

# Do information sharing platforms facilitate coordination?

## Evidence from the collapse of a cartel\*

Robert Clark  
University of Toronto

Jean-François Houde  
UW-Madison & NBER

September 10, 2026

### Abstract

Third-party platforms that aggregate sales data are under growing antitrust scrutiny, yet it remains empirically difficult to determine whether they facilitate collusion or instead generate legitimate efficiencies for firms (e.g., by enabling benchmarking). We study this question in Québec’s retail gasoline industry, where cartels operated in four cities between 2002 and 2006 before collapsing after the Competition Bureau of Canada executed search warrants. Throughout the period a market intelligence provider operated an information-sharing platform in this market, and we observe each station’s decision to supply its volumes to it. We first model a cartel that monitors its members through such a platform. The reports and direct communication are complements, and the arrangement is easier to sustain the more members report. This delivers a revealed-preference test: if the platform were integral to coordination, participation should fall when the cartel collapses. Using an event study design, we find that participation fell sharply in cartel cities relative to competitive markets following the investigation. It did not collapse entirely, which points to two components of the value of the reports: an efficiency value common to all stations, and a collusive value specific to cartel participants. We also find that volume moved toward the neighborhoods in which at least one station stopped sharing, consistent with information withdrawal triggering more aggressive local pricing. Our findings suggest that a passive platform with no coordinating intent can facilitate collusion when the data it shares are sufficiently granular and timely, with direct implications for the design of safe-harbor criteria for information-sharing arrangements.

---

\*We thank Michi Igami for excellent discussion, and participants at the IIOC, Greenline IO Day, MaCCi Summer Institute for Competition Policy, CRESSE Conference, including David Byrne, Daniel Chaves, Marco Duarte, Kei Kawai, Rob Porter, and Otto Toivanen for helpful suggestions.

# 1 Introduction

Third-party platforms that aggregate sales data are under growing antitrust scrutiny, yet whether they actually facilitate collusion or instead generate legitimate efficiencies for firms has never been directly tested. 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; Ellison, 1994). 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).

Separating the two accounts empirically requires direct observation of firms’ decisions to supply information and a source of variation in the value of coordination that does not also move the value of the service the platform sells. Existing empirical work typically has neither. Participation is 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 supplies both by leveraging unique data from a market research firm on stations’ decisions to voluntarily share their information with its platform, along with quasi-natural variation provided by an antitrust authority’s investigation into collusion to evaluate whether participation declined following the collapse of the cartel. We study Québec’s retail gasoline industry, where alleged cartels operated in Magog, Sherbrooke, Thetford Mines, and Victoriaville between 2002 and May 2006, when the Competition Bureau of Canada issued search warrants.<sup>1,2</sup> These cartels engaged in 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 the subsequent convictions caused a rapid collapse of collusive margins in the prosecuted markets. Throughout this period a market intelligence provider operated an information-sharing platform, which we call the *Platform*. The Platform surveyed stations for their sales volumes and sold the compiled reports back to subscribing firms. Participation was voluntary and we observe it directly, by station and by year, in the cartel cities and throughout the province.

The cartel arrangement creates a monitoring problem that prices alone cannot solve, since it fixed not only prices but also the division of volume through the timing of price adjustments. A station could therefore comply with the posted price while deviating by adjusting a few hours late, shifting volume without leaving any trace in posted prices. 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 are conclusive but which stations cannot observe at their rivals. The Bureau’s records show that participants used the Platform to monitor precisely this

---

<sup>1</sup>**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 leveled charges against any third-party information providers in this market.

<sup>2</sup>See Clark and Houde (2013, 2014) for analysis of these alleged cartels.

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 by publicly cutting price or, more subtly, by joining the last mover in the adjustment window and capturing market share without changing their posted price. The latter deviation is detectable only through 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 complementary. Second, the arrangement is easier to sustain when more members participate. Third, the deterrence value of monitoring is greatest for low-cost members, which have the strongest incentive to deviate. These results imply that the platform has a collusive value that varies across firms. This structure delivers our empirical prediction: when the cartel collapses, the value of reporting for collusive monitoring should fall, leading participation to decline, particularly among the firms most strongly constrained by the monitoring arrangement. 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 observed response of participation therefore provides a revealed-preference test that separates these two sources of value.

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 allegation of collusion was 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. It fell toward, but not below, the level observed in the competitive cities, which is the pattern the decomposition predicts and which rules out the possibility that the Platform was purely a collusive device.

The composition of the decline matches the cost ordering the model predicts. We use court records from the prosecution to classify stations in the cartel cities by ownership type. The cartel leaders, who organized the arrangement and used the Platform’s reports to police it, do not leave. Their sharing rate rises slightly and no leader station stops. The low-price guarantee chain that enforced price wars paid the largest fine and continued to share throughout. The stations that withdraw are instead the ones with genuine pricing discretion that were marginal participants in the arrangement: branded lessee dealers decline by 8 percentage points and unbranded lessee dealers by 12. This pattern is also casts doubt on the leading alternative explanation, that the

decline reflects legal exposure rather than the loss of collusive value. Legal exposure predicts that the prosecuted firms leave. We observe the opposite. Among lessee dealers participation falls by the same amount whether or not the station is named in the court documents.

Our second result concerns what withdrawal does to the market. Leaving the Platform is itself a competitive act, since a station that removes its volume from the reports makes its own deviations harder to detect and punish. The prediction is that stations that left following the investigation priced more aggressively afterward. We cannot test this directly because our pricing data are not sufficiently high frequency and because, by construction, we do not observe volumes for a station that stopped reporting. Instead, we test it on the neighbors that kept reporting. Competition in retail gasoline is local, so a station that starts pricing aggressively forces nearby rivals to follow, and the neighborhood as a whole gains volume from the rest of the city.<sup>3</sup> We find that always-sharing stations within 500 metres of a station that stopped sharing gained 0.226 log points of quarterly volume over 2006Q3 to 2008, relative to always-sharing stations elsewhere in the same market and quarter.

Finally, we ask whether the same pattern appears outside the four prosecuted cities. Clark and Houde (2014) document that a number of other Québec cities priced like the cartel cities before 2006, and that margins and price adjustments in those cities changed after the investigation. If those markets were coordinating in the same way, their stations should also have changed their sharing behavior. The evidence is mixed. In Québec City, the largest of these markets, we find no evidence that the investigation caused a decline in sharing. Participation there falls more or less continuously from the early 2000s, the decline is concentrated in two unbranded discount chains, and margins show no permanent break after 2006. In the eight other markets, closer in size to the cartel cities, sharing falls abruptly in 2007 and 2008, reaching  $-0.091$  by 2008, and the hazard of ceasing to share rises. These results are consistent with the investigation having destabilized coordination in the smaller potentially collusive markets, while in a market as large as Québec City the non-collusive value of the reports was enough to keep stations on the Platform.

**Related literature:** A broad literature examines information sharing among competitors. Early theoretical work by Vives (1984), Gal-Or (1985), and Raith (1996) asks whether oligopolists have an incentive to share private cost or demand information before competing, and shows that the answer depends on the mode of competition and the nature of the uncertainty. The literature closest to our model is the one 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 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 and are not evidence of a breakdown. 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

---

<sup>3</sup>Our approach is similar in spirit to Dafny (2009), who studies the impact of mergers by focusing on the response of rivals.

is sustainable. Our monitoring structure is the one in Green and Porter (1984), and our result that collusion is easier to sustain when every member reports (Proposition 2) is the comparative static of Abreu et al. (1991) and Kandori (1992) applied to the number of firms that agree to be measured. The institution imposes two modifications: the statistic is built from reports the members choose whether to file, and the deviation it is designed to catch is invisible in posted prices.

A second body of work studies cartels that exchange sales information directly, which is the closest analogue to the Platform. Harrington and Skrzypacz (2007) study a cartel whose members privately observe their own sales and report them to each other, and show that the arrangement is sustained by a reporting rule together with the threat of a price war, so that the reports are the instrument of enforcement rather than a by-product of it. Harrington and Skrzypacz (2011) extend the argument to arrangements in which firms exceeding their sales quotas compensate firms whose reported sales fall below quota, which is what makes truthful self-reporting incentive compatible. Our reporting decision plays the same role their reporting rule does. Awaya and Krishna (2016) show that cheap talk achieves nothing on its own but raises the collusive payoff when it is combined with noisy private monitoring, and Awaya and Krishna (2020) extend the result to information exchange in cartels. The complementarity between direct communication and the reports in our model (Proposition 1) is closest to their result. Marshall and Marx (2008) study aggregate data in a homogeneous-product market with large demand shocks: a firm with zero sales while aggregate sales are positive knows a rival cut its price, whereas zero aggregate sales is consistent with a negative demand shock and no deviation, so the usefulness of aggregate data depends on the size of demand shocks relative to the gains from deviation. Our setting differs from all of these in one respect that matters for the empirical design. The reports are not exchanged inside the cartel. They are filed with a commercial third party that sells them to colluding and non-colluding firms alike. Two consequences follow. The reports are station-level rather than aggregate, which makes detection more precise than in the models above. And the decision to supply a volume is a choice we observe directly, for colluding and non-colluding firms at the same time, which is what allows us to price the collusive value of the information.

Kühn and Vives (1995) and Kühn (2001) bring these insights to a policy framework, arguing that the key distinction is between sharing price information, which tends to be anticompetitive by facilitating coordination, and sharing quantity or cost information, which may be procompetitive. Bennett and Collins (2010) survey the legal and economic landscape and argue that authorities should focus on the nature and timing of what is shared rather than adopting per se rules.<sup>4</sup>

On the empirical side, the papers most closely related to ours are Byrne et al. (2025a), Byrne et al. (2025b), and Byrne et al. (2024), who study private subscription platforms that collect and redistribute price information among competing firms in retail gasoline markets. Like us, their central focus is on the role of information sharing in facilitating collusion; the key difference is that their platforms share price information, whereas ours shares information about sales volumes

---

<sup>4</sup>A separate line of work shows that some degree of strategic uncertainty can help sustain collusion by limiting the severity of punishment phases and reducing the temptation to deviate (Bernheim and Madsen, 2017; Sugaya and Wolitzky, 2018; Ortner et al., 2024; Kawai et al., 2026).

and transactions. To our knowledge, no existing empirical work examines platforms that provide information about sales volumes or transactions, which is the setting we study. A broader empirical literature examines mandatory public transparency regimes (Albæk et al., 1997; Luco, 2019; Ater and Rigbi, 2023; Rossi and Chintagunta, 2015). A common finding across this literature is that the competitive effects depend critically on who receives the information: sharing among firms tends to soften competition, while dissemination to consumers tends to intensify it. Aryal et al. (2022) study a related but distinct form of quantity communication, showing how public earnings announcements about capacity discipline influence pricing behavior in the airline industry.

There has also been work examining the role of trade associations in providing information that might facilitate collusion. In the Sugar Institute cartel, aggregated self-reported statistics on production, deliveries, and sales were shared with members while maintaining confidentiality, and the association’s Code of Ethics enforced rules that made price cuts more observable and punishable (Genesove and Mullin, 1998, 1999, 2001). Kawai et al. (2025) study a trade association that facilitated bid rigging, and Marshall and Marx (2012) provide other examples. Third-party information sharing platforms share features with the hub-and-spoke framework described in Sahuguet and Walckiers (2017), Garrod et al. (2021), and Clark et al. (2024), in that a central intermediary redistributes competitively sensitive information among competing firms. The key difference is that the platform we study is passive: it aggregates and disseminates sales data as a commercial service, without any deliberate intent to coordinate a collusive arrangement.

The economics literature on algorithmic pricing as a third-party information sharing platform has developed rapidly, with theoretical contributions examining the incentives of algorithm designers to facilitate coordination and the conditions under which efficiency and anticompetitive harm coexist (Harrington, 2022, 2024, 2025b,a; Calvano et al., 2025; Sugaya and Wolitzky, 2026; Aleksenko and Miklós-Thal, 2026; Ichihashi, 2026), and empirical work finding evidence of coordinated pricing in retail gasoline and rental housing markets (Assad et al., 2024; Calder-Wang and Kim, 2024).<sup>5</sup> A key theme is the tension between the legitimate efficiency value of information platforms and their potential to facilitate collusion. Our paper addresses this tension directly, showing that platform participation reflected both motives simultaneously and exploiting an exogenous cartel collapse to separately identify them.

In addition to the papers mentioned above, several others have examined collusion in retail gasoline, including Wang (2008, 2009) and Byrne and deRoos (2019) for Australia, Chaves and Duarte (2023) for Brazil, Sorgard and Steen (2024) for Sweden, and Martin and Verboven (2026) for Belgium. There are also papers on competition in retail gasoline more generally, including some focused on Québec’s market; see for instance Houde (2012) and Carranza et al. (2015).

---

<sup>5</sup>On tacit algorithmic collusion without a common platform, see for instance Kühn and Tadelis (2018), Miklós-Thal and Tucker (2019), Calvano et al. (2020), and Brown and MacKay (2023).

## 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. The section closes with a timeline that shows which data source covers which markets and which years.

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

Stations fall into two categories. Branded stations sell under a refiner’s trademark and operate under one of four contractual arrangements. Company-operated stations are owned and run by the refiner. Commission-operated stations are owned by the refiner, and the dealer earns a per-litre commission. Lessee dealers lease the site from the refiner and buy their supply from it. Independent branded retailers own the site and contract with a refiner to carry its brand. Unbranded stations operate on their own account: the larger ones hold supply contracts with one or more refiners, the smaller ones buy at the posted rack price, and all of them set their own retail prices. The contractual form determines who makes decisions at the station. At a company-operated or commission site, the retail price and the station’s dealings with third parties are set by corporate policy. A lessee dealer and an unbranded operator decide both for themselves. The same distinction separates the stations that could choose to stop supplying their sales volumes from those for which that decision was made at head office. Québec’s market contains all of these types.

The market also features a chain with a low-price guarantee program, which committed it to match the lowest posted price in each local market, and one large convenience-store chain that operated across the province. Big-box retailers whose primary business lies outside gasoline also operated stations, typically selling unbranded fuel at low 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.<sup>6</sup>

Figure 1 plots the timeline of the investigation and its overlap with our main sources of data, which we introduce below. 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

---

<sup>6</sup>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 as including Magog. The volume subsection below states where Magog is included and where it is not.




Figure 1: Timeline of the cartel, the investigation, and the data coverage

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.<sup>7</sup> 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).<sup>8</sup> 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.<sup>9</sup>

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

<sup>7</sup>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>.

<sup>8</sup>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.

<sup>9</sup>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.

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.<sup>10</sup> 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. Appendix Table 11 maps the classified cities into the Platform markets used in the information-sharing analysis.

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

|                                              | <i>Collusive</i> (2002–2006) |         |       | <i>Collapse</i> (2006–2007) |         |       | <i>Ex-post</i> (2007–2012) |         |       |
|----------------------------------------------|------------------------------|---------|-------|-----------------------------|---------|-------|----------------------------|---------|-------|
|                                              | Margin                       | P90-P10 | Ratio | Margin                      | P90-P10 | Ratio | Margin                     | P90-P10 | Ratio |
| <i>Panel A: Cartel cities</i>                |                              |         |       |                             |         |       |                            |         |       |
| 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  |
| <i>Panel B: Potentially collusive cities</i> |                              |         |       |                             |         |       |                            |         |       |
| Other coll. cities (n=18)                    | 7.44                         | 5.75    | 3.35  | 6.23                        | 7.99    | 2.35  | 6.98                       | 7.91    | 2.65  |
| <i>Panel C: Competitive control</i>          |                              |         |       |                             |         |       |                            |         |       |
| 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  $|\text{margin}| > 20$  are excluded as outliers. P90-P10 is the mean of city-level inter-decile ranges  $p_{90} - p_{10}$  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.<sup>11</sup> Panel A reports on the cartel cities, Panel B on the other collusive-cluster cities (18), and Panel C on the competitive controls (31). Appendix Figure 6 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

<sup>10</sup>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.

<sup>11</sup>This breakdown follows Clark and Houde (2014), which focused on the collusive and collapse periods.

change is in stability. The interdecile range in the cartel cities rises to between 8.60 and 9.20 cents per litre, roughly double its pre-investigation level, as margins begin to swing between near-zero during price wars and temporary spikes during failed attempts to restore coordination.<sup>12</sup> This instability persists beyond the investigation, but margins recover in two of the three cartel markets (Sherbrooke–Magog and Thetford Mines). 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.

## 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 do not disclose the name of the company and 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, spanning the cartel period, the investigation, and several years after it.<sup>13</sup>

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 19.5 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,

---

<sup>12</sup>Appendix C reports a city-level difference-in-differences version for completeness.

<sup>13</sup>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.

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

Table 12 reports 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 6,816 in Sherbrooke, and Sherbrooke has 61 active stations against 23 in Victoriaville and 12 in Thetford Mines. Stations are also larger in Sherbrooke, with 11.9 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 is an annual province-wide panel of stations, covering 2002 to 2013.<sup>14</sup> 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.<sup>15</sup> 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. Subscribers paid to participate and were contractually required 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 cumulative meter on each pump or on the storage tanks, and many did. During our sample a technological transition was under way. Stations were replacing analog pumps with digital systems that store volume data electronically. With the older equipment a surveyor could take the reading without assistance, and the default was to share. Digital systems require the

---

<sup>14</sup>The 2001 file duplicates the 2002 file, so we drop 2001 and the panel begins in 2002.

<sup>15</sup>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.

operator to retrieve the figure and hand it over, which makes sharing a deliberate act and moves the default toward non-participation. An operator could therefore withhold its volumes in two ways: by not subscribing, and by denying the surveyor access to the meters. Prices and station characteristics were collected either way, without the operator’s permission.

The result is that participation is voluntary for some stations and not for others. At vertically integrated stations and at the large retail networks that hold corporate subscriptions, participation is governed by head-office policy or by a contract with the Platform. For independent operators, branded lessee dealers, and convenience-store chains, the decision is made at the station by an operator who also sets the price. Our data do not record which stations subscribed to the reports. We observe whether a station supplied its volumes, not whether it purchased access to the compiled results, so we cannot measure how many stations supplied data without subscribing.

A volume can be missing for two kinds of reasons. Some are unrelated to the operator’s decision: the surveyor’s first visit to a new site, a power outage, a temporary closure for renovation, or an equipment failure such as a meter reset. Others are a decision by the operator: 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. Appendix E 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.

For the analysis of 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 is best understood as a marker of station type, a station that reveals itself to have been a marginal participant in the sharing arrangement, rather than as a change at a date.

Table 2 describes participation in the year before the investigation. The estimation sample is the 1,653 stations observed in 2006 across 29 markets. Column (1) reports the composition of the sample. Branded stations affiliated with the major refiners are the largest group at 44 percent, the low-price guarantee chain accounts for 24 percent, unbranded distributors for 16 percent, and retail chains for 7 percent. The remaining 9 percent are independent and unclassified stations, which form the omitted category in the regression. Only 5 percent of the 2006 cross-section entered after 2002, and 23 percent exited before 2013.

Participation is high. Ninety-one percent of stations supplied their volume in 2006. Column (2)

Table 2: Prevalence of Information Sharing by Station Type (2006)

|                              | (1)<br>Mean    | (2)<br>$Share_{it}$  |
|------------------------------|----------------|----------------------|
| Low-price guarantee chain    | 0.24<br>(0.43) | 0.342***<br>(0.026)  |
| Branded stations (excl. LPG) | 0.44<br>(0.50) | 0.337***<br>(0.024)  |
| Retail chain stations        | 0.07<br>(0.25) | 0.289***<br>(0.034)  |
| Unbranded distributor        | 0.16<br>(0.37) | 0.178***<br>(0.028)  |
| Entrant post 2002            | 0.05<br>(0.21) | -0.112***<br>(0.032) |
| Exit pre 2013                | 0.23<br>(0.42) | 0.076***<br>(0.016)  |
| Constant                     |                | 0.618***<br>(0.023)  |
| Observations                 | 1653           | 1653                 |
| Markets                      |                | 29                   |
| Mean dep. variable           |                | 0.908                |
| Std-deviation of market FE   |                | 0.098                |
| Fixed effects                |                | Market               |

*Notes:* Column (1) reports sample means and standard deviations of the covariates. Column (2) estimates a linear probability model of  $Share_{it}$  on station-type indicators with market fixed effects. The omitted station-type category is independent and other unclassified stations. Estimation sample restricted to 2006. Robust standard errors in parentheses. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1.

decomposes this with a linear probability model of  $Share_{it}$  on station-type indicators and market fixed effects. The baseline rate for independent and unclassified stations is 62 percent. Every organizational type shares at a higher rate. The two highest are the low-price guarantee chain, at 96 percent, and similarly for the other branded stations. Retail chains are at roughly 91 percent, and unbranded distributors at about 80 percent. This ordering is implied by the institutional description. 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.

Stations that entered after 2002 are 11.2 percentage points less likely to share, conditional on type and market. This is consistent with the technological transition described above. Newer stations have electronic pumps that require the operator to retrieve the figure, so the default moves toward non-participation. Stations that exit before 2013 share at a 7.6 percentage point higher rate, which we take to represent a survivorship pattern. Participation also varies across markets, though less than it varies within them. The standard deviation of the market fixed effect is 9.8 percentage points, about a third of the gap between branded stations 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. The documents record the phone calls made around each price adjustment: who called whom, at what time, and a summary or extract of what was said. They are close to a complete record of price changes in Victoriaville and Thetford Mines, and less complete for Sherbrooke and Magog, particularly for price decreases.<sup>16</sup> They describe a cartel that solved the coordination and enforcement problems of a price-posting market through two instruments: telephone communication before each adjustment, and a fixed order in which the adjustment was carried out. The order of moves is not specific to this cartel. Byrne and deRoos (2019) document price leadership with the same structure in Perth, where a dominant retailer initiated each increase and the other retailers followed in a stable sequence.<sup>17</sup>

The order of play assigned a distinct function to each type of participant. The leaders initiated the increases. They placed the calls, told the others what price to post and when, and served as the communication nexus through which the rest of the market learned the plan. A second group moved last on increases and first on decreases. Moving last means selling at the old, lower price after every other station has raised its price, which transfers volume to the last mover. The firms in this position were low-cost, including a chain operating a low-price guarantee that committed it to matching the lowest posted price in the market, and big-box retailers. The transfer is what bought their participation. A third group, the followers, adjusted in the prescribed order. Many were lessee dealers who set their own prices, so their compliance had to be obtained. One independent operator in Victoriaville, the dissident in the terminology of Clark and Houde (2013), was required to post the same price as the followers but often refused the leader’s instructions to raise. Because of its location, an increase could not hold without it. Two stations on the periphery were tolerated rather than included, and were allowed to post up to two cents per litre below the suggested price, with violations of that limit subject to retaliation (Clark and Houde, 2013).

---

<sup>16</sup>They contain the Bureau’s summaries of the recorded conversations and extracts from them, not the conversations themselves. The online appendix of Clark and Houde (2013) describes the coding of the calls, which measures the number and direction of calls, the duration of each price change, and the timing of each participant’s move.

<sup>17</sup>Byrne et al. (2026) formalize this arrangement as coordination through price signals. The difference with our setting is the evidence. They do not observe communication between the retailers and argue that the order of moves was established through posted prices alone. In Québec the court record shows the participants discussing the timing of each move by telephone.

As a result of the timing of adjustments, the agreement fixed not only a price but also a division of volume. Compliance had to be verified in quantities. A wrong posted price is seen by anyone and needs no data to be caught. A station that sets the right price but adjusts a few hours late moves volume without posting a price that differs from one already on a sign. The price record therefore cannot identify the deviation. Rivals may notice it at the pump and complain, but a complaint is not proof. Proof is in volumes, and rival volumes are the one competitively relevant variable a station cannot observe for itself. The Platform supplied them.

The Bureau’s records show that the participants used the Platform for exactly this purpose. The wiretaps ran in Victoriaville and Thetford Mines from March to June 2005, a period of open instability in these markets. Two extracts illustrate the practice.<sup>18</sup>

*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 the operator whose refusals to raise produced the two Victoriaville price wars before the investigation, in October 2004 and in December 2005 (Appendix Figure 7). The second episode kept margins near 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 could be seen from his sign. 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 supplied his volumes, so the reports showed his gain through the shares the members were losing. The case against the participants rests on those market shares, supported by complaints about delays. This is the sequence the model in Section 3.2 formalizes:

---

<sup>18</sup>Our translations of the Bureau’s summaries of recorded conversations submitted to the Superior Court of Québec. Firms and agents are identified by their position in the market rather than by name.




Figure 2: Market shares of non-cooperating stations and median margins in Victoriaville

*Notes:* Quarterly market shares (left axis) of the tolerated low-price station that reported its volumes, of the dissident, and of the low-cost (big-box) retailer, and median margins in cents per litre (right axis), 2001 to 2009. The tolerated station's share is measured from the survey. The dissident did not report, so his share is constructed from the surveyor's estimated volumes and is an approximation. The first vertical line marks the recorded conversations discussing the Platform's sales reports (March and April 2005). The second marks the investigation announcement. In the weekly Régie series, margins in Victoriaville fell from 8 to 9 cents per litre in September 2004 to between 1.0 and 1.9 cents over the three weeks from October 9, 2004, and from 6.8 cents in the week of December 3, 2005 to -0.1, -0.3, and -1.7 cents over the three weeks beginning December 17, below the regulated floor, recovering in mid-January 2006. The court record attributes both episodes to the dissident. The wiretap coverage cited for the second is the March to June 2005 campaign, and there is no Victoriaville recording from December 2005.

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, and each was allowed to post two cents below the collusive price. One reported its volumes and the other never did. 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 2 shows the pattern in the survey data. The tolerated station that reported its volumes and the dissident both gain share in the quarters before the recorded conversations, and both fall back after the investigation, when margins collapse.

The arrangement was already under strain before the warrants. Clark and Houde (2014) present structural break tests confirming that the May 2006 searches caused the collapse of collusive margins in the four markets, but the Bureau's records show unease among the participants in the preceding months. Some station owners and network managers were expressing concern about the legal implications of agreeing on prices by telephone, and at least two integrated firms sent letters to

their lessee operators instructing them not to discuss prices with competitors. The most likely trigger is the 2004 newspaper article that started the investigation. What matters for what follows is that the instrument under strain was communication. The reports themselves were a commercial service that no one had suggested was unlawful, and to which many stations continued to subscribe.

### 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 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  $\delta$ . Firm  $j$ 's volume in period  $t$  is

$$q_{jt} = s_{jt} m_t u_{jt}, \quad (1)$$

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 is independent across firms and periods. A firm observes its own volume and nothing else. The market-wide component is why it cannot determine a rival's conduct from its own sales, since a weak market and a rival stealing share look the same.

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  $\Delta_p$  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).<sup>19</sup> 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 and transfer less shares (Clark and Houde, 2013), so we assume that only increases matter. We take  $\tau$  and  $\Delta_p$  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 mechanically, and does not deviate.<sup>20,21</sup> If the arrangement breaks

<sup>19</sup>It could also be because  $L$  advertises a low-price guarantee and commits to move last.

<sup>20</sup>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  $\tau$ . This is not central to the paper and we set it aside. The questions we ask are about the followers.

<sup>21</sup>As mentioned above, the cartel leaders actually set the collusive price. For simplicity and to focus on the key mechanism, here we assign this role to  $L$ . The assumption being that  $L$  constrains the leaders when they set this

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.

**The two 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  $d \in (0, \bar{d}]$  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  $\Delta_p$  below the collusive one and gains in expectation

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

which we assume is positive. Appendix A.6 derives (2) and shows what it implies when followers have asymmetric costs. Posted prices pass through the same values under the delay, and legitimate dispersion in the timing of price changes makes a single late change unattributable, so the price record contains no admissible evidence. The deviation leaves traces only in volumes, and in what the members observe and communicate to each other.<sup>22</sup>

**The two signals.** The cartel observes two signals at the end of a period. The first is the Platform's report, which publishes the volume of every member that supplied one and prints the identity of the members whose volume is missing. It does not publish total market volume or say why a volume is missing. Let  $r$  be the number of reporting followers. The cartel compares  $L$ 's volume with the average of theirs,

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

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.

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

---

price.

<sup>22</sup>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.

the value “complaint” with probability  $\phi(d)$  after a delay of size  $d$ . We assume  $\phi$  is continuously differentiable and increasing on  $[0, \bar{d}]$ , with  $\phi(0) = \phi_0 \in (0, 1)$  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.

**The trigger, and why a deviator withholds.** The cartel’s monitoring problem is a one-sided test. Under the arrangement  $Y$  has mean  $\mu_0$ . 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  $\mu_0$ , with a critical value that may depend on the message:

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

Here  $\text{sd}(Y) = \sigma_u \sqrt{1 + 1/r}$  is the standard deviation of the statistic under compliance, and  $\kappa_C$  and  $\kappa_N$  are the critical values in standard-deviation units. The cartel chooses them so that no follower gains from a delay, at the smallest cost in wars that a compliant cartel sets off by accident.

Writing the critical value in standard deviations of the statistic the cartel actually computed captures an institutional feature: a single missing report is not itself an offense. The Platform indicates that a member’s volume is missing but not why. A volume can be missing because the member refused or because the surveyor could not take a reading. Punishing every missing volume would mean punishing accidents.<sup>23</sup> To simplify the model exposition we abstract away from the possibility that the number of reports is a random variable. This 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, and the statistic is computed on one fewer report. It is therefore noisier, and because the cutoff is written in standard deviations of the statistic actually computed, the cutoff in levels loosens with the noise. 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, but is part of the deviation.<sup>24</sup>

<sup>23</sup>We purchased from the survey firm a separate file giving the reason for each missing volume. The members did not have it.

<sup>24</sup>We assume that the cartel does not react to a missing report. No report is missing on the equilibrium path, so in the model as written a rule that also initiated a war on a missing report would cost nothing on path and would

**The two technologies are complements.** The cartel chooses the trigger to make price wars as rare as possible while still deterring both deviations. Write  $\alpha_0$  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. The Platform moves the argument of the normal distribution function inside each branch of (4), because a delay lowers the mean of  $Y$ . The message moves the weight between the two branches. Only a rule whose two cutoffs differ can use the second movement. 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. Both (B) and (C) are available to the cartel of regime (A).

**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  $\delta_{\text{crit}}(A) < \min\{\delta_{\text{crit}}(B), \delta_{\text{crit}}(C)\}$ .*

The optimal rule in regime (A) sets  $\kappa_C < \kappa_N$ . The intuition is that the cartel pays for deterrence with wars started by mistake, and a complaint tells it when a war is least likely to be a mistake. In the other two regimes the cartel conditions on a single signal, so it deters with more frequent price wars, and a smaller range of discount factors sustains collusion. The reason the cartel never ignores the reports is that a complaint is not proof. Because  $\phi_0 > 0$ , complaints also arrive when nobody has cheated, and a cartel that acts on complaints alone punishes the innocent at that rate. This is how we represent the calls the leader received: a member's complaint about a rival is a suspicion, not evidence. The report lets the cartel dismiss a complaint that the volumes do not support. Since a real delay lowers the statistic, dismissal rarely spares a deviator, so it removes accidental wars faster than it removes deterrence.

**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  $\kappa_C$  finite. Then  $\alpha_0$  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  $\delta_{\text{crit}}$  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.*

The statistic is the log of  $L$ 's share relative to the average reporting follower's share, which under the arrangement should equal  $\mu_0$ . The denominator is estimated from the reports, and each additional report averages out more of the demand noise, so the estimate becomes less precise as

---

make the deviator report. The cartel cannot write that rule for the reason just given:  $L$  does not observe why a volume is missing.

the number of reports falls.<sup>25</sup> A given theft of volume is then less likely to cross the standardized cutoff, and every remaining follower’s delay is harder to deter.

### 3.3 Predicted effects of a cartel investigation on information sharing

The test we run 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 pay for the reports but not the non-collusive component. This subsection makes that argument in two steps. The first says what the investigation does to the collusive value, and which stations it should move first. The second says why the other non-collusive value did not fall at the same time.

**The investigation is a shock, not a comparative static.** An investigation is unanticipated. The trigger, the length of a price war, and the collusive markup were all calibrated for the environment that preceded it, and the members had no occasion to re-solve for them. Retuning the arrangement would require exactly the kind of meeting the investigation made prosecutable. We therefore interpret the “collapse period” not as a new steady state of the repeated game but play against a rule that the cartel would no longer choose. The question is whether the constraint against the secret delay still holds at the old cutoffs and the old war length.

The investigation limits the members’ ability to communicate about cheating, which flattens  $\phi$ . It also raises the hazard that the arrangement ends altogether, which lowers  $\delta$ . Both increase the incentive to deviate while the rule stays where it was. The optimal gap between the two cutoffs is the log odds ratio of the message, so a flatter  $\phi$  calls for a smaller gap. A cartel holding the old gap makes both errors at once. It initiates wars on weak volume evidence after messages that no longer carry information, and it demands too much evidence when no message arrives. Deterrence fails, and the follower’s best response is the deviation the model already contains, which by Lemma 1 is delay and withhold together.

This gives two readings of the announcement. Under the first, the collusive value of the Platform was destroyed and the members reverted to the static Nash equilibrium. Subscriptions fall, and they fall unevenly, because the reports are worth something to a retailer for reasons unrelated to collusion. Under the second, the arrangement survives in degraded form and the withdrawals are deviations. Note that the second reading does not imply a higher on-path frequency of price wars. The trigger has not moved, so by construction neither has  $\alpha_0$ . Any increase in the frequency of wars, and any increase in the volatility of margins, comes from the deviations themselves.

The two readings therefore differ in what they predict for margins. Under static Nash, margins fall to the competitive level and stay there. Under the second reading, they move every time a deviation starts a war and every time an attempt to restore the collusive price breaks down. Margin volatility in the cartel cities rose after 2006, both in absolute terms (Table 4) and relative to the

---

<sup>25</sup>On the equilibrium path every follower reports, and under (4) a complying follower is indifferent about doing so, since  $\alpha_0$  does not depend on the number of reports. We do not model the decision to join because the question we ask is about participation after the arrangement is disrupted.

competitive cities (Appendix C). This is consistent with the second reading: at least initially the arrangement was not replaced by a settled competitive equilibrium but by repeated attempts to revive coordinated increases that failed.

**Which stations stop sharing.** With symmetric followers the incentive to deviate is uniform and the model delivers no cross-sectional prediction. We introduce asymmetry through the followers' costs, and replace the common cost  $c$  with costs  $c_i$  and assume  $\Delta_p < 1 - c_i$  for every  $i$ .

**Proposition 3** (The secret delay is hardest to deter for the efficient follower). *The constraint against the public price deviation is the same for every follower. The constraint against the secret delay is ordered by cost: if it holds for follower  $i$ , it holds for every follower with a higher cost. The critical discount factor is non-increasing in  $c_i$ , and strictly decreasing wherever the secret delay is the binding constraint, so the arrangement is hardest to sustain for the most efficient follower.*

Proposition 3 rests on a simple asymmetry. A secret delay sells the stolen volume at  $\Delta_p$  below the collusive price. That discount is the same number of cents for every follower, so it consumes a larger share of a high-cost follower's margin than of a low-cost follower's. A low-cost follower keeps most of its margin on the volume it takes; a high-cost follower keeps little. What each loses in a war is its collusive margin, so the gain from the delay relative to the profit lost in a war falls with cost.<sup>26</sup> Because the secret delay is the deviation that Lemma 1 bundles with withdrawal from the Platform, the ordering is a prediction about who stops sharing.<sup>27</sup>

We expect partial rather than full unravelling. Intuitively, deviations feed back on themselves, since a price war destroys the value of the arrangement and so weakens the threat that sustains it. The prediction is a chain of deviations rather than a single break, with the same stations deviating repeatedly. That the pattern is firm-specific is what distinguishes a chain of deviations from a common shock.

**The non-collusive value of the reports.** The revealed-preference argument needs one further element. We read the withdrawals as evidence that the collusive value of the reports fell, because monitoring a rival's volume is worth less once the arrangement it supported is no longer sustained. That is valid only if the non-collusive value did not fall at the same time. If both fell, the withdrawal

---

<sup>26</sup>That the public price deviation does not order the followers is specific to the model. With homogeneous demand a public undercut takes the whole market at the collusive margin, so the gain and the profit lost in a war are the same margin for every follower. With differentiated demand a low-cost follower chooses a deeper undercut and gains more from it. And with a regulated minimum price as the war price, which is the case in Québec, a low-cost follower earns a margin during a war while a high-cost follower earns almost none. Under either extension both deviations are more attractive to the low-cost follower, and the ordering by cost extends to both. We keep the simpler model because the secret delay alone delivers the ordering, and it is the deviation that matters for withdrawal.

<sup>27</sup>The other asymmetries in our application are not cost differences. A multi-site operator spreads a war across outlets that are not all in the same local market. A retailer that prices fuel against high-margin products in an adjoining store earns during a war on the complements rather than on the fuel. A lessee dealer near the end of its contract with a brand discounts the continuation value of the arrangement itself. Each has the same reduced form: it lowers what the firm stands to lose when the arrangement ends relative to what it gains by taking volume now, which in this model is a lower  $\delta_i$ , and the constraint against the secret delay fails first for the follower with the lowest  $\delta_i$ . The prediction is the same either way. The stations with the least to gain from the arrangement withdraw first.

is consistent with any division between them and says nothing about the cartel. We argue that the non-collusive value weakly rose. The assumption concerns the common change over time, not the cross-sectional dispersion in that value, which is what the test of the previous paragraph exploits.

The non-collusive value of the reports comes from benchmarking. A retailer observes its own sales but not its demand curve, and its sales move both with demand in its neighborhood and with its position relative to nearby stations. The Platform reports the volumes of the stations around it, which is the information that separates the two. That information is worth more when the environment is volatile, since own history is then a poor forecast of own sales. Appendix B 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.

By disrupting the arrangement, the investigation made the environment more volatile on both margins and shares, suggesting that the non-collusive value of subscribing to the Platform was weakly larger after 2006. Dispersion of station volumes within a market and quarter rises in all three prosecuted markets we observe over the full sample period, and the within-quarter standard deviation of the weekly retail margin rises by 1.1 to 1.8 cents per litre. Table 4 reports the comparison.

**Model summary.** The model explains why the arrangement is more stable the more of its members report, why the investigation lowered participation and lowered it first among the low-cost stations, for which the gain from a secret delay is largest relative to the collusive profit a war destroys, and why price wars followed. It also gives the reason the non-collusive value of the reports did not fall when the investigation began. The next section tests these predictions.

## 4 The effect of the cartel collapse on information sharing

The previous section formalizes the idea that a station’s willingness to pay for the Platform has both collusive and non-collusive components. The investigation reduced the former and weakly raised the latter, so the change in participation after May 2006 is a lower bound on the collusive value. We test this hypothesis by measuring the effect of the investigation on information sharing in the four prosecuted markets. We then discuss alternative interpretations and quantify the potential 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:

$$Share_{it} = \alpha_i + \delta_t + \sum_{\tau \neq 2002} \beta_{\tau} (Cartel_i \times \mathbf{1}\{t = \tau\}) + \varepsilon_{it}, \quad (5)$$

where  $Share_{it}$  is an indicator equal to one if station  $i$  shares volume data with the Platform in year  $t$ ,  $\alpha_i$  is a station (or market) fixed effect (depending on the specification), and  $\delta_t$  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 3 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.




Figure 3: Effect of the cartel collapse on information sharing: cartel cities

*Notes:* Each series plots the estimated coefficients  $\hat{\beta}_\tau$  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.

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. 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 gradual onset of the decline between 2006 and 2008 is also consistent with our understanding of the Platform’s contractual arrangements. Subscribing stations appear to have been bound by annual contracts that typically auto-renewed, so even stations with a clear desire to withdraw would have faced a cost of exiting immediately and may have continued sharing until their next renewal decision.

The magnitude of this effect is substantial given that the majority of stations in the sample are branded franchisees or chain affiliates whose participation is contractually required (Table 2). A 13 percentage point decline in the overall sharing rate implies a much 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 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 margins in Victoriaville remained at competitive levels through the end of the sample (Table 5).

**Who stops sharing?** The event-study estimates measure the average decline in participation. The composition carries 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-price guarantee chain consumed the Platform’s reports to police the arrangement (Section 3.1), while the followers’ were the ones being policed. If the loss of collusive value was the motive, withdrawal should be concentrated among policed members with pricing discretion, while the monitors, whose participation also carries intrinsic and contractual value, remain. Table 3 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).<sup>28</sup> 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 subscriber, increased its sharing rate from 0.75 to 1.00 with a zero stop rate. Both facts point the same way: where efficiency or contractual obligation dominates the participation decision, sharing is unchanged.

The convenience-store chain is the exception among the organizers. It repudiated the arrangement, fought the charges in court, and avoided conviction. Its sharing rate falls by 18 percentage points, and 25% of its stations stop sharing for at least one quarter after the investigation. Its incentives after the collapse are those of a firm that exited the agreement rather than those of a monitor, and its withdrawal is consistent with a deliberate change in competitive conduct.

The stations that do withdraw systematically are instead those with genuine pricing discretion that were marginal participants in the collusive arrangement. Branded lessee dealers, who retain pricing autonomy under their franchise agreements, show an 8 percentage point decline with a 12%

---

<sup>28</sup>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.

Table 3: Information Sharing by Station Ownership in Cartel Cities

|                                    | Num.<br>stations<br>(1) | Pre<br>(2) | Post<br>(3) | Fraction volume sharing<br>$\Delta$<br>(4) | Stop<br>sharing<br>(5) |
|------------------------------------|-------------------------|------------|-------------|--------------------------------------------|------------------------|
| <b>Panel A: Station ownership</b>  |                         |            |             |                                            |                        |
| 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                   |
| <b>Panel B: Post-2006 survival</b> |                         |            |             |                                            |                        |
| 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: Stations active in the cartel cities before and after the May 2006 announcement of the investigation. “Volume sharing” is an indicator variable equal to one for stations who participate in the market intelligence provider’s sales survey. The “post” period is defined after 2006Q1. “Fraction volume sharing” is the average of “Volume sharing” before and after the May 2006 announcement of the investigation. “Stop-sharing” is an indicator variable equal to one for stations who refused to share sales volume information for at least one quarter in the post period, conditional on sharing during all quarters in the pre-period.

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, we observe the same decline in participation for stations named in the court documents and for stations that are not.

**Robustness.** Three concerns bear on the estimates above, and we address each in turn. The first is the counterfactual. Under a difference-in-differences reading, the post-2006 coefficients must be read against what would have happened to sharing absent the investigation, and the standard counterfactual extrapolates the pre-period 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 7 in the Appendix summarizes the results. The pooled post-investigation effect is  $-0.079$  with a market-clustered standard error of 0.016. The estimated pre-period trend is  $-0.0042$  and is statistically indistinguishable from zero ( $p = 0.93$ ). 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. What binds these values is not a documented pre-trend violation but the precision of the post-period estimates. With four treated markets the confidence intervals are wide, so even a small hypothesized post-treatment violation of parallel trends moves the identified set across zero. These findings confirm that the post-2006 estimates are not driven by a pre-existing trend, although we acknowledge that the bands at the longer horizons are wider than they would be with more treated units.

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 8 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 hazard is 0.028 in 2007 and 0.067 in 2008, against a pre-period base rate of 0.012, 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 rather than 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 about 0.87 before the investigation to 0.72 by 2011, while the competitive control holds at about 0.90, consistent with stations that stop sharing in the cartel cities not resuming.

The size of the decline is also large relative to the normal fluctuations in the competitive markets. Appendix Table 6 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 Appendix E. 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 (Appendix Figure 10): 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 Appendix D, and the codings themselves are described in Appendix E.

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 it does 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 consequential because it targets the sign condition directly, 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 3 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 did withdraw, but it financed a full legal defence, so the subscription fee was not the margin on which it economized.

**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. A station would then withdraw to avoid legal discovery or the public appearance of cartel membership, whether or not the Platform retained any collusive value. Two institutional 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. The evidence assembled in the investigation was of three kinds: wiretap recordings in which station operators discussed price adjustments with one another, testimony of the independent operator whose complaint opened the file, price data gathered by the Régie de l’énergie, and volume data collected by the Platform. No charge was brought against the Platform or its operator, and sharing volumes through the survey was never alleged to be a facilitating practice. The Bureau used the survey’s data as evidence about the market while treating participation in it as unremarkable conduct. Second, participation is near universal in 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. Under this reading, the decline still reveals that participants regarded the Platform as part of the collusive apparatus. What changes is whether the decline 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 appearance story is also inconsistent with the fact that lessee dealers not named in the court documents withdrew at the same rate as those that were.<sup>29</sup>

**Regime-dependent intrinsic value.** The third explanation targets the sign condition directly. 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 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 way.

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 fell in 2003, three years before the warrants, when the entry of three stations destabilized the local arrangement and triggered a price war (Appendix Figure 9). 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 those markets also changed after the investigation (albeit by a smaller magnitude). If they 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 here 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.

---

<sup>29</sup>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.

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 4 reports Québec City in panel (a) and the eight other markets pooled in panel (b).




Figure 4: 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; panel (b) reports the station fixed-effects specification, which is our baseline. The reference year is 2002. The control group consists of the competitive-cluster cities. Panel (b) pools the eight potentially collusive markets outside Québec City. Grand-Mère is excluded throughout: it contributes a single station that never shares. Standard errors clustered by station.

**Québec City.** We find no evidence that the investigation caused a decline in sharing in Québec City. Panel (a) instead shows a gradual fall of about 2 percentage points between 2003 and 2005, stability through 2006, and a further decline afterwards reaching  $-0.10$  to  $-0.11$  by 2012–2013, so the series declines more or less continuously across the whole period rather than breaking at the investigation.

A further analysis of the identity of stations who stopped sharing reveal that the decline is not spread across the market but concentrated in two unbranded discount chains, while branded stations participate almost universally throughout (Appendix Table 10).

Except for a price-war episode following the investigation in 2006, the pricing data shows no decline in margins: annual mean survey margins are 6.7 cents per litre in 2006, 7.4 in 2007 and 7.6 in 2008, soften to between 6.6 and 6.8 over 2009 to 2011, and reach 7.8 in 2012 and 8.1 in 2013, remaining above competitive-market levels throughout.<sup>30</sup>

We therefore conclude that the investigation did not significantly destabilized coordination in

<sup>30</sup>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.

Québec, and branded stations continued to share volume information with the platform. The fact that Québec City is a large/urban market might have contributed to the stable participation rate.

**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, which tracks the cartel-city estimates in magnitude with a one-year delay. The stop hazards move at the same time:  $0.031$  in 2007 and  $0.048$  in 2008 (Appendix Table 8).

Panel (b) also shows a small drift in the pre-period, reaching  $-0.032$  in 2006, which is not statistically different from zero. The concern this raises is not really about these markets. If participation across the whole collusive cluster was already falling before the investigation, the same force could contaminate the cartel-city estimates. Appendix Table 9 presents the average participation rate for all potentially collusive cities other than the cartel markets and Québec City. The table shows that the pre-trend is concentrated in two of the eight markets, Drummondville and Rivière-du-Loup, which hold 33 and 14 stations in 2006, or 20 percent of the group. Station fixed effects mean that only stations which change status contribute to the identification of the drift term, and twenty-one of the 262 stations that appear in the group before 2006 change status in that window. The pre-period estimate is therefore driven by a small number of station-level events primarily in two markets.<sup>31</sup> In the sensitivity analysis of Appendix Table 14, dropping these two markets 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 volume should have moved toward the neighborhoods in which information sharing broke down during the price-war episodes. Evidence of this reallocation is consistent with stations that stopped sharing having had the least to gain from collusion, and therefore pricing most aggressively once the agreement collapsed.

---

<sup>31</sup> Although we do not have detailed data on prices, volumes and station characteristics for Rivière-du-Loup, our characteristics and volume data cover Drummondville. There 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 sit at or below zero for nine weeks between November 2003 and late January 2004, having averaged 6.6 cents in 2002. Participation falls from 0.94 in 2002 to 0.79 by 2006 and does not return to its 2002 level. We interpret this as a change in conduct triggered by a change in market structure.

## 5.1 Hypothesis and identification

Two results connect withdrawal to pricing and market shares. By Lemma 1, a station that delays its price adjustments and keeps reporting supplies the evidence that convicts it. By Proposition 3, the stations for which that act pays are the ones with the least to lose from a price war, which are the stations closest to indifferent between colluding and deviating before the collapse. We therefore treat sustained withdrawal after the investigation as a marker for those stations, and we measure the reallocation of market shares towards neighborhoods in which stations stopped sharing.

Two features of the setting dictate the form of the test. First, we do not observe volume for stations that stopped sharing, since reporting volume is the behavior they ceased. Second, neighbors of an aggressive station match its price rather than concede their traffic, so the price differential opens across neighborhoods rather than within them. The prediction is that the whole neighborhood gains, not that the aggressor’s closest rivals lose. We therefore 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.

Note that we cannot test the hypothesis directly with the cross-sectional price data available to us. Since price differences between stations are transitory, high-frequency prices and sales data would be required. The 2000 Québec City price war illustrates the point. Clark et al. (2024) show that maps of volume-share changes identify the origin of that war and the neighborhoods that gained and lost, while the same maps built from posted prices identify none of them. They also show why price differences cannot last. Adjustments occur within days, and a station that holds its price above the market during a war loses volume quickly. Consistent with all of this, Appendix Figure 13 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.

## 5.2 Estimation

The regression sample includes only stations that always shared with the survey throughout the sample period. Stations classified as  $StopShare_i = 1$  are excluded from the estimation sample and enter the regression only indirectly via the neighborhood indicator variables. The dependent variable,  $\log V_{it}$ , is the log of average daily volume at always-sharing station  $i$  in quarter  $t$ . The sample excludes the 2006Q1 survey round.<sup>32</sup>

---

<sup>32</sup>The estimates are not sensitive to the treatment of the transitional 2006 rounds. Including the 2006Q1 round yields a slightly smaller main interaction (0.220 at 500 m and 0.173 at 1 km, both  $p < 0.01$ ); excluding also the 2006Q2 round, whose cumulative reading straddles the May 29 warrants, yields a slightly larger one (0.237 and 0.192, both  $p < 0.01$ ). Both movements are as expected: readings taken around the announcement mix the two regimes and attenuate the contrast with the reference period.

We estimate the following specification:

$$\begin{aligned} \log V_{it} = & \alpha_i + \mu_b + \gamma_{mt} + \beta_1 StopNbr_{it} + \beta_2 StopNbr_{it} \times Post0608_t + \beta_3 StopNbr_{it} \times Post08_t \\ & + \delta_1 NonShareNbr_{it} + \delta_2 NonShareNbr_{it} \times Post0608_t + \delta_3 NonShareNbr_{it} \times Post08_t \quad (6) \\ & + W'_{it}\lambda + X'_i\theta + \varepsilon_{it}, \end{aligned}$$

where  $\alpha_i$  is a station fixed effect,  $\mu_b$  is a brand fixed effect, and  $\gamma_{mt}$  is a market $\times$ quarter fixed effect.  $StopNbr_{it}$  is an indicator equal to one if at least one station within distance  $d$  (500 m or 1 km) has  $StopShare_j = 1$ . It varies over time only through the entry and exit of such neighbors, so under station fixed effects the level coefficient  $\beta_1$  is identified from that variation alone.  $NonShareNbr_{it}$  is an indicator equal to one if at least one neighbor within  $d$  does not share during period  $t$ .  $Post0608_t$  equals one for quarters between 2006Q3 and the end of 2008;  $Post08_t$  equals one after 2008.

The market $\times$ quarter fixed effects  $\gamma_{mt}$  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 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 (Ultramar, Petro-Canada, Esso, Irving, Shell), and distributor-brand rivals (Sonic, Olco, Petro T, Crevier, Harnois). 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  $X_i$  includes station-level amenity controls (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 coefficients of interest are  $\beta_2$  and  $\beta_3$ , which capture the differential volume effect of having a stop-sharing neighbor in the two post-investigation periods, relative to the pre-period. The level coefficient  $\beta_1$  tests for pre-existing differences between neighborhoods with and without stop-sharing stations. The non-sharing neighbor indicators  $NonShareNbr_{it}$  and their interactions are included as controls rather than as objects of interest.<sup>33</sup> The coefficients  $\delta_1$ – $\delta_3$  absorb any baseline volume differential associated with being close to a non-sharing station, so that  $\beta_1$ – $\beta_3$  isolate the effect of proximity to a station that *changed* its sharing status, holding fixed the overall prevalence of non-sharing stations in the neighborhood.

We split the post-investigation period at 2008 for two reasons. First, the event study in Figure 3 shows that the sharing decline partially reverses after 2010, suggesting that the behavioral change

---

<sup>33</sup>Of the 111 stations observed in the three markets, 11 never shared with the Platform before the collapse and 10 switched from sharing to non-sharing afterwards.




Figure 5: Volume event study: effect of stop-sharing neighbors (500 m)

*Notes:* Coefficients from a regression of log average daily volume on year interactions with  $\mathbf{1}(\text{Neighbor stop sharing within 500 m})$ , controlling for non-sharing neighbor year interactions, neighborhood composition, station amenities, and station, brand, and market  $\times$  quarter fixed effects. The reference year is 2005. Standard errors clustered at the station level.

may have been temporary. Second, volume data after 2008 are available only for Sherbrooke, as the quarterly panel for the other cartel cities ends in 2008. Table 13 in the Appendix presents summary statistics for the main outcome and explanatory variables used in the analysis.

### 5.3 Results

Figure 5 presents the event-study version of the 500 m specification, interacting the stop-sharing neighbor indicator with year dummies (reference year 2005). Figure 12 in the Appendix estimates the same specification for 1 km radius. The pre-trend coefficients for 2001–2004 are close to zero, which is consistent with parallel trends. The effect appears in 2006 at 0.16 log points and reaches 0.25 in 2007 and 0.24 in 2008, and then falls back to between 0.14 and 0.17 over 2009–2011. This timing is consistent with the interpretation that the investigation triggered a behavioral change among stop-sharing stations.

Table 17 in the Appendix 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 0.226 log points more over the 2006Q3–2008 period than always-sharing stations in the same market and quarter with no stop-sharing neighbor ( $p < 0.01$ ). The level coefficient ( $\beta_1 = 0.095$ ) is identified only from entry and exit of stop-sharing neighbors and is not statistically significant ( $p > 0.10$ ). The post-2008 coefficient ( $\beta_3 = 0.152$ ) is smaller and only marginally significant ( $p < 0.10$ ). After 2008 the sample is Sherbrooke alone, so we do not read it as evidence that the reallocation continued beyond the period of instability. The results at the 1 km radius (column 2) follow the same pattern but with attenuated magnitudes: the 2006Q3–2008 interaction is 0.181 ( $p < 0.01$ ). 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 at roughly 20 log points. 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 20 log points better than neighborhoods that did not. Appendix Table 4 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.

This is consistent with two complementary interpretations. Under the first, neighborhoods containing a station that stopped sharing priced more aggressively during the price-war episodes that followed the investigation and drew volume from the rest of the city. Under the second, the always-sharing neighborhoods include stations that repeatedly attempted to restore the price cycle and were not followed, and those attempts cost them volume. The posted-price data cannot differentiate the two interpretations, for the reason given in Section 5.1. The survey records one price per station per quarter, and an attempt to restore the price that is not followed is often reversed within a day, so the read is as likely to fall outside such an episode as inside it.

Two further benchmarks put the magnitude in context. The first is the Victoriaville reallocation of Figure 2. Between the pre-investigation period and 2007–2008, the tolerated low-price station on the periphery lost about 50 log points of market share and the deviating station about 85, while the centrally located big-box retailer gained 11 to 14. Single stations moved by amounts of the same order as our neighborhood estimates, and in the same direction, away from the periphery, toward the center. The second is the 2000 Québec City price war, studied by Clark et al. (2024) with the same survey data and which reallocated volume across neighborhoods on a comparable scale.

The volume reallocation documented here links information exchange, conduct, and market performance. 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 Conclusions and policy discussion

This paper investigates whether a third-party platform that collects and redistributes sales volumes helps competing firms sustain a collusive arrangement. The setting is 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, and the market and station fixed-effects estimates agree, so the decline reflects incumbent stations changing behavior rather than entry and exit. The withdrawals also affected the distribution of sales volume across neighborhoods. Always-sharing stations with a neighbor that stopped sharing gained roughly 20 log points in volume relative to always-sharing stations without one, an effect concentrated in 2006 to 2008, the years of price instability that followed the announcement of the investigation. Stations paid to be observed by their rivals, and they stopped paying once the arrangement that made the observation valuable had collapsed.

The model in Section 3.2 explains why the two move together. The reports and direct communication are complements. The reports supply a statistic whose standard of proof the cartel can set, and communication supplies evidence whose informativeness it cannot. Neither alone deters cheating as cheaply as the two together. The court record puts the contribution of the Platform in perspective. The durability of these cartels rested primarily on the telephone, and the wiretap evidence shows the reports entering those conversations as the device used to verify volumes and police the arrangement. The decline in participation is therefore consistent with the Platform facilitating coordination. It does not imply that the Platform was the main facilitating tool. Direct communication was.

These findings have direct implications for the design of competition policy toward information-sharing arrangements. Until 2023 the DOJ and the FTC maintained a safe harbor under which exchanges of competitively sensitive information were presumptively legal if three conditions held: the exchange was 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 no participant accounting for more than 25 percent of any statistic. Both agencies withdrew the guidance in 2023 as overly permissive, without issuing replacement criteria. The Platform we study would not have qualified for protection because it reported station-level rather than aggregated volumes, and total sales for the immediately preceding quarter. Our findings suggest that these restrictions were economically meaningful: even information that was a few months old could facilitate collusion when reported at the individual-station level, allowing cartel members to monitor transfers of volume and detect deviations. More broadly, our evidence points to the importance of considering the interaction between the granularity and frequency of information when evaluating information-sharing arrangements, rather than treating either dimension in isolation.

In 2023 the DOJ filed a civil antitrust suit against Agri Stats, alleging that its reports on pricing, margins, inventories, and operational metrics for pork, chicken, and turkey processors facilitated collusion among subscribers. Litigation over RealPage raises the same question in rental housing. The Platform we study differs from both in one respect that the law treats as important: 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. The distinction between active and passive platforms is legally relevant, but our results suggest it is not the economically relevant one. What matters is whether the information a platform disseminates lets firms monitor each other precisely

enough to sustain an arrangement.

## 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.
- Bernheim, B. D. and E. Madsen (2017). Price cutting and business stealing in imperfect cartels. *American Economic Review* 107(2), 387–424.
- Brown, Z. Y. and A. MacKay (2023). Competition in pricing algorithms. *American Economic Journal: Microeconomics* 15(2), 109–156.
- Byrne, D., N. de Roos, A. R. Grinberg, and L. M. Marx (2024). The role of platforms for facilitating anticompetitive communication. In *Cartels Diagnosed: New Insight on Collusion*. Cambridge University Press.
- 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*. Forthcoming.
- 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>.
- Calvano, E., G. Calzolari, V. Nicolò, and S. Pastorello (2020). Artificial intelligence, algorithmic pricing, and collusion. *American Economic Review* 110(10), 3267–3297.
- 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 (2023). The inner workings of a hub-and-spoke cartel in the automotive fuel industry. Forthcoming, *American Economic Journal: Microeconomics*.
- Clark, R., M. Duarte, and J.-F. Houde (2024). Cartel instability and price wars: Retail gasoline in canada. In *Book chapter, June 2024 draft*.
- 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.
- Ellison, G. (1994, Spring). Theories of cartel stability and the joint executive committee. *The Rand Journal of Economics* 25(1), 37–57.

- 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 (1999). The sugar institute learns to organize information exchange. In N. R. Lamoreaux, D. M. G. Raff, and P. Temin (Eds.), *Learning by Doing in Markets, Firms, and Countries*, pp. 103–138. University of Chicago Press.
- 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.
- Harrington, J. (2004). Post-cartel post-cartel pricing during litigation. *Journal of Industrial Economics* 52.
- Harrington, J. E. (2022). The effect of outsourcing pricing algorithms on market competition. *Management Science* 68(9), 6889–6906.
- Harrington, J. E. (2024). The effect of demand variability on the adoption and design of a third party’s pricing algorithm. *Economics Letters* 244, 112011.
- Harrington, J. E. (2025a). An economic test for an unlawful agreement to adopt a third-party’s pricing algorithm. *Economic Policy* 40, 263–295.
- Harrington, J. E. (2025b). 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.
- Houde, J.-F. (2012). Spatial differentiation and vertical mergers in retail markets for gasoline. *American Economic Review* 102, 2147–2182.
- Ichihashi, S. (2026). Algorithmic pricing and consumer data. Working paper, February 2026.
- Kandori, M. (1992). Social norms and community enforcement. *Review of Economic Studies* 59(1), 63–80.

- Kawai, K., J. Nakabayashi, and J. Ortner (2025). The value of privacy in cartels: An analysis of the inner workings of a bidding ring. *Review of Economic Studies*. Forthcoming.
- 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*.
- Kühn, K.-U. and S. Tadelis (2018). The economics of algorithmic pricing: Is collusion really inevitable? Working Paper.
- 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 (2008). Explicit collusion and market share allocations. Working paper.
- 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.
- Miklós-Thal, J. and C. Tucker (2019). Collusion by algorithm: Does better demand prediction facilitate coordination between sellers? *Management Science* 65(4), 1552–1561.
- Ortner, J., T. Sugaya, and A. Wolitzky (2024). Mediated collusion. *Journal of Political Economy*.
- 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 (2015). Price transparency and retail prices: Evidence from fuel price signs in the italian motorway. *Journal of Marketing Research* 53, 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 (2018). Maintaining privacy in cartels. *Journal of Political Economy* 126(6), 2570–2607.
- 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).

## A Proofs

This appendix proves Lemma 1 and Propositions 1, 2 and 3. Section A.1 records the objects that Section 3.2 uses but does not display: a follower's value of the arrangement, the two incentive constraints, and the problem the cartel solves when it chooses the trigger. The remaining subsections take the results in the order in which they appear in the text. Section A.3 adds one lemma, which establishes the property of the optimal rule described after Proposition 1: the cutoff that applies after a complaint is lower than the cutoff that applies after none, and the gap is the log odds ratio of the message.

Two conventions apply throughout. We write  $\Phi$  for the standard normal distribution function and  $\varphi$  for its density. The letter  $\phi$  always denotes the message technology of Section 3.2 and never a normal density. The message takes the two values  $C$ , a complaint, and  $N$ , no complaint, which are the subscripts already attached to the cutoffs in (4).

### A.1 Value functions, incentive constraints, and the cartel's problem

We first fix the primitives that the propositions use. The idiosyncratic shock satisfies  $\log u_{jt} \sim N(-\sigma_u^2/2, \sigma_u^2)$ , so that  $u_{jt}$  has mean one, and the shocks are independent across firms and periods and independent of  $m_t$ . Within a period, firms first choose prices, including any deviation. Then  $m_t$  and the  $u_{jt}$  are realized and volumes are sold. Then the report and the message arrive. A firm therefore chooses its deviation before it knows the demand it will face, and every profit below is an expectation taken before that realization.

Under the arrangement a follower earns

$$\pi^{\text{col}} \equiv \bar{m} \left( s - \frac{\tau}{n-1} \right) (1-c) \quad (7)$$

per period. During a price war the price is  $c$ , firm  $L$  serves the market and every follower earns zero, so a war costs a follower its entire collusive profit and no parameter of the punishment enters the analysis beyond its length  $W$ .

The trigger (4) is stated in standard deviations of the statistic the cartel actually computed. It is convenient to name the two on-path probabilities it induces directly,

$$\alpha_{0,C} \equiv \Phi(-\kappa_C), \quad \alpha_{0,N} \equiv \Phi(-\kappa_N), \quad \alpha_0 \equiv \phi_0 \alpha_{0,C} + (1 - \phi_0) \alpha_{0,N}, \quad (8)$$

where  $\alpha_{0,C}$  is the probability of a war in a period in which nobody cheated but a complaint arrived anyway, and  $\alpha_{0,N}$  the probability in a period with neither. Their average  $\alpha_0$ , weighted by how often a complaint arrives on the equilibrium path, is the frequency of price wars the members live with. No standard deviation appears in (8), because the cutoff is standardized: the on-path probability of crossing it is the same whichever set of reports the statistic was computed on.

Let  $V^{\text{col}}$  be the value to a follower of being in the collusive phase at the start of a period. It earns  $\pi^{\text{col}}$  during the period. At the end of the period a war begins with probability  $\alpha_0$ , and a war is  $W$

periods of zero profit followed by a return to the collusive phase, so entering one is worth  $\delta^W V^{\text{col}}$ . With probability  $1 - \alpha_0$  the collusive phase continues. Hence  $V^{\text{col}} = \pi^{\text{col}} + \delta[\alpha_0 \delta^W V^{\text{col}} + (1 - \alpha_0) V^{\text{col}}]$ , and collecting terms,

$$V^{\text{col}} = \frac{\pi^{\text{col}}}{1 - \delta + \delta \alpha_0 (1 - \delta^W)}. \quad (9)$$

The denominator is the ordinary  $1 - \delta$  plus a term for the wars the cartel inflicts on itself by mistake. The more often the trigger fires on the equilibrium path, the less the arrangement is worth. Entering a war rather than continuing to collude is worth  $\delta^W V^{\text{col}}$  instead of  $V^{\text{col}}$ , a difference of  $(1 - \delta^W) V^{\text{col}}$ .

Consider the public price deviation. A follower that posts just below the collusive price earns  $\bar{m}(1 - c)$  this period instead of  $\pi^{\text{col}}$  and is caught with certainty, so  $V^P = \bar{m}(1 - c) + \delta^{W+1} V^{\text{col}}$ . Subtracting from (9) and using (7),

$$V^{\text{col}} - V^P = -\bar{m}(1 - c) \left(1 - s + \frac{\tau}{n - 1}\right) + \delta(1 - \alpha_0)(1 - \delta^W) V^{\text{col}}, \quad (10)$$

so the deviation is unprofitable when

$$\delta(1 - \alpha_0)(1 - \delta^W) V^{\text{col}} \geq \bar{m}(1 - c) \left(1 - s + \frac{\tau}{n - 1}\right). \quad (\text{IC-P})$$

The left side is the threat, the probability that the deviation is not already being punished times the discounted cost of a war. The right side is the one-period gain, the whole market less the share the firm already had. Substituting (9) and (7), the factor  $\bar{m}(1 - c)$  cancels from both sides and (IC-P) becomes a condition on the shares and the punishment alone:

$$\delta(1 - \alpha_0)(1 - \delta^W) \left(s - \frac{\tau}{n - 1}\right) \geq \left(1 - s + \frac{\tau}{n - 1}\right) [1 - \delta + \delta \alpha_0 (1 - \delta^W)]. \quad (\text{IC-P}')$$

Neither the level of demand nor the cost appears. Because a war costs a follower everything it was earning, the constraint is homogeneous in the margin.

Now the secret delay. Write  $\alpha(d)$  for the probability of a war after a delay of size  $d$ ; Section A.2 computes it once Lemma 1 has pinned down what the deviator reports. By (2) a follower that delays by  $d$  earns  $\pi^{\text{col}} + \bar{m}d(1 - \Delta_p - c)$  this period and faces a war with probability  $\alpha(d)$  rather than  $\alpha_0$ , so  $V^S(d) = \pi^{\text{col}} + \bar{m}d(1 - \Delta_p - c) + \delta[\alpha(d)\delta^W V^{\text{col}} + (1 - \alpha(d))V^{\text{col}}]$ . Subtracting from (9), the collusive profits cancel and the continuation terms differ only in the probability attached to the war branch, so the delay is unprofitable at every admissible size when

$$\delta[\alpha(d) - \alpha_0](1 - \delta^W) V^{\text{col}} \geq \bar{m}d(1 - \Delta_p - c) \quad \text{for every } d \in (0, \bar{d}]. \quad (\text{IC-S})$$

The object  $\alpha(d) - \alpha_0$  is the only place either signal appears. We call it the excess: the amount by which cheating of size  $d$  raises the probability of a war above the frequency the cartel runs anyway. Everything the two monitoring technologies do, they do through this one number. Substituting (9)

and (7) and cancelling  $\bar{m}(1 - c)$  again, and then dividing by  $d$ ,

$$\frac{\alpha(d) - \alpha_0}{d} \geq \frac{1 - \delta + \delta \alpha_0 (1 - \delta^W)}{\delta(1 - \delta^W)(s - \tau/(n - 1))} \cdot \frac{1 - \Delta_p - c}{1 - c} \quad \text{for every } d \in (0, \bar{d}]. \quad (\text{IC-S'})$$

The constant on the right is increasing in  $\alpha_0$  and decreasing in  $\delta$ , in  $W$  and in the follower's share. The limit of the left side as  $d \rightarrow 0$  is  $\alpha'(0)$ , so the requirement includes the marginal condition. Nothing in Propositions 1 and 2 depends on the value of the factor  $(1 - \Delta_p - c)/(1 - c)$ , only on its being a positive constant that involves neither  $d$  nor  $\delta$ , so both may be read at  $\Delta_p = 0$ . Proposition 3 is where the jump does work.

The cartel picks the trigger to make price wars as rare as possible while still deterring both deviations. Formally it minimizes  $\alpha_0$  subject to (IC-P') and (IC-S'), over the cutoffs available to it. Two features of this problem matter below. First, (IC-P') does not involve the signals, and it is easier to satisfy the smaller  $\alpha_0$  is, so it places an upper bound on the war frequency the cartel can afford and never conflicts with the objective. Second, (IC-S') is where the signals enter, and a smaller  $\alpha_0$  moves its two sides in opposite directions: a quieter trigger delivers a smaller excess, which tightens the constraint, but it also lowers the constant on the right, because fewer wars on the equilibrium path make the arrangement worth more and the threat of a war larger. The first effect dominates at the frequencies of interest, since a trigger that never fires deters nothing. We write  $\alpha_0^*$  for the infimum of the set of  $\alpha_0$  at which some rule satisfies (IC-S'), and the proofs below use only that definition. The infimum is attained. Allow each cutoff to take the values  $\pm\infty$ , so that a rule in regime (A) is a point of  $[-\infty, +\infty]^2$ , which is compact, and extend the left side of (IC-S') to  $d = 0$  by its limit  $\alpha'(0)$ . Both sides of (IC-S') are then continuous in the cutoffs and in  $d \in [0, \bar{d}]$ , the feasible set is closed, and  $\alpha_0$  attains its infimum on it. At the optimum (IC-S') binds at some  $d^* \in [0, \bar{d}]$ : were it slack at every  $d$ , the slack would have a positive minimum over  $[0, \bar{d}]$ , and raising both cutoffs slightly would preserve feasibility and lower  $\alpha_0$ . The same argument applies in regimes (B) and (C), whose rules are points of  $[0, 1]^2$  and of  $[-\infty, +\infty]$ . Collusion is sustainable if and only if  $\alpha_0^*$  also satisfies (IC-P').

## A.2 Proof of Lemma 1

**Lemma 1** (A deviator withholds). *Let the idiosyncratic shocks be log-normal and let the trigger be the lower-tail cutoff (4), standardized by the standard deviation of the statistic the cartel actually computed. Then a follower that takes a secret delay of size  $d > 0$  strictly prefers to withhold its volume from that period's report, for every pair of cutoffs in (4) at least one of which is finite.*

*Proof.* Compare the two options at a fixed trigger. If the deviator reports, all  $n - 1$  followers appear in the report, and the deviator's own inflated volume, whose share is  $s - \tau/(n - 1) + d$ , enters the average in (3) with weight  $1/(n - 1)$ . The mean of  $Y$  then falls by

$$-\log\left(1 - \frac{d}{s + \tau}\right) + \frac{1}{n - 1} \log\left(1 + \frac{d}{s - \tau/(n - 1)}\right), \quad (11)$$

the first term from  $L$ 's lower volume and the second from the deviator's higher one, and the standard deviation of  $Y$  is  $\sigma_u \sqrt{n/(n-1)}$ . If the deviator withholds, the report contains  $n-2$  followers, the mean falls by the first term only, and the standard deviation is  $\sigma_u \sqrt{(n-1)/(n-2)}$ , which is larger. The cutoff is restated in the standard deviation of the statistic actually computed, so withholding carries no penalty of its own. The standardized fall in the mean is therefore strictly smaller when the deviator withholds, on both counts. Because the shocks are log-normal, the standardized statistic is standard normal whichever set of followers it is computed on, so the probability of crossing a standardized cutoff depends only on the standardized fall, and it is strictly lower at every finite cutoff when the deviator withholds. Since  $\phi(d) \in (0, 1)$ , both branches of (4) carry weight, so one finite cutoff is enough for the preference to be strict.  $\square$

The last step is where log-normality is used. Under other shock distributions the statistic computed on  $n-2$  reports and the statistic computed on  $n-1$  share a mean but not a shape, and a cutoff set at a given quantile of one can have a different density at that quantile from the other. The lemma then need not hold at every cutoff. Proposition 1 does not depend on it.

The lemma pins down what the cartel observes after a deviation, and therefore  $\alpha(d)$ . Under a deviation the statistic is computed on  $n-2$  followers, its standard deviation is  $\sigma_u \sqrt{(n-1)/(n-2)}$ , and its mean is

$$\mu(d) = \mu_0 + \log\left(1 - \frac{d}{s + \tau}\right), \quad \mu(0) = \mu_0, \quad -\mu'(0) = \frac{1}{s + \tau} > 0, \quad (12)$$

which is below  $\mu_0$  and decreasing in  $d$ . This is where  $n \geq 3$  is used: at least one follower must still report after the deviator has withdrawn, so that the cartel retains a statistic in which the common component of demand cancels. Writing

$$z(d) \equiv \frac{\mu_0 - \mu(d)}{\sigma_u \sqrt{(n-1)/(n-2)}} \quad (13)$$

for the fall in the mean expressed in standard deviations of the statistic the cartel computes, the quantity  $(Y - \mu(d))/(\sigma_u \sqrt{(n-1)/(n-2)})$  is standard normal, and averaging over the two values of the message gives

$$\alpha(d) = \phi(d) \Phi(-\kappa_C + z(d)) + (1 - \phi(d)) \Phi(-\kappa_N + z(d)). \quad (14)$$

At  $d = 0$  this returns  $\alpha_0$ , as it must: a follower that does not cheat faces the on-path frequency. The two technologies enter differently. The Platform moves the argument of  $\Phi$  inside each branch. The message moves the weight between the two branches.

One restriction on the rule is worth stating, because it works against the results rather than for them. Since a deviator withholds, the statistic under a deviation has a larger variance as well as a lower mean, so a fully optimal test would also call a war when  $Y$  is extremely high. We exclude that. A report showing  $L$  with an unusually large share is not evidence that a follower has been stealing from  $L$ , and no cartel would write such a rule. The restriction costs the Platform some

power, so it makes Proposition 1 harder to prove, not easier.

### A.3 The optimal rule uses state-contingent cutoffs

**Lemma 2** (State-contingent cutoffs). *In regime (A), suppose the cartel's problem is solved with  $\kappa_C$  and  $\kappa_N$  both finite and with (IC-S') binding at a single deviation size  $d^* \in (0, \bar{d}]$ . Then  $\kappa_C < \kappa_N$  and*

$$\kappa_N - \kappa_C = \frac{1}{z(d^*)} \log \left[ \frac{\phi(d^*)}{\phi_0} \cdot \frac{1 - \phi_0}{1 - \phi(d^*)} \right]. \quad (15)$$

*Proof.* Write  $z \equiv z(d^*) > 0$ . Lowering  $\kappa_C$  by a small amount raises the on-path frequency at the rate  $\phi_0 \varphi(\kappa_C)$ , by (8), and raises the probability of catching the deviation at the rate  $\phi(d^*) \varphi(\kappa_C - z)$ , by (14). The ratio of what it buys to what it costs is therefore  $[\phi(d^*)/\phi_0] \cdot [\varphi(\kappa_C - z)/\varphi(\kappa_C)]$ , and the same computation in the other branch gives  $[(1 - \phi(d^*))/(1 - \phi_0)] \cdot [\varphi(\kappa_N - z)/\varphi(\kappa_N)]$ . At an interior optimum the two ratios are equal, since otherwise the cartel could relax the constraint at no cost by tightening the branch with the better ratio and loosening the other. This holds  $d^*$  fixed, which is legitimate to first order because the constraint binds there. A solution of the resulting equation exists because each ratio, as a function of its own cutoff, ranges over the whole positive line. Using  $\varphi(\kappa - z)/\varphi(\kappa) = \exp(\kappa z - z^2/2)$ , equality reads

$$\frac{\phi(d^*)}{\phi_0} e^{\kappa_C z} = \frac{1 - \phi(d^*)}{1 - \phi_0} e^{\kappa_N z}, \quad (16)$$

which rearranges to (15). The right side of (15) is strictly positive because  $\phi(d^*) > \phi_0$  makes the first factor exceed one and the second fall below one, so  $\kappa_C < \kappa_N$ .  $\square$

When the binding deviation is the marginal one the gap has a closed form. As  $d^* \rightarrow 0$ ,  $\mu_0 - \mu(d^*) \rightarrow d^*/(s + \tau)$  by (12) and the log odds ratio behaves like  $\phi'(0)d^*/[\phi_0(1 - \phi_0)]$ , so

$$\kappa_N - \kappa_C \longrightarrow \sigma_u \sqrt{\frac{n-1}{n-2}} (s + \tau) \frac{\phi'(0)}{\phi_0(1 - \phi_0)}. \quad (17)$$

The gap widens with the informativeness of the message at zero and narrows with the precision of the Platform statistic, since a smaller  $\sigma_u$  makes the volume evidence worth more relative to the message. A complaint is evidence, and evidence should lower the standard of proof the cartel demands of the volume report. How much it lowers it depends on how much more likely a complaint is when someone has cheated than when nobody has, which is the odds ratio, and on how precise the volume report is, since a precise report needs less help.

### A.4 Proof of Proposition 1

The three regimes of Section 3.2 differ only in the excess they deliver. In regime (A) it is (14) minus (8). In regime (B) there is no Platform, so the rule is a pair of probabilities, the cartel calls

a war with probability  $\alpha_{0,C}$  after a complaint and  $\alpha_{0,N}$  otherwise, and

$$\alpha(d) - \alpha_0 = [\phi(d) - \phi_0](\alpha_{0,C} - \alpha_{0,N}). \quad (18)$$

The excess is the informativeness of the message times the gap between the two war probabilities, and Lemma 1 has no content here. In regime (C) there is no message and a single cutoff  $\kappa$  applies, so  $\alpha_0 = \Phi(-\kappa)$  and

$$\alpha(d) - \alpha_0 = \Phi(-\kappa + z(d)) - \Phi(-\kappa). \quad (19)$$

Both rules are available to the cartel of regime (A): setting  $\kappa_C = \kappa_N$  in (4) gives (C), and making the war probability independent of  $Y$  within each branch gives (B).

One property of the optimum in regime (B) is used in the proof. There the cartel minimizes  $\phi_0\alpha_{0,C} + (1 - \phi_0)\alpha_{0,N}$  while (18) shows that only the gap  $\alpha_{0,C} - \alpha_{0,N}$  delivers deterrence. For a given gap the cheapest way to attain it is to set  $\alpha_{0,N} = 0$ , so the optimal rule never calls a war in the absence of a complaint and  $\alpha_0^*(B) = \phi_0\alpha_{0,C}^*$ , where  $\alpha_{0,C}^*$  is the smallest gap that satisfies (IC-S'). A cartel with communication alone pays for deterrence entirely in false accusations.

*Proof of Proposition 1. Step 1: regime (A) can do what regime (C) does.* Setting  $\kappa_C = \kappa_N = \kappa$  in (4) reproduces (19) exactly, so any  $\alpha_0$  feasible in (C) is feasible in (A) and  $\alpha_0^*(A) \leq \alpha_0^*(C)$ . The comparison with regime (B) is handled directly in Step 3, which exhibits for every rule of regime (B) a cutoff rule of regime (A) that is strictly better.

*Step 2: strictly better than the Platform alone.* Let  $\kappa$  solve the problem in regime (C), so that (IC-S') holds with equality at some binding  $d^*$ , and consider in regime (A) the pair  $\kappa_C = \kappa - \eta$  and  $\kappa_N = \kappa + \eta\phi_0/(1 - \phi_0)$ , for small  $\eta > 0$ . By (8) the on-path frequency changes by

$$\phi_0[\Phi(-\kappa + \eta) - \Phi(-\kappa)] + (1 - \phi_0)\left[\Phi\left(-\kappa - \frac{\phi_0}{1 - \phi_0}\eta\right) - \Phi(-\kappa)\right] = \phi_0\varphi(\kappa)\eta - \phi_0\varphi(\kappa)\eta + O(\eta^2) = O(\eta^2), \quad (20)$$

so  $\alpha_0$  is unchanged to first order. By (14) the probability of catching a deviation of size  $d$  changes by

$$\phi(d)\varphi(\kappa - z(d))\eta - (1 - \phi(d))\varphi(\kappa - z(d))\frac{\phi_0}{1 - \phi_0}\eta + O(\eta^2) = \varphi(\kappa - z(d))\frac{\phi(d) - \phi_0}{1 - \phi_0}\eta + O(\eta^2). \quad (21)$$

For every  $d > 0$  the coefficient of  $\eta$  in (21) is strictly positive, because  $\phi$  is increasing and  $\phi_0 < 1$ . Two further facts make the improvement uniform over  $(0, \bar{d}]$ . The coefficient divided by  $d$  has the strictly positive limit  $\varphi(\kappa)\phi'(0)/(1 - \phi_0)$  as  $d \rightarrow 0$ , and it is continuous and positive on  $(0, \bar{d}]$ , so its infimum over that interval is positive. And the  $O(\eta^2)$  terms remain uniform in  $d$  after division by  $d$ . The excess is a smooth function of  $\eta$  whose derivatives in  $\eta$  are continuously differentiable in  $d$  on  $[0, \bar{d}]$ , because  $\phi$  is, and it vanishes at  $d = 0$  for every  $\eta$ ; the excess divided by  $d$  therefore has the same properties, and its expansion in  $\eta$  has a remainder that is  $O(\eta^2)$  uniformly on  $[0, \bar{d}]$ . The left side of (IC-S') therefore rises by at least a fixed positive multiple of  $\eta$  at every  $d$ , while the right side moves by  $O(\eta^2)$  through  $\alpha_0$ . The cartel can then raise both cutoffs by a common  $\varepsilon$

proportional to  $\eta$ , which lowers  $\alpha_0$  by  $[\phi_0\varphi(\kappa_C) + (1 - \phi_0)\varphi(\kappa_N)]\varepsilon$  to first order and consumes only part of the slack. Hence  $\alpha_0^*(A) < \alpha_0^*(C)$ . This is where  $\phi'(0) > 0$  is used: without it the message would carry no information about small deviations and the limit above would be zero.

*Step 3: strictly better than communication alone.* Let  $(\alpha_{0,C}, \alpha_{0,N})$  solve the problem in regime (B), with  $\alpha_{0,N} = 0$ . There are two cases.

If  $\alpha_{0,C} < 1$ , consider in regime (A) the rule that keeps the  $N$  branch silent and replaces the constant probability  $\alpha_{0,C}$  by a cutoff of the same on-path size, that is  $\kappa_C = -\Phi^{-1}(\alpha_{0,C})$ , which is finite. The on-path frequency is unchanged, since  $\Phi(-\kappa_C) = \alpha_{0,C}$ . The probability of catching a deviation of size  $d$  becomes  $\phi(d)\Phi(-\kappa_C + z(d))$ , which strictly exceeds  $\phi(d)\Phi(-\kappa_C) = \phi(d)\alpha_{0,C}$  because  $\mu(d) < \mu_0$  for  $d > 0$  and  $\Phi$  is strictly increasing. The last expression is the regime (B) value, so the excess is strictly larger at every  $d > 0$  at an unchanged  $\alpha_0$ . The increase divided by  $d$  has the positive limit  $\phi_0\varphi(\kappa_C)z'(0)$  as  $d \rightarrow 0$ , so it is bounded away from zero on  $(0, \bar{d}]$  as in Step 2, and the cartel can raise  $\kappa_C$  slightly and lower  $\alpha_0$ .

If  $\alpha_{0,C} = 1$ , the regime (B) rule calls a war after every complaint, with  $\alpha_0 = \phi_0$  and excess  $\phi(d) - \phi_0$ . Consider in regime (A) the rule with  $\kappa_N = +\infty$  and  $\kappa_C = k$  for  $k$  very negative, so that a war follows a complaint unless the Platform statistic is extremely high. Its on-path frequency is  $\phi_0\Phi(-k) = \phi_0 - \phi_0\Phi(k)$ , strictly below  $\phi_0$ . Its excess at  $d$  is  $\phi(d)\Phi(-k + z(d)) - \phi_0\Phi(-k)$ , which exceeds the regime (B) excess by

$$\phi_0\Phi(k) - \phi(d)\Phi(k - z(d)). \quad (22)$$

Because  $\log \Phi$  is concave,  $\Phi(k - z) \leq \Phi(k) \exp(-z\varphi(k)/\Phi(k))$  for every  $z > 0$ , so (22) divided by  $\Phi(k)$  is at least

$$\phi_0 - \phi(d) \exp(-z(d)\varphi(k)/\Phi(k)). \quad (23)$$

This bound is zero at  $d = 0$ , and its derivative in  $d$  is  $\exp(-z(d)\varphi(k)/\Phi(k)) [\phi(d)z'(d)\varphi(k)/\Phi(k) - \phi'(d)]$ . Since  $\phi \geq \phi_0$  and  $z'$  is increasing, the derivative is strictly positive on all of  $[0, \bar{d}]$  as soon as  $\varphi(k)/\Phi(k)$  exceeds the maximum of  $\phi'$  on  $[0, \bar{d}]$  divided by  $\phi_0z'(0)$ , and  $\varphi(k)/\Phi(k)$  grows without bound as  $k \rightarrow -\infty$ . For  $k$  negative enough, (23) is therefore zero at  $d = 0$  with a derivative bounded away from zero on  $[0, \bar{d}]$ , so it is at least a fixed positive multiple of  $d$  at every  $d \in (0, \bar{d}]$ . The regime (A) rule then has a strictly lower on-path frequency than the regime (B) rule and an excess that exceeds the regime (B) excess by at least a positive multiple of  $d$  at every  $d \in (0, \bar{d}]$ , and the strict inequality follows. No proviso on regime (B) is needed.

*Step 4: from war frequencies to discount factors.* Fix a regime and a rule. Constraint (IC-S') reads

$$\frac{\alpha(d) - \alpha_0}{d} \geq \frac{1}{s - \tau/(n-1)} \cdot \frac{1 - \Delta_p - c}{1 - c} \left[ \frac{1 - \delta}{\delta(1 - \delta^W)} + \alpha_0 \right], \quad (24)$$

and  $(1 - \delta)/[\delta(1 - \delta^W)]$  is strictly decreasing in  $\delta$ , because its reciprocal  $\delta + \delta^2 + \dots + \delta^W$  is strictly increasing. The left side does not involve  $\delta$ , so for a fixed rule (IC-S') holds on an upper interval of

$\delta$ . Dividing (IC-P') by  $\delta(1 - \delta^W)$  in the same way gives

$$(1 - \alpha_0) \left( s - \frac{\tau}{n-1} \right) \geq \left( 1 - s + \frac{\tau}{n-1} \right) \left[ \frac{1 - \delta}{\delta(1 - \delta^W)} + \alpha_0 \right], \quad (25)$$

whose right side is decreasing in  $\delta$  by the same computation, so (IC-P') also holds on an upper interval of  $\delta$ . Collusion is therefore sustainable on an upper interval of  $\delta$  in each regime, which defines  $\delta_{\text{crit}}$  as the infimum of that interval. At any  $\delta$  at which regime (C) is sustainable, Step 2 gives a regime (A) rule with both constraints strictly slack, and by continuity in  $\delta$  that rule remains feasible at some smaller  $\delta$ , so regime (A) is sustainable on a strictly larger interval and  $\delta_{\text{crit}}(A) < \delta_{\text{crit}}(C)$ . Step 3 covers regime (B) in both of its cases, including the case in which regime (B) stops being sustainable because the war probability after a complaint has reached one, so the same argument gives  $\delta_{\text{crit}}(A) < \delta_{\text{crit}}(B)$ .  $\square$

One condition is worth recording, because it says when the comparison has content. Letting  $\delta \rightarrow 1$  in (25), and using  $(1 - \delta)/[\delta(1 - \delta^W)] \rightarrow 1/W$ , gives a necessary condition for collusion to be sustainable at any discount factor in any regime:

$$W(1 - \alpha_0) \left( s - \frac{\tau}{n-1} \right) \geq \left( 1 - s + \frac{\tau}{n-1} \right) (1 + W\alpha_0). \quad (26)$$

With  $\alpha_0$  near zero this says the punishment must last at least  $(1 - s + \tau/(n-1))/(s - \tau/(n-1))$  periods, a little more than  $n-1$ : a war must cost a follower at least as much as it can gain by taking the whole market for one period. Short punishments make every regime infeasible, and the proposition then compares nothing.

## A.5 Proof of Proposition 2

Let  $k \in \{0, 1, \dots, n-3\}$  be the number of followers who never supply a volume, so that  $n-1-k$  followers report under the arrangement and  $L$  reports always. Everything else is as before: the same trigger (4), the same message technology, and the same punishment of length  $W$ . Followers are symmetric here, because the question is about the number of reports and not about which firm supplies them.

Two observations fix the effect of  $k$  on the two sides of (IC-S'). The first is that  $k$  does not change the on-path frequency of price wars. On the equilibrium path the statistic (3) is computed on  $r = n-1-k$  followers and its standard deviation is  $\sigma_u \sqrt{1 + 1/r}$ , which is larger the larger is  $k$ . The cutoff is stated in standard deviations of the statistic the cartel actually computes, so  $\alpha_0$  in (8) is unchanged. The cartel does not call more wars by mistake when fewer members report. The second is that  $k$  does change the excess. Lemma 1 applies unchanged to a follower that reports, so a deviator of that kind withholds and the statistic is computed on  $n-2-k$  followers. This is where  $k \leq n-3$  is used: at least one follower must still report after the deviator has withdrawn, or the common component of demand does not cancel and there is no statistic. The mean under a

deviation,  $\mu(d)$  in (12), does not depend on  $k$ , because the fall comes from  $L$ 's volume alone and the reporting followers all hold the same share whatever their number. The standardized fall does. Writing

$$z_k(d) \equiv \frac{\mu_0 - \mu(d)}{\sigma_u \sqrt{1 + \frac{1}{n-2-k}}}, \quad \alpha_k(d) = \phi(d) \Phi(-\kappa_C + z_k(d)) + (1 - \phi(d)) \Phi(-\kappa_N + z_k(d)), \quad (27)$$

which is (14) with the number of reports made explicit, we can state the proposition and prove it.

**Proposition 2** (Participation and stability). *Fix the message technology and any trigger of the form (4) with  $\kappa_C$  finite, and let  $k \leq n-3$ . Then  $\alpha_0$  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  $\delta_{\text{crit}}$  is therefore strictly increasing in  $k$ .*

*Proof.* That  $\alpha_0$  does not depend on  $k$  is the first observation above. For the excess,  $1 + 1/(n-2-k)$  is strictly increasing in  $k$ , so  $z_k(d)$  in (27) is strictly decreasing in  $k$  for every  $d > 0$ , since  $\mu(d) < \mu_0$ . Both branches of  $\alpha_k(d)$  are strictly increasing in  $z_k(d)$  whenever the corresponding cutoff is finite, and  $\kappa_C$  is finite by hypothesis, so  $\alpha_k(d)$  is strictly decreasing in  $k$ . Subtracting  $\alpha_0$ , which does not move, gives the statement about the excess. For the discount factor, fix a rule and compare  $k$  with  $k' > k$ . The left side of (IC-S') is strictly smaller at  $k'$  at every  $d$ , and the right side is unchanged, since it involves only  $\alpha_0$ ,  $\delta$ ,  $W$ , the follower's share, and  $\Delta_p$ . By the argument of Step 4 of the proof of Proposition 1, the right side is strictly decreasing in  $\delta$  and the left side does not involve  $\delta$ , so each constraint holds on an upper interval of  $\delta$  and the interval at  $k'$  is a strict subset of the interval at  $k$ . Constraint (IC-P') does not involve the signals and is unaffected. Taking the union over rules preserves the strict inclusion. A rule that satisfies (IC-S') at  $(k', \delta)$  satisfies it at  $(k, \delta)$  with slack bounded away from zero on  $(0, \bar{d}]$ , by the argument of Step 2 of the proof of Proposition 1: the increase in the excess divided by  $d$  has a strictly positive limit at  $d = 0$  and is continuous and positive on  $(0, \bar{d}]$ . If (IC-P') binds at  $(k', \delta)$ , raise both cutoffs by a small common amount. This lowers  $\alpha_0$ , which relaxes (IC-P') strictly by (25), and it consumes only part of the slack in (IC-S'). Both constraints are then strictly slack at  $(k, \delta)$ , and by continuity in  $\delta$  they hold at some  $\delta' < \delta$ .  $\square$

The boundary case is worth naming. At  $k = n-2$  a single follower reports, and if that follower deviates and withholds, no follower reports at all. The cartel then has the message and nothing else, and once the rule is extended to periods with no report it is in regime (B). Proposition 2 therefore describes a continuous path from regime (A) to regime (B), travelled by the number of members willing to be measured rather than by the availability of the technology. Proposition 1 gives the two endpoints and Proposition 2 gives the interior.

## A.6 The value of a secret delay, and the proof of Proposition 3

This subsection does two things. It derives the gain from a secret delay stated in (2), which the body asserts without deriving, and it uses that derivation to prove Proposition 3. One change to the environment is needed for the proposition and only one. Follower  $i$  has constant marginal cost  $c_i$ , with  $0 < c_1 \leq c_2 \leq \dots \leq c_{n-1}$  and  $\Delta_p < 1 - c_{n-1}$ , while  $L$  still has cost zero. Every other object is the one used for Propositions 1 and 2: the same collusive price, the same shares, the same two deviations, the same signals, the same trigger, and the same punishment. Collusive shares are unchanged, so every follower holds  $s - \tau/(n - 1)$ , which we write  $s_f$  from here on, and we write  $D \equiv 1 - \delta + \delta\alpha_0(1 - \delta^W)$  for the denominator of (9).

**Where the gain comes from.** Within a quarter the collusive price is raised a fixed number of times. During each increase there is an interval in which the firms that have already raised post the collusive margin, normalized to one, and the firm that has not yet raised posts  $1 - \Delta_p$ . The product is homogeneous, so the whole flow of purchases in that interval goes to the low posting. Let  $\lambda$  be the fraction of the quarter's volume sold in those intervals. Under the arrangement  $L$  alone holds the low posting, so it takes all of  $\lambda$  and the remaining  $1 - \lambda$  is split equally, which gives  $L$  a share  $\lambda + s(1 - \lambda)$  and each follower  $s(1 - \lambda)$ . Matching these to the shares of Section 3.2 pins  $\lambda$ :

$$\lambda + s(1 - \lambda) = s + \tau \iff \lambda = \frac{\tau}{1 - s}, \quad \text{and then} \quad s(1 - \lambda) = s - \frac{\tau}{n - 1} = s_f \quad (28)$$

for each follower, so the transfer is funded equally by the followers with nothing further assumed. This requires  $\tau < 1 - s$ , which the model assumes. The number of increases does not appear. Shares aggregate linearly and every interval carries the same  $\Delta_p$ , so only the total transfer  $\tau$  and the size  $\Delta_p$  enter, and how many increases produce them is immaterial.

Two prices therefore appear in a follower's profit, and which one applies depends on whether the follower is in the interval. Under the arrangement it is not: it has already raised, it sells nothing while  $L$  holds the low posting, and its whole share  $s_f$  is sold at the collusive margin, so

$$\pi_i^{\text{col}} = \bar{m} s_f (1 - c_i), \quad (29)$$

and a price war costs it this amount in full. A follower that takes the secret delay is in the interval, because delaying is what puts it there. It posts  $1 - \Delta_p$  alongside  $L$ , takes a share  $d$  of the quarter's volume that  $L$  would otherwise have taken, and sells that volume at  $1 - \Delta_p$  while the share it already held is still sold at the collusive margin. Its profit is  $\bar{m}[s_f + d(1 - \Delta_p)] - \bar{m}(s_f + d)c_i$ , which is  $\pi_i^{\text{col}} + \bar{m}d(1 - \Delta_p - c_i)$ . The second term is (2). The stake and the prize are margins on two different prices, and that is what separates the followers.

**The index.** Repeating the derivations of Section A.1 with  $c_i$  in place of  $c$ , the two constraints for follower  $i$  are

$$\delta(1-\alpha_0)(1-\delta^W)\frac{\pi_i^{\text{col}}}{D} \geq \bar{m}(1-c_i)(1-s_f), \quad \delta[\alpha(d)-\alpha_0](1-\delta^W)\frac{\pi_i^{\text{col}}}{D} \geq \bar{m}d(1-\Delta_p-c_i), \quad (30)$$

the second for every  $d \in (0, \bar{d}]$ . Define the ratio of the prize from the secret delay to the stake it puts at risk,

$$R_i \equiv \frac{1-\Delta_p-c_i}{1-c_i} = 1 - \frac{\Delta_p}{1-c_i} \in (0, 1). \quad (31)$$

Substituting (29) into (30) and cancelling  $\bar{m}(1-c_i)$  from both sides of each constraint, they read

$$\delta(1-\alpha_0)(1-\delta^W)s_f \geq (1-s_f)D, \quad \frac{\alpha(d)-\alpha_0}{d} \geq \frac{DR_i}{\delta(1-\delta^W)s_f}. \quad (32)$$

The first is (IC-P') and contains no  $i$  at all. The second is (IC-S') with the factor  $R_i$  attached, and  $R_i$  is the only place the follower's identity appears anywhere in the model.

**Proposition 3** (The secret delay is hardest to deter for the efficient follower). *Let  $\Delta_p < 1 - c_{n-1}$ . Then the constraint against the public price deviation is the same for every follower. The index  $R_i$  is strictly decreasing in  $c_i$ , so the constraint against the secret delay is ordered by cost: if it holds for follower  $i$  it holds for every follower with a higher cost. The critical discount factor  $\delta_{\text{crit},i}$  is non-increasing in  $c_i$ , and strictly decreasing wherever the secret delay is the binding constraint.*

*Proof.* That the first constraint in (32) does not involve  $i$  is immediate: the deviator earns  $\bar{m}(1-c_i)$  by taking the whole market at the collusive price, it forgoes  $\pi_i^{\text{col}} = \bar{m}s_f(1-c_i)$ , and it loses that same amount in every period of the war, so the margin  $1-c_i$  is common to the gain and to the threat and cancels. The public price deviation ranks nobody, because it sells at the collusive price and its prize and its stake are margins on the same price. For the second constraint, differentiate (31):

$$\frac{\partial R_i}{\partial c_i} = -\frac{\Delta_p}{(1-c_i)^2} < 0, \quad (33)$$

strictly, because  $\Delta_p > 0$  and  $c_i < 1 - \Delta_p < 1$ . A follower with a smaller  $R_i$  faces a weaker requirement in the second constraint of (32), and nothing else in that constraint depends on  $i$ , so if it holds for follower  $i$  it holds for every follower with a higher cost.

For the discount factor, the quantity  $D/[\delta(1-\delta^W)] = (1-\delta)/[\delta(1-\delta^W)] + \alpha_0$  is strictly decreasing in  $\delta$  by the computation in Step 4 of the proof of Proposition 1, and neither side of either constraint in (32) involves  $\delta$  elsewhere. Each constraint therefore holds on an upper interval of  $\delta$ . Write  $\delta^P$  for the infimum of the interval defined by the first constraint, which is common to all followers, and  $\delta^S(R_i)$  for the infimum of the interval defined by the second, which is increasing in  $R_i$ . Then  $\delta_{\text{crit},i} = \max\{\delta^P, \delta^S(R_i)\}$ . Both infima are taken for a fixed rule. The union over rules preserves the ordering by the argument in the proof of Proposition 2, including the adjustment of the cutoffs when the first constraint binds. Since  $R_i$  is strictly decreasing in  $c_i$ , this is non-increasing

in  $c_i$ , and strictly decreasing at every  $c_i$  at which  $\delta^S(R_i) > \delta^P$ , which is the case in which the secret delay is the binding constraint.  $\square$

Two remarks close the argument. The first is that the jump is what creates the cross-section. By (31) the spread in  $R_i$  across followers is proportional to  $\Delta_p$ , and at  $\Delta_p = 0$  every  $R_i$  equals one and the two constraints in (32) collapse to (IC-P') and (IC-S'), which is the symmetric model. Cost cannot rank the followers when there is nothing to delay: a price war pays every follower the same amount, which is nothing, so what a war costs a firm carries no information about its cost, and if the prize and the stake were margins on the same price the ranking would disappear as it does for the public deviation. Larger jumps raise  $R_i$  for nobody, since  $R_i$  falls with  $\Delta_p$ , but they spread the followers further apart, so they make the arrangement easier to sustain and its cross-section sharper at the same time.

The second remark is about what the proposition is not. It ranks the followers by the return to the deviation that Lemma 1 bundles with withdrawal from the Platform. It does not say that the most efficient follower deviates in equilibrium, because in equilibrium nobody deviates. It says that if the environment moves against the arrangement, through a fall in  $\delta$  or a fall in the excess that the monitoring technologies deliver, the second constraint in (32) fails first for the follower with the lowest cost. Together with Lemma 1 that is the ranking of the order in which members stop supplying volumes to the Platform.

## B The non-collusive value of the reports

Section 3.3 asserts two things about the value of the reports to a station that is not colluding: that it is positive, and that it did not fall when the investigation began. This appendix gives one account of why, in three steps. The cost of carrying inventory rises with the variance of a station's forecast of its own volume. Subscribing lowers that variance. And the instability that followed the investigation raised the part of that variance which subscribing removes. We present this as one illustration and not as the mechanism. Any account in which stations benchmark against their rivals delivers the same comparative static. Symbols introduced here are local to this appendix; in particular,  $k$  below is a cost constant and not the number of non-reporting followers of Proposition 2.

### B.1 Setup

There are  $n$  stations in a local market, indexed by  $j$ , and a period is the interval between two reports. Volume is (1),  $q_{jt} = s_{jt}m_t u_{jt}$ , with  $\log u_{jt} \sim N(-\sigma_u^2/2, \sigma_u^2)$  independent across stations and periods. Two things are added. The first is persistence in the market component, without which there would be nothing worth forecasting. Writing  $x_t$  for  $\log m_t$  in deviation from its mean,

$$x_t = \rho x_{t-1} + \eta_t, \quad \eta_t \sim N(0, \sigma_\eta^2), \quad \rho \in (0, 1), \quad (34)$$

with stationary variance  $\sigma_x^2 = \sigma_\eta^2/(1 - \rho^2)$ . The second decomposes the share:

$$\log s_{jt} = \log \bar{s}_j + \zeta_{jt}, \quad (35)$$

where  $\bar{s}_j$  is the share the station holds under the prevailing conduct and  $\zeta_{jt}$  has mean zero and variance  $\sigma_\zeta^2$ . Stations are spatially differentiated, and that is what gives  $\zeta_{jt}$  its content. Were they perfect substitutes any price difference would send the whole market to the cheapest station and no interior allocation would survive. Because they are not, a price difference moves some volume and not all of it, and how much it moves depends on where the stations are. A station that undercuts takes volume from its immediate neighbours, and a neighbourhood that prices below the rest of the city takes volume from the other neighbourhoods. Both kinds of movement are  $\zeta_{jt}$ , since both are internal to the market the Platform surveys.

What moves  $\sigma_\zeta^2$  is the conduct. While the arrangement governs the allocation the members post nearly the same price, there is nothing to shop on, and shares are the ones it assigns. When it fails they no longer post the same price: some go to the regulated floor and others stop above it, the cross-section spreads out, and volume moves. The investigation raised the risk of such periods, and that is the only thing in the model it changes.

**Assumption 1** (Reallocation is zero sum).  $\sum_{j=1}^n \zeta_{jt} = 0$  in every period, and  $(\zeta_{1t}, \dots, \zeta_{nt})$  is exchangeable, independent across periods, and independent of  $m_t$  and of the  $u_{jt}$ .

Assumption 1 is the adding-up constraint and not a behavioural restriction: shares sum to one, so volume moving between stations inside the market cannot change what the market sells, while volume leaving for a neighbouring market is a movement in  $m_t$  and is already carried by (34). Exchangeability is what makes the unweighted sum the right form of adding up, since in levels the constraint is  $\sum_j \bar{s}_j \exp(\zeta_{jt}) = 1$ .

**Assumption 2** (Ordering and pricing are separable). The delivery is scheduled before the period's retail margin is chosen, and the fuel is in the tank when it is chosen.

Assumption 2 is what makes the value derived here non-collusive in the sense the paper needs: the ordering cost enters profit additively and not the first-order condition for price, so subscribing changes the level of a station's profit and not its conduct. Within a period the station orders for the next period using everything it has observed, prices are set and the next period's shocks are realized, volumes are sold, and the Platform publishes.

## B.2 The cost of ordering before demand is known

Write  $\hat{q}_{j,t+1}$  for the station's expectation of its next-period volume when it orders and

$$v \equiv \text{Var}(\log q_{j,t+1} \mid \mathcal{I}_{jt}) \quad (36)$$

for the variance of its forecast, where  $\mathcal{I}_{jt}$  is its information. The station orders  $a$  litres before the period's demand is realized and pays for the error. Ordering too much ties up working capital and

tank capacity; ordering too little means an emergency delivery or a sale lost to a rival. Both cost more the further the order is from what the station turns out to sell, and we take that cost to be proportional to the squared proportional miss,

$$k \mathbb{E}[(\log a - \log q_{j,t+1})^2 \mid \mathcal{I}_{jt}], \quad k > 0. \quad (37)$$

**Lemma 3** (The cost of ordering rises with the variance of the forecast). *The optimal order satisfies  $\log a^* = \mathbb{E}[\log q_{j,t+1} \mid \mathcal{I}_{jt}]$ , and the minimized cost is  $k v$ .*

*Proof.* Adding and subtracting  $\mathbb{E}[\log q_{j,t+1} \mid \mathcal{I}_{jt}]$  inside (37) gives  $k[(\log a - \mathbb{E}[\log q_{j,t+1} \mid \mathcal{I}_{jt}])^2 + v]$ , which is minimized at the stated order and takes the value  $k v$  there.  $\square$

A station cannot order its way out of uncertainty. The best it can do is centre the order, and what it then pays is the variance of its forecast. Richer models say the same thing: under the newsvendor, the standard model in this literature, the station orders a critical fractile of its forecast distribution and its cost is also strictly increasing in  $v$ . We use (37) because it is linear in  $v$ , which makes the comparison in Section B.4 exact. The constant  $k$  scales with the size of the station and drops out of that comparison.

### B.3 What the station forecasts, and what the reports remove

Substituting (35) into (1), next period's log volume is

$$\log q_{j,t+1} = \underbrace{\log \bar{s}_j + \mathbb{E} \log m_t}_{\text{known}} + \underbrace{x_{t+1}}_{\text{forecastable}} + \underbrace{\zeta_{j,t+1} + \log u_{j,t+1}}_{\text{not forecastable by anyone}}, \quad (38)$$

so that  $v = Q + \sigma_\zeta^2 + \sigma_u^2$ , where  $Q \equiv \text{Var}(x_{t+1} \mid \mathcal{I}_{jt})$ . The last two terms are the station's own share movement and its own idiosyncratic shock in the period it is ordering for. Neither is forecastable by the station, by the Platform, or by anyone else. Only  $Q$  is reducible by information, and everything the reports do, they do through it.

To reduce  $Q$  the station needs to know  $x_t$ , since  $x_{t+1} = \rho x_t + \eta_{t+1}$ . Its own volume this period tells it

$$\log q_{jt} - \log \bar{s}_j - \mathbb{E} \log m_t + \frac{\sigma_u^2}{2} = x_t + \zeta_{jt} + \log u_{jt} + \frac{\sigma_u^2}{2}, \quad (39)$$

a signal of  $x_t$  with noise variance  $\sigma_u^2 + \sigma_\zeta^2$ . The station cannot separate the two noise terms. A weak period is a weak market or a rival taking its customers, and from inside one station these look the same. This is the same confounding that makes the reports useful to the cartel in Section 3.2; here it is a forecasting nuisance rather than a monitoring problem.

A subscriber sees every station's volume and can ask where the volume it lost went. If a rival two blocks away gained what it lost, the market did not shrink and no revision to the forecast of  $x_t$  is called for. If nobody gained, the market did shrink. Formally, it averages the signals (39) across stations: by Assumption 1 the reallocation terms sum to zero and vanish from the average

exactly, while the idiosyncratic terms average down. The subscriber's signal therefore has noise variance  $\sigma_u^2/n$ , which does not depend on  $\sigma_\zeta^2$  at all. This is what the Platform contributes. It does not give a station more draws of the same thermometer. It gives it a thermometer immune to reallocation, because volume that moves from one station to another leaves the total unchanged, and reallocation is the one component of the noise that the collapse of the arrangement inflates.<sup>34</sup>

Writing  $\sigma^2$  for the noise variance of whichever signal the station has, and taking the stationary distribution of  $x_t$  as its prior, one round of Bayesian updating gives

$$Q(\sigma^2) = \rho^2 \frac{\sigma_x^2 \sigma^2}{\sigma_x^2 + \sigma^2} + \sigma_\eta^2, \quad Q'(\sigma^2) = \rho^2 \frac{\sigma_x^4}{(\sigma_x^2 + \sigma^2)^2} > 0. \quad (40)$$

A noisier signal leaves the station less sure of the market. Equation (40) also says when this account is empty: at  $\rho = 0$  the signal is irrelevant and at  $\sigma_\eta^2 = 0$  there is nothing to learn, and  $Q$  does not depend on  $\sigma^2$  in either case. The reports are worth something only because local demand moves and is persistent, which is a testable statement about the autocorrelation of market volume.

#### B.4 The value of subscribing, and why it rises with instability

That subscribing helps at all needs no argument. The subscriber can ignore the reports and use its own sales, so its information contains the non-subscriber's and its forecast variance is weakly smaller. By Lemma 3 its inventory cost is weakly lower.

The question the paper needs answered is what happens to that advantage when the arrangement stops governing the allocation. Write  $G$  for the saving from subscribing, which by Lemma 3 is  $k$  times the forecast variance the reports remove. The two stations differ only in the noise of their signal, so by (38) the terms in  $v$  that no information can touch are common to both and cancel in the difference:

$$G = k(v_N - v_S) = k \left[ Q(\sigma_u^2 + \sigma_\zeta^2) - Q(\sigma_u^2/n) \right]. \quad (41)$$

**Proposition 4** (The reports are worth more when shares move).  *$G > 0$ , and  $G$  is strictly increasing in  $\sigma_\zeta^2$ . Writing  $\sigma_{\zeta,c}^2$  for the reallocation variance while the arrangement governs the allocation,  $\sigma_{\zeta,w}^2 > \sigma_{\zeta,c}^2$  for its value when the arrangement does not, and  $\alpha$  for the rate at which the second occurs, the expected saving  $(1 - \alpha)G(\sigma_{\zeta,c}^2) + \alpha G(\sigma_{\zeta,w}^2)$  is strictly increasing in  $\alpha$ .*

*Proof.*  $\sigma_u^2 + \sigma_\zeta^2 > \sigma_u^2/n$  and  $Q' > 0$  by (40), so  $G > 0$ . In (41) the second term does not involve  $\sigma_\zeta^2$  and the first is strictly increasing in it, again by  $Q' > 0$ , so  $G$  is strictly increasing in  $\sigma_\zeta^2$ . The

<sup>34</sup>We assume the Platform covers the market. With  $r < n$  reporting, the noise variance of the subscriber's signal is  $\sigma_u^2/r + \sigma_\zeta^2(n-r)/(r(n-1))$ , so reallocation is removed in proportion to coverage and exactly at  $r = n$ . The subscriber's signal then also becomes noisier when  $\sigma_\zeta^2$  rises, and Proposition 4 needs a condition.  $G$  remains positive, but it is increasing in  $\sigma_\zeta^2$  only if

$$\frac{n-r}{r(n-1)} Q' \left( \frac{\sigma_u^2}{r} + \frac{(n-r)\sigma_\zeta^2}{r(n-1)} \right) < Q'(\sigma_u^2 + \sigma_\zeta^2),$$

which holds when coverage is close to full. It can fail at low coverage, because  $Q'$  is decreasing in its argument and the subscriber's argument is the smaller one.

expected saving is then a convex combination of two values of  $G$  with the larger weight  $\alpha$  on the larger one.  $\square$

Section 3.3 argues that the investigation raised  $\alpha$ , through deviations rather than through any change in the trigger, and that  $\alpha$  therefore exceeds the on-path frequency  $\alpha_0$  of Section 3.2. Proposition 4 says that the reports were worth more to a station after the investigation than before it, which is what the revealed-preference argument requires.

## B.5 Evidence

Table 4: Volatility before and after the May 2006 investigation

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

Each cell is the mean across quarters of a market-quarter volatility measure, before and after the May 2006 search warrants. The post period begins in 2006Q3 because 2006Q2 straddles the warrants.

*Panels A and B:* cross-sectional dispersion, within market-quarter, of the residual from a regression of log station volume on station and city-quarter fixed effects, in log points. Sample: stations reporting to the Platform with a direct volume reading. Station fixed effects absorb permanent size differences, so entry and exit do not move the measure.

*Panel C:* within-quarter standard deviation of the weekly retail margin, in cents per litre, from the Regie survey.

Window: 22 quarters before and 10 after, 2001Q1–2008Q4. Kent surveys these markets through 2008Q4.

Differences are from a regression of the quarterly series on an indicator for the post-warrant period, with Newey-West standard errors at four lags. \*, \*\* and \*\*\* denote significance at the 10, 5 and 1 percent level.

Table 4 is the empirical counterpart of Proposition 4. It compares the quarters before the warrants with the quarters after, in the three prosecuted markets we observe over the full sample period. Panels A and B report the dispersion of station volumes within a market and quarter, net of station and market-quarter fixed effects. Those two effects make the measure the analogue of  $\sigma_{\zeta}^2$ : the station effect removes  $\bar{s}_j$  and the market-quarter effect removes  $m_t$ , so the residual is reallocation between stations inside the same market and period. It rises in every market and on

Table 5: Effect of the Investigation on Margins and Price Adjustments

|                               | (1)<br>Margins (c/l) | (2)<br>Margin dispersion | (3)<br>Adjustment ratio |
|-------------------------------|----------------------|--------------------------|-------------------------|
| Cartel cities $\times$ Post   | -2.030**<br>(0.937)  | 2.793***<br>(0.911)      | -1.276**<br>(0.525)     |
| Other collusive $\times$ Post | -0.784*<br>(0.410)   | 0.708<br>(0.538)         | -1.151***<br>(0.355)    |
| Observations                  | 106                  | 106                      | 106                     |
| $R^2$                         | 0.940                | 0.900                    | 0.820                   |
| Period FE                     | Yes                  | Yes                      | Yes                     |
| City FE                       | Yes                  | Yes                      | Yes                     |

*Notes:* Each column estimates equation (42) with a different dependent variable. Margins are the median posted price minus the Régie price floor in cents per liter. Margin dispersion is the interdecile range (p90–p10) of weekly margins within a city-period, capturing the degree of price instability over time. The adjustment ratio is the ratio of median price increases to median price decreases, normalized by the corresponding ratio for the price floor. The estimation sample covers two periods: pre-investigation (January 2002 to May 2006) and post-investigation (May 2006 to March 2007). Robust standard errors clustered at the city level in parentheses. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ .

both measures, by 0.034 to 0.079 log points on the standard deviation and by 0.033 to 0.117 on the interquartile range. Panel C reports the within-quarter standard deviation of the weekly margin, which rises by 1.1 to 1.8 cents per litre, and the city-level difference in differences in Appendix C puts that increase at 2.79 cents relative to the competitive cities. The volume comparison in Panels A and B has no control group, since the survey covers only the prosecuted markets, so we read it as evidence on the direction of the change only.

## C The effect of the investigation on margins and price adjustments

This appendix reports the city-level replication of Clark and Houde (2014) summarized in Section 2.2.

To investigate the impact of the collapse of the cartels, we estimate the following difference-in-differences specification:

$$y_{ct} = \alpha_c + \delta_t + \beta_1 (\text{Cartel}_c \times \text{Post}_t) + \beta_2 (\text{OtherCollusive}_c \times \text{Post}_t) + \varepsilon_{ct}, \quad (42)$$

where  $y_{ct}$  is the outcome variable (margin, margin dispersion, or price adjustment ratio) for city  $c$  in period  $t$ ,  $\alpha_c$  is a city fixed effect,  $\delta_t$  is a period fixed effect,  $\text{Cartel}_c$  is an indicator for the four prosecuted cities, and  $\text{OtherCollusive}_c$  is an indicator for the remaining cities in the collusive cluster.  $\text{Post}_t$  equals one for the post-investigation period (May 2006 to March 2007). The estimation sample consists of two periods per city (pre-investigation and post-investigation), yielding 106 observations. Standard errors are clustered at the city level.

Table 5 reports the results. Column (1) shows that the investigation caused a significant decline in median retail margins. In the cartel cities, margins fell by 2.03 cents per liter ( $p < 0.05$ ) relative to competitive control markets, which is comparable to the 2.27 cent decline estimated at the station level in Clark and Houde (2014). The other potentially collusive cities also experienced a significant margin decline of 0.78 cents per liter ( $p < 0.10$ ), smaller than in the cartel cities but indicating that the disruption extended beyond the directly prosecuted markets. This could be because some of the potentially collusive cities featured elevated levels of instability prior to the collapse, or simply because they did not respond to the collapse, or even because our classification was imperfect and some of the cities are not in fact collusive.

Column (2) examines margin instability within each city, measured as the interdecile range (p90–p10) of weekly margins over the period. Under the cartel, margins in the prosecuted cities were remarkably stable: the coordinated pricing mechanism produced predictable, high margins week after week. The investigation shattered this stability. Cartel cities exhibit a significant increase in the interdecile range of 2.79 cents per liter ( $p < 0.01$ ), reflecting the onset of recurrent price wars in which margins swing between near-zero (during competitive episodes) and temporary spikes (during failed attempts to restore coordination). This instability is not a transitory adjustment: the descriptive statistics show that elevated margin dispersion persists into the 2007–2012 period, indicating that the cartel cities did not settle into a new stable equilibrium but instead experienced sustained competitive turbulence. The other collusive-cluster cities also show an increase in margin dispersion (0.71), although the effect is not statistically significant ( $p > 0.10$ ), suggesting that the transition to instability was most pronounced in the directly prosecuted markets.

Column (3) estimates the effect on the price adjustment asymmetry ratio. Cartel cities show a decline of 1.28 ( $p < 0.05$ ), indicating a shift toward more symmetric price adjustments as the coordinated Edgeworth cycle pattern breaks down. The other potentially collusive cities exhibit a decline of similar magnitude (1.15,  $p < 0.01$ ), consistent with the same breakdown of asymmetric pass-through in those markets. In Clark and Houde (2013), the corresponding station-level estimates were  $-1.63$  for cartel cities and  $-0.97$  for cyclical markets. The broadly similar magnitudes confirm that the city-level aggregate data replicate the main findings from the earlier station-level analysis.

Taken together, these results establish that the 2006 investigation caused a discrete disruption to collusive pricing in the cartel cities, with spillover effects on other markets exhibiting similar pre-investigation pricing patterns. This disruption is the necessary precondition for the information-sharing results that follow. If the investigation did not alter competitive conduct, there would be no reason to expect changes in stations' sharing behavior.

## D Additional tables and figures

Table 6: 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:* Cells report the mean of  $Share_{it}$  by survey year within each market group. Cartel cities: Magog, Sherbrooke, Thetford Mines, and Victoriaville. Montreal comprises the six Montreal-area markets. Other competitive: the remaining competitive-control cities. Other potentially collusive: the potentially collusive cluster excluding Québec City. Station counts are for 2006.

Table 7: Pre-trend test and HonestDiD sensitivity analysis: cartel cities

|                               | (1)<br>Cartel cities  |
|-------------------------------|-----------------------|
| <i>Treatment effects:</i>     |                       |
| DiD estimate (baseline)       | -0.079***<br>( 0.016) |
| <i>Pre-trend diagnostics:</i> |                       |
| Pre-investigation trend coef. | -0.0042<br>( 0.0064)  |
| Pre-trend test (p-value)      | 0.934                 |
| <i>Breakdown values:</i>      |                       |
| Short-run (2007)              | 0.20                  |
| Medium-run (2008–2010)        | 0.75                  |
| Long-run (2011–2013)          | 0.00 <sup>†</sup>     |

*Notes:* Standard errors clustered at the market level in parentheses. Significance: \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ . <sup>†</sup>: breakdown reported even though standard 95% CI marginally contains zero (upper bound within 0.01 of zero).

*Pre-investigation trend coef.:* estimate of  $\gamma$  from  $COOP_{it} = \alpha_i + \delta_t + \gamma(\text{year}_t - 2002) \cdot D_i + \varepsilon_{it}$ , fit on pre-period observations ( $\text{year} \leq 2006$ ).  $\alpha_i$ : station fixed effects;  $\delta_t$ : year fixed effects;  $D_i$ : treatment indicator.

*Breakdown  $M$ :* smallest value of  $\bar{M}$  under the relative-magnitudes restriction of Rambachan and Roth (2023) for which the 95% robust confidence interval for the post-period effect first contains zero. Larger values indicate greater robustness to post-treatment violations of parallel trends.

Table 8: 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<br>( 0.022)   | 0.032<br>( 0.022)   | 0.031*<br>( 0.016)   | 0.016<br>( 0.013)   | 0.012<br>( 0.013)  | 0.017<br>( 0.012)  |
| 2008                                   | 0.067**<br>( 0.030) | 0.070**<br>( 0.029) | 0.048**<br>( 0.019)  | 0.041**<br>( 0.018) | 0.001<br>( 0.011)  | 0.001<br>( 0.009)  |
| 2009                                   | -0.001<br>( 0.015)  | 0.001<br>( 0.015)   | 0.003<br>( 0.012)    | 0.001<br>( 0.011)   | -0.004<br>( 0.010) | -0.001<br>( 0.010) |
| 2010                                   | 0.002<br>( 0.015)   | -0.008<br>( 0.006)  | -0.000<br>( 0.009)   | -0.001<br>( 0.009)  | 0.000<br>( 0.010)  | 0.004<br>( 0.010)  |
| 2011                                   | 0.005<br>( 0.015)   | 0.008<br>( 0.015)   | 0.014<br>( 0.013)    | 0.008<br>( 0.011)   | 0.006<br>( 0.011)  | 0.005<br>( 0.009)  |
| 2012                                   | 0.005<br>( 0.016)   | 0.020<br>( 0.020)   | -0.010*<br>( 0.005)  | -0.005<br>( 0.007)  | -0.003<br>( 0.009) | 0.005<br>( 0.010)  |
| 2013                                   | 0.018<br>( 0.021)   | 0.019<br>( 0.021)   | 0.008<br>( 0.012)    | 0.006<br>( 0.012)   | -0.004<br>( 0.009) | -0.001<br>( 0.009) |
| Pre-2007 coefficients = 0 ( <i>p</i> ) | 0.543               | 0.726               | 0.393                | 0.335               | 0.589              | 0.858              |
| 2007–2008 = 0 ( <i>p</i> )             | 0.044               | 0.024               | 0.008                | 0.039               | 0.651              | 0.399              |
| 2007–2008 = 2009–2013 ( <i>p</i> )     | 0.033               | 0.025               | 0.005                | 0.024               | 0.386              | 0.441              |
| Coding of non-sharing                  | spec0               | spec1               | spec0                | spec1               | spec0              | spec1              |
| Observations                           | 15103               | 15248               | 15103                | 15248               | 15103              | 15248              |

*Notes:* Linear probability model for the stop-sharing hazard, with group by year interactions. The dependent variable equals one if a station that shared in the previous survey year does not share in the current year, and the sample conditions on sharing in the previous year. Because  $1 - \text{Share}_{it} = \text{Share}_{i,t-1} - \text{Share}_{it}$  on this subsample, the regression is a first-difference specification, and the market fixed effect absorbs a market-specific linear trend in the sharing level rather than a level difference. The pre-2007 test is therefore a test for non-linear pre-period movement, and is weaker than the pre-trend test on the levels reported in Section 4. Coefficients are measured relative to the pooled 2003–2006 pre-period, all of which precedes the May 2006 warrants because sharing is recorded in the first-quarter survey round. The omitted group is the competitive control, which pools Montreal with the other competitive cities, and Sherbrooke and Magog enter as a single market, as in the Bureau’s case. Columns 1, 3 and 5 use the baseline coding of non-sharing; columns 2, 4 and 6 use the strict coding, which counts explicit refusals only. The pre-2007 coefficients are estimated in an unreported specification that adds 2003, 2004 and 2006 deviations from the same pooled pre-period mean; only their joint test is reported here. Standard errors clustered by station in parentheses. \*, \*\*, and \*\*\* denote significance at the 10, 5, and 1 percent levels.

Table 9: Participation rates by potentially collusive market and year

| Market             | 2002 | 2003 | 2004 | 2005 | 2006 | 2007 | 2008 | 2009 | 2010 | 2011 | 2012 | 2013 | N    |
|--------------------|------|------|------|------|------|------|------|------|------|------|------|------|------|
| Alma               | 0.81 | 0.80 | 0.81 | 0.80 | 0.74 | 0.78 | 0.78 | 0.78 | 0.82 | 0.82 | 0.80 | 0.80 | 19   |
| Chicoutimi         | 0.84 | 0.90 | 0.91 | 0.89 | 0.88 | 0.84 | 0.81 | 0.79 | 0.82 | 0.82 | 0.83 | 0.82 | 86   |
| Drummondville      | 0.94 | 0.91 | 0.86 | 0.82 | 0.79 | 0.77 | 0.80 | 0.79 | 0.79 | 0.75 | 0.78 | 0.83 | 33   |
| Montmagny          | 0.86 | 0.86 | 0.86 | 0.86 | 0.86 | 0.86 | 0.71 | 0.71 | 0.71 | 0.71 | 0.71 | 0.83 | 7    |
| Rimouski           | 0.67 | 0.72 | 0.72 | 0.67 | 0.71 | 0.69 | 0.62 | 0.62 | 0.56 | 0.44 | 0.62 | 0.69 | 17   |
| Riviere-du-Loup    | 0.86 | 0.79 | 0.71 | 0.71 | 0.57 | 0.57 | 0.64 | 0.64 | 0.64 | 0.64 | 0.64 | 0.64 | 14   |
| Shawinigan         | 0.85 | 0.92 | 0.92 | 1.00 | 1.00 | 0.92 | 0.83 | 0.83 | 0.83 | 0.83 | 0.83 | 0.70 | 14   |
| Trois-Rivieres     | 0.96 | 0.96 | 0.96 | 0.96 | 0.96 | 0.92 | 0.80 | 0.84 | 0.93 | 0.93 | 0.95 | 0.95 | 49   |
| All eight markets  | 0.87 | 0.88 | 0.88 | 0.87 | 0.85 | 0.82 | 0.78 | 0.78 | 0.80 | 0.79 | 0.81 | 0.82 | 239  |
| Competitive cities | 0.89 | 0.91 | 0.92 | 0.92 | 0.92 | 0.92 | 0.92 | 0.91 | 0.91 | 0.91 | 0.91 | 0.91 | 1041 |

*Notes:* Cells report the mean of  $Share_{it}$  by survey year for each potentially collusive market outside Québec City. The last two rows pool the eight markets and report the competitive control group used in the event study of Figure 4. N is the number of stations observed in 2006. Grand-Mère is excluded throughout the analysis: it contributes a single station that never shares, so it carries no within-station variation. These raw series are reported in place of market-level event studies, which the market sizes cannot support.

Table 10: Québec City: participation by brand group

| Group                  | 2002 | 2004 | 2006 | 2008 | 2010 | 2012 | Stations (2006) |
|------------------------|------|------|------|------|------|------|-----------------|
| Discount chain 1       | 0.74 | 0.65 | 0.66 | 0.40 | 0.33 | 0.30 | 29              |
| Discount chain 2       | 0.85 | 0.80 | 0.76 | 0.62 | 0.53 | 0.58 | 17              |
| Major branded networks | 0.94 | 0.97 | 0.96 | 0.97 | 0.97 | 0.94 | 186             |
| All other stations     | 0.91 | 0.94 | 0.88 | 0.88 | 0.85 | 0.83 | 33              |

*Notes:* Cells report the mean of  $Share_{it}$  by survey year for stations in the Québec City market, by brand group. The two discount chains are unbranded distributor networks; major branded networks are the five vertically integrated or franchised brands. The last column reports the number of stations observed in 2006. Group names are withheld under the paper’s anonymity convention for the two discount chains.

Table 11: Crosswalk between Régie survey cities and Platform markets

| Platform market                    | Stations (2006) | Régie survey cities                                                                                                |
|------------------------------------|-----------------|--------------------------------------------------------------------------------------------------------------------|
| <i>Cartel</i>                      |                 |                                                                                                                    |
| Magog                              | 12              | Magog                                                                                                              |
| Sherbrooke                         | 62              | Sherbrooke                                                                                                         |
| Thetford Mines                     | 12              | Thetford Mines                                                                                                     |
| Victoriaville                      | 22              | Victoriaville                                                                                                      |
| <i>Québec City</i>                 |                 |                                                                                                                    |
| Québec City                        | 265             | Levis, Québec (Beauport), Québec (Charlesbourg), Québec (Québec), Québec (Sainte-Foy), Saint-Augustin-de-Desmaures |
| <i>Other potentially collusive</i> |                 |                                                                                                                    |
| Alma                               | 19              | Alma                                                                                                               |
| Chicoutimi                         | 86              | Saguenay (Chicoutimi), Saguenay (La Baie)                                                                          |
| Drummondville                      | 33              | Drummondville                                                                                                      |
| Grand-Mère                         | 1               | Shawinigan (Grand-Mère)                                                                                            |
| Montmagny                          | 7               | Montmagny                                                                                                          |
| Rimouski                           | 17              | Rimouski                                                                                                           |
| Riviere-du-Loup                    | 14              | Riviere-du-Loup                                                                                                    |
| Shawinigan                         | 14              | Shawinigan                                                                                                         |
| Trois-Rivieres                     | 49              | Trois-Rivieres                                                                                                     |
| <i>Montreal-area control</i>       |                 |                                                                                                                    |
| Laval                              | 92              | Laval                                                                                                              |
| Montreal (centre)                  | 236             | Mont-Royal, Montreal (Saint-Laurent), Montreal (Verdun), Montreal (centre)                                         |
| Montreal (east)                    | 118             | Montreal (Anjou), Montreal-Nord, Repentigny                                                                        |
| Montreal (north)                   | 90              | Blainville, Mirabel, Saint-Eustache, Terrebonne                                                                    |
| Montreal (south shore)             | 214             | Brossard, Chateauguay, Longueuil                                                                                   |
| Montreal (west)                    | 90              | Beaconsfield / Kirkland, Dorval / DDO, Vaudreuil-Dorion                                                            |
| <i>Other competitive control</i>   |                 |                                                                                                                    |
| Amos                               | 8               | Amos                                                                                                               |
| Berthierville                      | 8               | Berthierville                                                                                                      |
| Hull                               | 63              | Gatineau, Gatineau (Hull)                                                                                          |
| Joliette                           | 19              | Joliette                                                                                                           |
| Lachute                            | 10              | Lachute                                                                                                            |
| Rouyn-Noranda                      | 12              | Rouyn-Noranda                                                                                                      |
| Saint-Hyacinthe                    | 26              | Saint-Hyacinthe                                                                                                    |
| Saint-Jean                         | 16              | Saint-Jean-sur-Richelieu                                                                                           |
| Saint-Jerome                       | 28              | Saint-Jerome                                                                                                       |
| Val-d'Or                           | 11              | Val-d'Or                                                                                                           |

*Notes:* “City” denotes a Régie de l’énergie survey unit; “market” denotes a Platform survey unit. The classification of Section 2.2 is estimated on the Régie city panel: 53 cities with sufficient coverage, of which 4 are cartel cities, 18 potentially collusive, and 31 competitive (Table 1). Classified cities are assigned to Platform markets by name. The table lists the 30 matched markets, which form the sharing-analysis sample: 1,654 stations in 2006; Grand-Mère contributes a single station and is dropped as a singleton in estimation, which yields the 1,653 stations and 29 markets of Table 2. Several cities can map into one market, so the 18 potentially collusive cities correspond to Québec City plus nine other potentially collusive markets, with two cities (Maniwaki and Mont-Joli) unmatched. One of those nine, Grand-Mère, contains a single station and is excluded from the analysis of Section 4.3. Seventeen Platform markets have no classified Régie city and are excluded from the sharing analysis: Baie-Comeau, Chandler, Dolbeau, Forestville, Gaspé, Granby, Iberville, L’Assomption, Mascouche, Mistassini, Roberval, Sept-Îles, Sorel, Saint-Félicien, Saint-Georges-de-Beauce, Tracy, and Valleyfield. Four classified cities have no Platform market and enter the price analysis only: Maniwaki, Mont-Joli, Saint-Lin-Laurentides, and Ville-Marie.

Table 12: Summary statistics: quarterly volume panel, cartel cities

|                                         | Victoriaville |       | Thetford Mines |       | Sherbrooke |       |
|-----------------------------------------|---------------|-------|----------------|-------|------------|-------|
|                                         | Mean          | S.D.  | Mean           | S.D.  | Mean       | S.D.  |
| <i>Panel A: Sales volume</i>            |               |       |                |       |            |       |
| Daily gasoline volume (litres)          | 6,397         | 4,346 | 4,264          | 2,020 | 6,816      | 3,523 |
| Log daily volume                        | 8.54          | 0.69  | 8.25           | 0.46  | 8.66       | 0.66  |
| Volume reported (direct read)           | 0.80          | 0.40  | 0.86           | 0.35  | 0.80       | 0.40  |
| <i>Panel B: Station characteristics</i> |               |       |                |       |            |       |
| Gasoline pumps                          | 8.45          | 4.25  | 6.84           | 3.68  | 11.94      | 7.49  |
| Islands                                 | 2.03          | 0.88  | 1.95           | 1.01  | 2.51       | 1.15  |
| Convenience store                       | 0.80          | 0.40  | 0.50           | 0.50  | 0.68       | 0.47  |
| Car wash                                | 0.15          | 0.36  | 0.14           | 0.34  | 0.20       | 0.40  |
| Repair shop                             | 0.66          | 0.47  | 0.50           | 0.50  | 0.38       | 0.49  |
| Self-service                            | 0.35          | 0.48  | 0.50           | 0.50  | 0.60       | 0.49  |
| Stations in the same network            | 2.47          | 1.21  | 2.31           | 0.78  | 9.00       | 4.69  |
| <i>Panel C: Brand affiliation</i>       |               |       |                |       |            |       |
| Low-price guarantee chain               | 0.09          | 0.29  | 0.24           | 0.43  | 0.24       | 0.43  |
| Other major brand                       | 0.37          | 0.48  | 0.58           | 0.49  | 0.52       | 0.50  |
| Unbranded distributor                   | 0.27          | 0.44  | 0.02           | 0.12  | 0.18       | 0.38  |
| Retail chain                            | 0.04          | 0.21  | 0.08           | 0.27  | 0.05       | 0.21  |
| Independent and other                   | 0.22          | 0.42  | 0.08           | 0.27  | 0.01       | 0.12  |
| <i>Panel D: Local market structure</i>  |               |       |                |       |            |       |
| Rivals within 500 m                     | 1.18          | 1.09  | 0.88           | 0.68  | 1.04       | 1.08  |
| Rivals within 1 km                      | 3.22          | 2.26  | 2.11           | 0.76  | 2.11       | 1.51  |
| Stations in the market                  | 22.99         | 2.44  | 12.46          | 0.50  | 61.09      | 3.93  |
| Stations                                | 27            |       | 14             |       | 74         |       |
| Quarters                                | 32            |       | 32             |       | 44         |       |
| Station-quarters                        | 727           |       | 398            |       | 2,677      |       |
| Coverage                                | 2001–2008     |       | 2001–2008      |       | 2001–2011  |       |

*Notes:* Station-quarter observations from the Platform's quarterly survey of the three cartel markets with coverage over the full sample. Magog is excluded because it enters the survey in 2005q4. Daily gasoline volume is the cumulative volume of the four gasoline grades recorded between two meter readings, divided by the number of days elapsed; diesel is excluded. Volume is reported only when the surveyor obtained a direct reading, and the third row of Panel A gives the fraction of station-quarters for which this is the case. Means for volume are taken over reported observations. Convenience store, car wash, repair shop, self-service and the brand indicators are dummies. Rival counts are the numbers of competing stations within the stated radius, measured from station coordinates.

Table 13: Summary Statistics: Volume Event Study Sample

|                                                                | Mean   | Std. Dev. | Min   | Max       |
|----------------------------------------------------------------|--------|-----------|-------|-----------|
| <i>Panel A: Dependent Variable and Station Characteristics</i> |        |           |       |           |
| Log Volume                                                     | 8.580  | 0.664     | 5.572 | 9.969     |
| Number of Pumps                                                | 10.217 | 6.986     | 2.000 | 32.000    |
| Number of Islands                                              | 2.292  | 1.077     | 1.000 | 6.000     |
| Convenience Store                                              | 1.434  | 1.258     | 0.000 | 3.000     |
| Car Wash                                                       | 0.167  | 0.373     | 0.000 | 1.000     |
| Repair Shop                                                    | 0.462  | 0.499     | 0.000 | 1.000     |
| Network size                                                   | 6.848  | 4.996     | 1.000 | 17.000    |
|                                                                | 500 m  |           | 1 km  |           |
|                                                                | Mean   | Std. Dev. | Mean  | Std. Dev. |
| <i>Panel B: Information Sharing Variables</i>                  |        |           |       |           |
| <b>1</b> (Neighbor Stopped Sharing)                            | 0.122  | 0.327     | 0.202 | 0.402     |
| <b>1</b> (Neighbor Non-Sharing)                                | 0.134  | 0.341     | 0.256 | 0.437     |
| <i>Panel C: Neighborhood Composition</i>                       |        |           |       |           |
| Number of Rivals                                               | 1.063  | 1.104     | 2.428 | 1.708     |
| Number of Same-Brand Stations                                  | 0.020  | 0.140     | 0.134 | 0.340     |
| Number of Distributor Brands                                   | 0.188  | 0.430     | 0.465 | 0.693     |
| Number of Major Brands                                         | 0.707  | 0.822     | 1.615 | 1.190     |

*Notes:* Sample includes cartel-market stations with valid volume data, excluding Magog and stations that stopped cooperating with the provider. Neighborhood variables are measured within 500m and 1 km radii. The sample contains 2747 station-quarter observations, 96 unique stations, and 43 quarters. “Neighbor Stopped Sharing” indicates the presence of at least one neighbor that refused to share volume data with the provider for at least one quarter after 2006. “Neighbor Non-Sharing” indicates the presence of neighboring stations that refused to share data in the current period.

Table 14: Pre-trend test and HonestDiD sensitivity analysis: potentially collusive markets

|                               | (1)<br>Québec City <sup>‡</sup> | (2)<br>Other PC<br>excl. QC | (3)<br>Other PC<br>excl. QC, Drum. & RdL |
|-------------------------------|---------------------------------|-----------------------------|------------------------------------------|
| <i>Treatment effects:</i>     |                                 |                             |                                          |
| DiD estimate (baseline)       | -0.037***<br>( 0.003)           | -0.052***<br>( 0.006)       | -0.050***<br>( 0.005)                    |
| <i>Pre-trend diagnostics:</i> |                                 |                             |                                          |
| Pre-investigation trend coef. | -0.0026<br>( 0.0029)            | -0.0083<br>( 0.0080)        | -0.0007<br>( 0.0052)                     |
| Pre-trend test (p-value)      | 0.063                           | 0.683                       | 0.896                                    |
| <i>Breakdown values:</i>      |                                 |                             |                                          |
| Short-run (2007)              | 0.10                            | 0.30                        | 0.55                                     |
| Medium-run (2008–2010)        | 0.30                            | 0.10                        | 0.45                                     |
| Long-run (2011–2013)          | 0.35                            | 0.00                        | 0.00 <sup>†</sup>                        |

*Notes:* Standard errors clustered at the market level in parentheses. Significance: \* p<0.10, \*\* p<0.05, \*\*\* p<0.01. †: breakdown reported even though standard 95% CI marginally contains zero (upper bound within 0.01 of zero). ‡: Québec City sample restricted to 2004 onward; see Chaves et al. (2025) on the documented transition from price wars to uniform pricing in 2003–2004.

*Pre-investigation trend coef.:* estimate of  $\gamma$  from  $\text{COOP}_{it} = \alpha_i + \delta_t + \gamma(\text{year}_t - 2002) \cdot D_i + \varepsilon_{it}$ , fit on pre-period observations ( $\text{year} \leq 2006$ ).  $\alpha_i$ : station fixed effects;  $\delta_t$ : year fixed effects;  $D_i$ : treatment indicator.

*Breakdown  $M$ :* smallest value of  $\bar{M}$  under the relative-magnitudes restriction of Rambachan and Roth (2023) for which the 95% robust confidence interval for the post-period effect first contains zero. Larger values indicate greater robustness to post-treatment violations of parallel trends.


Table 15: Pre-trend test and HonestDiD analysis under alternative information sharing classification (Specification 1): identified set under bounded violations of parallel trends

|                               | (1)<br>Cartel cities  | (2)<br>Québec City <sup>‡</sup> | (3)<br>Other PC<br>excl. QC | (4)<br>Other PC<br>excl. QC & Drum. | (5)<br>Other PC<br>excl. QC, Drum. & RdL |
|-------------------------------|-----------------------|---------------------------------|-----------------------------|-------------------------------------|------------------------------------------|
| <i>Treatment effects:</i>     |                       |                                 |                             |                                     |                                          |
| DiD estimate (baseline)       | -0.073***<br>( 0.019) | -0.041***<br>( 0.003)           | -0.037***<br>( 0.008)       | -0.037***<br>( 0.009)               | -0.032***<br>( 0.006)                    |
| <i>Pre-trend diagnostics:</i> |                       |                                 |                             |                                     |                                          |
| Pre-investigation trend coef. | -0.0049<br>( 0.0047)  | -0.0051<br>( 0.0029)            | -0.0077<br>( 0.0064)        | -0.0058<br>( 0.0065)                | -0.0022<br>( 0.0045)                     |
| Pre-trend test (p-value)      | 0.123                 | 0.020                           | 0.345                       | 0.317                               | 0.402                                    |
| <i>Breakdown values:</i>      |                       |                                 |                             |                                     |                                          |
| Short-run (2007)              | 0.25                  | 0.00                            | 0.00 <sup>†</sup>           | 0.00 <sup>†</sup>                   | 0.00 <sup>†</sup>                        |
| Medium-run (2008–2010)        | 0.60                  | 0.40                            | 0.15                        | 0.25                                | 0.40                                     |
| Long-run (2011–2013)          | 0.00                  | 0.55                            | 0.00                        | 0.00 <sup>†</sup>                   | 0.00 <sup>†</sup>                        |

*Notes:* Standard errors clustered at the market level in parentheses. Significance: \* p<0.10, \*\* p<0.05, \*\*\* p<0.01. †: breakdown reported even though standard 95% CI marginally contains zero (upper bound within 0.01 of zero). ‡: Québec City sample restricted to 2004 onward; see Chaves et al. (2025) on the documented transition from price wars to uniform pricing in 2003–2004.

*Pre-investigation trend coef.:* estimate of  $\gamma$  from  $\text{COOP}_{it} = \alpha_i + \delta_t + \gamma(\text{year}_t - 2002) \cdot D_i + \varepsilon_{it}$ , fit on pre-period observations ( $\text{year} \leq 2006$ ).  $\alpha_i$ : station fixed effects;  $\delta_t$ : year fixed effects;  $D_i$ : treatment indicator.

*Breakdown  $M$ :* smallest value of  $\bar{M}$  under the relative-magnitudes restriction of Rambachan and Roth (2023) for which the 95% robust confidence interval for the post-period effect first contains zero. Larger values indicate greater robustness to post-treatment violations of parallel trends.



Notes: Markets classified via k-medians clustering on pre-investigation (2002–2006) median margins, price adjustment asymmetry, and price stickiness. Montréal includes five sub-markets. Three Abitibi-Témiscamingue markets shown in inset. Boundaries: Statistics Canada, 2021 Census boundary files.

Figure 6: Québec retail gasoline markets by competitive classification

Table 16: Pre-trend test and HonestDiD analysis under alternative information sharing classification (Specification 2): identified set under bounded violations of parallel trends

|                               | (1)<br>Cartel cities  | (2)<br>Québec City <sup>‡</sup> | (3)<br>Other PC<br>excl. QC | (4)<br>Other PC<br>excl. QC & Drum. | (5)<br>Other PC<br>excl. QC, Drum. & RdL |
|-------------------------------|-----------------------|---------------------------------|-----------------------------|-------------------------------------|------------------------------------------|
| <i>Treatment effects:</i>     |                       |                                 |                             |                                     |                                          |
| DiD estimate (baseline)       | -0.079***<br>( 0.016) | -0.037***<br>( 0.003)           | -0.052***<br>( 0.006)       | -0.054***<br>( 0.007)               | -0.050***<br>( 0.005)                    |
| <i>Pre-trend diagnostics:</i> |                       |                                 |                             |                                     |                                          |
| Pre-investigation trend coef. | -0.0042<br>( 0.0064)  | -0.0026<br>( 0.0029)            | -0.0083<br>( 0.0080)        | -0.0058<br>( 0.0081)                | -0.0007<br>( 0.0052)                     |
| Pre-trend test (p-value)      | 0.934                 | 0.063                           | 0.683                       | 0.728                               | 0.896                                    |
| <i>Breakdown values:</i>      |                       |                                 |                             |                                     |                                          |
| Short-run (2007)              | 0.20                  | 0.10                            | 0.30                        | 0.25                                | 0.55                                     |
| Medium-run (2008–2010)        | 0.75                  | 0.30                            | 0.10                        | 0.20                                | 0.45                                     |
| Long-run (2011–2013)          | 0.00 <sup>†</sup>     | 0.35                            | 0.00                        | 0.00 <sup>†</sup>                   | 0.00 <sup>†</sup>                        |


Notes: Standard errors clustered at the market level in parentheses. Significance: \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ . <sup>†</sup>: breakdown reported even though standard 95% CI marginally contains zero (upper bound within 0.01 of zero). <sup>‡</sup>: Québec City sample restricted to 2004 onward; see Chaves et al. (2025) on the documented transition from price wars to uniform pricing in 2003–2004.

*Pre-investigation trend coef.*: estimate of  $\gamma$  from  $\text{COOP}_{it} = \alpha_i + \delta_t + \gamma(\text{year}_t - 2002) \cdot D_i + \varepsilon_{it}$ , fit on pre-period observations ( $\text{year} \leq 2006$ ).  $\alpha_i$ : station fixed effects;  $\delta_t$ : year fixed effects;  $D_i$ : treatment indicator.


*Breakdown  $M$* : smallest value of  $\bar{M}$  under the relative-magnitudes restriction of Rambachan and Roth (2023) for which the 95% robust confidence interval for the post-period effect first contains zero. Larger values indicate greater robustness to post-treatment violations of parallel trends.




(a) Québec City



(b) Sherbrooke



(c) Victoriaville



(d) Thetford Mines

Figure 7: Weekly retail margins and price-war episodes

*Notes:* Weekly retail margins (cents per liter) defined as the posted retail price minus the Régie price floor. Shaded regions identify price-war episodes using a two-step classification: seed periods of at least two consecutive weeks with margins  $\leq 2$  cents per liter, extended forward to include subsequent weeks with margins  $\leq 4$  cents per liter.




Figure 8: Raw participation rates by market group

Notes: Mean of  $Share_{it}$  by survey year for the five market groups. The dashed vertical line marks the May 2006 search warrants; the survey is a first-quarter snapshot, so 2007 is the first post-investigation observation.




Figure 9: Effect of the cartel collapse on information sharing: Drummondville

Notes: Each series plots the estimated coefficients  $\hat{\beta}_\tau$  from equation (5), estimated separately with market fixed effects 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.




Figure 10: Effect of the cartel collapse on information sharing under alternative information sharing classification (Specification 1): cartel cities

*Notes:* Each series plots the estimated coefficients  $\hat{\beta}_\tau$  from equation (5), estimated separately with market fixed effects 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.




Figure 11: Effect of the cartel collapse on information sharing under alternative information sharing classification (Specification 2): cartel cities

*Notes:* Each series plots the estimated coefficients  $\hat{\beta}_\tau$  from equation (5), estimated separately with market fixed effects 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.

Table 17: Effect of Stop-Sharing Neighbors on Station Volume

|                                                                    | (1)<br>500 m        | (2)<br>1 km         |
|--------------------------------------------------------------------|---------------------|---------------------|
| 1(Neighbor stop sharing)                                           | 0.095<br>(0.093)    | 0.058<br>(0.076)    |
| 1(Neighbor stop sharing) $\times$ 1(2006Q3 $\leq$ t $\leq$ 2008Q4) | 0.226***<br>(0.070) | 0.181***<br>(0.066) |
| 1(Neighbor stop sharing) $\times$ 1(Year > 2008)                   | 0.152*<br>(0.087)   | 0.162<br>(0.120)    |
| Observations                                                       | 2,746               | 2,746               |
| $R^2$                                                              | 0.945               | 0.948               |
| Station FE                                                         | Yes                 | Yes                 |
| Brand FE                                                           | Yes                 | Yes                 |
| Market $\times$ Quarter FE                                         | Yes                 | Yes                 |

*Notes:* Dependent variable is log quarterly volume at always-sharing stations in cartel cities (excluding Magog). StopNbr equals one if at least one neighbor within the specified radius stopped sharing after 2006. The specification also includes NonShareNbr, which equals one if at least one neighbor within the radius does not share in quarter t, together with its two period interactions. Those coefficients are not reported: the timing of never-sharing status is not tied to the investigation, so they do not admit a causal reading, and including them means the stop-sharing coefficients are estimated relative to a non-sharing baseline. Additional controls include dummies for the number of competitors, same-brand stations, major-brand stations, and distributor-brand stations within the radius, as well as station amenity controls (convenience store type, car wash, repair, service type, islands, pumps with quadratic and cubic terms). Robust standard errors clustered at the station level in parentheses. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1.




Figure 12: Volume event study: effect of stop-sharing neighbors (1 km)

*Notes:* Coefficients from a regression of log quarterly volume on year interactions with  $\mathbf{1}(\text{Neighbor stop sharing within 1 km})$ , controlling for non-sharing neighbor year interactions, neighborhood composition, station amenities, and station, brand, and market $\times$ quarter fixed effects. The reference year is 2005. Standard errors clustered at the station level.




Figure 13: Price event study: effect of stop-sharing neighbors on the posted-price snapshot

*Notes:* Coefficients from the same specification as Figures 5 and 12, with the posted price of unleaded gasoline on the survey date (cents per litre) replacing log volume as the dependent variable: year interactions with  $\mathbf{1}(\text{Neighbor stop sharing})$ , controlling for non-sharing neighbor year interactions, neighborhood composition, station amenities, and station, brand, and market $\times$ quarter fixed effects. The reference year is 2005. Standard errors clustered at the station level.

## E Construction of the information-sharing variables

This appendix documents the construction of the station-year sharing indicator  $Share_{it}$  and the cross-sectional indicator  $StopShare_i$  introduced in Section 2.4. It lists the survey codes that record why a volume reading was not obtained, states the rules that map them into the indicators, and describes two alternative codings used in the sensitivity analysis. Specification 0 is the baseline used in the main text. Specification 1 replaces the inferential rules used to classify ambiguous survey responses with a strict literal-text rule. Specification 2 retains the baseline classification of  $Share_{it}$  but widens the post-period window used to construct  $StopShare_i$ .

### E.1 Survey responses

Table 18: Volume-status codes in the survey, by period

| Volume-status code     | 2002–2006 |       | 2007–2013 |       |
|------------------------|-----------|-------|-----------|-------|
|                        | N         | Share | N         | Share |
| BLANK (direct reading) | 8399      | 0.885 | 9998      | 0.854 |
| REFUSED                | 753       | 0.079 | 1242      | 0.106 |
| RESETS                 | 200       | 0.021 | 200       | 0.017 |
| SOME ACTUAL VOLUMES    | 92        | 0.010 | 139       | 0.012 |
| PUMP PROBLEMS          | 29        | 0.003 | 71        | 0.006 |
| PUMPS SEIZED           | 13        | 0.001 | 46        | 0.004 |
| R                      | 1         | 0.000 | 3         | 0.000 |
| SOME REFUSED           | 0         | 0.000 | 2         | 0.000 |
| POWER OUTAGE           | 0         | 0.000 | 1         | 0.000 |
| FIRST READS            | 0         | 0.000 | 1         | 0.000 |
| SOME ACTUAL VOLUME     | 2         | 0.000 | 0         | 0.000 |
| SITE JUST PICKED UP    | 3         | 0.000 | 0         | 0.000 |
| FIRST READ             | 1         | 0.000 | 0         | 0.000 |
| Total                  | 9493      |       | 11703     |       |

*Notes:* Station-year counts of the volume-status code recorded by the provider’s surveyors, pooled across all markets. A blank code indicates a valid direct meter reading. The share of station-years coded REFUSED rises from 7.9 percent before the investigation to 10.6 percent after.

For each station–year observation, the Platform records a status code (**st1cds**) and a free-text explanation string. The status code takes four values: blank (a valid volume reading was obtained), **A** and **E** (no reading was obtained; the file does not document the distinction between the two, and we treat them identically), and **T** (station temporarily closed, typically for renovation). A fifth value, **C**, marks a permanent closure, and these station–years are dropped before classification. The explanation field is populated mainly when the status is non-blank, although a small number of valid reads carry quality annotations (**M** or **S**); these annotations are stripped prior to classification and do not affect any indicator. The substantive explanation strings that drive classification fall into three groups:

- **Explicit refusal:** REFUSED, SOME (partial or selective refusal), and the one-letter shorthand R.
- **Non-strategic:** FIRST READ, FIRST READS, SITE JUST PICKED UP, POWER OUTAGE.
- **Ambiguous mechanical:** RESETS, PUMP PROBLEMS, PUMPS SEIZED.

The first group documents an explicit refusal to share volume data. We treat **SOME** as a refusal because the two codes are used by the same stations: of the 927 station-years at stations that carry a **SOME** flag in some year, 640, or 69 percent, belong to stations that also carry a **REFUSED** flag in another year. The second group identifies reasons for a missing record that are unrelated to the operator’s decision. The third group describes equipment failures. An isolated failure is plausibly genuine, but a mechanical excuse that repeats across annual survey visits, or that appears next to a documented refusal, is more likely to cover a decision to withhold. The rules below treat the third group accordingly.

## E.2 Construction of $Share_{it}$

Let  $NonShare_{it} = 1 - Share_{it}$ . A station-year is classified as  $NonShare_{it} = 1$  if the station can be plausibly inferred to have withheld its volume data, and  $NonShare_{it} = 0$  otherwise. The three specifications differ only in how the ambiguous mechanical codes are classified.

**Specification 0 (baseline).** A station-year is assigned  $NonShare_{it} = 1$  if either of the following holds.

1. *Explicit refusal.* The explanation contains **REFUSED** or **SOME**, or is the shorthand **R**, and no reading was obtained.
2. *Ambiguous mechanical with adjacency.* The explanation is one of **RESETS**, **PUMP PROBLEMS**, or **PUMPS SEIZED**, no reading was obtained, and at least one of the following adjacency conditions holds:
  - (a) *Two-year run rule.* The same station’s prior or next observation in the panel also has one of these three ambiguous codes.
  - (b) *Confirmed-refusal adjacency rule.* The same station’s prior or next observation in the panel is itself classified as  $NonShare = 1$  under rule (1) or rule (2a). Rule (2b) is evaluated after rule (2a), so that any reclassification from the two-year run rule propagates one step further.

All other station-years are assigned  $NonShare_{it} = 0$ . This includes valid reads, the four non-strategic codes, temporary closures (status **T**), and isolated single-year occurrences of the ambiguous mechanical codes.

The rationale for the adjacency logic is that a single-year **RESETS** (or related) entry, surrounded on both sides by valid reports, is most naturally read as a genuine equipment issue. Equipment

failures are unlikely to persist across annual survey visits, so a run of two or more such years is coded as non-sharing. An ambiguous year adjacent to a documented refusal is coded as non-sharing as well, which captures stations that alternate between an outright refusal and a mechanical excuse. Any ambiguous year that survives both rules is an isolated one surrounded by clean reads, and it is coded as sharing.

**Specification 1 (explicit refusals only).** A station–year is assigned  $NonShare_{it} = 1$  if the explanation contains REFUSED or SOME, or is the shorthand R, and no reading was obtained. All other station–years, including every occurrence of the ambiguous mechanical codes, are assigned  $NonShare_{it} = 0$ .

Specification 1 thus drops the inferential layer entirely and uses only the literal content of the explanation field. It is conservative in the sense that it counts only stations whose non-cooperation is explicitly documented in the survey response.

**Specification 2 (extended post window).** The  $NonShare_{it}$  classification is identical to specification 0. Specifications 0 and 2 differ only in the post-period window used to construct  $StopShare_i$  (described below).

Table 19: Mapping of survey responses to  $NonShare_{it}$  across specifications

| Survey response                                             | Spec 0 | Spec 1 | Spec 2 |
|-------------------------------------------------------------|--------|--------|--------|
| Valid read ( <code>st1cds</code> blank)                     | 0      | 0      | 0      |
| REFUSED / SOME / R, no reading obtained                     | 1      | 1      | 1      |
| FIRST READ / SITE JUST PICKED UP / POWER OUTAGE             | 0      | 0      | 0      |
| <code>st1cds</code> = T (temporarily closed)                | 0      | 0      | 0      |
| RESETS / PUMP PROBLEMS / PUMPS SEIZED, no reading obtained: |        |        |        |
| isolated single year                                        | 0      | 0      | 0      |
| two or more consecutive years                               | 1      | 0      | 1      |
| adjacent to a confirmed refusal                             | 1      | 0      | 1      |

### E.3 Construction of $StopShare_i$

The cross-sectional indicator  $StopShare_i$  identifies stations that cooperated in every pre-period observation but failed to cooperate in at least one post-period observation. The construction rule is identical across the three specifications:

$$StopShare_i = \mathbf{1} \left[ \overline{Share}_i^{\text{pre}} = 1 \wedge \overline{Share}_i^{\text{post}} < 1 \right],$$

where  $\overline{Share}_i^{\text{pre}}$  and  $\overline{Share}_i^{\text{post}}$  denote the mean of  $Share_{it}$  over the station’s pre-period and post-period observations, respectively. Two auxiliary rules apply in every specification: stations whose last observed year is before 2006 are assigned  $StopShare_i = 0$ , and stations for which either the pre- or post-period mean cannot be computed (because the station has no observations in that window)

are assigned  $StopShare_i = \cdot$  (missing). Under the baseline specification the indicator is defined for 2,028 stations and equals one for 86 of them, including 13 of the 121 stations in the cartel cities.

The pre-period is the same in all three specifications: every observation in year  $\leq 2006$ . The specifications differ in the post-period:

- **Specifications 0 and 1:** post-period = year  $\in \{2007, 2008, 2009\}$ .
- **Specification 2:** post-period = year  $> 2006$ , with no upper bound.

The three-year window used in specifications 0 and 1 corresponds to the period immediately following the May 2006 disclosure of the Bureau’s investigation, before the partial recovery of cooperation documented in Section 4.1. Specification 2 widens the window to include the medium and long run. A station that resumes cooperation in 2007–2009 but withdraws again in 2010 or later is therefore classified as  $StopShare_i = 1$  under specification 2 but as  $StopShare_i = 0$  under specifications 0 and 1. Table 20 summarizes.

Table 20: Pre- and post-period windows for  $StopShare_i$

| Specification                   | Pre-period       | Post-period            |
|---------------------------------|------------------|------------------------|
| Spec 0 (baseline)               | year $\leq 2006$ | $\{2007, 2008, 2009\}$ |
| Spec 1 (explicit refusals only) | year $\leq 2006$ | $\{2007, 2008, 2009\}$ |
| Spec 2 (extended post window)   | year $\leq 2006$ | year $> 2006$          |