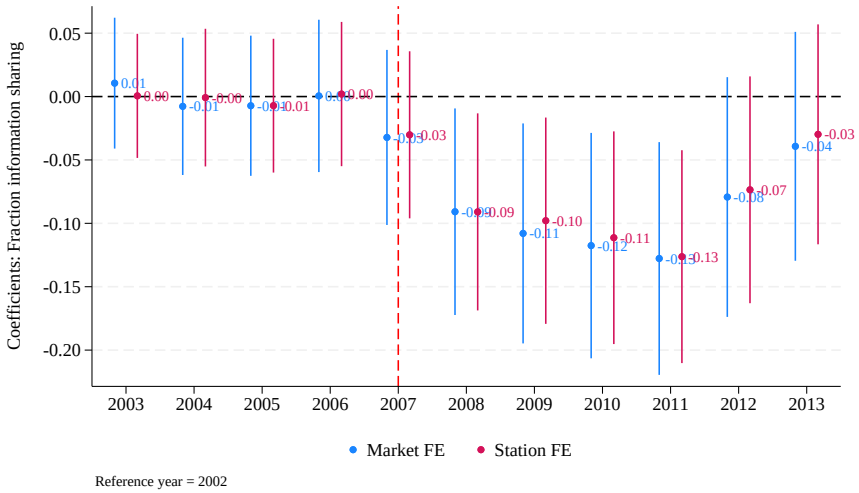


Figure 2: Effect of the cartel collapse on information sharing: cartel cities

Notes: Each series plots the estimated coefficients $\hat{\beta}_y$ from equation (5), estimated separately with market and station fixed effects. The reference year is 2002. The dashed vertical line marks the year of the investigation (2006). The control group consists of non-collusive markets. Standard errors clustered at the station level.

that the decline is driven by behavioral change. Incumbent stations that were previously sharing switched to non-sharing after the investigation, and this is not the result of entry or exit patterns.

The decline partially reverses after 2011, with sharing rates recovering toward pre-investigation levels by 2013. It is not the result of compositional change, since stations that exit the sample are more likely to share than continuing stations, so attrition alone would push sharing rates down, not up. We do not take a position on what this implies about conduct, because we do not observe coordination directly and the pricing evidence is mixed across markets. Margins recovered in Sherbrooke by 2012, though with higher dispersion than under collusion, while Victoriaville remained in recurrent price wars until the end of the sample (Online Appendix Table F.1). The years after 2008 are outside the period our design speaks to, which is the response to the investigation of the 2002 to 2006 conduct.

Who stops sharing? The event-study estimates measure the average decline in participation. The composition provides additional information. We use court records from the cartel prosecution to classify stations in the cartel cities by ownership type. Recall that monitoring was centralized. The cartel leaders and the low-cost firms examined the Platform's reports to enforce the arrangement (Section 3.1), while the followers were the ones being monitored. If the loss of collusive value was the motive, withdrawal should be concentrated among followers with pricing discretion, while the monitors, whose participation also carries intrinsic and contractual value, remain. The convenience-store chain is a special case among the organizers. The court record shows that it fought the charges and was not found guilty, so we treat it as a firm that exited the agreement rather than as a monitor.

Table 2: Information Sharing by Station Ownership in Cartel Cities

	Num. stations	Fraction volume sharing			Stop sharing
		Pre	Post	Δ	
	(1)	(2)	(3)	(4)	(5)
Panel A: Station ownership					
Cartel leaders	7	0.79	0.86	0.07	0.00
Convenience-store chain	17	0.96	0.78	-0.18	0.25
Hardware chain	4	0.75	1.00	0.25	0.00
Lessee dealers (branded)	32	0.84	0.76	-0.08	0.12
Lessee dealers (unbranded)	20	0.78	0.67	-0.12	0.13
Low-price guarantee chain	16	0.98	0.98	0.00	0.07
Panel B: Post-2006 survival					
Continuing stations	75	0.87	0.77	-0.09	0.16
Exiting stations	21	0.87	0.89	0.04	0.00
Observations	96	96	94	94	77

Sample: the 96 stations of the quarterly panel for Victoriaville, Thetford Mines and Sherbrooke that were active in the first survey round of 2006 and are observed in the annual sharing panel both before and after the May 2006 announcement of the investigation. "Volume sharing" is $Share_{it}$, the annual indicator equal to one when the station supplied its volume to the Platform. "Fraction volume sharing" is its mean over the survey years up to 2006 (pre) and over 2007 to 2009 (post). "Stop sharing" is $StopShare_i$ as defined in Online Appendix H: the station shared in every year up to 2006 and did not share in at least one of 2007 to 2009; column (5) is its mean among the stations that shared in every pre-period year. Columns (3) and (4) exclude the two stations with no survey observation in 2007 to 2009, so column (4) need not equal column (3) minus column (2). Panel B splits the sample by whether the station is still observed in 2013.

Table 2 reports sharing rates before and after the investigation for each ownership group among stations active in the cartel cities (excluding Magog, for which the quarterly panel starts in 2006). The pattern of withdrawal does not track legal exposure. The cartel leaders show a zero stop-sharing rate (column 5) and a slight increase in sharing from 0.79 to 0.86 (column 4).[24] The low-price guarantee chain incurred the largest fine, but continued sharing volume information. This chain behaved as if it were vertically integrated, with its participation governed by corporate policy, and shows no change in sharing (0.98 in both periods) and a stop rate of only 7%, associated with a single switch. The hardware chain, another corporate network, increased its sharing rate from 0.75 to 1.00 with a zero stop rate. Both facts point in the same direction. In cases where efficiency or contractual obligation dominate the participation decision, sharing is unchanged.

The convenience-store chain behaves as a firm that exited the agreement. Its sharing rate falls by 18 percentage points, and 25% of its stations stop sharing in at least one year after the investigation. Stations with large convenience stores also have less to gain from collusion, since cross-selling gives them a lower effective marginal cost of selling gasoline (Houde, 2012).

[24]The leaders' pre-period rate below one reflects two stations of one leadership firm in Sherbrooke: one did not share in 2002–2003, the other did not share before the investigation. The remaining leader stations shared in every pre-period year.

The stations that do withdraw systematically are instead those with genuine pricing discretion that were marginal participants in the collusive arrangement. Unlike the leaders, who ran several stations each, most lessee dealers operate a single site. Branded lessee dealers, who retain pricing autonomy under their franchise agreements, show an 8 percentage point decline with a 12% stop rate, and unbranded lessee dealers exhibit a 12 percentage point decline and a 13% stop rate. It is for these stations that the collusive monitoring value of the platform was likely decisive, and for which the collapse of the cartel removed the primary reason to participate. Moreover, among lessee dealers, the decline is at least as large for the 10 stations not named in the court documents (13 percentage points) as for the 42 that are named (8 points).

Robustness. We consider three potential concerns about the estimates and address each in turn. The first is the counterfactual. In a difference-in-differences framework, the post-2006 coefficients must be interpreted relative to what sharing would have looked like absent the investigation, typically extrapolating the pre-investigation trend. If that trend is not zero, the post-2006 estimate mixes the treatment effect with its continuation. We test the joint hypothesis that the pre-2006 event-study coefficients are zero, and we apply the relative magnitudes approach of Rambachan and Roth (2023) to compute breakdown values $\bar{M}$, the smallest deviation from linear parallel trends, expressed as a multiple of the largest pre-period violation, at which the 95 percent robust confidence interval on the post-period effect first contains zero. We report three horizons, the short run (2007), the medium run (2008–2010), and the long run (2011–2013).

Table G.1 in the Online Appendix summarizes the results. The pooled post-investigation effect is -0.079 with a market-clustered standard error of 0.016. These estimates use 2006 as the reference year and cluster by market, as the relative-magnitudes method requires. The estimated pre-period trend is -0.0042 (standard error 0.0064), and the joint test that the 2002 to 2004 event-study coefficients are zero has $p = 0.93$. The breakdown values are 0.20 in the short run, 0.75 in the medium run, and 0.00 in the long run, the last flagged as marginal because its standard 95 percent confidence interval marginally contains zero. The pre-period diagnostics show no trend. The breakdown values then say how large a post-2006 violation of parallel trends, relative to the largest pre-period violation, would be needed to overturn each estimate. The medium-run estimate, where the effect is largest, survives a violation three quarters that size. The short-run value is low because the response had only begun in 2007, and the long-run value is zero because participation recovers after 2011, so the 2011 to 2013 effect is itself not distinguishable from zero.

The second concern is that the decline reflects noisy fluctuations in participation rather than a response to the investigation. If the decline in the level of participation is a response to the investigation, it should be produced by a change in the rate at which stations stop sharing during the investigation period, between the May 2006 warrants and the June 2008 charges, and not by elevated churn in other years. Table B.4 in Appendix B reports year-by-year stop hazards (the probability that a station stops sharing in a given year, conditional on having shared in the previous one). On that subsample the stop indicator is the negative first difference of the sharing indicator,

so the specification is analogous to a first-difference regression and the market fixed effect absorbs a market-specific linear trend in participation. In the cartel cities the stop hazard exceeds that of the competitive control by 0.028 in 2007 and 0.067 in 2008, against a pre-period stop rate of 0.012 per year in the cartel cities over 2003–2006 (bottom rows of the table), and the 2007 and 2008 coefficients are jointly different from zero ($p = 0.04$), while the pre-2007 coefficients are jointly zero. The 2008 survey round precedes the June 2008 charges, so the 2008 estimate measures a response to the investigation itself rather than to the charges.

The hazard falls back to its pre-period level from 2009 onward. This is a property of the risk set as opposed to a reversal of behavior. The estimate conditions on sharing in the previous year, so the stations that stopped in 2007 or 2008 leave the sample, and those still sharing in 2009 are increasingly the ones whose participation is contractual or corporate rather than discretionary. The share of cartel-city stations remaining in the risk set falls from 0.85 to 0.89 before the investigation to 0.75 by 2010 and 2011, while the pooled competitive control stays near 0.90 (Table B.3).

The size of the decline is also large relative to the normal fluctuations in the competitive markets. Table B.3 in Appendix B reports raw sharing rates by year and market group. Rates never leave the 0.90 to 0.94 band in Montréal or the 0.84 to 0.89 band in the other competitive cities, over twelve years. Against that, the cartel cities fall from 0.89 in 2006 to 0.75 in 2011. The decline is several times larger than any movement the control markets produce.

The third concern is the coding of non-sharing. Our baseline classification combines the explicit refusal codes with the adjacency rules for ambiguous mechanical codes described in Online Appendix H. We re-estimate the event study and the sensitivity analysis under a strict alternative that counts only explicit refusals and treats every ambiguous mechanical entry as sharing, whatever its temporal pattern. The cartel-city event study is visually unchanged as can be seen in Figure G.6 in the Online Appendix. The pre-period coefficients remain close to zero and the decline still reaches about -0.13 by 2010–2011. The sensitivity bounds are somewhat weaker under the strict coding, but the null of no pre-trend in the cartel cities is still not rejected. The full set of figures and tables under the alternative codings is reported in Online Appendix G, and the codings themselves are described in Online Appendix H.

Taken together, the three exercises support the same conclusion. The May 2006 investigation caused a substantial decline in information sharing in the four prosecuted cities. The decline is not an artifact of compositional turnover, since the market and station fixed-effects estimates agree. It is not the continuation of a pre-existing trend, since the pre-period coefficients are jointly zero and the estimated pre-trend is economically negligible. It is concentrated in the two years following the warrants, which is when the hazard of ceasing to share rises. And the event-study estimates do not depend on the discretionary part of our coding rule.

4.2 Alternative interpretations

In this section we consider three alternative explanations for the empirical observations. The first is that stations cancelled the subscription to fund their legal defenses. The second is that participation

in the survey itself became a liability once the case was open. The third, and the most important, is that the value of the reports to a station that does not collude fell when the pricing regime changed.

Legal cost-cutting and financial constraints. Stations facing prosecution may have cancelled their subscriptions to redirect funds toward legal expenses. The decomposition in Table 2 is difficult to reconcile with this explanation. Cost-cutting predicts that the stations with the greatest legal exposure, and therefore the highest legal bills, should be the most likely to stop sharing. Those are precisely the stations that remain. The convenience-store chain that did withdraw had sufficiently deep pockets to mount a full legal defense and ultimately avoided conviction. For a firm of its size, subscription costs are trivially small relative to legal expenditures.

Reputational and discovery exposure. A second explanation holds that after May 2006 being documented as a participant in a volume-sharing arrangement was itself a liability, so that stations withdrew to avoid legal discovery or the appearance of cartel membership, whether or not the Platform retained any collusive value. Two facts weigh against this story. First, the Bureau's case was built on direct communication between competitors, and participation in the survey formed no part of it. No charge was brought against the Platform or its operator, sharing volumes was never alleged to be a facilitating practice, and the Bureau used the survey's data as evidence about the market. Second, participation is high and stable in the competitive markets across the province throughout the period. If survey participation by itself signalled cartel membership, this pattern would be difficult to understand.

The reputation story survives these facts, but it does not overturn the conclusion. If stations distanced themselves from the Platform because association with it looked incriminating, they did so because they perceived participation to be linked to the collusive conduct. The decline then still reveals that participants regarded the Platform as part of the collusive apparatus. What changes is whether it measures the Platform's monitoring value or the participants' perception of it. Either reading is inconsistent with a purely efficiency-motivated demand for the data. The story is also inconsistent with the fact that lessee dealers not named in the court documents withdrew at least as much as those that were.[25]

Regime-dependent efficiency value. The third explanation is that the non-collusive value of the reports fell after the investigation. The value of the reports to a station that does not collude need not be constant across conduct regimes. If a quarterly volume report is worth less in the competitive regime, part of what we attribute to lost collusive value is lost non-collusive value under a different regime, and the decline overstates the collusive component. Section 3.3

[25] We do not use the leaders' behaviour as evidence on this point. That the cartel leaders, the main targets of the investigation, stayed on the Platform is consistent with the appearance story if their lawyers advised that withdrawing after the raid would look like an admission that the Platform was part of the conduct. The argument is similar to Harrington (2004) on keeping post-cartel prices high to limit damages.

argued that the direction runs the other way, because benchmarking is worth more when the local environment is less predictable, and reported the increases in margin and share volatility that followed the investigation. Two further pieces of evidence, which do not rely on that argument, point the same direction.

The first is that participation in the competitive control cities is high throughout the sample and does not decline around the investigation. If the intrinsic value of a volume report collapsed whenever prices became volatile, participation should have fallen in these markets as well. By the dispersion measures of Table 1, the competitive cities have the widest week-to-week margin dispersion in the province in the collusive and ex-post periods, wider than the cartel cities before 2006, and participation there is the highest and flattest in the sample. Volatility does not depress demand for the reports where conduct is constant.

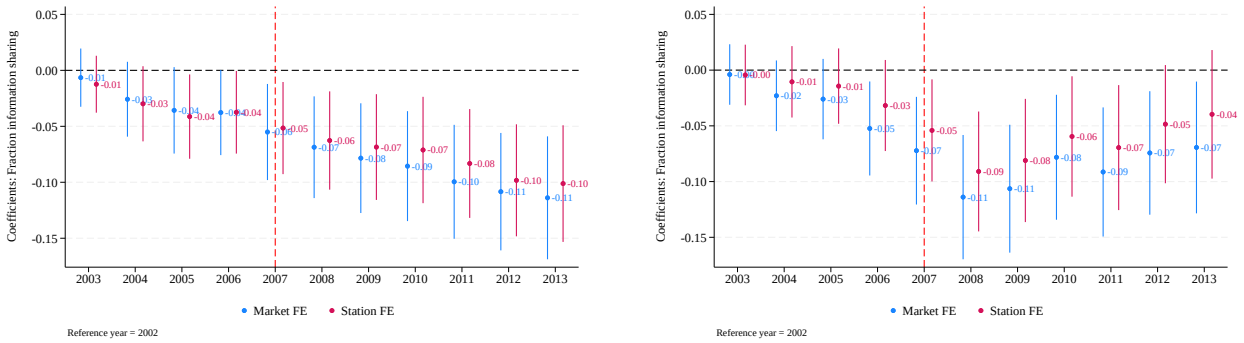
The second is a single potentially collusive market in which conduct changed without the investigation. Participation in Drummondville began to fall in 2003, three years before the warrants, when the entry of a discount distributor and a big-box grocery retailer destabilized the local arrangement and triggered a price war (Online Appendix Figure G.5). That was a shock to conduct and not to the informativeness of the data, and participation tracked the stability of the arrangement rather than the predictability of the environment.

4.3 Spillover effects in potentially collusive cities

The clustering analysis of Section 2.2 identified a second group of cities that priced like the cartel cities before 2006, and Clark and Houde (2014) document that margins and price adjustments in these markets also changed after the investigation in a manner similar to the cartel cities (albeit by a smaller magnitude). If firms in these cities were coordinating in the same way and with a similar ownership structure, their stations should have changed their sharing behavior too. We test this hypothesis with one important caveat. These markets were never investigated, so we have no wiretap or court record describing what was in place or what other shocks preceded May 2006.

We estimate the equivalent of equation (5), replacing $Cartel_i$ with an indicator PC_i for membership in the potentially collusive cluster, with the competitive-cluster cities as the control group and standard errors clustered by station. We examine Québec City separately, for two reasons. It is the largest market in the cluster, and it is the only one for which we have independent evidence on conduct. Chaves et al. (2025) document its transition from recurrent price wars to uniform pricing between 2000 and 2003. Figure 3 reports Québec City in panel (a) and the eight other markets pooled in panel (b).

Québec City. We find no evidence of a break in sharing at the investigation in Québec City. Panel (a) shows a secular decline that starts after 2002: about 4 percentage points by 2005, stability through 2006, and -0.10 to -0.11 by 2012–2013. In the survey years before our sample, 1991 to 2001, the share of stations with a direct volume reading fluctuates between 0.87 and 0.95 with



(a) Québec City

(b) Other potentially collusive (excluding Québec City)

Figure 3: Event studies for the potentially collusive cities

Notes: Each panel plots event-study coefficients from a regression of $Share_{it}$ on year interactions with the treatment-group indicator. Panel (a) reports market and station fixed-effects estimates for Québec City; panel (b) reports the same two specifications for the eight potentially collusive markets outside Québec City, pooled. Reference year is 2002. The control group consists of the competitive-cluster cities. (Grand-Mère is excluded throughout, since it has a single station that never shares). Standard errors clustered by station.

no trend,[26] so the decline starts with the transition from recurrent price wars to uniform pricing documented by Chaves et al. (2025). It is concentrated in two unbranded discount chains, which held 46 of the 265 stations in 2006, while the branded networks participate almost universally throughout (Online Appendix Table G.3). Margins show no decline apart from a price-war episode in 2006. Annual mean survey margins stay between 6.6 and 8.1 cents per litre from 2006 to 2013, above competitive-market levels throughout.[27] A possible interpretation is that the city-wide arrangement was solid enough to accommodate defections by a few independent stations, since the major branded networks, which hold most of the market, continued to share throughout. In the model the precision of the monitoring statistic depends on the number of reports through the factor $\sqrt{1+1/r}$, so the loss of a few reports matters little when r is large. The evidence does not indicate that the investigation destabilized coordination in Québec City.

The other markets. The picture in the eight remaining markets is different, and it resembles the cartel cities. Panel (b) shows an abrupt decline beginning in 2007 and reaching -0.091 by 2008, the same size and the same year as in the cartel cities. The stop hazards move at the same time: 0.031 in 2007 and 0.048 in 2008 (Table B.4 in Appendix B).

Panel (b) also shows a small drift in the pre-period, reaching -0.032 in 2006, which is not statistically different from zero. A drift common to the whole collusive cluster would also bias the cartel-city estimates. The cartel cities show no such pre-trend, and the drift here is not spread across the cluster. Table G.2 in the Online Appendix presents the average participation rate for

[26] The 1991–2001 survey records one summer read per station and year. A direct reading is the closest available measure to our sharing indicator.

[27] Margins are the weekly Régie survey price net of the price floor, averaged within the year across the Québec City survey municipalities. The competitive-market average over the same years ranges from 3.9 to 6.9 cents per litre.

each of the eight potentially collusive cities other than the cartel markets and Québec City. The pre-trend is concentrated in Drummondville and Rivière-du-Loup, which in 2006 had 33 and 14 stations, or 20 percent of the stations in these cities. With station fixed effects, only stations that change status contribute to the identification of the drift term, and twenty-one of the 262 stations that appear in the eight cities before 2006 change status in that window. The pre-period estimate is therefore driven by a small number of station-level events, most of them in Drummondville and Rivière-du-Loup.[28] In the sensitivity analysis of Table G.7 of the Online Appendix, dropping those two cities raises the pre-trend p-value from 0.68 to 0.90 and leaves the post-2006 estimate essentially unchanged, -0.050 against -0.052.

The bottom line is that the potentially collusive markets outside Québec City behave collectively like the prosecuted cities. This is consistent with the investigation having destabilized coordination in those markets as well.

5 Market consequences of information withdrawal

The announcement of the investigation had two consequences in the cartel cities: margins collapsed in price wars (Section 2.2), and information sharing declined (Section 4). Both point to a breakdown of the agreement. This section tests a third implication, that stations that stopped sharing had the least to gain from collusion, and therefore priced most aggressively once the agreement collapsed.

5.1 Hypothesis and identification

Section 3.3 predicts that the stations that stop sharing after the investigation are the ones that deviate from the arrangement, and that they price more aggressively than the stations that keep sharing. A station that delays its price adjustments and keeps reporting supplies the evidence that convicts it (Lemma 1). We cannot test this mechanism directly for two reasons. First, naturally we do not observe the sales volumes of stations that stopped reporting. Second, since price differences between stations are transitory, high-frequency prices and sales data would be required.[29]

To overcome these challenges, we examine whether volume shifts toward neighborhoods where a station stops sharing. Since nearby stations tend to match price cuts rather than concede sales, we expect a price differential to open across neighborhoods rather than within them, causing volume to increase across the neighborhood as a whole. This is not a prediction of the model, which has no spatial dimension. It is a property of retail gasoline markets with spatially differentiated

[28] In Drummondville, which our characteristics and volume data cover, the pre-2006 decline coincided with the entry of a discount distributor and a big-box grocery retailer in 2003, which triggered a city-level price war. Weekly margins sat at or below zero for nine weeks between November 2003 and January 2004, against an average of 6.6 cents in 2002, and participation fell from 0.94 in 2002 to 0.79 by 2006. We interpret this as a change in conduct triggered by a change in market structure.

[29] Evidence from the 2000 Québec City price war shows why price differences cannot last (Clark et al., 2025). Adjustments occur within days, and a station that holds its price above the market during a war loses volume quickly. Consistent with this, Figure G.10 in the Online Appendix re-estimates our event study with the posted-price snapshot as the outcome. The estimates stay within 0.15 cents per litre of zero at 500 m and within 0.3 at 1 km, and no post-investigation coefficient is statistically distinguishable from zero.

stations (a price cut is matched within hours or days by the stations that compete for the same traffic, and a war spreads from a few stations to the whole city faster than the weekly price survey can record (Clark et al., 2025)). The prediction is therefore that the whole neighborhood gains, not that the aggressor's closest rivals lose. We measure the spillover onto nearby stations. Always-sharing stations with a stop-sharing neighbor should gain volume relative to always-sharing stations without one, in the same city and quarter. This reallocation effect can be caused by isolated price wars following a deviation, or failed attempts to raise prices by stations in neighborhoods with no stop-sharing station. Both explanations attribute the reallocation to the breakdown of coordination that followed the investigation.

5.2 Estimation

The regression sample excludes the stop-sharing stations, which enter the regression only indirectly through the neighborhood indicators, and keeps every other station in the quarters for which the surveyor obtained a direct reading of its volume. We refer to these stations as always-sharing stations. The dependent variable, $\log V_{it}$, is the log of average daily volume at station i in quarter t. The sample excludes the 2006Q1 survey round.[30]

We estimate the following specification:

$$\begin{aligned}\log V_{it} = {} & \alpha_i + \mu_b + \gamma_{mt} + \beta_1 \mathit{StopNbr}_{it} + \beta_2 \mathit{StopNbr}_{it} \times \mathit{Post0608}_t \\ & + \beta_3 \mathit{NeverShareNbr}_{it} + \beta_4 \mathit{NeverShareNbr}_{it} \times \mathit{Post0608}_t \\ & + W_{it}'\lambda + X_{it}'\theta + \varepsilon_{it},\end{aligned} \tag{6}$$

where α_i is a station fixed effect, μ_b is a brand fixed effect, identified from stations that change brand, and γ_{mt} is a market×quarter fixed effect. $\mathit{StopNbr}_{it}$ is an indicator equal to one if at least one station within the radius (500 m or 1 km) has $\mathit{StopShare}_j = 1$. It varies over time only through the entry and exit of such neighbors, so under station fixed effects the level coefficient β_1 is identified from that variation alone. $\mathit{NeverShareNbr}_{it}$ is an indicator equal to one if at least one neighbor within the radius does not share in the survey year and is not a stop-sharing station. Most such neighbors never shared, and the others have mixed histories. $\mathit{Post0608}_t$ equals one for quarters between 2006Q3 and 2008Q4, the end of the sample.

The market×quarter fixed effects γ_{mt} absorb all market-level demand shocks, gasoline price movements, and seasonal patterns. Identification therefore comes entirely from within-market, within-quarter variation. The comparison is between always-sharing stations in the same market and quarter that differ in whether they have a stop-sharing neighbor.

The vector W_{it} captures the local competitive environment with full sets of indicator variables

[30]The estimates are not sensitive to the treatment of the transitional 2006 rounds. Including the 2006Q1 round yields a slightly smaller main interaction (0.239 at 500 m and 0.163 at 1 km); excluding also the 2006Q2 round, whose cumulative reading straddles the May 29 warrants, yields a slightly larger one (0.262 and 0.176). In both cases the estimate is significant at the 1 percent level at 500 m and at the 5 percent level at 1 km. Both movements are as expected: readings taken around the announcement mix the two regimes and attenuate the contrast with the reference period.

for the number and types of competing stations. The counts are the total number of rivals within the radius and, separately, the numbers of same-brand rivals, major-brand rivals, and distributor-brand rivals. Because these enter as indicators for each count, the specification absorbs the effect of any change in the number or type of nearby stations nonparametrically. This addresses the compositional concern with the design. If a station exits, its former customers would mechanically appear at neighboring stations. Exit and entry change the neighborhood counts and are absorbed by W_{it}, so the coefficients of interest are identified from stations whose local market structure is held fixed. The vector X_{it} includes the station amenities recorded at each survey visit (convenience store type, car wash, repair shop, service type, number of islands, and cubic polynomials in the number of pumps and in car capacity). Standard errors are clustered at the station level.

The coefficient of interest is β_2, which captures the differential volume effect of having a stop-sharing neighbor over 2006Q3 to 2008, relative to the pre-period. Pre-existing differences between neighborhoods with and without stop-sharing stations are absorbed by the station fixed effects, and the event-study version in Figure 4 tests for differential pre-trends. The never-sharing neighbor indicator $NeverShareNbr_{it}$ and its interaction are included as controls rather than as objects of interest.[31] The coefficients β_3 and β_4 hold fixed the neighborhoods that contain a never-sharing station, so that β_1 and β_2 compare neighborhoods with a station that *changed* its sharing status to neighborhoods in which every station shares.

The sample ends in 2008Q4. Volume data after 2008 are available only for Sherbrooke, as the quarterly panel for the other cartel cities ends in 2008, and the years after 2008 are outside the period our design speaks to (Section 4.1). The post-investigation period is therefore a single window, 2006Q3 to 2008Q4, which covers the investigation and the June 2008 charges. Table G.6 in the Online Appendix presents summary statistics for the main outcome and explanatory variables used in the analysis.

5.3 Results

Figure 4 presents the event-study version of the 500 m specification, interacting the stop-sharing neighbor indicator with year dummies (reference year 2005). Figure G.9 in the Online Appendix estimates the same specification for a 1 km radius. The pre-trend coefficients for 2001–2004 are small, between -0.02 and -0.06, and none is statistically different from zero, which is consistent with parallel trends. The effect appears in 2006 at 0.16 log points (the 2006 estimate pools the second to fourth quarters, the first being excluded) and reaches 0.24 in 2007 and 0.26 in 2008. This timing is consistent with the interpretation that the investigation triggered a behavioral change among stop-sharing stations.

Table B.5 in Appendix B reports the point estimates for both distance bands. At the 500 m radius (column 1), always-sharing stations located near a stop-sharing neighbor sell 0.247 log points

[31] The quarterly panel for the three markets over 2001 to 2008 contains 112 stations (Table G.5). Of these, 84 shared with the Platform in every survey year, 10 are stop-sharing stations, 11 never shared before 2006, and 7 have mixed histories, 3 of them because they are not observed on both sides of the investigation.

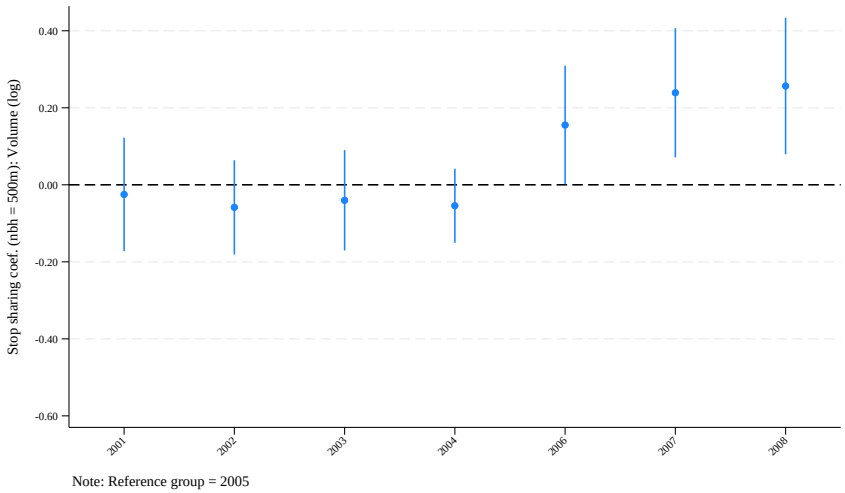


Figure 4: Volume event study: effect of stop-sharing neighbors (500 m)

Notes: Coefficients from a regression of log average daily volume on year interactions with **1**(Neighbor stop sharing within 500 m), controlling for never-sharing neighbor year interactions, neighborhood composition, station amenities, and station, brand, and market×quarter fixed effects. The sample runs from 2001 to 2008 and the reference year is 2005. Standard errors clustered at the station level.

more over the 2006Q3–2008 period than always-sharing stations in the same market and quarter whose neighbors all share ($p < 0.01$). The level coefficient ($\beta_1 = 0.013$) is identified only from entry and exit of stop-sharing neighbors and is not statistically different from zero. The results at the 1 km radius (column 2) follow the same pattern with an attenuated magnitude (the 2006Q3–2008 interaction is 0.166, $p < 0.05$). The estimate is identified from Sherbrooke and Victoriaville, since no station in Thetford Mines stopped sharing. Excluding Victoriaville, the 2006Q3–2008 interaction is 0.185 (0.058) at 500 m and 0.122 (0.053) at 1 km, so the result does not rest on one market. The spatial decay from 500 m to 1 km is consistent with local price competition as the mechanism. The competitive effect of a station that starts pricing aggressively dissipates with distance as consumers substitute among more proximate alternatives.

The reallocation is large: about 25 log points at 500 m. The specification includes market by quarter fixed effects, so the coefficient is a reallocation among always-sharing stations within a city and quarter rather than growth. What the estimate says is that neighborhoods containing a station that stopped sharing did about 25 log points better than neighborhoods that did not. Appendix Table B.2 reports the interquartile range of the log-volume residual within a market and quarter after the investigation: 0.167 in Sherbrooke, 0.284 in Thetford Mines, and 0.338 in Victoriaville. The gap is of the same order as the dispersion these markets display.

The reallocation can arise from two kinds of episode, and both rest on the same behavior. In the first, a station that stopped sharing cuts its price, its neighbors match, and the neighborhood draws volume from the rest of the city until the cut is matched city-wide. In the second, stations elsewhere in the city attempt to restore the price and are not followed by the stations near the stop-sharer, so the restoring stations lose volume to that neighborhood until they reverse the increase. The court

documents describe non-cooperation in exactly these terms, as a refusal to follow an increase. In both episodes the neighborhood of the stop-sharer is the one where the price stays low, which is what the estimate measures. The posted-price data cannot separate the two, for the reason given in Section 5.1.

The size of the estimate is in line with the reallocations that one price war produces in these markets. Until 2005 the larger tolerated station on the periphery of Victoriaville posted one to three cents below the city median and sold 12 to 15 percent of the market's volume (Figure 1). The price war of December 2005 ended the discount. From the first read of 2006 onward the station posts the median price, and its share falls by about half, to 7 to 8 percent of the market, without recovering. The second tolerated station follows the same path on a smaller scale. A movement of this size in response to a differential of a few cents is what the demand estimates of Clark and Houde (2013) imply for these markets, and it is the order of magnitude of our estimate. The episode also affects the estimate directly. Neither station has a stop-sharing neighbor, so both belong to the comparison group, and their loss contributes to the estimated gain of neighborhoods with a stop-sharing station. Excluding the two stations reduces the 2006Q3–2008 interaction from 0.247 to 0.192 at 500 m ($p < 0.01$) and from 0.166 to 0.100 at 1 km, where it is no longer statistically different from zero ($p = 0.12$).

The estimate measures the gain of the stop-sharing station's neighbors, not of the station itself. We do not observe the stop-sharing station's volume, since the surveyor imputes a constant rate for stations that refuse to report. If its neighbors matched its prices with a delay, the stop-sharing station gained more than they did, and the estimate is a lower bound on its own gain. Before the investigation the stop-sharing stations were not in decline. Relative to the always-sharing stations in the same market and quarter, their volumes in 2001 to 2004 differ from 2005 by 0.03 to 0.07 log points, none statistically different from zero, and in Sherbrooke they were larger than the average station, with 16 percent of the market's volume in 2004 and 2005. Whether they lost volume after they stopped reporting we cannot observe.

The volume reallocation connects the withdrawal from the Platform to pricing and market shares. Stations that withdrew from the information-sharing arrangement after the cartel collapsed are located in neighborhoods that subsequently experienced relative gains in sales volume. Combined with the evidence from Section 2.2 that price wars became a regular feature of the cartel cities after 2006, the volume reallocation supports the interpretation that stop-sharing stations were the ones driving the post-cartel competitive dynamics either by refusing to raise prices or by initiating price cuts.

6 Conclusion

This paper investigates whether a third-party platform that collects and redistributes sales volumes helps competing firms sustain a collusive arrangement, in retail gasoline in Québec. It leverages unique data on platform participation and a revealed-preference test developed around the collapse

of four city-level cartels after the Competition Bureau executed search warrants in May 2006. Two results follow. Participation in the Platform falls by 13 percentage points in the prosecuted cities relative to the control cities, from a flat pre-period. The withdrawals also affected the distribution of sales volume across neighborhoods. Always-sharing stations with a neighbor that stopped sharing gained roughly 25 log points in volume relative to always-sharing stations without one over 2006 to 2008, the years of price instability that followed the announcement of the investigation. Stations let their rivals observe their volumes, and they stopped once the arrangement that made the observation valuable had collapsed.

The model explains these findings. The cartel uses two kinds of evidence to monitor. Complaints are soft, and the Platform's reports provide hard evidence of lost volume that a central office could act on before authorizing retaliation. The two are complements, and the precision of the hard evidence depends on how many members agree to be measured. A member that cheats prefers not to be measured, so withdrawal from the Platform is part of the deviation, and the stations that stopped sharing after the investigation are the ones whose pricing shifted volume toward their neighborhoods. The court record also puts the Platform in perspective. The cartels relied heavily on direct communication, and the reports entered those conversations as the device used to verify volumes. The Platform facilitated coordination, but direct communication was the main tool.

Our identification relies on the response to the investigation between 2006 and 2008. The subsequent years provide useful context, but do not identify the role of the Platform in sustaining collusion. We describe these longer-run patterns without interpreting them as evidence about that mechanism. The prosecuted markets did not follow the same path after the collapse. In Victoriaville, margins stayed low. In the Régie series, which we extend to 2016 for this paragraph, the median weekly margin was below two cents in more than half of the weeks between 2009 and 2013, and it remained well below its pre-2006 level through 2016. Participation nevertheless returned to its pre-2005 level by 2013. None of the stations that stopped sharing after the investigation resumed. The increase comes from three stations that had not shared since 2004 or earlier, two of which had changed banner in the meantime. In Sherbrooke and Magog, price wars became rare after 2011 and margins stayed above their 2006 to 2010 levels through 2016, although the premium over the competitive cities remained smaller than before 2006. Participation also recovered, and most of the stations that resumed were stations that had stopped sharing after the investigation. The Sherbrooke pattern is consistent with stations having found other ways to coordinate without explicit communication. We do not observe conduct after 2006, so we cannot say whether the Platform played any part in it. In Victoriaville, the market appears to have become competitive.

These findings speak to how competition policy treats information-sharing arrangements. Until 2023 the DOJ and the FTC maintained a safe harbor, in their health-care enforcement statements, under which exchanges of competitively sensitive information were presumptively legal if they were managed by an independent third party, the information was more than three months old, and the data were aggregated across at least five participants with none accounting for more than 25 percent of any statistic. Both agencies withdrew the guidance in 2023 as overly permissive.

The Platform we study would not have qualified, because it reported station-level volumes for the preceding quarter.

The Platform also differs from those at issue in the Agri Stats and RealPage litigation in one important legal aspect, namely that it is passive. It sets no prices, makes no recommendations, and there is no evidence that it intended to coordinate anyone. It facilitated collusion nonetheless. What matters economically is whether the information a platform disseminates lets firms monitor each other precisely enough to sustain an arrangement. This is the criterion that Bennett and Collins (2010) propose, the nature and timing of what is shared rather than a per se rule. The Platform shared quantities only, the case that the usual distinction between sharing prices and sharing quantities (Kühn and Vives, 1995; Kühn, 2001) treats as the less harmful one.

Canadian authorities have since acted on this concern in the market we study. In September 2026 the Competition Bureau registered a consent agreement with the company that now operates the Platform. It concluded that distributing station-level data on sales volumes, prices, and station characteristics was an abuse of a dominant position, because it allowed competing retailers to obtain competitively sensitive information about one another and increased the risk of coordination. It did not find that coordination had occurred. Under the agreement, data on an individual retailer can no longer be given to another retailer, only data aggregated so that individual retailers cannot be identified can be distributed, pricing information cannot be distributed, and data can be released only for periods of at least one month and no sooner than ten business days after the period ends. The remedy limits the two dimensions our results point to, the granularity of the reports and their timeliness, and they matter together. Under its terms, the station-level reports that the cartel members consulted in 2005 could no longer be produced.

A Proofs

This appendix proves Lemma 1 and Propositions 1 and 2. Φ and φ are the standard normal distribution function and density, ϕ the message technology, with messages C (complaint) and N (none). The shock satisfies $\log u_{jt} \sim N(-\sigma_u^2/2, \sigma_u^2)$, independent across firms and periods and of m_t. The report and the message arrive at the end of the period, and a war, if triggered, starts the next period. A follower's collusive share is $s_f \equiv s - \tau/(n-1)$, it earns $\pi^{\text{col}} = \bar{m}\, s_f(1-c)$ per period of collusion and nothing during a war, and its gain from a delay of size d is (2), derived in Online Appendix C.

Incentive constraints. On the equilibrium path the trigger (4) starts a war with probability $\alpha_0 = \phi_0\Phi(-\kappa_C) + (1-\phi_0)\Phi(-\kappa_N)$, which does not depend on the set of reports the statistic is computed on, because the cutoffs are standardized. A war is W periods of zero profit, so $V^{\text{col}} = \pi^{\text{col}} + \delta[\alpha_0\delta^W + (1-\alpha_0)]V^{\text{col}}$. A public cut earns $\bar{m}(1-c)$ for one period and is punished with certainty; a delay of size d earns $\pi^{\text{col}} + \bar{m}d(1-\Delta_p-c)$ and is punished with probability $\alpha(d)$. Let $\rho(\delta) \equiv (1-\delta)/[\delta(1-\delta^W)] = 1/(\delta+\cdots+\delta^W)$, strictly decreasing in δ. Writing out the two deviation values, substituting V^{col}, cancelling $\bar{m}(1-c)$, and

dividing both by $\delta(1-\delta^W)$ and the second also by $s_f d$, the no-deviation conditions for every $d \in (0, \bar{d}]$ read

$$(1-\alpha_0)\, s_f \;\geq\; (1-s_f)\big[\rho(\delta)+\alpha_0\big], \tag{IC-P}$$

$$\frac{\alpha(d)-\alpha_0}{d} \;\geq\; \frac{R}{s_f}\big[\rho(\delta)+\alpha_0\big], \qquad R \equiv \frac{1-\Delta_p-c}{1-c} \in (0,1). \tag{IC-S}$$

The left side of (IC-S) is the *excess*, the only place the signals enter. The cartel minimizes α_0 subject to both constraints. Since (IC-P) is easier to satisfy the smaller is α_0, collusion is sustainable if and only if (IC-P) holds at α_0^*, the smallest α_0 at which some rule satisfies (IC-S); the minimum exists because, with cutoffs in $[-\infty, +\infty]$ and the left side of (IC-S) extended to $d=0$ by $\alpha'(0)$, each regime's feasible set is compact. Only the right sides involve δ, and they decrease in it, so δ_{crit} is the infimum of the δ at which some rule satisfies both. The set of such δ is closed, by the same compactness, so δ_{crit} is attained. A rule satisfies (IC-S) *with slack of order* d if $\alpha(d)-\alpha_0$ exceeds d times the right side of (IC-S) by at least γd on $(0, \bar{d}]$ for some $\gamma > 0$, that is, if the two sides of (IC-S) differ by at least γ.

Lemma A.1 (Perturbation)**.** *Let a regime (A) rule with a finite cutoff satisfy* (IC-S) *with slack of order* d*. A small common increase in both cutoffs, proportional to the slack, preserves* (IC-S)*, with slack of order* d *if the shift is small enough, and lowers* α_0 *by an amount of the same order. Hence (i) if the rule also satisfies* (IC-P)*, both constraints hold at some* $\delta' < \delta$*; (ii) if its* α_0 *is no larger than the optimal value of another problem (another regime, or another number of reporting followers), the optimal value of its own problem is strictly smaller.*

Proof. A common shift ε lowers α_0 at the rate $\phi_0\varphi(\kappa_C) + (1-\phi_0)\varphi(\kappa_N) > 0$. The excess vanishes at $d=0$ for every shift, and it and its derivative in the shift are C^1 in d, so its change is at most $B\varepsilon d$ for some $B < \infty$, and (IC-S) survives any $\varepsilon \leq \gamma/B$, with slack of order d if $\varepsilon < \gamma/B$. Both sides of (IC-P) then move in the right direction; continuity of the right sides in δ gives (i), and (ii) is immediate. Nothing changes when Y is computed on fewer followers, with z_k in place of z below. □

Proof of Lemma 1. Let $r \geq 2$ followers report, the deviator among them. If the deviator reports, Y is computed on r followers, its standard deviation is $\sigma_u\sqrt{1+1/r}$, and its mean falls by $-\log(1-d/(s+\tau)) + \frac{1}{r}\log(1+d/s_f)$ (L's volume falls, the deviator's rises). If it withholds, Y is computed on $r-1$ followers, the mean falls by the first term only, and the standard deviation $\sigma_u\sqrt{1+1/(r-1)}$ is larger, so the standardized fall z is strictly smaller. Since Y is normal on any set of reports, a branch of (4) with a finite cutoff κ fires with probability $\Phi(-\kappa+z)$, increasing in z, and each branch has positive probability because $\phi(d) \in (0,1)$. Withholding therefore strictly lowers the war probability at the same gain. □

Under full participation Y is computed on $n-2 \geq 1$ followers after a deviation, so m_t cancels. Its mean is $\mu(d) = \mu_0 + \log(1-d/(s+\tau))$; with $z(d) \equiv (\mu_0 - \mu(d))/[\sigma_u\sqrt{(n-1)/(n-2)}]$, increasing and convex with $z(0) = 0 < z'(0)$,

$$\alpha(d) = \phi(d)\,\Phi\big(-\kappa_C + z(d)\big) + \big(1-\phi(d)\big)\,\Phi\big(-\kappa_N + z(d)\big), \qquad \alpha(0) = \alpha_0. \tag{A.1}$$

Proof of Proposition 1. In regime (B) the rule is a pair of war probabilities $(\alpha_{0,C}, \alpha_{0,N})$, one per message, so $\alpha_0 = \phi_0\alpha_{0,C} + (1-\phi_0)\alpha_{0,N}$ and the excess is $[\phi(d)-\phi_0](\alpha_{0,C}-\alpha_{0,N})$; only the gap enters while α_0 rises in both, so the optimum sets $\alpha_{0,N} = 0$. In regime (C) a single cutoff κ applies, $\alpha_0 = \Phi(-\kappa)$, and the excess is $\Phi(-\kappa+z(d)) - \Phi(-\kappa)$.

Regime (C). Let κ, necessarily finite, solve the regime (C) problem, which regime (A) reproduces at $\kappa_C = \kappa_N = \kappa$, and consider $\kappa_C = \kappa - \eta$, $\kappa_N = \kappa + \eta\phi_0/(1-\phi_0)$ for small $\eta > 0$. The first-order terms in α_0 cancel, so α_0 changes by $O(\eta^2)$, while by (A.1) the excess at d rises by $\varphi(\kappa - z(d))[\phi(d) - \phi_0]\eta/(1-\phi_0) + O(\eta^2)$, uniformly in d after division by d since the excess vanishes at $d = 0$. The coefficient divided by d is continuous and positive on $(0, \bar{d}]$ with limit $\varphi(\kappa)\phi'(0)/(1-\phi_0) > 0$, which uses $\phi'(0) > 0$, so the perturbed rule has slack of order ηd. Lemma A.1 then lowers α_0 to first order in η, which dominates the $O(\eta^2)$ increase, so $\alpha_0^*(A) < \alpha_0^*(C)$.

Regime (B). Let $(\alpha_{0,C}, 0)$ solve its problem. If $\alpha_{0,C} < 1$, set $\kappa_N = +\infty$ and $\kappa_C = -\Phi^{-1}(\alpha_{0,C})$, which is finite. Then $\alpha_0 = \phi_0\alpha_{0,C}$ is unchanged and the excess exceeds the regime (B) excess by $\phi(d)[\Phi(-\kappa_C + z(d)) - \alpha_{0,C}] > 0$, which divided by d tends to $\phi_0\varphi(\kappa_C)z'(0) > 0$: slack of order d, and Lemma A.1(ii) applies. If $\alpha_{0,C} = 1$, then $\alpha_0 = \phi_0$ and the excess is $\phi(d) - \phi_0$. Set $\kappa_N = +\infty$ and κ_C finite, so that $\alpha_0 = \phi_0\Phi(-\kappa_C) < \phi_0$ and the excess exceeds the regime (B) excess by $\phi_0\Phi(\kappa_C) - \phi(d)\Phi(\kappa_C - z(d))$. By log-concavity of Φ, $\Phi(\kappa_C - z) \leq \Phi(\kappa_C)e^{-zh(\kappa_C)}$ with $h \equiv \varphi/\Phi$, so this is at least $\Phi(\kappa_C)\, b(d)$ with $b(d) \equiv \phi_0 - \phi(d)e^{-z(d)h(\kappa_C)}$, $b(0) = 0$ and $b'(d) = e^{-z(d)h(\kappa_C)}[\phi(d)z'(d)h(\kappa_C) - \phi'(d)]$. Since $\phi(d)z'(d) \geq \phi_0 z'(0)$ and $h(\kappa_C) \to \infty$ as $\kappa_C \to -\infty$, for κ_C negative enough b' is bounded below by a positive constant on $[0, \bar{d}]$, so (IC-S) holds at a strictly lower α_0 and $\alpha_0^*(A) < \alpha_0^*(B)$.

Discount factors. Wherever regime (B) or (C) is sustainable, the two steps give a regime (A) rule with slack of order d at a strictly lower α_0, so (IC-P) holds strictly and Lemma A.1(i) gives feasibility at a smaller δ: $\delta_{\text{crit}}(A) < \min\{\delta_{\text{crit}}(B), \delta_{\text{crit}}(C)\}$. □

Proof of Proposition 2. With $k \in \{0, \ldots, n-3\}$ non-reporting followers, α_0 does not depend on k. By Lemma 1 a reporting deviator withholds, so after a deviation Y is computed on $n - 2 - k \geq 1$ followers, $\mu(d)$ is unchanged, and the war probability $\alpha_k(d)$ is (A.1) with z replaced by $z_k(d) \equiv (\mu_0 - \mu(d))/[\sigma_u\sqrt{1 + 1/(n-2-k)}]$, strictly decreasing in k for $d > 0$. Each term is strictly increasing in $z_k(d)$ when its cutoff is finite, and one is, so the excess is strictly decreasing in k. For $k < k'$, a rule feasible at (k', δ) has at k the same right sides and an excess larger by an amount that, divided by d, is continuous and positive on $(0, \bar{d}]$ with a positive limit at 0: slack of order d. Lemma A.1(i) gives feasibility at some $\delta' < \delta$, so δ_{crit} increases strictly in k whenever collusion is sustainable at k' for some δ; applied to a rule attaining α_0^* at k', which has a finite cutoff since the excess vanishes otherwise, Lemma A.1(ii) gives a strictly smaller feasible α_0 at k. □

B Additional tables

This appendix collects the tables whose estimates are discussed in the text.

Table B.1: Prevalence of Information Sharing by Station Type (2006)

	(1)	(2)	(3)
	Mean	$Share_{it}$	$Share_{it}$
Low-price guarantee chain	0.24	0.342***	0.283***
	(0.43)	(0.039)	(0.041)
Branded stations (excl. LPG)	0.44	0.337***	0.284***
	(0.50)	(0.039)	(0.041)
Retail chain stations	0.07	0.289***	0.273***
	(0.25)	(0.045)	(0.045)
Unbranded distributor	0.16	0.178***	0.184***
	(0.37)	(0.044)	(0.044)
Entrant post 2002	0.05	-0.112***	-0.091**
	(0.21)	(0.040)	(0.042)
Exit pre 2013	0.23	0.076***	0.074***
	(0.42)	(0.016)	(0.016)
Site owned by a refiner	0.47		0.060***
	(0.50)		(0.018)
Site owned by a regional distributor	0.08		-0.062*
	(0.28)		(0.037)
Site owner unknown	0.02		-0.103
	(0.13)		(0.078)
Constant		0.618***	0.634***
		(0.039)	(0.039)
Observations	1653	1653	1653
Markets		29	29
Mean dep. variable		0.908	0.908
Std-deviation of market FE		0.098	0.098
Fixed effects		Market	Market

Notes: Column (1) reports means and standard deviations. Columns (2) and (3) are linear probability models of $Share_{it}$ in 2006 with market fixed effects. The omitted type is independent and unclassified stations. Column (3) adds the owner of the site from the 2003 (else 2013) tank-permit registry. The omitted owner is a dealer, a co-operative, or a retail chain. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table B.2: Volatility before and after the May 2006 investigation

Stations per quarter	Sherbrooke 51	Thetford Mines 11	Victoriaville 18
Panel A: Volume residual, standard deviation			
Before	0.120	0.135	0.182
After	0.154	0.181	0.260
Difference	0.034**	0.045***	0.079***
Panel B: Volume residual, interquartile range			
Before	0.134	0.175	0.221
After	0.167	0.284	0.338
Difference	0.033*	0.109***	0.117***
Panel C: Weekly margin, standard deviation			
Before	1.620	1.903	2.052
After	3.175	3.675	3.175
Difference	1.554***	1.772***	1.123***

Means across quarters of market-quarter volatility measures, 22 quarters before and 10 after the May 2006 warrants (2001Q1–2008Q4; the post period starts in 2006Q3). *Panels A and B*: dispersion within market-quarter of the residual of log station volume on station and city-quarter fixed effects, in log points, for stations with a direct reading. *Panel C*: within-quarter standard deviation of the weekly retail margin, in cents per litre. Differences are from a regression on a post-warrant indicator with Newey-West standard errors at four lags. *, ** and *** denote significance at the 10, 5 and 1 percent level.

Table B.3: Information-sharing rates by year and market group

Year	Cartel cities	Montreal	Other competitive	Québec City	Other pot. collusive
2002	0.852	0.900	0.845	0.908	0.863
2003	0.885	0.926	0.854	0.924	0.881
2004	0.879	0.938	0.866	0.917	0.873
2005	0.872	0.928	0.867	0.900	0.864
2006	0.889	0.933	0.886	0.906	0.846
2007	0.849	0.929	0.862	0.880	0.818
2008	0.790	0.930	0.859	0.867	0.775
2009	0.765	0.924	0.844	0.849	0.774
2010	0.755	0.917	0.867	0.841	0.806
2011	0.745	0.922	0.847	0.827	0.787
2012	0.800	0.931	0.842	0.825	0.810
2013	0.839	0.932	0.839	0.819	0.818
Stations (2006)	108	840	201	265	240

Notes: Mean of $Share_{it}$ by survey year and market group. Montreal comprises the six Montreal-area markets, and the other potentially collusive group excludes Québec City. Station counts are for 2006.

Table B.4: Stop-sharing hazard by year, market group and coding of non-sharing

	Cartel cities		Other pot. collusive		Québec City	
Year	(1)	(2)	(3)	(4)	(5)	(6)
2007	0.028	0.032	0.031*	0.016	0.012	0.017
	(0.022)	(0.022)	(0.016)	(0.013)	(0.013)	(0.012)
2008	0.067**	0.070**	0.048**	0.041**	0.001	0.001
	(0.030)	(0.029)	(0.019)	(0.018)	(0.011)	(0.009)
2009	-0.001	0.001	0.003	0.001	-0.004	-0.001
	(0.015)	(0.015)	(0.012)	(0.011)	(0.010)	(0.010)
2010	0.002	-0.008	-0.000	-0.001	0.000	0.004
	(0.015)	(0.006)	(0.009)	(0.009)	(0.010)	(0.010)
2011	0.005	0.008	0.014	0.008	0.006	0.005
	(0.015)	(0.015)	(0.013)	(0.011)	(0.011)	(0.009)
2012	0.005	0.020	-0.010*	-0.005	-0.003	0.005
	(0.016)	(0.020)	(0.005)	(0.007)	(0.009)	(0.010)
2013	0.018	0.019	0.008	0.006	-0.004	-0.001
	(0.021)	(0.021)	(0.012)	(0.012)	(0.009)	(0.009)
Pre-2007 coefficients = 0 (p)	0.543	0.726	0.393	0.335	0.589	0.858
2007–2008 = 0 (p)	0.044	0.024	0.008	0.039	0.651	0.399
2007–2008 = 2009–2013 (p)	0.033	0.025	0.005	0.024	0.386	0.441
Pre-2007 stop rate: group	0.012	0.009	0.013	0.013	0.016	0.012
Pre-2007 stop rate: control	0.009	0.008	0.009	0.008	0.009	0.008
Coding of non-sharing	spec0	spec1	spec0	spec1	spec0	spec1
Observations	15103	15248	15103	15248	15103	15248

Notes: Linear probability model of the stop-sharing hazard with group-by-year interactions, on stations that shared in the previous survey year, so the market fixed effect absorbs a market-specific linear trend in sharing. Coefficients are relative to the pooled 2003–2006 pre-period. The omitted group is the competitive control, and Sherbrooke and Magog form one market. Columns 1, 3 and 5 use the baseline coding of non-sharing, columns 2, 4 and 6 the strict coding (explicit refusals only). The pre-2007 coefficients come from an unreported specification, and only their joint test is shown. The last two rows give the 2003–2006 stop rate in each group and in the control. Standard errors clustered by station in parentheses. *, ** and *** denote significance at the 10, 5 and 1 percent levels.

Table B.5: Effect of Stop-Sharing Neighbors on Station Volume

	(1) 500 m	(2) 1 km
1(Neighbor stop sharing)	0.013	0.044
	(0.101)	(0.074)
1(Neighbor stop sharing)×**1**(2006Q3 ≤ t ≤ 2008Q4)	0.247***	0.166**
	(0.077)	(0.067)
Observations	2,279	2,279
R^2	0.951	0.953
Station FE	Yes	Yes
Brand FE	Yes	Yes
Market × Quarter FE	Yes	Yes

Notes: Dependent variable is log quarterly volume at always-sharing stations in the cartel cities (excluding Magog), 2001Q1 to 2008Q4, excluding the 2006Q1 round. The neighbor indicator equals one if a station within the radius is a stop-sharing station ($StopShare_j = 1$, Online Appendix H). Both columns also include an indicator for a never-sharing neighbor and its interaction with the 2006Q3 to 2008Q4 indicator (not reported), local competition controls, and station amenities. Robust standard errors clustered by station in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

References

Abreu, D., P. Milgrom, and D. Pearce (1991). Information and timing in repeated partnerships. *Econometrica 59*(6), 1713–1733.

Albæk, S., P. Møllgaard, and P. B. Overgaard (1997). Government-assisted oligopoly coordination? a concrete case. *The Journal of Industrial Economics 45*(4), 429–443.

Aleksenko, S. and J. Miklós-Thal (2026). Third-party pricing algorithms and information sharing. Working paper.

Aryal, G., F. Ciliberto, and B. T. Leyden (2022, November). Coordinated capacity reductions and public communication in the airline industry. *The Review of Economic Studies 89*(6), 3055–3084.

Assad, S., R. Clark, D. Ershov, and L. Xu (2024). Algorithmic pricing and competition: Empirical evidence from the German retail gasoline market. *Journal of Political Economy 132*(3), 723–771.

Ater, I. and O. Rigbi (2023, February). Price transparency, media, and informative advertising. *American Economic Journal: Microeconomics 15*(1), 1–29.

Awaya, Y. and V. Krishna (2016). On communication and collusion. *American Economic Review 106*(2), 285–315.

Awaya, Y. and V. Krishna (2020). Information exchange in cartels. *RAND Journal of Economics 51*(2), 421–446.

Bennett, M. and P. Collins (2010). The law and economics of information sharing: The good, the bad and the ugly. *European Competition Journal 6*(2), 311–337.

Byrne, D. and N. deRoos (2019). Learning to coordinate: A study in retail gasoline. *American Economic Review 109*, 591–619.

Byrne, D. P., N. de Roos, M. S. Lewis, L. M. Marx, and X. Wu (2025a). Asymmetric information sharing in oligopoly: A natural experiment in retail gasoline. *Journal of Political Economy 133*(7), 2031–2088.

Byrne, D. P., N. de Roos, M. S. Lewis, L. M. Marx, and X. Wu (2025b). Price coordination with asymmetric information sharing: Theory and evidence. *Review of Economics and Statistics*.

Byrne, D. P., N. de Roos, M. S. Lewis, L. M. Marx, and X. Wu (2026). Negotiating coordination: A study in retail gasoline. Working paper, University of Melbourne.

Calder-Wang, S. and G. H. Kim (2024). Algorithmic pricing in multifamily rentals: Efficiency gains or price coordination? Working paper.

Calvano, E., G. Calzolari, V. Denicolò, and S. Pastorello (2025, February). Artificial intelligence, algorithmic recommendations and competition. Available at SSRN: https://ssrn.com/abstract=4448010.

Carranza, J. E., R. Clark, and J.-F. Houde (2015). Price controls and market structure: Evidence from gasoline retail markets. *The Journal of Industrial Economics 63*(1), 152–198.

Chaves, D., R. Clark, M. Duarte, and J.-F. Houde (2025). Misallocation in differentiated product markets: Evidence from collusion under uniform pricing. Working paper, December 2025.

Chaves, D. and M. Duarte (2025). The inner workings of a hub-and-spoke cartel in the automotive fuel industry. *American Economic Journal: Microeconomics 17*(1), 41–65.

Clark, R., M. Duarte, and J.-F. Houde (2025). Cartel instability and price wars. In *Cartels Diagnosed: New Insight on Collusion*. Cambridge University Press.

Clark, R., I. Horstmann, and J.-F. Houde (2024). Hub-and-spoke cartels: Theory and evidence from the grocery industry. *American Economic Review 114*(3), 783–814.

Clark, R. and J.-F. Houde (2013). Collusion with asymmetric retailers: Evidence from a gasoline price-fixing case. *American Economic Journal: Microeconomics 5*, 97–123.

Clark, R. and J.-F. Houde (2014). The effect of explicit communication on pricing: Evidence from the collapse of a gasoline cartel. *The Journal of Industrial Economics 62*, 191–228.

Dafny, L. (2009). Estimation and identification of merger effects: An application to hospital mergers. *The Journal of Law and Economics 52*(3), 523–550.

Duso, T., J. Harrington, Joseph E., C. Kreuzberg, and G. Sapi (2026). Public communication and collusion: New screening tools for competition authorities. *Journal of Competition Law & Economics*.

Gal-Or, E. (1985). Information sharing in oligopoly. *Econometrica 53*(2), 329–343.

Garrod, L., J. E. Harrington, and M. Olczak (2021). *Hub-and-Spoke Cartels: Why They Form, How They Operate, and How to Prosecute Them*. MIT Press.

Genesove, D. and W. P. Mullin (1998, Summer). Testing static oligopoly models: Conduct and cost in the sugar industry, 1890-1914. *The Rand Journal of Economics 29*(2), 355–377.

Genesove, D. and W. P. Mullin (2001, Jun.). Rules, communication, and collusion: Narrative evidence from the sugar institute case. *American Economic Review 91*(3), 379–398.

Green, E. J. and R. H. Porter (1984, Jan.). Noncooperative collusion under imperfect price information. *Econometrica: Journal of the Econometric Society 52*(1), 87–100.

Grennan, M. and A. Swanson (2020). Transparency and negotiated prices: The value of information in hospital-supplier bargaining. *Journal of Political Economy 128*(4), 1234–1268.

Harrington, J. (2004). Post-cartel pricing during litigation. *Journal of Industrial Economics 52*(4), 517–533.

Harrington, J. E. (2022). The effect of outsourcing pricing algorithms on market competition. *Management Science 68*(9), 6889–6906.

Harrington, J. E. (2025). Hub-and-spoke collusion with a third-party pricing algorithm. *Journal of Industrial Economics*. Forthcoming.

Harrington, J. E. and A. Skrzypacz (2007). Collusion under monitoring of sales. *RAND Journal of Economics 38*(2), 314–331.

Harrington, J. E. and A. Skrzypacz (2011, October). Private monitoring and communication in cartels: Explaining recent collusive practices. *American Economic Review 101*(6), 2425–2449.

Hastings, J. S. (2004). Vertical relationships and competition in retail gasoline markets: Empirical evidence from contract changes in Southern California. *American Economic Review 94*(1), 317–328.

Houde, J.-F. (2012). Spatial differentiation and vertical mergers in retail markets for gasoline. *American Economic Review 102*, 2147–2182.

Kandori, M. (1992). The use of information in repeated games with imperfect monitoring. *Review of Economic Studies 59*(3), 581–593.

Kawai, K., J. Nakabayashi, and J. Ortner (2026). The value of privacy in cartels: An analysis of the inner workings of a bidding ring. *Review of Economic Studies 93*(1), 600–628.

Kreps, A. A. (2026). Information sharing and collusion: Theory and evidence from poultry processing. Job market paper, Northwestern University, October 2026.

Kühn, K.-U. and X. Vives (1995). *Information Exchanges Among Firms and their Impact on Competition*. Luxembourg: Office for Official Publications of the European Community.

Kühn, K. (2001). Fighting collusion by regulating communication between firms. *Economic Policy 16*, 168–204.

Luco, F. (2019). Who benefits from information disclosure? The case of retail gasoline. *American Economic Journal: Microeconomics 11*, 277–305.

Marshall, R. C. and L. M. Marx (2012). *The Economics of Collusion: Cartels and Bidding Rings*. The MIT Press.

Martin, S. and F. Verboven (2026, February). Price caps and coordinated effects: Evidence from retail gasoline. Working paper.

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.

Raith, M. (1996). A general model of information sharing in oligopoly. *Journal of Economic Theory 71*(1), 260–288.

Rambachan, A. and J. Roth (2023). A more credible approach to parallel trends. *Review of Economic Studies 90*(5), 2555–2591.

Rossi, F. and P. Chintagunta (2016). Price transparency and retail prices: Evidence from fuel price signs in the Italian motorway. *Journal of Marketing Research 53*(3), 407–423.

Sahuguet, N. and A. Walckiers (2017). A theory of hub-and-spoke collusion. *International Journal of Industrial Organization 53*, 353–370.

Sorgard, L. and F. Steen (2024). Coordinated rebate reductions and semi-collusion: Retail gasoline in Sweden. Forthcoming in *Cartels Diagnosed: New Insight on Collusion*.

Stigler, G. J. (1964). A theory of oligopoly. *Journal of Political Economy 72*, 42–61.

Sugaya, T. and A. Wolitzky (2026). Collusion with optimal information disclosure. Forthcoming in *Quarterly Journal of Economics*.

Vives, X. (1984). Duopoly information equilibrium: Cournot and Bertrand. *Journal of Economic Theory 34*(1), 71–94.

Wang, Z. (2008). Collusive communication and pricing coordination in a retail gasoline market. *Review of Industrial Organization 32*(1), 35–52.

Wang, Z. (2009). (mixed) strategy in oligopoly pricing: Evidence from gasoline price cycles before and under a timing regulation. *Journal of Political Economy 117*(6).