Price Coordination with Asymmetric Information Sharing: Theory and Evidence*

David P. Byrne† Nicolas de Roos‡ Matthew S. Lewis§
Leslie M. Marx¶ Xiaosong Wu‖

October 19, 2024

Abstract

Platform-based information sharing among ostensibly competing firms presents a challenge for antitrust authorities, as evidenced by, for example, issues surrounding the *Informed Sources* antitrust case and its price-sharing platform for retail gasoline stations. Even when anticompetitive effects of information sharing are clear, effective remedies may not be. We provide a theoretical model that offers guidance for disrupting anticompetitive coordination facilitated through a price-sharing platform. Our analysis highlights the potential ineffectiveness of removing only one participant from the platform and points to benefits from the combination of ensuring that (i) at least two substantive market participants do not have ready access to each other's prices and (ii) the costs of price leadership are sufficiently high.

JEL Classification: D22, D43, D83, L13

Keywords: information sharing, price-sharing platform, collusion

*We thank Joe Harrington, Toshimasa Maruta, and seminar audiences at Duke University, 2024 MaCCI Annual Conference, and CRESSE 2024 for helpful feedback. Funding from the Australian Research Council (DP200103574 and DP210102321) is gratefully acknowledged. Disclosures: David P. Byrne and Leslie M. Marx were retained by the ACCC for the Informed Sources matter, while 7-Eleven retained Nicolas de Roos. Disclaimers: The interpretations of all results are those of the authors and do not necessarily represent those of the ACCC, Australian Competition Tribunal, Australian Government, or any parties involved with the Informed Sources matter. Any references to collusion in the paper are strictly in the economic and not legal sense.

†Department of Economics, University of Melbourne, byrned@unimelb.edu.au.

‡School of Management, University of Liverpool, nicderoos@gmail.com.

§John E. Walker Department of Economics, Clemson University, mslewis@clemson.edu.

¶Fuqua School of Business, Duke University, marx@duke.edu.

‖Department of Economics, University of Melbourne, andy.wu1@unimelb.edu.au.

1 Introduction

Information sharing and oligopolistic coordination is a longstanding issue in Industrial Organization, and it is a growing issue with the emergence of digital information-sharing platforms. Policymakers have raised concerns that platform-based information sharing facilitates price leadership and coordination to the detriment of consumers. Concerns have come up related to a variety of price-sharing platforms, including, for example, Informed Sources' PriceWatch service for retail gasoline,[1] Agri Stats' service for agricultural products,[2] and CoStar's STR service for hotels.[3] Antitrust economists and policymakers argue that competition is likely to soften when competing firms share contemporaneous price information.[4] Such concerns have reached new heights with the emergence of information-sharing platforms with near-immediate dissemination of price information among ostensibly competing firms.

Although authorities have raised concerns about such platforms, the best approach to remedies may not be obvious. In some cases, it may be feasible to ban information-sharing platforms,[5] restrict the nature of information sharing,[6] or limit platform membership.[7] Competition authorities could also potentially restrict a platform's ability to actually set prices or limit the ability of firms to set prices based on rival's prices (Brown and MacKay, 2022).

[1] Informed Sources' service provides "accurate, reliable, timely data" because "to make decisions with confidence, you need a complete view of the market" (https://informedsources.com/). Informed Sources and major gasoline retailers entered into a settlement with the Australian Competition and Consumer Commission (ACCC) ACCC (2014).

[2] Agri Stats, whose mission includes to "improve the bottom line profitability for our participants by proving accurate and timely comparative data" (https://www.agristats.com/history), is accused of facilitating cartels for slaughtering chickens, pigs, and turkeys (United States v. Agri Stats, Complaint, D MN, September 28, 2023).

[3] Hilton, Hyatt, Marriott Ritz-Carlton, InterContinental, Loews, and Accor are accused of "continuously sharing sensitive information about pricing, revenue, and occupancy" through CoStar's STR and a "give-to-get" policy and providing shared data that "have replaced uncertainty with comfort, reducing competitors' inclination to lower prices or engage in competitive actions" ("CoStar, Hotel Giants Accused Of Data-Driven Price-Fixing," *Law360*, February 21, 2024; Class Action Complaint in Portillo et al. v. CoStar et al., WD WA, February 20, 2024, para. 17).

[4] "Other things being equal, the sharing of information relating to price, output, costs, or strategic planning is more likely to raise competitive concern than the sharing of information relating to less competitively sensitive variables. Similarly, other things being equal, the sharing of information on current operating and future business plans is more likely to raise concerns than the sharing of historical information" (U.S. FTC, 2000, p. 15).

[5] In United States v. Cargill Meat Solutions Corp., the consent decree included an industry ban on information exchange (Civil Action No.: 1:22-cv-1821, Final Judgement, June 2, 2023). Under the proposed "Preventing the Algorithmic Facilitation of Rental Housing Cartels Act" (S.3692, 2024), it would be illegal in the United States for a rental property owner to subscribe to a software or data analytics service that performs a "coordinating function."

[6] In the *Airline Tariff Publishing Company* (ATPCO) matter, the settlement limited airlines from pre-announcing future airline fares for individual routes on ATPCO's platform before they became "live" in the market and available to consumers.

[7] The settlement in the *Informed Sources* matter excluded a major gasoline retailer from the platform.

Restrictions on the frequency of price changes are possible, but this could simplify the task of establishing focal points for price coordination (Byrne and de Roos, 2019). In other cases, it might make sense to require a platform to facilitate demand-side information sharing (Luco, 2019; Brown and MacKay, 2022). Proposed legislation in the United States aims to facilitate prosecution of antitrust violations when platforms are involved.[8]

Yet, despite these mounting concerns, we know little about how information sharing affects coordination and price leadership and the associated effects of remedies. Indeed, our models typically take the identity of price leaders and the timing of their play as exogenous (see, e.g., Miller et al., 2021), so these models cannot be used to understand how limiting information affects oligopolistic coordination. Longstanding research on information sharing in oligopoly focuses instead on detecting cheating and the scope for collusion,[9] but, in light of emerging digital price-sharing cases, this seems a second-order concern in the policy sphere relative to the coordination aspect of price-sharing platforms.

In this paper, we develop a model of oligopolistic coordination that allows for alternative information-sharing regimes and apply the model to evaluate the impacts of different remedies in the recent *Informed Sources* antitrust case in retail gasoline. In the model, each firm decides when to raise its price in a finite time interval. Firms face a trade-off: a first-mover disadvantage discourages early price increases, but each firm receives higher payoffs once all firms have raised their prices. The information-sharing structure in the model can be altered by adjusting the set of firms that can monitor and quickly respond to their rivals' price increase decisions perfectly. The framework, therefore, allows us to examine how different information-sharing structures affect firms' trade-offs in deciding when to raise prices and associated equilibrium outcomes.

The retail gasoline market that we study exhibits repeated price cycles and, within them, repeated coordination games with firms timing price increases.[10] Within our sample period, there is also an exogenous shock to information sharing among the firms associated with the *Informed Sources* antitrust case. These two features of the market allow us to test our theory of coordination and examine the impacts of alternative information-sharing structures using observed coordination outcomes.

The insights that we obtain, including the possibility of rapid coordination of price in-

[8]Under the "Preventing Algorithmic Collusion Act" (S.3686, 2024), there would be presumed "agreement" whenever companies share competitively sensitive information through a pricing algorithm to raise prices.

[9]There is a large related literature on this point including, for example, Luco (2019). See also Kuhn (2001) and Kuhn and Vives (1995).

[10]There is a substantive literature on Edgeworth price cycles (Edgeworth, 1925; Maskin and Tirole, 1988; de Roos, 2012, 2017) in retail gasoline with contributions by, among others, Castanias and Johnson (1993); de Roos and Katayama (2013); Eckert (2003); Bachmeier and Griffin (2003); Eckert and West (2004); Noel (2007a,b, 2008, 2009); Lewis and Noel (2011).

creases, potentially extend to any market in which firms interact repeatedly in an environment with common cost or demand shocks, frequent and observable price changes, and a high cost of price leadership.[11] In such cases, cost or demand increases can give rise to a coordination game in which firms would like to raise prices but do not want to be the first to do so. Opportunities for coordination may be particularly prevalent in markets with platform-enabled price sharing because such information sharing can enhance the observability of price changes and, if prices are made available to consumers, can reinforce the cost of price leadership.[12]

Our key results are as follows. First, with full information sharing, corresponding to the information-sharing environment that prevailed before the *Informed Sources* case, we show that a *domino effect* operates whereby all firms coordinate on price increases with minimal delay. Intuitively, equilibrium strategies cannot involve a long delay, or each firm would have an incentive to increase its price immediately to precipitate a cascade of rival price increases.

Our second result relates to the information-sharing environment after the settlement of the case, in which the antitrust authority (the ACCC) removed one participant from the platform. This remedy induced the platform to manually collect prices for that firm, creating a setting with full information sharing except for one blinded firm. In this case, our theory predicts price increases across firms with minimal delay will persist because the excluded firm can become the price leader, essentially knocking down the dominoes. Using high-frequency, station-level pricing data, we empirically confirm both of these results by studying the coordination outcomes surrounding the case's settlement.

For our third result, we consider a counterfactual intervention. Suppose that two firms are removed from the information sharing platform and face a bilateral *information blockage* whereby neither firm can readily observe the prices of the other. With two uninformed firms, neither firm can lead in increasing prices and know that the other firm will quickly observe this action. We show that, in the unique globally risk-dominant equilibrium, firms do not increase their prices if and only if the cost of price leadership is high. Thus, from the perspective of antitrust action, sufficient off-platform competition and costly price leadership are complements. Interventions that increase the cost of price leadership without facilitating information sharing can produce competitive benefits.

[11] These conditions suggest a wide range of possible applications, including the industries with price-sharing platforms already mentioned. Other possible applications include in banking—on the response of mortgage rates to changes in monetary policy and the transmission to household behavior, see, e.g., Kashyap and Stein (2000); Maggio et al. (2017).

[12] Industries with platform-enabled price sharing include the retail gasoline industry studied here, agricultural products (see footnote 2), hotels (see footnote 3), and products sold on digital marketplaces that offer price-sharing capability (see Musolff, 2024).

Related literature

While the value of price information sharing among explicitly colluding firms for monitoring purposes is well studied in the academic literature[13] and well recognized in antitrust practice,[14] the literature on the role of price information sharing for tacit coordination is more limited. Empirical studies of tacit coordination aided by price information sharing include Borenstein (2004), Byrne and de Roos (2019), Luco (2019), and Byrne et al. (2024a).

The literature typically assumes that firms are *symmetric* in their ability to share and monitor information, which means that it offers limited insights regarding the potential benefits of remedies that disrupt the symmetry of information sharing. Recent exceptions include the theoretical analysis of Brown and MacKay (2023), showing that *asymmetry* in pricing technology can contribute to higher equilibrium prices, and the empirical analysis of asymmetric information sharing in retail gasoline by Byrne et al. (2024b). In Byrne et al. (2024b), we also study the *Informed Sources* case settlement, and evaluate the effect of removing one firm from the platform on firms' prices, margins, and profit levels. The results in that study align with predictions from a wide range of models of oligopoly pricing under strategic complementarity and the idea that the case settlement endowed the excluded firm with short-run price commitment due to its informational disadvantage.

In contrast to this previous work, the current paper focuses on the *coordination problem* among firms in timing price increases. Through our analysis, we develop a theoretical model that allows for the consideration of counterfactual remedies involving asymmetric information sharing among firms to offer guidance for antitrust agencies in information-sharing cases with an eye to disrupting anticompetitive price coordination. Our further empirical examination of the *Informed Sources* case settlement motivates this new model and helps validate its predictions.

In Section 2, we provide background on the *Informed Sources* case that provides foundations for our theoretical model. In Section 3, we develop a model of asymmetric price information sharing and use it to analyze the actual interventions in the case in Section 4. In Section 5, we use the model to consider counterfactual interventions. In Section 6, we discuss policy implications for antitrust. Section 7 concludes the paper.

[13] See, e.g., Friedman (1971); Green and Porter (1984); Scherer and Ross (1990); Kuhn and Vives (1995); Ivaldi et al. (2007); Harrington and Skrzypacz (2011); Luco (2019); Asker et al. (2020); Kubitz and Woodward (2020)

[14] See, e.g., Vives (2007); U.S. FTC (2000); OECD (2011); U.S. FTC (2014); European Commission (2023); DOJ (2023)

2 Market context

We begin by providing background on pricing and information sharing in the retail gasoline market at the center of the *Informed Sources* case, the price-sharing platform, and the associated legal settlement. In developing this background, we draw on a unique dataset described in Byrne et al. (2024b), which consists of complete, daily station-level price data from Informed Sources for retailers in the market. The primary sample in the current paper covers July 1, 2014 to December 31, 2018. Over this period, there are no major shocks to the market structure other than the *Informed Sources* antitrust case.[15] See Byrne et al. (2024b) for a detailed description of the data. Here, we describe key contextual features of the market, data, and case to motivate our theory of oligopolistic coordination and information sharing.

2.1 Market structure

The case was alleged in Melbourne, Australia, a city of approximately 5 million people in 2014. The five leading gasoline retailers in the market are BP, Caltex, Coles, Woolworths, and 7-Eleven, and at the time of the case, they dominated the market. We illustrate their station numbers and shares in 2015 (i.e., before the case) in Table 1. BP and Caltex are vertically integrated distributors and importers of gasoline, Coles and Woolworths are dominant retailers in the grocery market, and 7-Eleven operates an extensive chain of convenience stores.

Table 1: Gasoline Retailer Station Counts and Shares in Melbourne

Retailer	**Stations**	
BP	127	(19%)
Caltex	92	(13%)
Coles	147	(22%)
Woolworths	93	(14%)
7-Eleven	148	(22%)
Other	75	(11%)
Total	682	(100%)

[15] We confirm the stability of all other aspects of the market structure from detailed annual industry monitoring reports from the ACCC. These are available at https://www.accc.gov.au/by-industry/petrol-and-fuel/fuel-and-petrol-monitoring.

Figure 1: Daily average retailer-level prices over the full sample

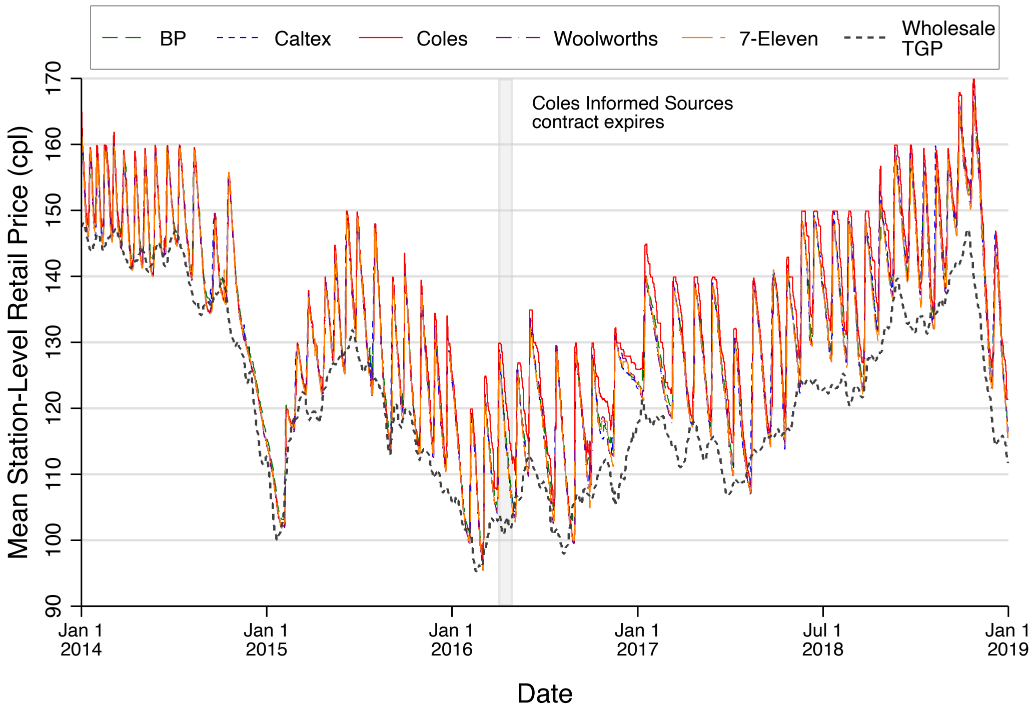

Notes: The figure shows for each major retailer the daily average of its station-level retail prices computed from daily, station-level prices from Informed Sources for January 1, 2014, to December 31, 2018. The wholesale Terminal Gate Price (TGP) corresponds to the daily wholesale price for unleaded 91 gasoline from Melbourne's local terminal gate from the Australian Institute of Petroleum.

2.2 Retail pricing

As with many gasoline markets, Melbourne has regular asymmetric cycles in retail prices.[16] We illustrate this in Figure 1, which plots daily average prices for each retailer over the primary sample period. The plot also contains the wholesale terminal gate price (TGP) for gasoline, which is the main time-varying component of the retailers' marginal cost. The TGP fluctuates with world oil prices, and Figure 1 shows retail price levels move with the TGP, but with higher frequency periodic jumps.

Figure 2 zooms in and plots typical price cycles from Figure 1 over two months in 2015.

[16] Retail price cycles have been extensively documented in the United States (Lewis, 2012), Europe (Foros and Steen, 2013; Assad et al., 2024), Canada (Noel, 2007b; Byrne et al., 2015), and Australia (Wang, 2009; de Roos and Katayama, 2013; Byrne et al., 2024a,b). Musolff (2024) uncovers high-frequency (within-day) price cycles across thousands of products for a range of product categories on Amazon. Cycles in all settings involve "restoration" and "undercutting" phases, as described in this paper. Further, Musolff (2024) shows that the prevalence of price cycles increased notably after September 2013, when Amazon introduced the MWS Subscription API, which allows re-pricers to subscribe to pricing change notifications.

Figure 2: Price cycle examples from 2015

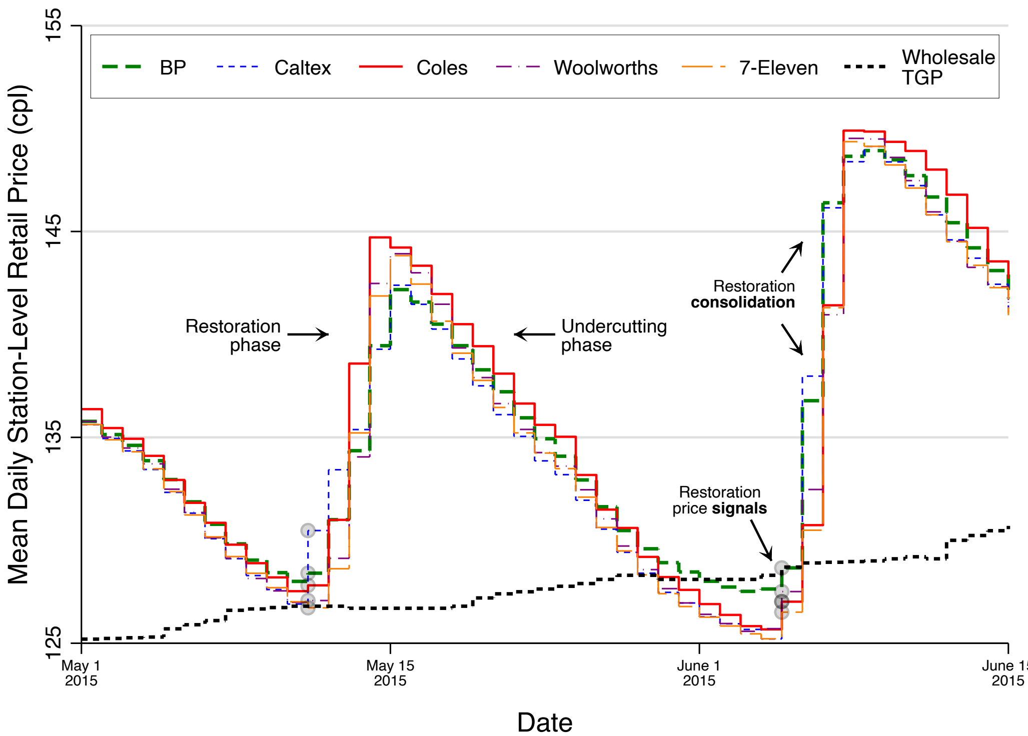

Notes: The figure shows for each major retailer the daily average of its station-level retail prices computed from daily, station-level prices from Informed Sources for September 15 to November 7, 2015. The wholesale Terminal Gate Price (TGP) corresponds to the daily wholesale price for unleaded 91 gasoline from Melbourne's local terminal gate from the Australian Institute of Petroleum.

The figure shows that prices jump rapidly in what we call the *restoration* phase and then gradually fall towards the wholesale price of petrol in the *undercutting* (or discounting) phase before a new cycle begins. Figures 1 and 2 illustrate how undercutting continues until retail prices approach daily wholesale TGPs, which triggers the start of a restoration phase. From this point, stations raise prices substantially, typically by way of a single price increase on the order of fifteen percent, as the retailers collectively restore their price-cost margins.

Figure 2 further shows that the price restoration phases involves a two-part *signal*-then-*consolidate* structure. In particular, the retailers implement complete price restorations at a small subset of stations ahead of large market-wide price jumps, which *signals* the start of a restoration phase and a target restoration price. These signals lead to small increases in average daily prices from the bottom of the cycle. After the signals are sent, retailers then *consolidate* restorations by rapidly implementing price restorations across their station networks, leading to large increases in their average daily prices.[17] In total, the two-part

[17]The stepwise increases in average prices during the restoration phase typically reflect the full restoration

signal-then-consolidate restoration process in Figure 2 occurs over approximately 5 to 7 days, as prices go from a local minimum at the bottom of the cycle to a local maximum at the top.

In what follows, we use *price leadership* to refer to rapid network-wide price increases that lead the market to consolidate restorations after price signals have been sent. The Australian Competition and Consumer Commission (ACCC), in its 2010 market study of retail gasoline, summarizes the pricing structure as follows

> *"During the discounting phase, retailers adjust prices on a site-by-site basis in response to local competition. During the restoration phase, retailers adjust prices across a broad number of sites at the same time. In the vast majority of cases, prices are increased to the same level at each of those sites across multiple local retail markets. The restoration phase involves a sharp increase in prices of around 10-17 cents per litre. In a given city, the price increase is initiated by one retailer (one of the refiner-marketers) and other retailers follow by increasing their prices to the same or very similar level."* (ACCC, 2010, paras. 74–75).

Notably, the quote highlights restoration consolidation in the last sentence, emphasizing how one of the major retailers tends to engage in price leadership in consolidating restorations.

Coordination problem with restorations

After price signals are sent, the retailers face a trade-off in determining whether to lead the market in consolidating restorations. Being a leader is costly in terms of revenue because of large station-level demand elasticities in retail gasoline markets. Houde (2012); Clark and Houde (2013); Wu et al. (2024) obtain estimates ranging from -10 to -30.[18] Ideally, a retailer would prefer to follow a rival in consolidating restorations. However, without at least one retailer taking the lead to initiate restoration consolidation, retailers can remain "stuck" at the bottom of the cycle, earning low profit margins. The faster a leader emerges to consolidate restorations, the faster all retailers are able to realize higher profits collectively.

of prices at an expanding set of stations rather than, for example, gradual increases in prices at all stations. In Byrne et al. (2024c), we examine the origins of the *signal*-then-*consolidate* pricing structure, which dates to 2012. Historically, and in our sample, all major retailers tend to send price signals ahead of market restorations and do so with different stations and in different orders over time. For the current study, what matters is simply that price signals are sent ahead of restoration consolidation, meaning that retailers are aware of candidate restoration prices as they solve the coordination problem of determining who leads in consolidating a restoration.

[18] Further evidence of the cost of price leadership is that retailers take steps to mitigate those costs, including within-brand randomization of leadership across stations (Wang, 2009; Byrne and de Roos, 2019), short-lived price "flares" to minimize temporal exposure and the use of geographically remote stations (Byrne et al., 2024a).

Within the larger dynamic game underlying price cycles, the retailers repeatedly face this trade-off while solving the coordination problem inherent to restoration consolidation cycle-by-cycle. Our model and analysis in Sections 3 and 4 focuses on this trade-off.

2.3 Role of Informed Sources

At the start of our sample period in 2014, the five largest gasoline retailers, BP, Caltex, Coles, Woolworths, and 7-Eleven, were subscribers to Informed Sources' price-sharing service (see https://informedsources.com/). Subscribers share high-frequency and disaggregated price data through a digital platform that allows them to: (1) provide station-level price data to the platform every 15 or 30 minutes; and (2) gain on-demand access to all price information from all other subscribers on the platform.

Beyond automated, digital price information sharing, Informed Sources uses a workforce of price spotters who drive around the market and manually collect non-subscriber station-level price data at daily, weekly, or other frequencies. As noted in historical ACCC market studies, this information sharing is central to the retailers' willingness and ability to lead restoration consolidation

> *"The quality and quantity of the information available to the major players also effectively reduces the risk encountered by retailers who seek to lead prices up in a market."* (ACCC, 2007, p. 135).

The ACCC again notes another key feature of the coordination problem with restoration consolidation: how information sharing among firms moderates the *risk* of price leadership with restorations. We will return to the interaction of information sharing and the risk of engaging price leadership in Section 5 below.

Informed Sources case

In August 2014, the ACCC brought an antitrust case against Informed Sources and the five major gasoline retailers.[19] In December 2015 (16 months later), the involved parties settled the case with a settlement containing two provisions: (1) Informed Sources would make its data available "on reasonable commercial terms" to third parties, including researchers and

[19] The ACCC alleged that Informed Sources and Australia's five major gasoline retailers violated section 45 of the Competition and Consumer Act of 2010, which makes illegal "contracts, arrangements or understandings that have the purpose, effect, or likely effect of substantially lessening competition" (ACCC, 2015a,b). The ACCC alleged that "the price information exchange service allowed those retailers to communicate with each other about their prices, and had the effect or likely effect of substantially lessening competition for the sale of petrol in Melbourne."

app developers; and (2) Coles would withdraw from the platform in April 2016 on the expiry of their contract.

In line with best practices by antitrust agencies internationally at the time of the case (OECD, 2011; European Commission, 2011; U.S. FTC, 2014), the settlement provisions were intended to make coordination on price increases more difficult and price leadership more costly. In particular, they were meant to increase the cost of price leadership by enhancing demand-side search behavior (provision 1)[20] and reducing the speed and accuracy with which firms can receive and respond to strategic price data (provision 2).[21] In particular, the expectation was that after the settlement, other retailers could not respond as effectively to attempts by Coles to signal or consolidate a restoration, and Coles could not respond as effectively to attempts by other retailers in restoration signaling and consolidation.

Altogether, the ACCC expected that the settlement would disrupt price restorations by exacerbating the coordination problem with restorations and making restorations riskier for retailers. This was expected to provide consumer benefits:

> *"The ACCC considers that even a few additional failed restorations in a year may have a significant impact on the average retail price for petrol over the year."* (ACCC, 2010, para. 89).

In addition, a relatively more drawn-out restoration phase would potentially benefit consumers because it would allow them to observe stations at both the top and bottom of the cycle and intertemporarily substitute towards low-price stations before the market fully consolidates a restoration. Making the price data publicly available per provision 1 was intended to enable such consumer search.

2.4 Effects of the settlement

However, despite the settlement provision requiring Informed Sources to make its data available on reasonable terms, no app developers subsequently used the Informed Sources data, either before or after the settlement.[22] Thus, the platform has only facilitated price information sharing on the supply side of the market.

[20] "Making [Informed Sources] pricing information available to consumers will allow consumers to make better-informed purchasing decisions and therefore create greater competition in petrol pricing" (ACCC, 2015b, p. 1).

[21] Chairman of the ACCC, Rod Sims, stated: "I welcome and appreciate the decision of Coles Express to cease using the Informed Sources information-sharing service at the earliest available opportunity .... The ACCC considers this to be an extremely positive step towards increasing competition in the petrol market ...." (ACCC, 2015a, p. 1).

[22] We document the history of app entry and exit in Byrne et al. (2024b) and offer various potential explanations for a lack of app entry after the *Informed Sources* case in Melbourne.

Figure 3: Manual price uploads before and after Coles' exit

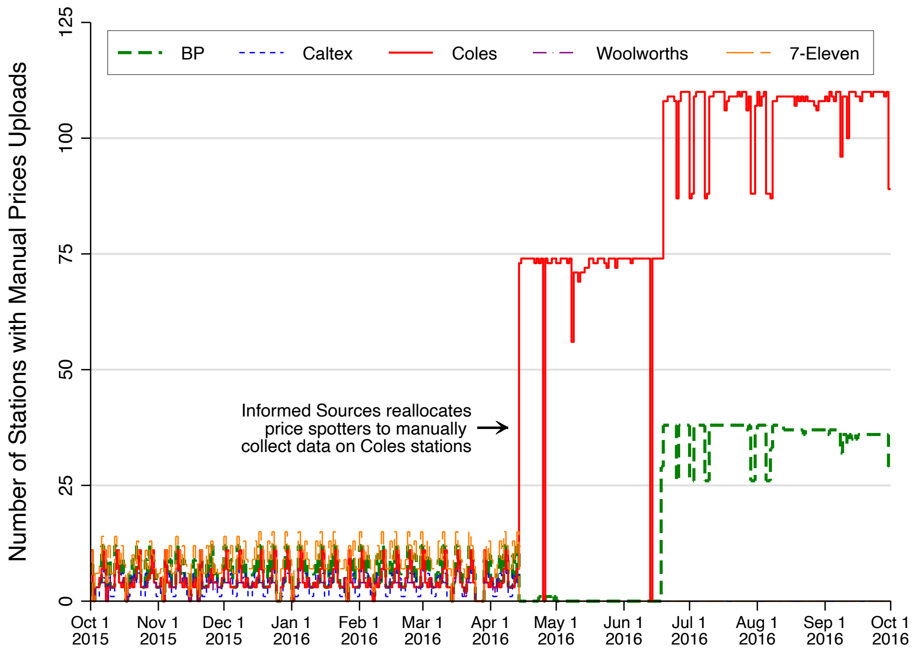

Notes: The figure shows for each major retailer the daily number of stations with manual price updates for the six months before and six months after Coles' exit from the platform. The red line represents Coles and the green line represents BP. The days in April and June with zero uploads are national holidays (ANZAC Day and Queen's Birthday).

Moreover, Informed Sources unexpectedly responded to the settlement by employing price spotters to manually collect Coles' prices daily following their exit the platform. We illustrate this unanticipated response in Figure 3, which plots the number of stations that Informed Sources manually collects data on.[23] Thus, while the Informed Sources settlement did have the effect of reducing Coles' ability to respond effectively to attempts by other retailers to signal or consolidate a restoration, it did *not* have the expected effect of reducing the ability of the remaining Informed Sources subscribers to observe Coles' prices.

Various factors support the conclusion that Coles faced an informational disadvantage following its exit from the Informed Sources platform. First, stations on the platform had access to station-level data from Informed Sources subscribers every 15 to 30 minutes, and they had access to the platform's manually collected data on Coles' stations. In contrast, to

[23]Byrne et al. (2024b) provide additional details on manual data collection by Informed Sources. By July 2016, Informed Sources was manually collecting data on 111 of 147 (76%) Coles stations strictly from the urban core of the market and not the remote outer suburbs. In mobilizing this manual data collection, Informed Sources shifts away from manually collecting data on all retailers before Coles exits the platform, as seen in Figure 3. Such manual data collection before April 2016 would have enabled data verification checks on its platform.

have comparable data, Coles would have had to manually collect data on all of their major rivals' stations at frequent intervals and quickly and reliably aggregate that information.[24] There is no evidence that Coles engaged in any systematic price collection.

Second, while manual collection would have been possible by observing the price boards of rival stations from the perspective of Coles locations, for the information to be valuable for decision-making, station managers would need to accurately input that information into a central database available to pricing managers at Coles.[25] In addition, the prohibition on "pay at the pump" in Australia means that station managers must actively service customer payments, increasing the opportunity cost of price spotting activity. Given the size of the Melbourne metropolitan area, the number of stations involved, and the at-times substantial price dispersion across stations within a major brand,[26] an approach based on station employee price spotting seems unlikely to yield information at a level comparable to that available through Informed Sources.

Lastly, empirical evidence in Byrne et al. (2024b) is consistent with Coles facing an informational disadvantage when it exited the Informed Sources platform. This includes asymmetry in the price adjustment frequency between Coles and its rivals, with Coles having a comparatively large decrease in its price adjustment frequency after its exit, and a reduction in the responsiveness of Coles' stations' prices to local rival stations' price adjustments.

Regarding the settlement's effect on price cycles, as shown in Figure 1, price restorations continued unabated. Thus, removing Coles from the Informed Sources platform did not have the desired effect of disrupting price coordination.[27] Motivated by this observation, in Section 3, we develop a model of asymmetric information sharing and coordination and use the model in Section 4 to explain the continuation of successful restoration coordination after Coles exits the platform. Then, in Section 5, we use the model to provide insights into the potential effectiveness of counterfactual interventions aimed at changing the nature of

[24]Other platforms were providing retail gasoline pricing, including the MotorMouth website and the PetrolSpy search app, which were established in July 2013 and September 2014, respectively, but their coverage was incomplete and their data often had multiple-day lags.

[25]Reliable manual collection can be challenging. According to "How Much is Poor Retail Pricing Data Costing Your Business?," available at: https://informedsources.com/how-much-is-poor-retail-pricing-data-costing-your-business/, field studies reveal nonnegligible differences between automated and manual price reports by local station managers for a given station.

[26]For example, the daily within-brand standard deviation of prices across stations regularly reaches 12 cents per liter (cpl), which is 10% of the sample average daily price of 124 cpl.

[27]The settlement also resulted in changes to overall margin levels during the undercutting phase, which we explore in Byrne et al. (2024b). There, we examine how Coles' strategic ignorance gives rise to short-lived price commitment, leading to an overall price increase. However, we do not examine how changes in information sharing affect the microdynamics of price coordination around restorations. As alluded to in the introduction, our previous study thus complements our analysis below, which abstracts from case effects on overall margin levels to focus on how asymmetric information sharing affects how firms coordinate price restorations.

information sharing among retailers to disrupt price coordination.

3 A model of restoration timing

In this section, we develop a model of price restorations, allowing some firms to be informed about their rivals' prices and some not. The model focuses on the coordination problem inherent to restoration consolidation described above and firms' dynamic trade-offs in determining when to consolidate restorations in light of restoration price signals that define restoration price and profit levels outside the model. To simplify our exposition, throughout, we use "restore prices" or "coordinate restorations" to mean "engage in restoration consolidation" per the two-part *signal*-then-*consolidate* restoration process described above. As defined above, a firm engages in price leadership if it engages in restoration consolidation before its rivals.

3.1 Setup

To examine the role of information flows in coordinating price restorations, we modify the technology diffusion model of Reinganum (1981). In this model, two firms choose when to adopt a new technology, and payoffs are such that there is a first-mover advantage. In our adaptation, there are n firms that choose when to restore prices, and there is a first-mover disadvantage.

In our model, each firm chooses a time in the interval $[0, 1]$ to restore its price. The unit interval is divided into T periods of length $\varepsilon = 1/T$ for a positive integer T that is large. The set of possible restoration times for each firm is $\mathcal{T} \equiv \{0, \varepsilon, 2\varepsilon, \ldots, 1\}$.

Starting at time 0, as long as no firm has restored, each firm earns profit of π_0 (*baseline* profit) per unit of time. Once a firm has restored its price, if at least one of its rivals has not, then the restoring firm earns profit of π_1 (*leader* profit) per unit of time. Following a restoration by at least one firm, any firms that have not yet restored earn profit of π_2 (*follower* profit) per unit of time. Once all firms have restored, each firm earns profit π_3 (*coordination* profit) per unit of time. Because the game ends at time 1, restoration at time 1 is equivalent to not restoring at all.

Our assumptions on the relative size of the profit flows reflect the retail gasoline market. A firm that restores before all its competitors loses price-sensitive customers to firms that delay their restoration. Thus, the leader profit flow is less than the baseline profit flow, and the follower profit flow is greater than the baseline profit flow: $\pi_1 < \pi_0 < \pi_2$. As mentioned above, reasons to expect that price leadership is costly include high station-level gasoline price elasticity and previous evidence of actions taken by retailers to mitigate leadership costs.[28]

However, once all firms have restored, each firm enjoys elevated profits at the restored price: $\pi_3 > \pi_2$. This condition is likely to hold in our gasoline market application primarily because splitting the market at the high post-restoration price is typically more profitable than attracting a larger share of the market at the much lower pre-restoration price.

Accordingly, we assume that:[29]

$$\underset{\text{leader}}{\pi_1} < \underset{\text{baseline}}{\pi_0} < \underset{\text{follower}}{\pi_2} < \underset{\text{coordination}}{\pi_3}. \tag{1}$$

This setup captures the problem that gasoline firms repeatedly face when trying to coordinate the timing of their price restorations at the beginning of each price cycle.[30] In our attempt to shed light on the impact of information sharing on oligopoly coordination, we focus our model on an individual coordination event and abstract away from dynamic considerations that may arise across coordination events. In practice, it is possible that early coordination could have additional (unmodeled) benefits.[31] For example, firms may anticipate that early coordination in the current restoration will increase the likelihood of swift coordination in future restorations.[32] While not directly modeled, this could offer an additional justification for our assumption that the profits of completing a restoration are greater than those of a follower remaining at the bottom of the cycle (i.e., $\pi_3 > \pi_2$).

Further, the market contains a fringe of small independent firms. As is common in retail petrol markets characterised by price cycles (Noel, 2007b; de Roos and Katayama, 2013), these firms consistently act as followers in the restoration phase. The presence of such firms provides an additional incentive to complete a restoration to speed up the possibility that

[28]See Wang (2009), Byrne and de Roos (2019), and Byrne et al. (2024a) for evidence of the retailers taking steps to mitigate restoration consolidation leadership costs.

[29]By way of contrast, Reinganum (1981) assumes in her Assumption 2 that $\pi_2 < \pi_3 < \pi_1$ and $\pi_2 < \pi_0 < \pi_1$. She further assumes in her Assumption 4 that the cost savings from non-adoption, $\pi_1 - \pi_0$, cannot continue forever, which rules out firms' postponing adoption forever.

[30]For simplicity, we assume that the leader and follower payoffs are the same regardless of the number of firms that have restored. This is not essential for our results. For intuition, note that as long as the follower payoff is less than the coordinated payoff, i.e., $\pi_2 < \pi_3$, then we have the result discussed below that the "dominoes fall" once only informed retailers remain unrestored, although the upper bound on the time of full restoration will be affected.

[31]We also implicitly assume that firms have a common understanding of π_0, π_1, π_2, and π_3. In practice, π_0 corresponds to near-zero profits at the bottom of the cycle. The restoration consolidation-related profits, π_1, π_2, and π_3, come from the firms' (unmodeled) restoration price signals before the restoration consolidation timing game begins, along with the cross-price elasticity of demand. We will return to demand elasticity in Section 5 below.

[32]Appendix B.2 empirically confirms that restoration consolidation speed across cycles exhibits non-neglible persistence. Fast restoration speed in the current cycle predicts fast restoration speed in the next cycle. Figure 4 below visually highlights this temporal persistence in restoration consolidation speed.

independents observe the completed restoration and then also restore their prices. As we illustrate in Appendix B.1, the independent retails reliably act as price followers, and only once all of the major firms have coordinated a restoration will these firms restore their prices.

Informed and uninformed firms

Consistent with a price-sharing platform being active in a market and an *Informed Sources* case history that pursued removing firms from it, we allow the possibility that some firms are informed about rival price restorations while others are uninformed. We refer to a firm as *informed* if at each time t it observes the history of restorations by the other firms (i.e., on the platform), including other informed and uninformed firms. Such a firm chooses whether to restore at any given time t as a function of the observed history of restorations up to that point. In contrast, we refer to a firm as *uninformed* if it does not observe the restoration choices of the other firms (i.e., as it would at high frequency on the platform) and must choose and commit to its restoration time independently of the other firms at time $t = 0$. There is thus an asymmetry in the model between informed and uninformed firms in terms of their ability to observe and respond to rivals' price adjustments (restorations) at high frequency, in line with our discussion of the effects of the *Informed Sources* case settlement in Section 2.4 above.

We now describe the firms' strategies, beginning with uninformed firms. Let $\mathcal{N}_U$ be the set of uninformed firms. A pure strategy for uninformed firm i is a restoration time $t_i \in \mathcal{T}$. A mixed strategy is a cdf $G_i \in \Delta(\mathcal{T})$, where $G_i(t)$ indicates the probability that a firm will restore at or before time t, with $G_i(1) = 1$, i.e., uninformed firms must restore at time 1 if they have not done so prior to that. The requirement of restoration at time 1 is without loss of generality because the game ends at time 1 and so there are no payoff implications of actions taken at time 1. Thus, we sometimes refer to restoration at time 1 as the choice of not restoring at all.

Let $\mathcal{N}_I$ be the set of informed firms. An informed firm can observe the history of restoration decisions of its rivals. The state at the beginning of time $t \in \mathcal{T}$, denoted by $\mathbf{h}_t \in \{0,1\}^n$, indicates the status of each firm i as not restored, $h_{t,i} = 0$, or restored, $h_{t,i} = 1$, prior to decisions made at time t, where $\mathbf{h}_0 \equiv (0,\ldots,0)$. Partial history $\mathbf{H}_t = (\mathbf{h}_0, \mathbf{h}_1, \ldots, \mathbf{h}_t)$ is feasible if for any $t' > t$, $\mathbf{h}_{t'} \geq \mathbf{h}_t$. Let $\mathcal{H}_t$ be the set of all feasible partial histories in period t and define $\mathcal{H} \equiv \{\mathbf{H}_t\}_{t\in\mathcal{T}}$. A strategy for informed firm i is a mapping from feasible partial histories to a binary restoration decision, $\sigma_i : \mathcal{H} \to \{0,1\}$, where $\sigma_i(\mathbf{H}_t) = 1$ if $h_{t,i} = 1$, i.e., restoration is irreversible,[33] and at time 1, we have $\sigma_i(\mathbf{H}_1) = 1$, i.e., informed firms must

[33]Although "failed restorations" are possible, they are extremely rare and do not occur in Melbourne in the time period shown in Figure 1.

restore at time 1 if they have not done so prior to that.

The payoff to firm i per unit of time while state $\mathbf{h} \in \{0,1\}^n$ prevails is

$$u_i(\mathbf{h}) \equiv \begin{cases} \pi_0 & \text{if } \max_{j\in\mathcal{N}} h_j = 0, \\ \pi_1 & \text{if } h_i = 1 \text{ and } \min_{j\in\mathcal{N}} h_j = 0, \\ \pi_2 & \text{if } h_i = 0 \text{ and } \max_{j\in\mathcal{N}} h_j = 1, \\ \pi_3 & \text{if } \min_{j\in\mathcal{N}} h_j = 1. \end{cases}$$

Given the partial history in period t, $\mathbf{H}_t$, the strategy profile $(\mathbf{G}, \boldsymbol{\sigma})$ implies a probability distribution over the sequence of continuation states $\mathbf{h}_{t+1}, \ldots, \mathbf{h}_T$ and thus an expected continuation payoff to each firm. Specifically, for firm i the expected continuation payoff from period t given partial history $\mathbf{H}_t$ is

$$\sum_{t'\in\{t,\ldots,1-\varepsilon\}} \sum_{\mathbf{h}\in\{0,1\}^n} \varepsilon u_i(\mathbf{h}) \Pr(\mathbf{h} \text{ at time } t' \mid \mathbf{H}_t, \mathbf{G}, \boldsymbol{\sigma}).$$

As indicated by this expression, although all firms must restore by time $t = 1$, the final time interval over which payoffs accrue is from time $1 - \varepsilon$ to 1.

3.2 Equilibrium

Given the game just defined, with its players, strategies, and payoffs, we now define a notion of equilibrium that accommodates different assumptions about the numbers of informed and uninformed firms.

For uninformed firm $i \in \mathcal{N}_U$, uninformed strategy G_i is a best response to $(\mathbf{G}_{-i}, \boldsymbol{\sigma})$ if G_i maximizes firm i's expected payoff (from the perspective of time 0) given $\mathbf{G}_{-i}$ and $\boldsymbol{\sigma}$. For $i \in \mathcal{N}_I$, informed strategy $\sigma_i : \mathcal{H} \to [0,1]$ is a best response to $(\mathbf{G}, \boldsymbol{\sigma}_{-i})$ at every partial history if for every $t \in \mathcal{T}$ and $\mathbf{H}_t \in \mathcal{H}_t$, σ_i maximizes firm i's expected continuation payoff from period t given $\mathbf{H}_t$, $\mathbf{G}$, and $\boldsymbol{\sigma}_{-i}$.

If we have mutual best responses, then we have a market equilibrium:

Definition 1. The strategy profile $(\mathbf{G}, \boldsymbol{\sigma})$, where $\mathbf{G} = \{G_i\}_{i\in\mathcal{N}_U}$ and $\boldsymbol{\sigma} = \{\sigma_i\}_{i\in\mathcal{N}_I}$, is a *market equilibrium* if:

(i) for each $i \in \mathcal{N}_U$, G_i is a best response to $(\mathbf{G}_{-i}, \boldsymbol{\sigma})$;

(ii) for each $i \in \mathcal{N}_I$, σ_i specifies a best response to $(\mathbf{G}, \boldsymbol{\sigma}_{-i})$ at every partial history.

3.3 Discussion

In this section, we discuss two questions that naturally arise from our model setup. First, how does the model's focus on a single restoration phase square with the reality that restoration phases occur repeatedly? Second, given our presumption that firms are able to ensure compliance with a target price through repeated interaction, could they not also rely on repeated interaction to ensure compliance with the timing of price increases?

On the first question, the firms in our sample repeatedly consolidate price restorations across cycles with highly visible prices on Informed Sources' platform and physical price boards. The repeated games literature (e.g., Friedman, 1971) emphasizes the value of such information to ensure compliance among firms with supra-competitive prices. Given price signaling ahead of restoration consolidation and a high degree of price observability, any short-lived benefits from deviating from rivals' prices would likely not be worth risking future restoration disruption, coordination breakdown, and cycle collapse.[34] Thus, by specifying a model focused on a single restoration phase in which firms make binary, irreversible restoration decisions, we effectively focus on the coordination problem that firms face within a larger dynamic game involving price cycles, presuming that the monitoring and compliance problem is resolved. Further, we presume that any minor miscoordination on restoration prices cycle-by-cycle can be rectified through a tâtonnement process of adjustment towards a common price among the major firms. In separate work (Byrne and de Roos, 2019; Byrne et al., 2024a), we find evidence of such price adjustments among the major firms.

Turning to the second question, as documented in Byrne and de Roos (2019) and Byrne et al. (2024a), signal-then-consolidate pricing can establish a target restoration price. Given a target and high-frequency price information, firms pricing below the target may be relatively straightforward to observe and verify and, therefore, deter. By contrast, it is not obvious how firms could enforce the *timing* of restorations. Any deviation that involves restoration delay may be difficult to immediately punish and deter without a mutual understanding among firms about the appropriate time window within which a firm should follow a rival's restoration. Therefore, in the game we study, firms cannot easily rely on a regime of punishments and rewards to enforce the timing of restorations.

[34] Byrne and de Roos (2019) find evidence of a price war between BP and Caltex and cycle collapse in 2008-09 in Perth, Australia, after a period of repeated and significant restoration delay by Caltex that leaves BP as an exposed price leader across multiple consecutive restoration phases. Following a price war period where BP abandons restorations altogether, the firms take steps to re-initiate the cycle.

Price signals ahead of restoration consolidation reduce uncertainty among the firms over their collective willingness to raise prices and what level to raise them to. In this way, the signals make it easier for firms to infer a rivals' intentions to steal business through abnormally delayed restoration consolidation, which also facilitates firms' ability to swiftly punish cheating, per Genesove and Mullin (2001).

4 Tipping dominoes: outcomes with zero or one uninformed firm

We are now in a position to examine equilibrium outcomes. We examine the cases of zero or one uninformed firm, corresponding to the market setting before and after the *Informed Sources* case settlement, respectively. In Section 5, we consider the counterfactual case of two uninformed firms.

As we show in this section, if all firms are informed, restoration by one firm creates a *domino effect* whereby all remaining firms restore shortly thereafter. For intuition, note that if there is only one remaining unrestored firm, then the uniquely optimal continuation strategy for that firm is to restore in the next period because $\pi_3 > \pi_2$. If there are two unrestored firms, then continuation strategies cannot result in too long of a delay before both restore. Otherwise, one would restore in the next period to induce the other to restore immediately after that. This domino result extends to any number of unrestored informed firms. It is then straightforward to extend the result to allow one of the firms to be uninformed, in which case the uninformed firm has the incentive to restore immediately to induce a cascade of restorations by the informed firms.

4.1 The domino effect in coordination

Consider the case of n firms, where $n-1$ are informed and allow for the remaining firm to be either informed or uninformed. We begin with the following lemma, which shows that restoration by the uninformed firm, if there is one, creates the domino effect just described.

To state the lemma below, we define $\tau(1) \equiv 1$ as the maximum number of periods of delay before full restoration when there is only one remaining unrestored informed firm. Then for $m \in \{2, 3, \dots\}$, we define recursively the maximum number of periods of delay before full restoration when there are m remaining unrestored informed firms

$$\tau(m) \equiv 1 + \frac{\pi_2 - \pi_1}{\pi_3 - \pi_2} \sum_{i=1}^{m-1} \tau(i),$$

which implies, for example, that $\tau(2) = 1 + \frac{\pi_2-\pi_1}{\pi_3-\pi_2}$, $\tau(3) = 1 + (1 + \tau(2)) \frac{\pi_2-\pi_1}{\pi_3-\pi_2}$, and $\tau(4) = 1 + (1 + \tau(2) + \tau(3)) \frac{\pi_2-\pi_1}{\pi_3-\pi_2}$. Defining $\overline{m}_\varepsilon(t)$ to be the largest integer such that

$$t + \varepsilon \sum_{i=1}^{\overline{m}_\varepsilon(t)} \tau(i) \leq 1, \tag{2}$$

we have the following result:

Lemma 1. If in period t^* all firms are restored other than $m \in \{1, \ldots, \overline{m}_\varepsilon(t^*)\}$ informed firms, then in any market equilibrium all firms restore at or before time $t^* + \varepsilon \sum_{i=1}^{m} \tau(i)$.

Proof. See Appendix A.1.

As Lemma 1 shows, there is a sense in which the coordination problem faced by a group of unrestored informed firms is more severe the larger the number of such firms. In particular, the upper bound on achieving full restoration increases as the number of firms involved in the coordination increases. With more firms, there are more ways in which other firms might be the ones to launch the restoration, and then more ways in which other firms might be the ones to move second, and so on. Nevertheless, as $\varepsilon \to 0$, if all the unrestored firms are informed, they all restore immediately.

We now use Lemma 1 to derive an upper bound on the time of having all firms restored:

Lemma 2. In any market equilibrium with $n \in \{1, \ldots, \overline{m}_\varepsilon(0)\}$ firms, if all n firms are informed, then all firms restore at or before time $\varepsilon \sum_{i=1}^{n} \tau(i)$, and if exactly one firm is uninformed, then all firms restore at or before time $\varepsilon \sum_{i=1}^{n} \tau(i) - \varepsilon$.

Proof. See Appendix A.2.

Although the upper bounds for the time of full restoration provided in Lemma 2 are not tight, it is notable that the bound is slightly lower when there is one uninformed than when there are no uninformed firms. In a sense, having one uninformed firm provides a focal firm for initiating the restoration process—the uninformed firm takes the lead in tipping the dominoes.

As long as $\varepsilon > 0$ is sufficiently small such that $n < \overline{m}_\varepsilon(0)$, it follows from inequality (2) that $\varepsilon \sum_{i=1}^{n} \tau(i) \leq 1$, and so Lemma 2 implies that with zero or one informed firm, in any market equilibrium all firms restore before time 1. Further, as ε tends towards zero, which corresponds to T tending to infinity, the upper bound on the time of all firms being restored also tends towards zero. This gives us the following result:

Proposition 1. In a market equilibrium with n firms with zero or one uninformed firm, the upper bound on the time at which all firms are restored goes to zero as ε goes to zero.

As shown in Proposition 1, for small ε, the presence of one uninformed firm does not meaningfully disrupt the ability of firms to restore prices. A domino effect operates among the informed firms, which prevents them from delaying too long before restoration. Further, because early restoration by the uninformed firm triggers restorations by the informed firms, an uninformed firm's inability to observe and respond to the actions of the other firms does not disrupt the price restoration dynamic nor the phase length with which the firms coordinate a restoration.

Our model provides an additional result that when there is an uninformed firm, in equilibrium, the informed firms essentially wait until after the uninformed firm has restored. We formalize this in the following proposition:

Proposition 2. With 1 uninformed firm and $n-1$ informed firms, in any pure-strategy market equilibrium, the maximum time by which the first informed firm's restoration can precede the uninformed firm's restoration is $\frac{\pi_3-\pi_1}{\pi_0-\pi_1}\varepsilon\sum_{i=1}^{n-2}\tau(i)$, which goes to zero as ε goes to zero.

Proof. See Appendix A.3.

Proposition 2 is informative about the underlying mechanism behind the irrelevance result of Proposition 1. It implies that when there is one uninformed firm, the burden of price leadership falls largely on the uninformed firm. Once the uninformed firm restores, the informed firms observe this and follow in quick succession. In the limit, as ε goes to zero, it cannot be the case that informed firms restore before the uninformed firm.

4.2 Understanding how case settlement affects coordination

We now bring Propositions 1 and 2 to the Informed Sources data. In particular, we study whether, after case settlement: (1) restoration phase length is unchanged; and (2) the uninformed firm (Coles) shifts toward leading restoration consolidation. In short, the evidence aligns with the theoretical predictions from the Propositions 1 and 2.

Restoration speed

We first examine the evolution of the speed of restoration consolidation over time. Throughout our sample period, price cycles are stable for all retailers with 30 to 40-day cycle lengths, as illustrated in Figure 1. As we show in Byrne et al. (2024b), given the completeness of our price data and stability of the price cycles, it is straightforward to empirically identify local minima (bottom of the cycle) and maxima (top of the cycle) in average retail prices across the entire market or for each retailer. As foreshadowed by our discussion of the example cycles in Figure 2, for a given cycle, we measure restoration *phase length* as the number of days it takes the market to transition from the bottom to the top of the cycle. The shorter the phase length, the faster restoration consolidation occurs.

Figure 4: Restoration phase length with Coles on and off the platform

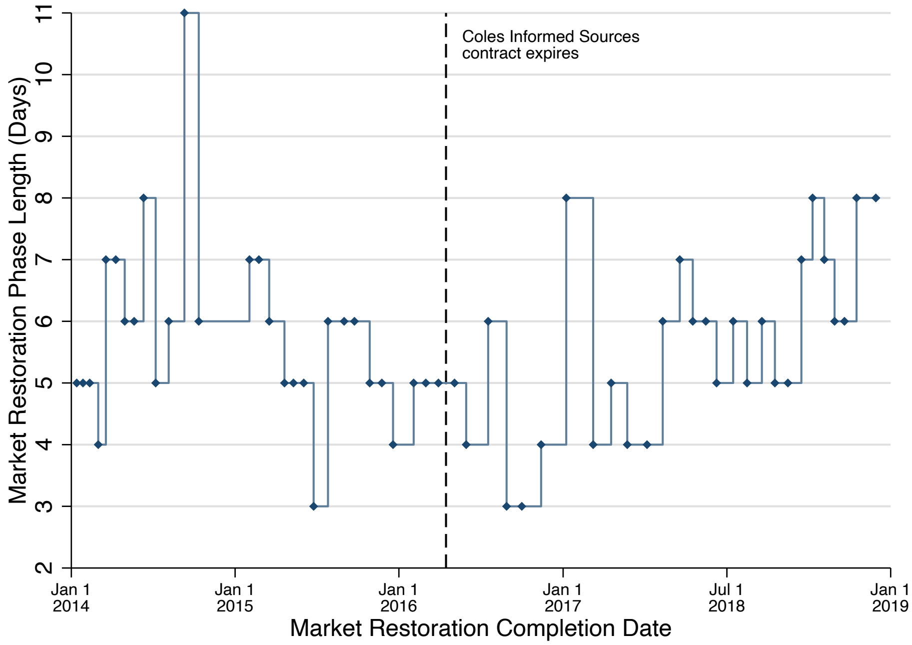

Figure 4 plots restoration phase lengths for each market cycle over our five-year sample period.[35] Restoration phases tend to vary between 4 and 8 days, with no discernible shifts or trends after Coles exits the platform. The mean restoration phase lengths are 5.57 days before and 5.60 days after, which is not a statistically significant difference (p-value $= 0.93$).[36] In sum, the rate of restoration coordination remains unchanged after Coles loses access to rival price information from the platform, consistent with the theoretical prediction from Proposition 1 of unchanged restoration speed with zero or one uninformed firm.

Restoration leadership

The key prediction from Proposition 2 is a shift by the uninformed firm (i.e., Coles) toward price leadership with restoration consolidation after it exits the platform. To understand how we can test this prediction empirically, we return to our example cycles from Figure 2. Rather than plotting daily average prices, we can plot *counts* of stations that restore their prices

[35]For each market-level cycle, we plot its restoration phase length and the date that a given cycle's restoration phase finishes, and prices are at a local maximum (i.e., the date when the "top of the cycle" is reached). See Definition 1 from Byrne et al. (2024b) for formally classifying dates where the cycle reaches a local maximum.

[36]This test of the difference in mean restoration phase length is based on $N_1 = 28$ market restoration phases before Coles exits and $N_2 = 30$ after Coles exits.

by station and date.[37] Figure 5 plots these station-level restoration counts for our example cycles. The plots highlight the two-part restoration process with a small number of station-level restorations (*signals*) precipitating waves of restorations (*consolidation*). Empirically, we seek to test whether the timing of the uninformed firm's (i.e., Coles) wave of restorations shifts from occurring with its rivals when on the platform to leading its rivals when off.

Figure 5: Station-level restoration counts for price cycle examples from 2015

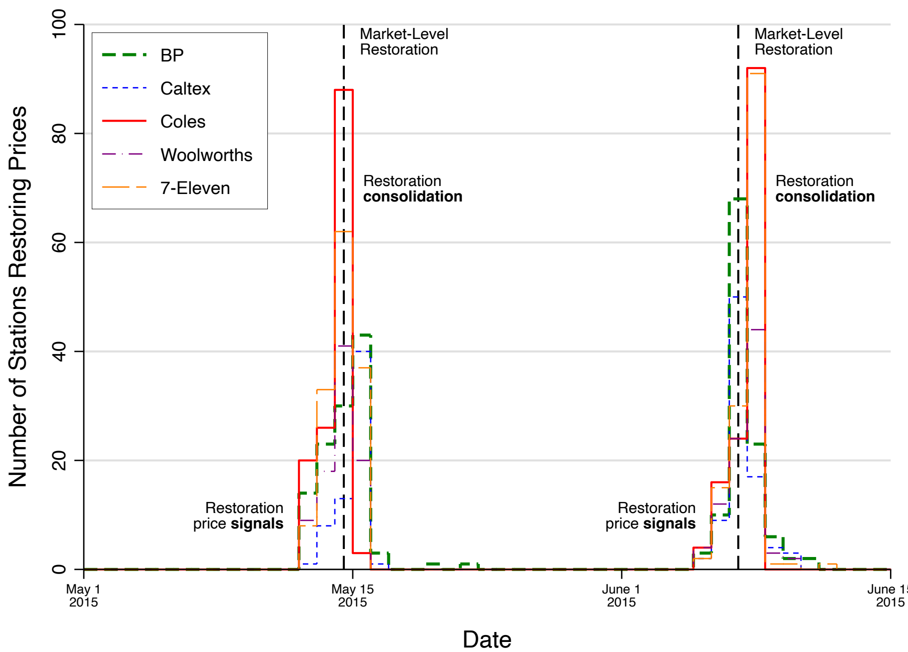

Building from Figure 5, we implement our test of a shift in restoration consolidation leadership using a linear-in-probability model of the following form

$$\text{Rest}_{it} = \alpha_0 + \sum_{j=-3}^{3} \left[\beta_j \times \text{MktRest}_{jt} + \gamma_j \times \text{ColesOff}_t \times \text{MktRest}_{jt}\right] + \epsilon_{ijt}, \tag{3}$$

where Rest_{it} equals 1 if station i restores its price on date t and 0 otherwise, MktRest_{jt} equals 1 if date t is j days away from a market restoration and 0 otherwise, and ColesOff_t equals 1 if t is after when Coles exits the Informed Sources platform (April 15, 2016) and 0 otherwise. Given a set of market restoration dates, we construct the MktRest_{jt} dummy variables in (3) based on dates that fall within a 7-day window of a market restoration date,

[37] Formally, we classify *station-level restorations* from Definition 1 in Byrne et al. (2024b) by identifying dates where a given station's price reaches a local maximum. We likewise classify *market-level restorations* by finding dates where average daily prices across all stations in the market reach a local maximum. The stability of price cycles at the station and market level allows us to use this simple yet robust restoration classification approach.

including the market restoration date itself.[38]

Interpreting the coefficients in (3), β_j quantifies a station's propensity to restore its prices within $j = -3, \ldots, 3$ days of a market restoration event when Coles is *on* the platform with symmetric, complete information sharing among all the firms. For $j < 0$, the coefficients measure the propensity for stations to *lead* market restorations. For $j = 0$ and $j > 0$, the coefficients respectively measure the propensity for stations to *go with* and *follow* market restorations. Likewise, the sum $\beta_j + \gamma_j$ quantifies a station's propensity to lead ($j < 0$), go with ($j = 0$), or follow ($j > 0$) market restorations after Coles is *off* the platform and Coles becomes asymmetrically 'uninformed' about rival prices compared to its rivals (BP, Caltex, Woolworths, 7-Eleven) who remain on the platform. Therefore, by estimating separate β_j and γ_j coefficients for Coles and its rivals, we can flexibly test whether there is a shift toward Coles price leadership during restoration phases after Coles exits the platform.[39] To this end, we modify (3) as follows

$$\begin{aligned}\text{Rest}_{it} = \alpha_0 + \alpha_1^{col} \times \text{Coles}_i + \sum_{j=-3}^{3} [\beta_j \times \text{MktRest}_{jt} + \beta_j^{col} \times \text{MktRest}_{jt} \times \text{Coles}_i \\ + \gamma_j \times \text{ColesOff}_t \times \text{MktRest}_{jt} + \gamma_j^{col} \times \text{ColesOff}_t \times \text{MktRest}_{jt} \times \text{Coles}_i] + \epsilon_{ijt}, \quad (4)\end{aligned}$$

where Coles_i equals 1 if station i is a Coles station. We estimate (4) using our sample of 891,294 daily station-level price observations between July 1, 2014 and December 31, 2018, spanning 645 unique stations among the major retailers, BP, Caltex, Coles, Woolworths, and 7-Eleven. We report standard errors clustered by station and date to account for cross-section and temporal persistence in station-level restorations.

Panels (a) and (b) of Figure 6 graphically present our restoration leadership findings for Coles and its rivals. Respectively, the lightly and darkly shaded bars in each panel report restoration probabilities (in terms of percentage points or 'pp') from (4) and their 95% confidence intervals when Coles is on and off the platform. The figures illustrate the *signal*-then-*consolidate* pricing structure with market restorations. With Coles on and off the platform, Coles' and its rivals' stations exhibit between 4 and 12 pp restoration probabilities two to three days before market restorations. These small probabilities correspond to the

[38] All of the market-level restoration windows in our sample consist of mutually exclusive dates. This feature of the data reflects the typical 20 to 30-day length of market-level price cycles, as mentioned above. We have checked extending the window around market restorations to +/− five days and found minimal restoration activity outside of the +/− three-day window. In other words, virtually all market restoration events are initiated and consolidated within 7-days.

[39] As in Byrne et al. (2024b), we pool price data for Coles' rivals in our analysis. Appendix B.3 presents results in the leader-follower setup before and after Coles exits the platform separately for each rival. These auxiliary results reveal similar results across retailers, which motivates our pooling of their data in estimating (4).

firms having a handful of stations restore their prices to signal an intent to restore prices and candidate price levels, as was illustrated in our example cycles in Figures 1 and 5 above.

With restoration consolidation, which our theoretical model focuses on, there is a large and asymmetric change in the timing of restorations after Coles exits the platform. In particular, Figure 6(a) shows that before Coles' exit, there is a 58 pp probability that a Coles station restores its price with the market restoration. However, after Coles exits the platform, this probability drops to just 22 pp. Additionally, Coles' stations shift from having a 19 pp restoration probability the day before market restorations when on the platform to a 41 pp restoration probability after exiting the platform. In sum, Coles shifts its primary restoration consolidation day to *leading* the market restoration by one day after it exits the platform.[40]

In contrast, Figure 6(b) reveals minimal change in Coles' rivals' restoration timing profile after Coles exits the platform. In particular, all the changes in the restoration probabilities are small in magnitude. Formally, we fail to reject the null H_0: $\gamma_{-3} = 0$ and $\ldots$ and $\gamma_3 = 0$ versus the alternative H_1: at least one of $\gamma_j \neq 0$ for $j \in \{-3, \ldots, 3\}$ with $p > 0.1$, confirming, statistically, that there is no change in Coles' rivals' restoration timing profile around market restorations. In sum, after Coles exits the platform, it shifts toward price leadership with restoration consolidation by one day, with its rivals following the following day. These results confirm the main prediction in Proposition 2 that the uninformed firm takes on a leadership role in coordinating restorations.[41]

Summary

The analysis in this section explains why asymmetric information-sharing failed to hinder price coordination after the *Informed Sources* case settlement. Before settlement, when all

[40]Formally, three tests from (4) confirm the changes in Coles' restoration timing profiles in Figure 6(a):

H_0: $\gamma_0^{col} = 0$ vs. H_1: $\gamma_0^{col} \neq 0$ rejects the null with $p < 0.01$ for Coles' stations, confirming a statistically significant change in Coles' restoration probability the *day of* market restorations.

H_0: $\gamma_{-1}^{col} = 0$ vs. H_1: $\gamma_{-1}^{col} \neq 0$ rejects the null with $p < 0.01$ for Coles' stations, confirming a statistically significant change in Coles' restoration probability the *day before* market restorations.

H_0: $\gamma_{-3}^{col} = 0$ and $\ldots$ and $\gamma_3^{col} = 0$ vs. H_1: at least one of $\gamma_{-j}^{col} \neq 0$ for $j \in \{-3, \ldots, 3\}$ rejects the null with $p < 0.01$, confirming a statistically significant change in Coles' stations' restoration timing profile around market restorations after Coles exits the platform.

[41]Appendix B.4 provides auxiliary results that examine temporal variation in station-level restoration shares around market restoration events underlying Figure 5. These results highlight a sharp, persistent shift toward restoration leadership unique to Coles' stations after Coles exits the Informed Sources platform. For instance, within our estimation sample, Coles has the largest share of stations leading a market restoration for just 1 of 11 (9%) market restorations while on the platform but does so for 9 of 16 (56%) market restorations when off the platform.

Figure 6: Station-Level Restoration Timing Before and After Coles Exits the Informed Sources Platform

(a) Coles

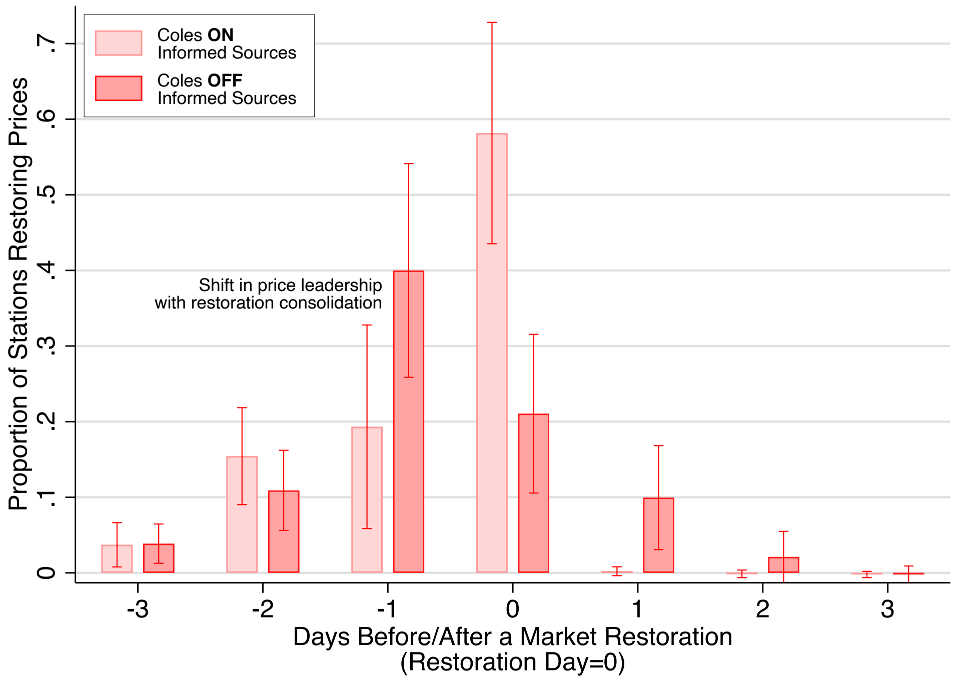

(b) Other major retailers

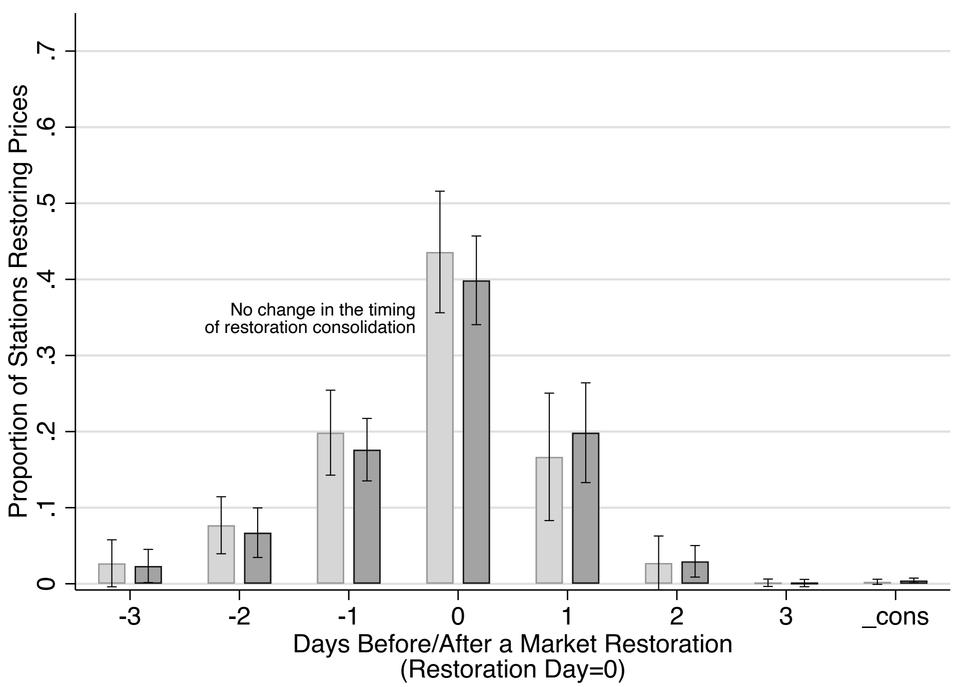

Notes: The figures contain estimates of the proportion of stations restoring prices on different days relative to the market-level restoration and their 95% confidence intervals derived from equation (3). Panels (a) and (b) respectively contain results based on Coles' and its rivals' (pooling BP, Caltex, Woolworths, 7-Eleven) stations.

major retailers subscribed to the platform and were informed firms in the language used here, all firms could be expected to restore "early" arbitrarily close to time 0 as ε goes to zero. After settlement, even if one of the firms was removed from the platform and so became uninformed, it remains the case that all firms could be expected to restore early in that same sense, with the uninformed firm becoming a price leader that "tips the dominoes." Thus, removing one firm from the platform, combined with the platform's response to the case settlement to actively collect data on the removed firm, does little to disrupt coordination or reduce average prices and these theoretical predictions were borne out empirically.[42]

An unfortunate implication of this theory and evidence is that the ACCC's imposed remedy to anticompetitive harms from price information sharing requiring one firm to exit the price-sharing platform was ineffective. Had it been anticipated that the platform would devote resources to manual observation of Coles' prices and had the theoretical results presented here been available at that time, then the ineffectiveness of the remedy might have been anticipated. This case history, however, begs the question of what interventions might be expected to meaningfully disrupt price restorations. One possibility is the exclusion of two firms from the price-sharing platform. We consider this possibility in the next section.

5 Creating informational blockage: outcomes with two uninformed firms

Armed with a theory of information sharing and price coordination, we are now in a position to consider counterfactual remedies. Given that the irrelevance result above relies on the uninformed firm leading restorations and on the informed firms observing that leadership—thereby initiating the fall of all the dominoes—what would happen if we removed *two* firms from the platform? As we now show, having two off-platform firms can have a potentially large effect by introducing a coordination problem between them. We show that if the cost of leadership is sufficiently high, it may be that neither of the uninformed firms wants to *risk* being the first to restore, thereby hindering coordination.

[42] In some sense, our theoretical and empirical results in this section rationalize Informed Sources' decision to actively collect data on Coles' stations after it exited the platform. Moreover, our findings from Byrne et al. (2024b), which highlight an overall increase in margin levels, further rationalize Informed Sources' decision. In both our analysis here and our prior work, Coles emerges as a natural price leader, which maintains coordination and leads to higher profit levels.

5.1 Preliminaries: a static game with two uninformed firms

The intuition above and the policy discourse in Section 2.3 from the ACCC surrounding the *Informed Sources* case suggest an important interaction between information sharing and price leadership risk. As we will demonstrate, risk concerns may play a role in equilibrium selection in the coordination game between two uninformed firms. It is instructive to begin by considering how those considerations play out in the case of just two uninformed firms and no informed firms.

Suppose we have just two uninformed firms with a single decision point (i.e., $\varepsilon = 1$). Then, the firms simultaneously decide to restore at time 0 or not restore, and the payoffs are as in Figure 7.

Figure 7: Illustration with two players and two actions

		Firm 2: do not restore	Firm 2: restore at time 0
Firm 1	do not restore	π_0, π_0	π_2, π_1
	restore at time 0	π_1, π_2	π_3, π_3

If both firms restore at time zero, each gets payoff π_3 for the whole period. If neither restores, each gets payoff π_0 for the whole period. If one restores and the other does not, the restorer gets payoff π_1, and the other gets payoff π_2.

This setting corresponds to a classic coordination game with two pure-strategy Nash equilibria: one in which both firms choose "restore at time 0," and one in which both firms choose "do not restore." Looking at our two pure-strategy Nash equilibria, both choosing "restore at time 0" is Pareto efficient (payoff dominant). However, if $\pi_0 - \pi_1 > \pi_3 - \pi_2$, then both choosing "do not restore" is the unique risk dominant equilibrium in the sense of Harsanyi-Selten.

The notion of *risk dominance* applies to symmetric 2×2 games with two symmetric pure-strategy Nash equilibria. For example, in the game in Figure 7, the Nash equilibrium (do not restore, do not restore) is risk dominant if "do not restore" is the unique best response for firm 1 to every mixed strategy for firm 2 that places weight at least 1/2 on "do not restore." This holds if and only if

$$\pi_0 - \pi_1 > \pi_3 - \pi_2, \tag{5}$$

which we refer to as the case with a *high cost of price leadership*. In this case, the profit loss from leading relative to the baseline is greater than that from not achieving full coordination relative to following. We refer to the opposite case, with $\pi_0 - \pi_1 > \pi_3 - \pi_2$, to be one with a *low cost of price leadership.*

With a high cost of price leadership, the price leadership payoff of π_1 is relatively small, and so restoring at time 0 is particularly risky—the player leading the restoration suffers a large loss if it is incorrect in its guess that the other firm will also restore at time 0. This case is salient in the case of retail gasoline, where, as mentioned, cross-price demand elasticities are large. In such a case, one might expect that in a setting that allows learning, it would be unlikely for play to converge to such a risky equilibrium with immediate restoration, even if the equilibrium is Pareto efficient (Fudenberg and Levine, 1998, pp. 20–21). The Harsanyi-Selten criterion of risk dominance captures this notion of risk (Harsanyi and Selten, 1988). The literature provides several arguments and foundations for why, in a case like this, one might expect coordination on the risk dominant equilibrium, typically involving evolutionary or stochastic dynamics. For example, the risk dominant equilibrium is the unique stochastically stable equilibrium in the sense of Young (1993).[43]

These preliminaries foreshadow our main theoretical findings from our general dynamic restoration game. Removing two firms from an information-sharing platform introduces a coordination problem over who leads restoration price increases between uninformed firms. Under risk dominance, if the cost of price leadership is sufficiently high (e.g., if cross-price demand elasticity is sufficiently large), the unique Nash Equilibrium is for neither firm to restore prices, which, through the lens of our model, implies coordination breakdown. We coin this remedy of limiting platform participation among firms to introduce coordination problems as *creating informational blockage*.

We now return to our dynamic coordination game and formally characterize conditions under which introducing two uninformed firms in a market with (other) informed firms leads to coordination breakdown. A reader looking to skip to our policy discussion in Section 6 can do so with the basic intuition from our preliminary analysis here.

5.2 Global risk dominance

Assuming a high cost of price leadership, the unique risk dominant equilibrium of the game in Figure 7 captures the notion that restoration exposes a retailer to the risk of receiving the low price leader payoff and provides a potentially relevant equilibrium selection criterion for our setup. However, risk dominance is only defined for two-player games where each player has two strategies. To address this, we consider the notion of *global risk dominance*

[43] In addition, the risk dominant equilibrium has the largest basin of attraction and is the unique outcome of iterated eliminated of dominated strategies in a global game setup in which the game is randomly drawn and there is noisy observation of the selected game (Carlsson and van Damme, 1993a,b; Morris and Shin, 1993). It is also the outcome of adaptive dynamics (Foster and Young, 1990; Blume, 1993) and stochastic evolutionary dynamics (Kandori et al., 1993; Kandori and Rob, 1995, 1998).

introduced by Maruta (1997) to extend risk dominance to general two-person strategic-form games:

Definition 2. (Maruta, 1997) Given a two-player strategic-form game, pure-strategy Nash equilibrium (a_1, a_2) is a *globally risk dominant Nash equilibrium* if: a_1 is the unique best response to every mixed strategy for player 2 that places weight at least $1/2$ on a_2, and a_2 is the unique best response to every mixed strategy for player 1 that places weight at least $1/2$ on a_1.

As noted in Maruta (1997), a globally risk dominant Nash equilibrium is unique whenever it exists.

In focusing on globally risk dominant equilibria, we follow a long tradition of using risk dominance and related notions as an equilibrium selection criterion. While risk dominance is typically used in two-by-two games, global risk dominance provides a generalization to any two-player strategic-form game. In that context, global risk dominance is motivated by its relation to the stochastic equilibrium selection process of Young (1993).[44]

To make this notion of equilibrium selection precise, define a game to be *weakly acyclic* if, from every strategy profile, there is a pure best response path that leads to a strict Nash equilibrium. Now, note that our game with two uninformed firms is weakly acyclic: given t_2, firm 1's best response is to also restore at time t_2, and (t_2, t_2) is a strict Nash equilibrium for any t_2, and symmetrically for firm 2. As shown by Maruta (1997), if a weakly acyclic two-person game has a globally risk-dominant equilibrium, then this is the one that is selected by the stochastic equilibrium selection process of Young (1993). In this stochastic process, the game is played repeatedly by players chosen randomly from a large finite population and who have access to a sample of information regarding past play. This process converges almost surely to a pure-strategy Nash equilibrium, and equilibria that have positive probability in the limit as the probability of mistakes goes to zero are known as *stochastically stable* equilibria. Thus, in our game with two uninformed firms and a high cost of price leadership, the unique stochastically stable equilibrium has no restoration by either firm.

5.3 Two uninformed and no informed firms

As a stepping stone to analyzing the information-sharing structure of direct interest for our counterfactual, we return to our dynamic coordination model and first consider a setup with just two uninformed firms.

[44] One could alternatively apply Ellison's (2000) notion of $\frac{1}{2}$-dominance, which also provides a generalization of risk dominance to any two-player strategic-form game. Ellison motivates $\frac{1}{2}$-dominance by its relation to the (Darwinian) stochastic equilibrium selection process of Kandori et al. (1993).

Denoting the restoration times of the first and second movers by t_1 and t_2, with $t_1 \leq t_2$, the payoffs to the leader (ℓ) and follower (f), respectively, are as follows:

$$\begin{aligned}\ell(t_1, t_2) &\equiv \pi_0 t_1 + \pi_1(t_2 - t_1) + \pi_3(1 - t_2) \\ &= t_1(\pi_0 - \pi_1) + t_2(\pi_1 - \pi_3) + \pi_3\end{aligned}$$

and

$$\begin{aligned}f(t_1, t_2) &\equiv \pi_0 t_1 + \pi_2(t_2 - t_1) + \pi_3(1 - t_2), \\ &= t_1(\pi_0 - \pi_2) + t_2(\pi_2 - \pi_3) + \pi_3.\end{aligned}$$

Given the profile (t_1, t_2), define the profits for firm 1 as

$$\Pi_1(t_1, t_2) = \begin{cases} \ell(t_1, t_2) & \text{if } t_1 \leq t_2, \\ f(t_1, t_2) & \text{if } t_1 > t_2, \end{cases}$$

and symmetrically for firm 2,

$$\Pi_2(t_1, t_2) = \begin{cases} \ell(t_1, t_2) & \text{if } t_1 \geq t_2, \\ f(t_1, t_2) & \text{if } t_1 < t_2. \end{cases}$$

A profile $\mathbf{t} = (t_1, t_2)$ constitutes a Nash equilibrium if, for each i and for all $t_i' \in \mathcal{T}$,

$$\Pi_i(\mathbf{t}) \geq \Pi_i(t_i', t_{-i}).$$

Taking firm 2's strategy to be specified according to cdf G_2 with (discrete) probability density $g_2(t)$ over $t \in \mathcal{T}$, by restoring at time $t_1 \in \mathcal{T}$, firm 1 earns expected payoff of

$$\Pi_1(t_1) \equiv \mathbb{E}_{t_2}[\Pi_1(t_1, t_2)] = \sum_{\tau \in \{0, \varepsilon, \ldots, 1-\varepsilon, 1\}} \Pi_1(t_1, \tau) g_2(\tau).$$

Analogously, taking firm 1's strategy to be specified according to the discrete probability density $g_1(t)$ over $t \in \mathcal{T}$, by restoring at time $t_2 \in \mathcal{T}$, firm 2 earns expected payoff of

$$\Pi_2(t_2) \equiv \mathbb{E}_{t_1}[\Pi_2(t_1, t_2)] = \sum_{\tau \in \{0, \varepsilon, \ldots, 1-\varepsilon, 1\}} \Pi_2(\tau, t_2) g_1(\tau).$$

To test for global risk dominance, we define for pdf g_i, time s, and $\lambda \in [0, 1]$, the mixed strategy that modifies g_i by placing weight λ on restoration at time s: $\hat{g}_i(t; \lambda, s) \equiv \lambda \cdot \mathbf{1}_{t=s} +$

$(1-\lambda)g_i(t)$. This allows us to check for global risk dominance by setting $\lambda \geq 0.5$. Specifically, pure-strategy Nash equilibrium (t_1, t_2) is a globally risk dominant Nash equilibrium if for all pdfs g over $\mathcal{T}$ and all $\lambda \geq 1/2$, t_1 is the unique best response to the mixed strategy for firm 2 of $\hat{g}(\cdot; \lambda, t_2)$ and t_2 is the unique best response to the mixed strategy for firm 1 of $\hat{g}(\cdot; \lambda, t_1)$.

We are then able to show that no restoration by either firm uniquely satisfies the criteria for being globally risk dominant when there is a high cost of price leadership and that restoration at time 0 by both firms uniquely satisfies the criteria for global risk dominance when there is a low cost of price leadership:

Proposition 3. With two uninformed firms, the unique globally risk dominant equilibrium has no restoration by either firm if $\pi_0 - \pi_1 > \pi_3 - \pi_2$, and the unique globally risk dominant equilibrium has restoration at time 0 by both firms if $\pi_0 - \pi_1 < \pi_3 - \pi_2$.

Proof. See Appendix A.4.

Using the notion of stochastically stable equilibria discussed above and the results of Maruta (1997), we have the following corollary (with the analogous result holding for the case of a low cost of price leadership):

Corollary 1. With two uninformed firms and a high cost of price leadership, the unique stochastically stable equilibrium has no restoration by either firm.

Proof. Noting that our game with two uninformed firms is weakly acyclic (the pure best response by i to any t_j is t_j, and (t_j, t_j) is a strict Nash equilibrium) and that uninformed firms observe a history of length zero ($k = 0$ in the notation of Maruta (1997)), the Theorem of Maruta (1997) applies. □

What is key for the potential disruption of restorations described in Proposition 3 and Corollary 1 is that the two firms in the market are characterized by what we call *bilateral information blockage*—neither firm can readily observe a restoration by the other. Such a blockage introduces a coordination problem that limits the firms' ability to earn high profits quickly.

5.4 Two uninformed and one or more informed firms

Of course, removing two retailers from the platform would have left multiple retailers still on the platform. Therefore, we now incorporate on-platform retailers into our model with two off-platform uninformed retailers. As described above, we assume that the on-platform retailers are informed, such as when the platform collects information on restorations by the off-platform retailers and distributes that information to one or more informed on-platform retailers.

We show here that the addition of informed firms does not affect the globally risk dominant strategies of the two uninformed firms. This is intuitive because, as we know from Lemma 1, once the second of the two uninformed firms restore, all informed firms will restore shortly thereafter. Thus, the game faced by the two uninformed firms is essentially unchanged when the informed firms are added—once the second of the two uninformed firms restores, or shortly thereafter, the uninformed firms will start receiving payoff π_3 per unit of time. The remainder of this section formalizes these results. A reader looking to go straight to the policy implications of these findings can jump to Section 6.

To formalize these results, we first extend the notion of global risk dominance to accommodate the presence of informed firms. Given a strategy profile $\boldsymbol{\sigma}$ for the informed firms, we define the *restricted strategic-form game given* $\boldsymbol{\sigma}$ for the uninformed firms to be the strategic-form game in which the payoffs for any pure-strategy profile $\mathbf{t}$ for the uninformed firms are the payoffs associated with the strategy profile $(\mathbf{t}, \boldsymbol{\sigma})$. Given this, we can extend the notion of a globally risk dominant Nash equilibrium to our setup as follows:

Definition 3. A market equilibrium $(\mathbf{t}, \boldsymbol{\sigma})$ is a *globally risk dominant market equilibrium* if $\mathbf{t}$ is a globally risk dominant Nash equilibrium for the uninformed firms in the restricted strategic-form game given $\boldsymbol{\sigma}$.

To define equilibrium strategies, we must consider off-equilibrium behavior. Off the equilibrium path, if an informed firm observes that all other firms have restored, then its best response is to restore at the first opportunity. Relevant to other off-equilibrium cases, we assume that the observation by the informed firm of a deviation by one of the uninformed firms does not affect the informed firm's belief about that strategy being played by the nondeviating uninformed firm, akin to the "passive beliefs" of, e.g., McAfee and Schwartz (1994). Further, if an uninformed firm deviates by not restoring when expected, we assume that the informed firms believe that the deviating uninformed firm will restore at time 1. This latter assumption implies that the failure of an uninformed firm to restore when expected results in any remaining unrestored informed firms following the continuation strategy of waiting until time 1 to restore, unless it is observed that all other firms have restored. We refer to these beliefs as *passive-aggressive beliefs* and restrict attention to such off-equilibrium beliefs for the remainder of the paper.

We now examine globally risk dominant market equilibria. As one might expect, the results are easiest to see with only one informed firm, so we begin there. However, we show below that our results extend to any number of informed firms. With only one informed firm, in any best response for the single informed firm to pure strategies by the uninformed firms, the informed firm must restore in the same period as when the second of the two uninformed firms restores. It follows that the only market equilibrium involving pure strategies, which is required for a globally risk dominant market equilibrium, is for all three firms to restore in the same period t.

Consistent with this, letting firm 3 be the informed firm, we can define a strategy $\sigma^*_{t^*}$ that is a best response for firm 3 to both uninformed firms restoring at time t^* as follows. Given state $\mathbf{h} \in \{0,1\}^3$ with $h_3 = 0$, i.e., firm 3 has not yet restored, the strategy $\sigma^*_{t^*}$ for firm 3 in period t given state $\mathbf{h}$ is given by $\sigma^*_{t^*}(\mathbf{h}, t) = 1$ if $h_1 = h_2 = 1$ or $t = t^*$, and $\sigma^*(\mathbf{h}) = 0$ otherwise.

We first observe that $\{1, 1, \sigma^*_1\}$ is a globally risk dominant market equilibrium if and only if $\pi_0 - \pi_1 > \pi_3 - \pi_2$ (i.e., there is a high cost of price leadership). To see this, consider a mixed strategy by firm 2 of $\hat{g}_2(\cdot; \lambda, 1)$. Firm 1's expected payoff from restoring at time 1 is then

$$\Pi_1(1) = \lambda\pi_0 + (1-\lambda)\sum_{\tau\in\mathcal{T}}(\tau\pi_0 + (1-\tau)\pi_2)g_2(\tau), \tag{6}$$

and firm 1's expected payoff from restoring at time $t_1 \in \mathcal{T}\backslash\{1\}$ is

$$\begin{aligned}\Pi_1(t_1) = {} & \lambda\left(t_1\pi_0 + (1-t_1)\pi_1\right) \\ & + (1-\lambda)\sum_{\tau\in\{0,\varepsilon,\ldots,t_1\}}\left(\tau\pi_0 + (t_1-\tau)\pi_2 + \varepsilon\pi_1 + (1-t_1-\varepsilon)\pi_3\right)g_2(\tau) \\ & + (1-\lambda)\sum_{\tau\in\{t_1+\varepsilon,\ldots,1-\varepsilon\}}\left(t_1\pi_0 + (\tau+\varepsilon-t_1)\pi_1 + (1-\tau-\varepsilon)\pi_3\right)g_2(\tau) \\ & + (1-\lambda)(t_1\pi_0 + (1-t_1)\pi_1)g_2(1),\end{aligned} \tag{7}$$

where the expression for $\Pi_1(t_1)$ accounts for the restoration of firm 3 under σ^*_1 one period after the restoration of the second of the uninformed firms to restore (or at time 1).

Proceeding as in the proof of Proposition 3, we can show that:

Lemma 3. With two uninformed and one informed firm, for $\varepsilon > 0$ sufficiently small, $\Pi_1(1) > \Pi_1(t_1)$ for all $t_1 \in \mathcal{T}\backslash\{1\}$, g_2, and $\lambda \geq 1/2$ if and only if there is a high cost of price leadership.

Proof. See Appendix A.5.

It follows from Lemma 3 that with two uninformed and one informed firm, $(1, 1, \sigma^*_1)$ is a globally risk dominant market equilibrium if and only if there is a high cost of price leadership, given ε sufficiently close to zero.

The result of Lemma 3 extends straightforwardly to the case of more than two informed firms. In the case of $m + 2$ total firms, with two uninformed and $m \geq 2$ informed, as before, for any t, there is an equilibrium in which all firms restore in period t. That is, there exist strategies for the informed firms $\sigma^*_{3,t}, \ldots, \sigma^*_{m,t}$ such that on the equilibrium path, firms $3, \ldots, m$ restore at time t, and such that $(t, t, \sigma^*_{3,t}, \ldots, \sigma^*_{m,t})$ is a market equilibrium.

Define σ^*_1 to be a strategy for an informed firm that specifies that, on the equilibrium path, the firm restores at time 1 and that, off the equilibrium path, it restores optimally given the other firms' equilibrium strategies and passive-aggressive beliefs. We first show that $(1, 1, \sigma^*_1, \ldots, \sigma^*_1)$ is a globally risk dominant market equilibrium if and only if there is a high cost of price leadership with $\pi_0 - \pi_1 > \pi_3 - \pi_2$.

Lemma 4. With two uninformed firms and $m \geq 1$ informed firms and $\varepsilon > 0$ sufficiently small, $(1, 1, \sigma^*_1, \ldots, \sigma^*_1)$ is a globally risk dominant market equilibrium if and only if there is a high cost of price leadership.

Proof. See Appendix A.6

In order to establish uniqueness, suppose that (5) holds and that for some $t_1 \in \mathcal{T}\backslash\{1\}$, $(t_1, t_1, \sigma^*_3, \ldots, \sigma^*_n)$ is a globally risk dominant market equilibrium. From Lemma 1, in equilibrium, all firms restore at or before time $t_2 \equiv \min\{1, t_1 + \varepsilon \sum_{i=1}^{n-2} \tau(i)\}$. Denote by Y firm 1's expected payoff on the equilibrium path associated with the time interval $[0, t_1]$. Denote by X firm 1's expected payoff on the equilibrium path associated with the time interval $[t_1, t_2]$, where $X \leq (t_2 - t_1)\pi_3$.

Consider a mixed strategy by firm 2 with weight $1/2$ on t_1 and $1/2$ on 1. If firm 2 restores at t_1, then all firms get their equilibrium payoffs. If firm 2 does not restore at t_1, firm 1's expected payoff under passive-aggressive beliefs is π_1 between time t_1 and time 1. Thus, firm 1's expected payoff from restoring at time t_1 is

$$\hat{\Pi}_1(t_1) = Y + \frac{1}{2}\left(X + (1 - t_2)\pi_3\right) + \frac{1}{2}(1 - t_1)\pi_1,$$

and firm 1's expected payoff from restoring at time 1 is

$$\hat{\Pi}_1(1) = Y + \frac{1}{2}(1 - t_1)\pi_2 + \frac{1}{2}\left(1 - t_1\right)\left(w\pi_2 + (1 - w)\pi_0\right).$$

where $w \in [0, 1]$ is the probability that an informed firm restores at or before t_1.

We then have

$$\begin{aligned}\hat{\Pi}_1(1) - \hat{\Pi}_1(t_1) &= \frac{1}{2}(1-t_1)\left(\pi_2 + w\pi_2 + (1-w)\pi_0 - \pi_1\right) - \frac{1}{2}\left(X + (1-t_2)\pi_3\right) \\ &\geq \frac{1}{2}(1-t_1)\left(\pi_2 + \pi_0 - \pi_1\right) - \frac{1}{2}(1-t_1)\pi_3 \\ &= \frac{1}{2}(1-t_1)((\pi_0 - \pi_1) - (\pi_3 - \pi_2)) \\ &> 0,\end{aligned}$$

where the first inequality uses $w\pi_2 + (1-w)\pi_0 \geq \pi_0$ and $X \leq (t_2 - t_1)\pi_3$ and the final inequality uses (5). This implies that t_1 is not a best response for firm 1. Thus, $(1, 1, \sigma_1^*, \ldots, \sigma_1^*)$ is the essentially unique globally risk dominant equilibrium (up to off-equilibrium strategies for the informed firms).

This completes the proof of the following proposition for the case of a high cost of price leadership, with the result for a low cost of price leadership following by analogous arguments. However, with a low cost of price leadership, we no longer have the uniqueness result because, for example, there could be a globally risk dominant market equilibrium in which all firms restore at time ε.

Proposition 4. With two uninformed firms and any number of informed firms and $\varepsilon > 0$ sufficiently small, if there is a high cost of price leadership, then the unique outcome of any globally risk dominant market equilibrium has no restoration by any firm, and if there is a low cost of price leadership, then there is a globally risk dominant market equilibrium in which all firms restore at time 0.

As Proposition 4 shows, the low-price outcome with no price restoration is a globally risk dominant market equilibrium when there are two uninformed firms, regardless of the number of informed firms, as long as there is a high cost of price leadership.

6 Antitrust policy

Turning to antitrust policy, our model has allowed us to consider equilibria under different information-sharing structures: the one that emerged after the *Informed Sources* case and others that did not emerge but were possible. Table 2 summarizes our main findings regarding the ease of coordination under different information-sharing structures. In sum, disrupting price coordination with high-frequency information sharing can be hard, but it can be done, at least in theory. However, it may require both: (i) multiple uninformed firms and (ii) a high cost of price leadership. Sufficient off-platform competition and high cost of price leadership are thus *complements*: one may require, e.g., a high demand elasticity to induce coordination-disrupting informational blockage via off-platform firms.

Table 2: Information-Sharing Structures and Coordination Outcomes

	Cost of price leadership	
Uninformed Firms	**High**	**Low**
0 or 1	Coordination occurs quickly (all firms restore near $t = 0$ in every equilibrium)	
2	Coordination fails (no firms restore in any GRD-E)	Coordination can occur quickly (all firms restore near $t = 0$ in a GRD-E)

Notes: GRD-E is short for Globally Risk Dominant Equilibrium per Definition 3.

Digging deeper, our analysis illustrates how small changes to a market's information-sharing structure can yield drastically different competitive outcomes. Under the observed case settlement, having one uninformed firm yields an uninformed price leader and successful coordination. However, under a counterfactual settlement with *just one* additional uninformed firm, a coordination problem arises through bilateral informational blockage, and coordination breaks down in the globally risk dominant equilibrium if price leadership is sufficiently costly. In the context of the *Informed Sources* case, given the large magnitude of cross-price demand elasticities in retail gasoline, removing two firms from the platform could potentially have disrupted coordination.

These results suggest two complementary components of a policy that restricts information sharing to disrupt price coordination. The policy would need to ensure (1) the creation of two significant firms in a market with bilateral information blockage, which, at a minimum, requires that those two firms do not participate in the price-sharing platform, and (2) that the cost of price leadership is high. In the *Informed Sources* case settlement, the requirement that Informed Sources' price data be made available to providers of consumer price-comparison apps might have enhanced the cost of price leadership and complemented an exclusion of two firms from the platform (had that been part of the settlement).

However, as mentioned in Section 2.4, no consumer apps were developed. We describe various reasons why in Byrne et al. (2024b), including announcements around the time of the case settlement by the state of New South Wales (where Sydney is located) that they were set to launch government-mandated price disclosure on consumer search apps with complete data coverage for all stations at all times. Even if apps had been launched with full-price disclosure to consumers, they represent a challenge for competition authorities in regulating digital information sharing. While they can raise demand elasticity and the cost of price leadership, they enable monitoring and price coordination among firms that can ultimately defeat any attempts to disrupt coordination.

Data processing may help thread this policy needle. For instance, through an app-enabled mandatory price disclosure policy, a government could implement informational nudges through an app using internally available high-frequency seller-level data that gives consumers useful information on where and when to make purchases.[45] Processing and publishing data in this way, rather than publishing the raw data itself, can potentially enhance demand elasticity while restricting supply-side information sharing that facilitates coordination. Further study is needed on how data processing can be leveraged to enable consumer search while undermining monitoring and coordination in pursuit of informational blockage, particularly as industries increasingly use information-sharing platforms and algorithms for price setting.

A final new insight from our analysis is that increasing demand elasticity potentially has multiple effects. As suggested by the prior literature, an increase in the elasticity increases gains from cheating and reduces the scope for equilibria at supra-competitive prices (Stigler, 1964; Friedman, 1971; Luco, 2019). In that case, a smaller range of discount factors can support monopoly prices. In this paper, we have shown that an increase in elasticity also increases the cost of price leadership and can reduce the scope for getting coordination off the ground in the first place. Thus, an increase in the elasticity pays a *double dividend* in the sense that it has the potential to handicap both the implementation and the initiation of coordination. These twin benefits for competition emphasize the potential benefits of, for example, consumer apps and demand-side algorithms that enable consumers to respond better to price differentials.

7 Conclusion

The theoretical model developed in this paper reveals how asymmetric access to information on rivals' actions can affect firms' ability to coordinate price increases in a timely manner. While we elucidated implications from specific antitrust case—*Information Sources*—the

[45]Such information-sharing platforms exist worldwide in retail gasoline and other important goods and services markets, including health insurance, retirement savings, real estate, retail electricity, or grocery. For example, the FuelCheck platform and app (https://www.fuelcheck.nsw.gov.au/) in New South Wales and Australian Capital Territory, launched in August 2016, provides real-time information sharing from all stations as part of its mandatory price disclosure policy. As of 2023, 33% of households have downloaded the app (ACCC, 2024).

findings are relevant to any market in which firms interact repeatedly in an environment with common cost or demand shocks, frequent and observable price changes, and a high cost of price leadership. For example, access to a price-sharing platform may allow hotels in a local market to increase prices more rapidly in response to a temporary reduction in demand elasticity, whereas this may be more difficult in markets where multiple hotel operators do not have access to such timely price information.

With regards to the *Informed Sources* case, our analysis provides a foundation for the rationale that a settlement excluding two firms, rather than just one, from the information-sharing platform at issue might have had more success achieving procompetitive benefits by delaying or disrupting coordination on price increases, particularly when combined with efforts to improve consumer access to price information. The results from the case and our theory highlight the need for additional research on how information structures shape coordination because that understanding can play a pivotal role in informing remedies.

Price-sharing platforms can facilitate price coordination, but we have now found multiple cautionary tales for competition authorities trying to prevent that. We have seen strategic responses by platforms to case settlements where Informed Sources implemented the manual collection of Coles' prices after the firm exited the platform. We have also seen that firms can be clever in the retailers' recognition that Coles could still lead restoration coordination after its exit from the platform (this paper) and achieve higher profit margins (Byrne et al., 2024b). And we have seen that limiting platform membership may not disrupt coordination—in particular, excluding only one firm may not be enough. On a positive note, however, we can offer some guidance for emerging digital information sharing based on our theoretical results and, in this way, learn from the *Informed Sources* case. In particular, we show that there may be benefits from excluding multiple firms from the platform if price leadership is costly and that policies that increase the cost of price leadership may be beneficial for preventing coordination from emerging.

Topics for future research abound with the continued rise of price-sharing platforms and algorithms in the economy. Of particular note are sharpening distinctions between price-sharing platforms, price-setting platforms, and algorithmic pricing and their implications for coordination and market power. Price-setting platforms go further than price-sharing platforms by recommending or setting prices for subscribers.[46] Price recommendations may be accompanied by requirements that they be followed at least for some share of the sub-

[46]Examples include RealPage, which advertises "moving from manual pricing to refining revenue" (https://www.realpage.com/case-studies/manual-pricing-to-refining-revenue/, https://www.justice.gov/d9/2023-11/418053.pdf) and OPIS PricePro, which advertises "predictable and explainable proposed pricing" and "capable of delivering your determined price to your customer through your point-of-sale system" (https://pricepro.opisnet.com/).

scriber's business. Algorithmic pricing, which Brown and MacKay (2022) define as the use of "a computer program that autonomously adjusts prices based on current and past data related to demand, cost, or rivals' prices," provides another variant on platform-based price setting with as-yet incompletely understood effects.[47] In addition, future research establishing the extent to which the presence of consumer apps affects price elasticity in different market settings would help unveil the micro-foundations of the cost of price leadership and the feasibility of price coordination.

[47]As an example, "The vast majority of hotels on the Las Vegas Strip use Rainmaker Group's pricing algorithms" (https://www.reddit.com/r/law/comments/10qs5qb/the_vast_majority_of_hotels_on_the_las_vegas/).

References

ACCC (2007): "Petrol Prices and Australian Consumers," Report of the ACCC Inquiry into the Price of Unleaded Petrol.

——— (2010): "Monitoring of the Australian Petroleum Industry: Report of the ACCC into the Prices, Costs and Profits of Unleaded Petrol in Australia," December 2010.

——— (2014): "ACCC Takes Action Against Informed Sources and Petrol Retailers for Price Information Sharing," Press Release, August 20, 2014.

——— (2015a): "ACCC and Coles Express Resolve Petrol Price Information Sharing Proceedings," Press Release, December 16, 2015.

——— (2015b): "Petrol Price Information Sharing Proceedings Resolved," Press Release, December 23, 2015.

——— (2024): "Making the Most of Fuel Price Apps and Websites," .

Asker, J., C. Fershtman, J. Jeon, and A. Pakes (2020): "A Computational Framework for Analyzing Dynamic Auctions: The Market Impact of Information Sharing," *RAND Journal of Economics*, 51, 805–839.

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.

Bachmeier, L. J. and J. M. Griffin (2003): "New Evidence on Asymmetric Gasoline Price Responses," *Review of Economics and Statistics*, 85, 772–776.

Blume, L. E. (1993): "The Statistical Mechanics of Strategic Interaction," *Games and Economic Behavior*, 5, 387–424.

Borenstein, S. (2004): "Rapid Price Communication and Coordination: The Airline Tariff Publishing Case (1994)," in *The Antitrust Revolution: Economics, Competition and Policy*, ed. by J. Kwoka Jr. and L. White, Oxford University Press, 233–251.

Brown, Z. Y. and A. MacKay (2023): "Competition in Pricing Algorithms," *American Economic Journal: Microeconomics*, 15, 109–156.

Brown, Z. Y. and A. J. MacKay (2022): "Are Online Prices Higher Because of Pricing Algorithms?" *Brookings Series: The Economics and Regulation of Artificial Intelligence and Emerging Technologies*.

Byrne, D. P. and N. de Roos (2019): "Learning to Coordinate: A Study in Retail Gasoline," *American Economic Review*, 109, 591–619.

Byrne, D. P., N. de Roos, A. R. Grinberg, and L. M. Marx (2024a): "Informed Sources and the Role of Platforms for Facilitating Anticompetitive Communication," in *Cartels Diagnosed: New Insight on Collusion*, ed. by J. E. Harrington and M. P. Schinkel, Cambridge.

Byrne, D. P., N. de Roos, M. S. Lewis, L. M. Marx, and X. Wu (2024b): "Asymmetric Information Sharing in Oligopoly: A Natural Experiment in Retail Gasoline," *Journal of Political Economy*, forthcoming.

——— (2024c): "Negotiating Price Coordination: A Study in Retail Gasoline," Working Paper, Duke.

Byrne, D. P., G. Leslie, and R. Ware (2015): "How do Consumers Respond to Gasoline Price Cycles?" *The Energy Journal*, 36, 115–147.

Carlsson, H. and E. van Damme (1993a): "Equilibrium Selection in Stag Hunt Games," in *Frontiers of Game Theory*, ed. by K. Binmore, A. Kirman, and P. Tani, Cambridge, MA: MIT Press, vol. 4, 237–253.

——— (1993b): "Global Games and Equilibrium Selection," *Econometrica*, 61, 989–1018.

Castanias, R. and H. Johnson (1993): "Gas Wars: Retail Gasoline Price Fluctuations," *Review of Economics and Statistics*, 75, 171–174.

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.

de Roos, N. (2012): "Static Models of the Edgeworth Cycle," *Economics Letters*, 117, 881–882.

——— (2017): "Edgeworth cycles with partial price commitment," *Economics Letters*, 150, 122–125.

de Roos, N. and H. Katayama (2013): "Gasoline Price Cycles Under Discrete Time Pricing," *Economic Record*, 89, 175–193.

DOJ, U. (2023): "Principal Deputy Assistant Attorney General Doha Mekki of the Antitrust Division Delivers Remarks at GCR Live: Law Leaders Global 2023," Feburary 2, 2023, Miami, Florida.

Eckert, A. (2003): "Price Cycles and the Presence of Small Firms," *International Journal of Industrial Organization*, 21, 151–170.

Eckert, A. and D. West (2004): "Retail Gasoline Price Cycles Across Spatially Dispersed Gasoline Stations," *Journal of Law and Economics*, 22, 889–1015.

Edgeworth, F. Y. (1925): *Papers Relating to Political Economy. Volume I, The Pure Theory of Monopoly*, Macmillan, London.

Ellison, G. (2000): "Basins of Attraction, Long-Run Stochastic Stability, and the Speed of Step-by-Step Evolution," *Review of Economic Studies*, 67, 17–45.

European Commission (2011): "Guidelines on the Applicability of Article 101 of the Treaty on the Functioning of the European Union to Horizontal Co-operation Agreements," Document 52011XC0114(04).

——— (2023): "Guidelines on the Applicability of Article 101 of the Treaty on the Functioning of the European Union to Horizontal Co-operation Agreements," Official Journal of the European Union 2023/C 259/01.

Foros, O. and F. Steen (2013): "Vertical Control and Price Cycles in Gasoline Retailing," *Scandinavian Journal of Economics*, 115, 640–661.

Foster, D. and P. Young (1990): "Stochastic Evolutionary Game Dynamics," *Theoretical Population Biology*, 38, 219–232.

Friedman, J. W. (1971): "A Non-Cooperative Equilibrium for Supergames," *The Review of Economic Studies*, 38, 1–12.

Fudenberg, D. and D. K. Levine (1998): *The Theory of Learning in Games*, MIT Press, Cambridge, Massachusetts.

Genesove, D. and W. Mullin (2001): "Rules, Communication, and Collusion: Narrative Evidence from the Sugar Institute Case," *American Economic Review*, 91, 379–398.

Green, E. J. and R. H. Porter (1984): "Noncooperative Collusion under Imperfect Price Information," *Econometrica*, 52, 87–100.

Harrington, J. E. and A. Skrzypacz (2011): "Private Monitoring and Communication in Cartels: Explaining Recent Collusive Practices," *American Economic Review*, 101, 2425–49.

Harsanyi, J. C. and R. Selten (1988): *A General Theory of Equilibrium Selection in Games*, MIT Press, Cambridge, Massachusetts.

Houde, J.-F. (2012): "Spatial Differentiation and Vertical Mergers in Retail Markets for Gasoline," *American Economic Review*, 105, 2147–2182.

Ivaldi, M., B. Jullien, P. Rey, P. Seabright, and J. Tirole (2007): "The Economics of Tacit Collusion: Implications for Merger Control," in *The Political Economy of Antitrust*, ed. by V. Ghosal and J. Stennek, Elsevier, 217–239.

Kandori, M., G. J. Mailath, and R. Rob (1993): "Learning, Mutation, and Long Run Equilibria in Games," *Econometrica*, 61, 29–56.

Kandori, M. and R. Rob (1995): "Evolution of Equilibria in the Long Run: A General Theory and Applications," *Journal of Economic Theory*, 65, 383–414.

——— (1998): "Bandwagon Effects and Long Run Technology Choice," *Games and Economic Behavior*, 22, 30–60.

Kashyap, A. K. and J. C. Stein (2000): "What Do a Million Observations on Banks Say About the Transmission of Monetary Policy?" *American Economic Review*, 90, 407–428.

Kubitz, G. and K. Woodward (2020): "Sharing Cost Information in Dynamic Oligopoly," Working Paper, Queensland University of Technology.

Kuhn, K.-U. (2001): "Fighting Collusion by Regulating Communication between Firms," *Economic Policy*, 16, 169–204.

Kuhn, K.-U. and X. Vives (1995): "Information Exchanges Among Firms and Their Impact on Competition," European Commission Document, Luxembourg: Office for Official Publications of the European Communities.

Lewis, M. and M. Noel (2011): "The Speed of Gasoline Price Response in Markets with and without Edgeworth Cycles," *The Review of Economics and Statistics*, 93, 672–682.

Lewis, M. S. (2012): "Price Leadership and Coordination in Retail Gasoline Markets with Price Cycles," *International Journal of Industrial Organization*, 30, 342–351.

Luco, F. (2019): "Who Benefits from Information Disclosure? The Case of Retail Gasoline," *American Economic Journal: Microeconomics*, 11, 277–305.

Maggio, M. D., A. Kermani, B. J. Keys, T. Piskorski, R. Ramcharan, A. Seru, and V. Yao (2017): "Interest Rate Pass-Through: Mortgage Rates, Household Consumption, and Voluntary Deleveraging," *American Economic Review*, 107, 3550–3588.

Maruta, T. (1997): "On the Relationship between Risk-Dominance and Stochastic Stability," *Games and Economic Behavior*, 19, 221–234.

Maskin, E. and J. Tirole (1988): "A Theory of Dynamic Oligopoly, II: Price Competition, Kinked Demand Curves, and Edgeworth Cycles," *Econometrica*, 56, 571–99.

McAfee, R. P. and M. Schwartz (1994): "Opportunism in Multilateral Vertical Contracting: Nondiscrimination, Exclusivity, and Uniformity," *American Economic Review*, 84, 210–230.

Miller, N. H., G. Sheu, and M. Weinberg (2021): "Oligopolistic Price Leadership and Mergers: The United States Beer Industry," *American Economic Review*, 111, 3123–3159.

Morris, S. and H. S. Shin (1993): "Global Games: Theory and Applications," in *Advances in*

Economics and Econometrics: The Eighth World Congress, ed. by M. Dewatripont, L. P. Hansen, and S. J. Turnovsky, Cambridge: Cambridge University Press, 56–114.

Musolff, L. (2024): "Algorithmic Pricing, Price Wars and Tacit Collusion: Evidence from E-Commerce," Working Paper, Wharton School of the University of Pennsylvania.

Noel, M. D. (2007a): "Edgeworth Price Cycles, Cost-Based Pricing, and Sticky Pricing in Retail Gasoline Markets," *Review of Economics and Statistics*, 89, 324–334.

——— (2007b): "Edgeworth Price Cycles: Evidence from the Toronto Retail Gasoline Market," *The Journal of Industrial Economics*, 55, 69–92.

——— (2008): "Edgeworth Price Cycles and Focal Prices: Computational dynamic Markov Equilibria," *Journal of Economics and Management Strategy*, 17, 345–377.

——— (2009): "Do Gasoline Prices Respond Asymmetrically to Cost Shocks? The Effect of Edgeworth Cycles," *RAND Journal of Economics*, 40, 582–595.

OECD (2011): "Information Exchanges Between Competitors under Competition Law," *OECD Policy Roundtables*, Directorate for Financial and Enterprise Affairs Competition Committee, DAF/COMP(2010)37.

Reinganum, J. F. (1981): "On the Diffusion of New Technology: A Game Theoretic Approach," *Review of Economic Studies*, 48, 395–405.

Scherer, F. M. and D. R. Ross (1990): *Industrial Market Structure and Economic Performance*, Houghton Mifflin.

Stigler, G. J. (1964): "A Theory of Oligopoly," *Journal of Political Economy*, 72, 44–61.

U.S. FTC (2000): "Antitrust Guidelines for Collaborations Among Competitors," Issued by the Federal Trade Commission and the U.S. Department of Justice.

——— (2014): "Information Exchange: Be Reasonable," *Competition Matters*, by Michael Bloom, December 2014.

Vives, X. (2007): "Information Sharing: Economics and Antitrust," IESE Business School Occasional Paper 07/11, University of Navarra.

Wang, Z. (2009): "Station-Level Gasoline Demand in an Australian Market with Regular Price Cycles," *Agricultural and Resource Economics*, 53, 467–483.

Wu, X., M. S. Lewis, and F. A. Wolak (2024): "Search with Learning in the Retail Gasoline Market," *RAND Journal of Economics*, 55, 292–323, https://doi.org/10.1111/1756-2171.12466.

Young, H. P. (1993): "The Evolution of Conventions," *Econometrica*, 61, 57–84.

A Proofs

A.1 Proof of Lemma 1

Proof. Suppose that at time t^* there remain $m \in \{1, 2, \ldots, \overline{m}(t^*)\}$ unrestored informed firms. If $m = 1$, then, in equilibrium, the remaining unrestored informed firm restores at time $t^* + \varepsilon = t^* + \varepsilon\tau(1)$, where the equality uses the definition of $\tau(1)$.

Let $m \in \{2, \ldots, \overline{m}(t^*) - 1\}$ be given and assume, in an induction step, that if all firms are restored in period t^* other than m informed firms, then in equilibrium all firms restore at or before time $t^* + \varepsilon \sum_{i=1}^{m} \tau(i)$.

Consider the case of $m + 1$ unrestored informed firms at time t^*. Suppose that the equilibrium continuation strategies by the $m + 1$ informed firms following time t^* result in at least one of the $m + 1$ informed firms not restoring until after time $t^* + \varepsilon \sum_{i=1}^{m+1} \tau(i)$. Let firm 1 be such a firm. From the perspective of time $t^* + \varepsilon$, firm 1's continuation payoff is less than $\pi_2(t^* + \varepsilon \sum_{i=1}^{m+1} \tau(i) - (t^* + \varepsilon)) + (1 - t^* - \varepsilon \sum_{i=1}^{m+1} \tau(i))\pi_3$ (strictly less than because full restoration occurs by supposition strictly after $t^* + \varepsilon \sum_{i=1}^{m+1} \tau(i)$). If firm 1 deviates from its equilibrium continuation strategy and restores at time $t^* + \varepsilon$, then by the induction assumption, that induces the remaining m informed firms to all restore by time $t^* + \varepsilon + \varepsilon \sum_{i=1}^{m} \tau(i)$, and so gives firm 1 a continuation payoff from the perspective of time $t^* + \varepsilon$ of at least $\pi_1 \varepsilon \sum_{i=1}^{m} \tau(i) + (1 - t^* - \varepsilon - \varepsilon \sum_{i=1}^{m} \tau(i))\pi_3$. This is strictly profitable for firm 1 if

$$\pi_2(t^* + \varepsilon \sum_{i=1}^{m+1} \tau(i) - (t^* + \varepsilon)) + (1 - t^* - \varepsilon \sum_{i=1}^{m+1} \tau(i))\pi_3 \leq \pi_1 \varepsilon \sum_{i=1}^{m} \tau(i) + (1 - t^* - \varepsilon \sum_{i=1}^{m} \tau(i) - \varepsilon)\pi_3,$$

which we can rewrite as

$$\tau(m+1) \geq 1 + \frac{\pi_2 - \pi_1}{\pi_3 - \pi_2} \sum_{i=1}^{m} \tau(i),$$

which is satisfied by our definition of $\tau(\cdot)$. This contradicts the supposition that equilibrium continuation strategies by the $m+1$ informed firms following time t^* result in at least one of the $m + 1$ unrestored informed firms not restoring until after time $t^* + \varepsilon \sum_{i=1}^{m+1} \tau(i)$, giving us the result that all of the remaining $m + 1$ unrestored informed firms restore by time $t^* + \varepsilon \sum_{i=1}^{m+1} \tau(i)$. This completes the induction argument and proves that if all firms are restored in period t^* other than $m \in \{1, \ldots, \overline{m}(t^*) - 1\}$ informed firms, then in equilibrium all firms restore at or before time $t^* + \varepsilon \sum_{i=1}^{m} \tau(i)$. $\square$

A.2 Proof of Lemma 2

Proof. If all n firms are informed, then by Lemma 1, all firms restore at or before time $\varepsilon \sum_{i=1}^{n} \tau(i)$. Suppose that there is exactly one uninformed firm. Suppose, by way of contradiction, that there exists an equilibrium in which the earliest time with all n firms restored is $t' > \varepsilon \frac{\pi_3 - \pi_1}{\pi_3 - \pi_2} \sum_{i=1}^{n-1} \tau(i)$. Then the uninformed firm's payoff is at most $t'\pi_2 + (1 - t')\pi_3$. Consider a deviation by the uninformed firm to restore at time 0. In that case, by Lemma 1, all firms restore by time $\varepsilon \sum_{i=1}^{n-1} \tau(i)$. Thus, the uninformed firm's payoff from the deviation to restore at time 0 is at least $\varepsilon \sum_{i=1}^{n-1} \tau(i)\pi_1 + (1 - \varepsilon \sum_{i=1}^{n-1} \tau(i))\pi_3$. This is greater than the uninformed firm's supposed equilibrium payoff if

$$t'\pi_2 + (1 - t')\pi_3 < \varepsilon \sum_{i=1}^{n-1} \tau(i)\pi_1 + (1 - \varepsilon \sum_{i=1}^{n-1} \tau(i))\pi_3,$$

which we can rewrite as

$$t' > \varepsilon \frac{\pi_3 - \pi_1}{\pi_3 - \pi_2} \sum_{i=1}^{n-1} \tau(i),$$

which holds by our supposition that $t' > \varepsilon \frac{\pi_3 - \pi_1}{\pi_3 - \pi_2} \sum_{i=1}^{n-1} \tau(i)$. Because the uninformed firm's deviation is thus profitable, we have a contradiction and conclude that with one uninformed firm, there is no equilibrium in which the earliest time of having all firms restored is $t' > \varepsilon \frac{\pi_3 - \pi_1}{\pi_3 - \pi_2} \sum_{i=1}^{n-1} \tau(i) = \varepsilon \sum_{i=1}^{n} \tau(i) - \varepsilon$, which completes the proof. □

A.3 Proof of Proposition 2

Proof. Consider a setup with 1 uninformed firm and $n - 1$ informed firms. Suppose an equilibrium in which the first restoration is by informed firm 1 at time t_1 and the uninformed restores at time t_2, where $t_1 < t_2$. In this case, firm 1's equilibrium payoff is no more than $\pi_0 t_1 + (t_2 - t_1)\pi_1 + (1 - t_2)\pi_3$, which occurs when all other informed firms restore by time t_2. If firm 1 deviates by not restoring at time t_1, and instead restoring at time t_2, then firm 1's payoff is at least $\pi_0 t_2 + \pi_1 \varepsilon \sum_{i=1}^{n-2} \tau(i) + \pi_3 \left(1 - t_2 - \varepsilon \sum_{i=1}^{n-2} \tau(i)\right)$, where by Lemma 1, $t_2 + \varepsilon \sum_{i=1}^{n-2} \tau(i)$ is the upper bound on the time of restoration by the other $n - 2$ informed firms. This deviation is profitable if

$$\pi_0 t_2 + \pi_1 \varepsilon \sum_{i=1}^{n-2} \tau(i) + \pi_3 \left(1 - t_2 - \varepsilon \sum_{i=1}^{n-2} \tau(i)\right) > \pi_0 t_1 + (t_2 - t_1)\pi_1 + (1 - t_2)\pi_3,$$

which we can rewrite as

$$t_2 - t_1 > \frac{\pi_3 - \pi_1}{\pi_0 - \pi_1}\varepsilon \sum_{i=1}^{n-2} \tau(i).$$

Thus, in any pure-strategy equilibrium, the time between the first restoration by an informed firm and the restoration by the uninformed firm is bounded above by $\frac{\pi_3-\pi_1}{\pi_0-\pi_1}\varepsilon \sum_{i=1}^{n-2} \tau(i)$, which goes to zero as ε goes to zero. This completes the proof. $\square$

A.4 Proof of Proposition 3

Proof. Observe that, for any $s \in \mathcal{T}$, the profile (s, s) is a pure-strategy Nash equilibrium. Consider the Nash equilibrium $(1, 1)$. Given pdf g_2 on $\mathcal{T}$ and the mixed strategy for firm 2 with pdf $\hat{g}_2(\cdot; \lambda, 1)$, which has point mass of $\lambda + g_2(1)$ on time 1, firm 1's expected profit from restoring at time 1 is

$$\Pi_1(1) = \lambda\pi_0 + (1-\lambda)\sum_{\tau\in\mathcal{T}}(\pi_0\tau + (1-\tau)\pi_2)g_2(\tau),$$

and from restoring at time $t_1 \in \mathcal{T}\backslash\{1\}$ is

$$\begin{aligned}\Pi_1(t_1) = {} & \lambda\left(t_1\pi_0 + (1-t_1)\pi_1\right) \\ & + (1-\lambda)\sum_{\tau\in\{0,\varepsilon,\ldots,t_1\}}(\pi_0\tau + (t_1-\tau)\pi_2 + (1-t_1)\pi_3)g_2(\tau) \\ & + (1-\lambda)\sum_{\tau\in\{t_1+\varepsilon,t_1+2\varepsilon,\ldots,1\}}(\pi_0 t_1 + (\tau-t_1)\pi_1 + (1-\tau)\pi_3)g_2(\tau).\end{aligned}$$

For sufficiency, we need to show that if $\pi_0 - \pi_1 > \pi_3 - \pi_2$, then for any $\tau \in \mathcal{T}\backslash\{1\}$, $\Pi_1(1) > \Pi_1(\tau)$ for all g_2 and $\lambda \geq 1/2$. And for necessity, we need to show that if $\pi_0 - \pi_1 \leq \pi_3 - \pi_2$, then there exists $t_1 \in \mathcal{T}\backslash\{1\}$, g_2, and $\lambda \geq 1/2$ such that $\Pi_1(1) \leq \Pi_1(t_1)$. Define the mean restoration time under g_2 by

$$\mu_{g_2} \equiv \sum_{\tau\in\mathcal{T}} \tau g_2(\tau) \in [0, 1],$$

and, for $t \in \mathcal{T}\backslash\{1\}$ and $G_2(t) < 1$, define the mean restoration time under g_2, conditional on restoration occurring after time t by

$$\mu_{g_2}^t \equiv \frac{\sum_{\tau\in\{t+\varepsilon,t+2\varepsilon,\ldots,1\}} \tau g_2(\tau)}{1 - G_2(t)} \in (t, 1].$$

Let $t_1 \in \mathcal{T}\backslash\{1\}$ be given.

Case 1. $G_2(t_1) = 1$. In this case, $\Pi_1(1) > \Pi_1(t_1)$ if and only if

$$\begin{aligned} 0 &< \lambda\pi_0 - \lambda\left(t_1\pi_0 + (1-t_1)\pi_1\right) + (1-\lambda)\sum_{\tau\in\mathcal{T}}(\pi_0\tau + (1-\tau)\pi_2)g_2(\tau) \\ & \quad -(1-\lambda)\sum_{\tau\in\{0,\varepsilon,\ldots,t_1\}}(\pi_0\tau + (t_1-\tau)\pi_2 + (1-t_1)\pi_3)g_2(\tau), \end{aligned}$$

which we can write as

$$\lambda(\pi_0 - \pi_1) > (1-\lambda)(\pi_3 - \pi_2),$$

which holds for $\pi_0 - \pi_1 > \pi_3 - \pi_2$ and $\lambda \geq 1/2$.

Case 2. $G_2(t_1) < 1$. In this case, $\Pi_1(1) > \Pi_1(t_1)$ if and only if

$$\frac{\pi_0 - \pi_1}{\pi_3 - \pi_2} > \frac{(1-\lambda)(1-t_1) - (1-\lambda)(1-G_2(t_1))(\mu_{g_2}^{t_1} - t_1)}{\lambda(1-t_1) + (1-\lambda)(1-G_2(t_1))(\mu_{g_2}^{t_1} - t_1)}. \tag{A.1}$$

Using $\mu_{g_2}^{t_1} \in (t_1, 1]$, if $\pi_0 - \pi_1 > \pi_3 - \pi_2$ and $\lambda \geq 1/2$, then left slide is greater than 1 and the expression on the right is less than or equal to 1, and so (A.1) holds.

Combining the two cases, we have shown that $\Pi_1(1) > \Pi_1(t_1)$ for all $t_1 \in \mathcal{T}\backslash\{0,1\}$ and $\lambda \geq 1/2$ if $\pi_0 - \pi_1 > \pi_3 - \pi_2$. It follows that restoring at time 1 is a unique best response for firm 1 to firm 2's mixed strategy of $\hat{g}_2(\cdot;\lambda,1)$ for all g_2 and $\lambda \geq 1/2$, establishing (using symmetry) that $(1,1)$ is a globally risk dominant equilibrium. Further, if $\pi_0 - \pi_1 \leq \pi_3 - \pi_2$, then for $t-1 = 0$, $\lambda = 1/2$, and g_2 with all weight on 0 so that $G_2(0) = 1$, using the analysis of Case 1 above, $\Pi_1(1) > \Pi_1(t_1)$ does not hold. Combining these establishes that $(1,1)$ is globally risk dominant if and only if $\pi_0 - \pi_1 > \pi_3 - \pi_2$.

It remains to establish uniqueness. Suppose that $\pi_0 - \pi_1 > \pi_3 - \pi_2$ and that for some $\hat{t} \in \mathcal{T}\backslash\{1\}$, $(\hat{t},\hat{t})$ is a globally risk dominant equilibrium. Consider a mixed strategy by firm 2 with weight 1/2 on $\hat{t}$ and 1/2 on 1. Then firm 1's expected payoff from restoring at time $\hat{t}$ is

$$\hat{\Pi}_1(\hat{t}) = \pi_0\hat{t} + \frac{1}{2}(1-\hat{t})\pi_3 + \frac{1}{2}(1-\hat{t})\pi_1,$$

and its expected payoff from restoring at time 1 is

$$\hat{\Pi}_1(1) = \pi_0\hat{t} + \frac{1}{2}(1-\hat{t})\pi_2 + \frac{1}{2}(1-\hat{t})\pi_0,$$

where

$$\hat{\Pi}_1(1) - \hat{\Pi}_1(\hat{t}) = \frac{1}{2}(1-\hat{t})\left[(\pi_0 - \pi_1) - (\pi_3 - \pi_2)\right] > 0,$$

which implies that $\hat{t}$ is not a best response for firm 1, contradicting the supposition that

$(\hat{t},\hat{t})$ is globally risk dominant. Thus, $(1,1)$ is the unique globally risk dominant equilibrium, completing the proof for the case of a high cost of price leadership.

The argument for the case of a low cost of price leadership is analogous. Take as given a probability density function g_2 on $\mathcal{T}$ and the mixed strategy for firm 2 with pdf $\hat{g}_2(\cdot;\lambda,0)$, which has point mass of $\lambda + g_2(0)$ on time 0. Then firm 1's expected profit from restoring at time 0 is

$$\Pi_1(0) = \lambda\pi_3 + (1-\lambda)\sum_{\tau\in\mathcal{T}}(\pi_1\tau + (1-\tau)\pi_3)g_2(\tau),$$

and from restoring at time $t_1 \in \mathcal{T}\backslash\{0\}$ is

$$\begin{aligned}\Pi_1(t_1) = {} & \lambda\left(t_1\pi_2 + (1-t_1)\pi_3\right) \\ & + (1-\lambda)\sum_{\tau\in\{0,\varepsilon,\ldots,t_1\}}(\pi_0\tau + (t_1-\tau)\pi_2 + (1-t_1)\pi_3)g_2(\tau) \\ & + (1-\lambda)\sum_{\tau\in\{t_1+\varepsilon,t_1+2\varepsilon,\ldots,1\}}(\pi_0 t_1 + (\tau-t_1)\pi_1 + (1-\tau)\pi_3)g_2(\tau).\end{aligned}$$

Let $t_1 \in \mathcal{T}\backslash\{0\}$ be given. If $G_2(t_1) = 0$, then $\Pi_1(0) > \Pi_1(t_1)$ if and only if

$$(1-\lambda)\left(\pi_0 - \pi_1\right) < \lambda\left(\pi_3 - \pi_2\right),$$

which holds for $\lambda \geq 1/2$. If $G_2(t_1) > 0$, then define $\hat{\mu}_{g_2}^{t_1} \equiv \sum_{\tau\in\{0,\varepsilon,\ldots,t_1\}}\tau g_2(\tau)/G_2(t_1) \leq t_1$. In this case, $\Pi_1(0) > \Pi_1(t_1)$ if and only if

$$(\pi_0-\pi_1)\left((1-\lambda)t_1 - (1-\lambda)G_2(t_1)(t_1-\hat{\mu}_{g_2}^{t_1})\right) < (\pi_3-\pi_2)\left(\lambda t_1 + (1-\lambda)G_2(t_1)(t_1-\hat{\mu}_{g_2}^{t_1}))\right),$$

which holds for $\lambda \geq 1/2$. To see this, note that the left side is less than

$$(\pi_3-\pi_2)\left((1-\lambda)t_1 - (1-\lambda)G_2(t_1)(t_1-\hat{\mu}_{g_2}^{t_1})\right),$$

which is less than $(\pi_3-\pi_2)\left(\lambda t_1 - (1-\lambda)G_2(t_1)(t_1-\hat{\mu}_{g_2}^{t_1})\right)$, which is less than the right side. This establishes that $(0,0)$ is globally risk dominant.

To establish uniqueness, suppose that $\pi_0 - \pi_1 < \pi_3 - \pi_2$ and that for some $\hat{t} \in \mathcal{T}\backslash\{0\}$, $(\hat{t},\hat{t})$ is a globally risk dominant equilibrium. Consider a mixed strategy by firm 2 with weight 1/2 on $\hat{t}$ and weight 1/2 on 0. Then firm 1's expected payoff from restoring at time $\hat{t}$ is

$$\hat{\Pi}_1(\hat{t}) = \frac{1}{2}\hat{t}\pi_0 + \frac{1}{2}\hat{t}\pi_2 + \frac{1}{2}(1-\hat{t})\pi_3,$$

and its expected payoff from restoring at time 0 is

$$\hat{\Pi}_1(0) = \frac{1}{2}\hat{t}\pi_1 + \frac{1}{2}\hat{t}\pi_3 + \frac{1}{2}(1-\hat{t})\pi_3,$$

where

$$\hat{\Pi}_1(0) - \hat{\Pi}_1(\hat{t}) = -\frac{1}{2}\hat{t}\left[(\pi_0 - \pi_1) - (\pi_3 - \pi_2)\right] > 0,$$

which implies that $\hat{t}$ is not a best response for firm 1, contradicting the supposition that $(\hat{t}, \hat{t})$ is globally risk dominant. Thus, $(0, 0)$ is the unique globally risk dominant equilibrium, completing the proof. □

A.5 Proof of Lemma 3

Proof. Define the mean restoration time under g_2, μ_{g_2}, and the mean restoration time under g_2, conditional on restoration occurring after time t, $\mu^t_{g_2}$, as in the proof of Proposition 3. Define $\Pi_1(1)$ and $\Pi_1(t_1)$ as in (6) and (7). Note that for $t_1 \in \mathcal{T}\backslash\{1\}$, $\Pi_1(1) > \Pi_1(t_1)$ if and only if

$$\begin{aligned}
0 < {} & \lambda\pi_0 + (1-\lambda)\left(\mu_{g_2}(\pi_0 - \pi_2) + \pi_2\right) - \lambda\left(t_1\pi_0 + (1-t_1)\pi_1\right) \\
& - (1-\lambda)\sum_{\tau\in\mathcal{T}}\left(\tau\pi_0 + (t_1-\tau)\pi_2 + \varepsilon\pi_1 + (1-t_1-\varepsilon)\pi_3\right)g_2(\tau) \\
& + (1-\lambda)\sum_{\tau\in\{t_1+\varepsilon,\ldots,1-\varepsilon\}}\left(\tau\pi_0 + (t_1-\tau)\pi_2 + \varepsilon\pi_1 + (1-t_1-\varepsilon)\pi_3\right)g_2(\tau) \\
& + (1-\lambda)\left(\pi_0 + (t_1-1)\pi_2 + \varepsilon\pi_1 + (1-t_1-\varepsilon)\pi_3\right)g_2(1) \\
& - (1-\lambda)\sum_{\tau\in\{t_1+\varepsilon,\ldots,1-\varepsilon\}}\left(t_1\pi_0 + (\tau+\varepsilon-t_1)\pi_1 + (1-\tau-\varepsilon)\pi_3\right)g_2(\tau) \\
& - (1-\lambda)(t_1\pi_0 + (1-t_1)\pi_1)g_2(1).
\end{aligned}$$

Rewriting the sums to be inclusive of 1, we have

$$
\begin{aligned}
0 &< \lambda\pi_0 + (1-\lambda)\left(\mu_{g_2}(\pi_0-\pi_2)+\pi_2\right) - \lambda\left(t_1\pi_0+(1-t_1)\pi_1\right) \\
&- (1-\lambda)\sum_{\tau\in\mathcal{T}}\left(\tau\pi_0+(t_1-\tau)\pi_2+\varepsilon\pi_1+(1-t_1-\varepsilon)\pi_3\right)g_2(\tau) \\
&+ (1-\lambda)\sum_{\tau\in\{t_1+\varepsilon,\ldots,1\}}\left(\tau\pi_0+(t_1-\tau)\pi_2+\varepsilon\pi_1+(1-t_1-\varepsilon)\pi_3\right)g_2(\tau) \\
&- (1-\lambda)\left(\pi_0+(t_1-1)\pi_2+\varepsilon\pi_1+(1-t_1-\varepsilon)\pi_3\right)g_2(1) \\
&+ (1-\lambda)\left(\pi_0+(t_1-1)\pi_2+\varepsilon\pi_1+(1-t_1-\varepsilon)\pi_3\right)g_2(1) \\
&- (1-\lambda)\sum_{\tau\in\{t_1+\varepsilon,\ldots,1\}}\left(t_1\pi_0+(\tau+\varepsilon-t_1)\pi_1+(1-\tau-\varepsilon)\pi_3\right)g_2(\tau) \\
&+ (1-\lambda)\left(t_1\pi_0+(1+\varepsilon-t_1)\pi_1-\varepsilon\pi_3\right)g_2(1) \\
&- (1-\lambda)(t_1\pi_0+(1-t_1)\pi_1)g_2(1).
\end{aligned}
$$

Combining the two sums that run from $t_1+\varepsilon$ to 1, using $\mu_{g_2}=\sum_{\tau\in\mathcal{T}}\tau g_2(\tau)$, and simplifying, we have

$$
\begin{aligned}
0 &< \lambda(1-t_1)(\pi_0-\pi_1)-(1-\lambda)(1-t_1)(\pi_3-\pi_2)+\varepsilon(1-\lambda)(\pi_3-\pi_1)\left(1-g_2(1)\right) \\
&+ (1-\lambda)\sum_{\tau\in\{t_1+\varepsilon,\ldots,1\}}\left(\pi_0-\pi_1+\pi_3-\pi_2\right)(\tau-t_1)g_2(\tau)
\end{aligned}
\tag{A.2}
$$

Let $t_1\in\mathcal{T}\backslash\{1\}$, be given.

Case 1. $G_2(t_1)=1$. In this case, for all $\tau\in\{t_1+\varepsilon,\ldots,1\}$, $g_2(\tau)=0$, and so inequality (A.2) reduces to

$$
\lambda\frac{\pi_0-\pi_1}{\pi_3-\pi_2}+\varepsilon\frac{1-\lambda}{1-t_1}\frac{\pi_3-\pi_1}{\pi_3-\pi_2}>1-\lambda,
$$

which holds if $\pi_0-\pi_1>\pi_3-\pi_2$ and $\lambda\geq 1/2$ because then the left side is greater than λ and the right side is less than or equal to λ.

Case 2. $G_2(t_1)<1$. In this case, using $\mu_{g_2}^t(1-G_2(t))=\sum_{\tau\in\{t+\varepsilon,\ldots,1\}}\tau g_2(\tau)$ and rearranging, including dividing through by $\pi_3-\pi_2$, which is positive, we can rewrite (A.2) as

$$
\begin{aligned}
0 &< \left[\lambda(1-t_1)+(1-\lambda)(\mu_{g_2}^{t_1}-t_1)(1-G_2(t_1))\right]\frac{\pi_0-\pi_1}{\pi_3-\pi_2} \\
&- (1-\lambda)(1-t_1)+(1-\lambda)(\mu_{g_2}^{t_1}-t_1)(1-G_2(t_1))+(1-\lambda)\frac{\pi_3-\pi_1}{\pi_3-\pi_2}\left(1-g_2(1)\right)\varepsilon.
\end{aligned}
$$

Using $t_1<1$ and $\mu_{g_2}^{t_1}>t_1$, the expression above in square brackets is positive for all $\lambda\geq 1/2$, so we can rearrange the inequality above as

$$\frac{\pi_0-\pi_1}{\pi_3-\pi_2} + \varepsilon \frac{\frac{\pi_3-\pi_1}{\pi_3-\pi_2}(1-\lambda)(1-g_2(1))}{\lambda(1-t_1)+(1-\lambda)(\mu_{g_2}^{t_1}-t_1)(1-G_2(t_1))} > \frac{(1-\lambda)(1-t_1)+(1-\lambda)(\mu_{g_2}^{t_1}-t_1)(1-G_2(t_1))}{\lambda(1-t_1)+(1-\lambda)(\mu_{g_2}^{t_1}-t_1)(1-G_2(t_1))},$$

which holds if $\pi_0 - \pi_1 > \pi_3 - \pi_2$ and $\lambda \geq 1/2$ because in that case the left side is greater than 1 and the right side is less than or equal to 1 (and equal to 1 when $\lambda = 1/2$).

Combining the two cases above, we have shown that if $\pi_0 - \pi_1 > \pi_3 - \pi_2$, then $\Pi_1(1) > \Pi_1(t_1)$ holds for all $t_1 \in \mathcal{T}\backslash\{1\}$, g_2, and $\lambda \geq 1/2$, establishing that $(1, 1)$ is globally risk dominant. Further, if $\pi_0 - \pi_1 < \pi_3 - \pi_2$, then for $\lambda = 1/2$ and g_2 with all weight on 0, so that $G_2(0) = 1$, then we are in Case 1 above, but with $\frac{\pi_0-\pi_1}{\pi_3-\pi_2} < 1$, and so the inequality $\Pi_1(1) > \Pi_1(t_1)$ does not hold for ε sufficiently close to 0. $\square$

A.6 Proof of Lemma 4

Proof. Let firms 1 and 2 be uninformed and firms $3, \ldots, m$ be informed. Consider a mixed strategy by firm 2 of $\hat{g}_2(\cdot; \lambda, 1)$. Firm 1's expected payoff from restoring at time 1 is then $\Pi_1(1)$ as given in (6). Now consider what happens if firm 1 restores at time $t_1 \in \mathcal{T}\backslash\{1\}$. Because the informed firms expect both uninformed firms to restore at time 1, they optimally do not restore until time 1 or until after observing that both uninformed firms have restored. Thus, firm 1's expected payoff from restoring at time $t_1 \in \mathcal{T}\backslash\{1\}$ is bounded above by its payoff when all of the informed firms restore one period after the restoration of the second of the uninformed firms to restore (or at time 1), which is given by $\Pi_1(t_1)$ in (7). It then follows from Lemma 3 that for $\varepsilon > 0$ sufficiently small, $\Pi_1(1) > \Pi_1(t_1)$ for all $t_1 \in \mathcal{T}\backslash\{1\}$, g_2, and $\lambda \geq 1/2$ if and only if $\pi_0 - \pi_1 > \pi_3 - \pi_2$, completing the proof. $\square$

B Supplemental empirics

B.1 Independents follow majors during restoration phases

Figure B.1 reproduces Figure 2 from the paper but also includes daily average prices among non-major retailers in the market. In line with our discussion of payoffs and interpretation of our model in Section 3.1, we see that independents tend to restore their prices *after* the major retailers consolidate a restoration, thus providing an incentive for major retailers to complete a restoration.

Figure B.1: Price cycle examples from 2015 including independent stations

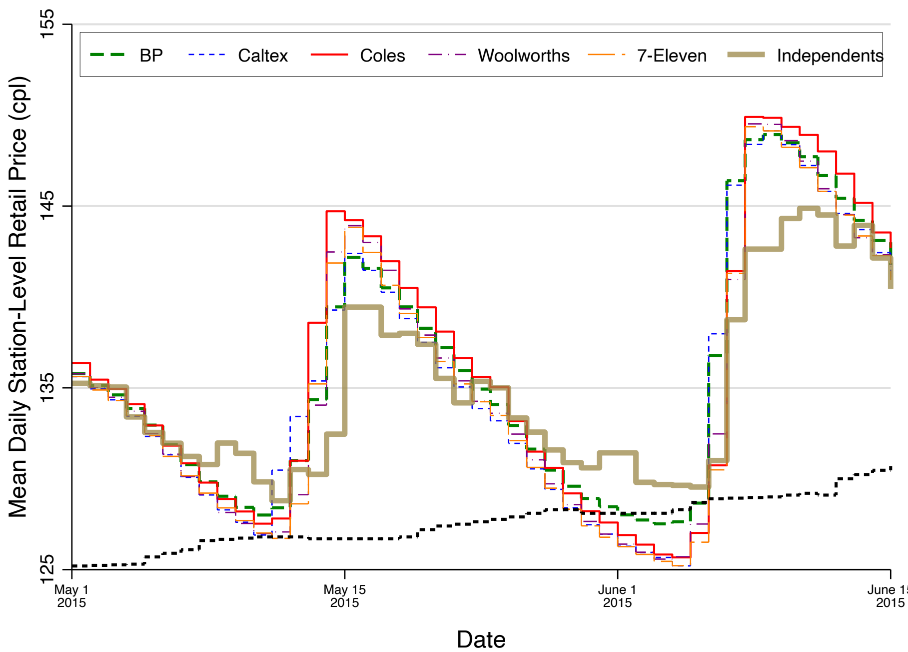

B.2 Persistence in restoration consolidation speed

Figure 4 illustrates persistence in restoration consolidation speed over time between cycles. Here, we formally test for temporal dependence by estimating an AR(1) model,

$$speed_t = \alpha_0 + \alpha_1 speed_{t-1} + \epsilon_t,$$

where $speed_t$ is our measure of restoration consolidation speed from Figure 4, which recall is the number of days it takes market restoration phase t to transition from market-level average prices being at a local minimum to a local maximum. Here, t enumerates, over time, the count of market restorations in our sample (see Section 4.2 in the paper for how we classify market restorations). We estimate $\hat{\alpha}_1 = 0.28$ with a standard error of 0.10, implying statistically significant persistence in restoration speed ($p = 0.016$). Quantitatively, the model estimates imply that a one-standard-deviation reduction in $speed_t$ of 1.5 days predicts a 0.42-day reduction in $speed_{t+1}$, which compares to a sample mean for $speed_t$ of 5.5 days.

B.3 Retailer-specific restoration timing estimates

Figure B.2 provides retailer-specific plots for BP, Caltex, Coles, and Woolworths, highlighting how their station-level restoration probabilities evolve before and after Coles exits the Informed Sources platform. Specifically, they report the coefficient estimates from a variation on (4), where we allow for retailer-specific heterogeneity in the β and γ coefficient through an appropriate combination of interaction terms in the regression. These graphs show minimal change in these retailers' restoration profiles in contrast to Coles' large shift toward leading restoration consolidation in Figure 5(a) after Coles exits the platform.

Figure B.2: Restoration Timing for Rivals Before and After the Case Settlement

(a) BP

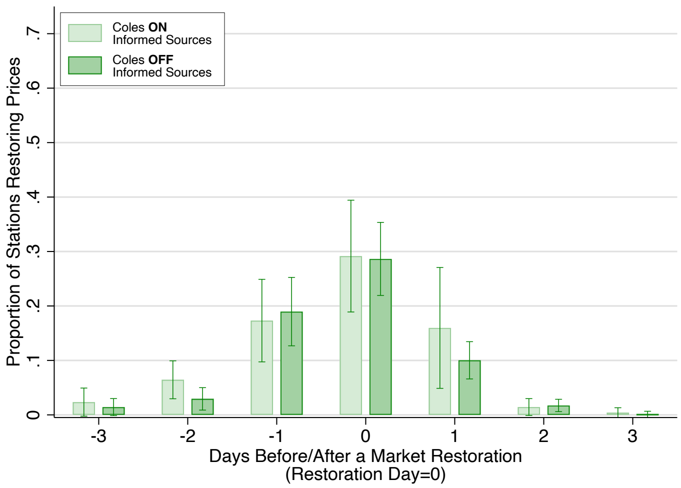

(b) Caltex

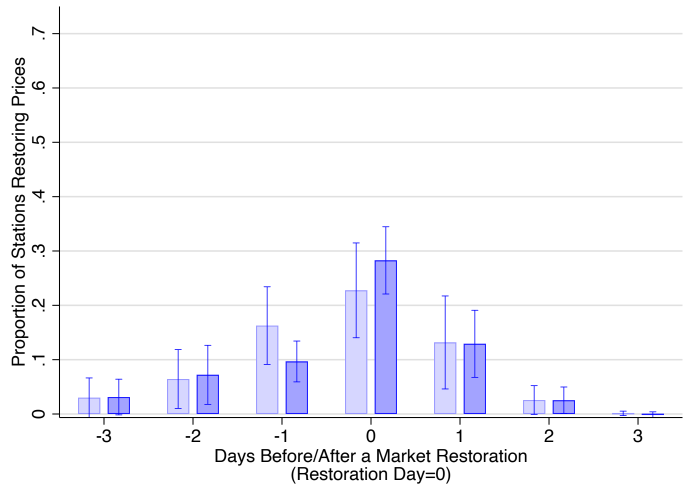

(c) Woolworths

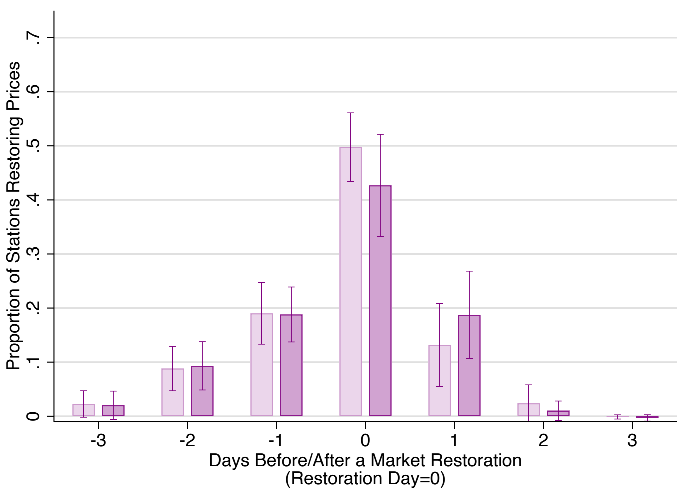

(d) 7-Eleven

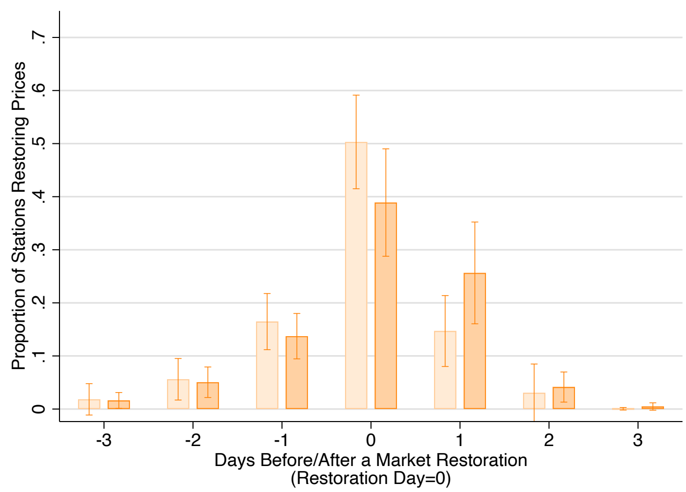

Note: The figures contain β_j (light shade) and $\beta_j + \gamma_k$ (dark share) estimates and their 95% confidence intervals. Standard errors are two-way clustered by station and date. Panels (a)–(d), respectively, show results for BP, Caltex, Woolworths, and 7-Eleven stations. The overall lower probability that a given BP or Caltex station engages in a restoration within a market restoration window reflects that independent stations in the market also operate these brand names that do not always cycle their prices.

B.4 Temporal Variation in Restoration Consolidation Leadership

This appendix explores temporal variation in the station-level restoration activity across retailers that underlies Figure 5. Figure B.3 examines the evolution in restoration consolidation leadership across retailers over time. Specifically, panels (a)–(e) in the figure plot, by retailer, the share of stations that have restored prices on dates that are either the *day before*, *day of*, or *day after* market restorations. The key comparison in each panel is between the solid black line with triangles (shares of stations "going with the market" in consolidating restorations) and the dashed magenta line with squares (shares of stations "leading the market" in consolidating restorations), and how these lines evolve before and after Coles exits the Informed Sources platform (indicated by the vertical dashed lines in each of panels (a)–(e)). Looking across the panels, only Coles exhibits opposing level shifts in these two lines, with the rise in the magenta dashed line after Coles exits the platform illustrating a systematic shift toward restoration consolidation leadership. All other retailers' plots of restoration timing are comparably stable. This difference-in-difference in the timing of station-level restoration shares underlies the estimates in the paper in Figure 5.[48]

Figure B.4 provides a complementary look at restoration consolidation leadership by *ranking* the retailers in terms of their share of station-level restorations around market restoration dates. Panel (a) plots this share rank on the *days before* market restoration dates, with the vertical dashed line again indicating when Coles exits the Informed Sources platform. We find a clear permanent shift from a rank of 5 to a rank of 1 among Coles stations before and after the retailer exits the platform. Further, panel (b) reveals a simultaneous shift by Coles stations away from "going with the market" in consolidating restorations after Coles exits the platform. The other retailers do not exhibit such dual persistent shifts in their restoration share ranks before and after Coles exits the platform, underlining a particular shift toward restoration leadership by Coles after it exits the platform.

[48] Panels (a)–(e) also highlight non-negligible variation in restoration shares on days around market restorations over time. While retailers exhibit a degree of systematic behavior in leading or going with market restorations over time, it is far from deterministic and involves noise. Such randomness in the timing of restorations can reflect, among other things, wholesale cost shocks as firms reach the bottom of the cycle. Positive shocks can induce retailers to quickly restore prices, while negative shocks can induce retailers to delay restorations, and the degree to which different retailers respond differently to these shocks over time can introduce noise into the retailers' relative timing in consolidating restorations.

Figure B.3: Station Restoration Shares Around Market Restoration Dates by Retailer

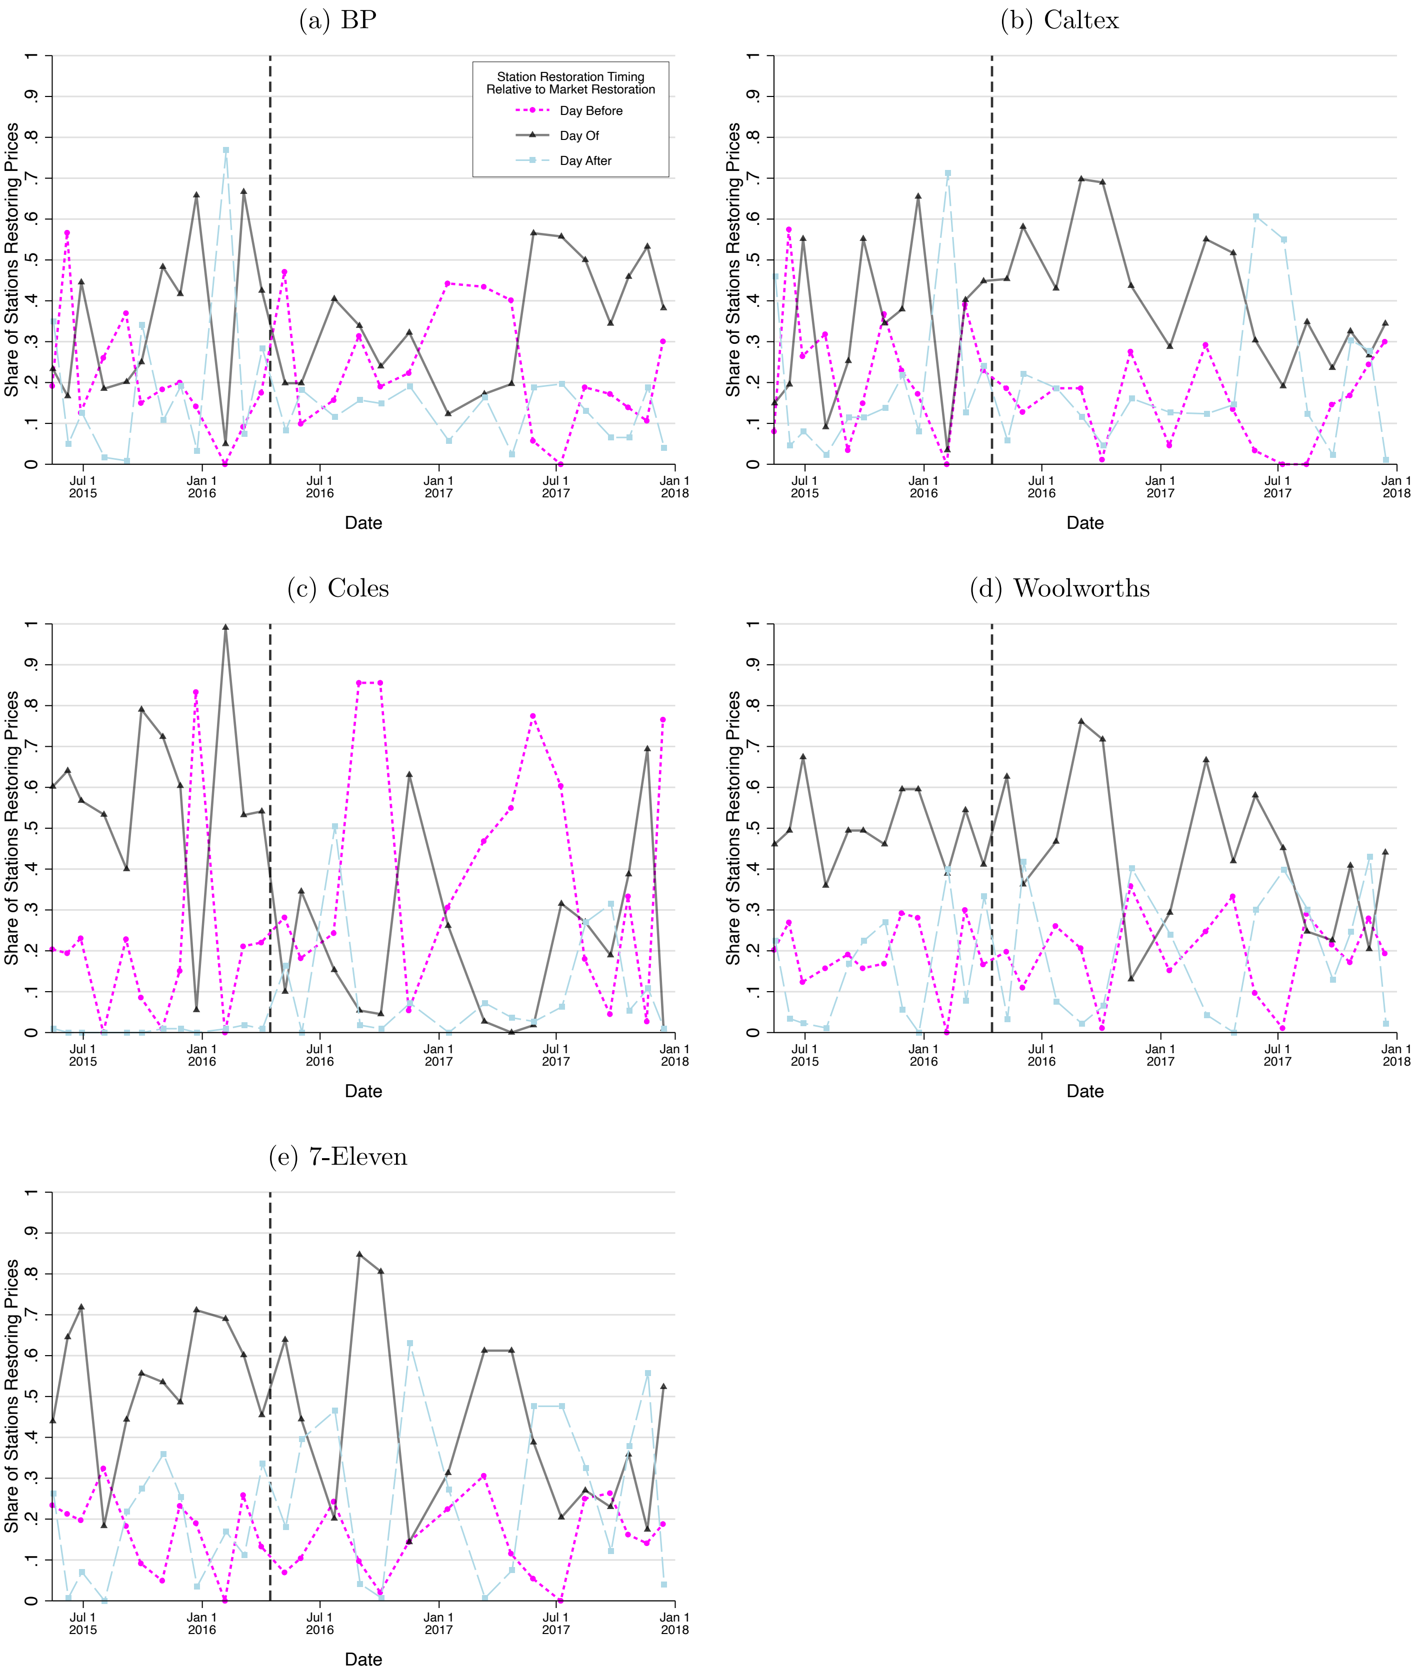

Note: The figures contain, by market restoration episode and retailer, the share of stations that have restored their prices one day before, the day of, and the day after market restorations. The horizontal dashed line in each figure marks April 15, 2016, when Coles exits the Informed Sources platform.

Figure B.4: Station Restoration Share Rank by Retailer Around Market Restorations Dates

(a) **Day Before** Market Restorations

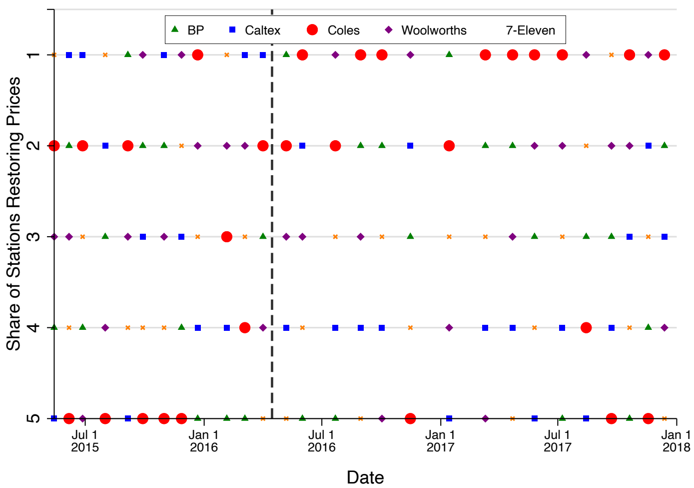

(b) **Day Of** Market Restorations

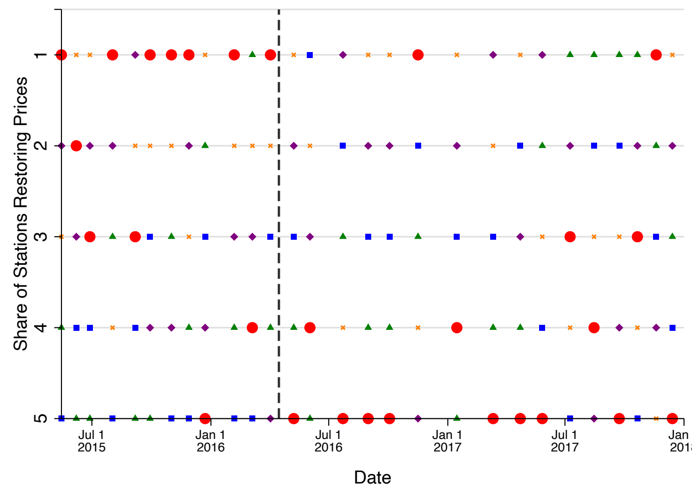

(c) **Day After** Market Restorations

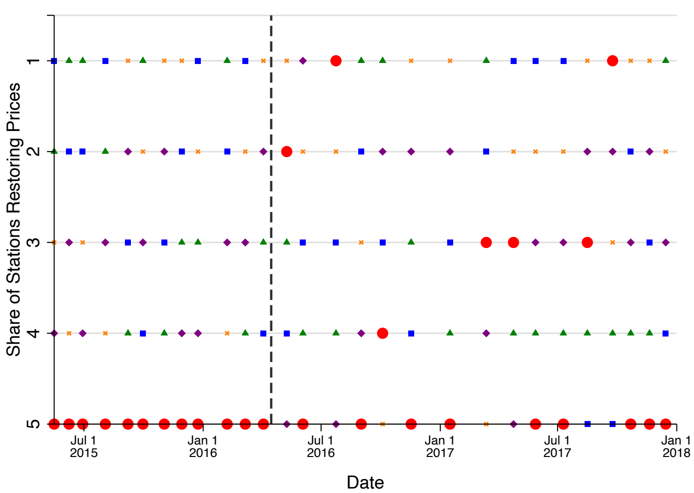

Note: The figures contain, by market restoration episode, the rank order of major retailers in terms of their share of stations having restored prices. Respectively, panels (a)–(c) plot this rank order the day before, day of, and day after market restorations. The horizontal dashed line in each figure marks April 15, 2016, when Coles exits the Informed Sources platform.