

Rockets, Feathers, and Delays: Gasoline Price Cycles under Imperfect Monitoring

Ayush Gupta

30 August 2025

[Please click here for the latest version](#)

Abstract

This paper studies the “rockets-and-feathers” price cycles which have been extensively documented in retail gasoline markets worldwide. We model a repeated game between two price-setting firms who monitor each other via randomly updated price signals. We introduce an Ex-Post Incentive Compatibility (XIC) condition which strengthens sequential rationality by requiring each period’s equilibrium action to be optimal for every on-path signal realization. XIC eliminates ex-post regret and ensures that there are no on-path punishments. We provide a full characterization of XIC-consistent price paths: either prices become stationary after finitely many periods, or they exhibit a cyclic pattern which mirrors observed gasoline cycles. Using a constructive example, we show that cycles expand the set of feasible equilibrium payoffs relative to any stationary cartel price. Price cycles therefore function as an implicit transfer mechanism, allowing intertemporal redistribution of surplus when explicit side payments are not possible. Our results are supported by evidence that cartels use adjustment delays to transfer surplus, and suggest a broader role for XIC as a behaviourally grounded selection criterion in dynamic environments with imperfect monitoring and accountability concerns.

1 Introduction

Retail gasoline prices in many cities exhibit a distinctive and persistent “rockets and feathers” pattern. A large and synchronous increase in pump prices is followed by many small and asynchronous decreases until prices approach marginal cost and the pattern repeats. Often called price cycles, this phenomenon has been documented in Perth, Vienna, Toronto, and many other markets despite substantial heterogeneity in taxation, regulation, and market concentration. Existing theories of competition and collusion struggle to offer a satisfactory explanation for these price cycles. Text-book Bertrand competition predicts marginal cost pricing, whereas collusion suggests a stationary cartel price. Competitive behaviour in the tradition of Maskin and Tirole (1988) can generate cyclical prices, but only in knife-edge cases with carefully chosen parameters. More general folk theorem constructions accommodate almost any dynamic, leaving unanswered the critical question of why this particular trajectory emerges.

This paper develops a model of collusion in a repeated game with imperfect monitoring to demonstrate why price cycles may emerge in equilibrium. We impose a behavioural refinement which ensures that there is no ex-post regret and rules out equilibria with on-path punishment. This refinement constrains equilibrium outcomes to two possibilities: either the same prices are observed in every period or prices exhibit a cyclic pattern like the one observed in retail gasoline markets. We show that price cycle equilibria allow the players to achieve payoff vectors which cannot be achieved in any fixed price equilibrium. Consequently, price cycles expand the set of feasible payoff vectors because they serve as a mechanism for transferring surplus from one player to the other. We therefore posit that price cycles emerge precisely when players wish to redistribute the surplus from collusion but cannot make explicit side payments.

We call our equilibrium refinement Ex-Post Incentive Compatibility (XIC). It requires that each player’s equilibrium action in period t is optimal for every possible realization of the imperfect signal which is observed at the end of period t . In other words, even if players are told the future realization of the public signal, no player will ever want to deviate from their equilibrium action. The XIC

refinement ensures that, on the path of play, players never wish to go back in time and change their past actions. Intuitively, this means that players never regret playing their equilibrium action. This refinement is motivated by the fact that anticipated regret is often an important consideration when human decision makers are facing uncertainty. Filiz-Ozbay and Ozbay (2007); Strack and Viefers (2021); Jhunjhunwala (2021); Fioretti et al. (2022) and many other studies (in economics and other fields) have documented that agents may prioritize regret minimization over expected utility maximization. In our context, this seems especially likely since the decision maker may be an employee, rather than the owner of the gas station. Employees often have to justify their choices ex-post, so they will naturally prefer equilibria under which their action is certain to appear optimal in hindsight.

The XIC refinement significantly strengthens standard sequential rationality which only requires that the expected payoff from deviating is less than the *expected* payoff on-path. In comparison, we require that the expected payoff from deviating is less than the *worst-case* payoff on-path. Consequently, only a small subset of Perfect Public Equilibria (PPE) satisfy the XIC refinement. In addition to no regret, these equilibria have many properties which are intuitively appealing when players are trying to coordinate in an uncertain environment.

First, equilibria which satisfy the XIC refinement have no punishment on the path of play. As a result, the players never engage in costly price wars similar to the ones documented by Green and Porter (1984). Avoiding such price wars not only increases payoffs but also helps ensure the long-term stability of collusive arrangements. Second, XIC equilibria guarantee that, even in the worst-case scenario, players are compensated for any payoff which they sacrifice by not deviating from their equilibrium action. This helps ensure continued cooperation when the gains from cooperating are uncertain and the players have different and / or unknown subjective beliefs about the underlying uncertainty. Third, any equilibrium which satisfies the XIC refinement is robust to adversarial realizations of the imperfect signal. In particular, unlike in standard folk theorem constructions, XIC equilibria do not require the punishments to be finely tuned for the specific monitoring technology.

To model the retail gasoline market, we use a standard repeated game with two heterogeneous firms that set prices. We introduce a realistic form of imperfect monitoring under which each firm observes its rival's price through a randomly updated channel. Formally, at the end of period t station i observes a signal y_j^t which equals the rival's actual price p_j^t with probability $1 - \gamma$ and repeats y_j^{t-1} with probability $\gamma \in (0, 1)$. The monitoring technology therefore obscures timing but not actions – stations are certain that every observed price was posted at some point in the past but they do not know when. We argue that this information structure is a good approximation of the informational constraints faced by actual gas stations since gas station managers often monitor their rivals using web-scraping applications and online data sources which are typically not updated in real time.

Under this monitoring technology, the XIC refinement requires that delaying a price change is never profitable for either firm (Theorem 1). More specifically, if firm i is required to change its price between periods t and $t + 1$, then the new price must be more profitable than the old price for firm i . Intuitively, this is because the price signal might fail to update and display the old price even when firm i has in fact changed its price. If firm i has an incentive to maintain the old price, incentive compatibility requires firm j to initiate punishment whenever the signal does not update. This results in on-path punishment and violates the XIC refinement.

Lemma 1 shows that, while price decreases are generally compatible with the XIC refinement, price increases are possible only when the old price is sufficiently close to marginal cost. That is because the demand facing an individual gas station is very sensitive to the price. Even a small increase typically results in a large decrease in demand. Consequently, price increases are profitable only when the old margin is so close to zero that the increased margin compensates for the decreased demand.

Lemma 2 establishes a sharp restriction on continuation play whenever the observed prices are high. Either the prices will remain fixed at this high level for all subsequent periods or at least one firm will reduce its price. Holding prices constant preserves the existing payoff profile indefinitely

so the only way to alter the distribution of surplus is a price cut. As a result, in every sub-game following high observed prices, play must take one of two forms: a stationary price equilibrium in which prices never change, or a phase of decreasing prices initiated by at least one firm.

Lemmas 1 and 2 thus rationalize the key features of the price cycles observed in retail gasoline markets. Asynchronous price cuts are necessary to alter the distribution of flow payoffs – otherwise both firms will receive the same payoff in each period. Avoiding regret or on-path punishment requires that the firms raise their price only when the old price is sufficiently close to marginal cost. The observed price cycles are therefore not coincidental. They emerge as the only way for firms to reallocate surplus over time while satisfying the behavioural restrictions imposed by the XIC refinement.

Theorem 2 uses these results to further characterize the entire set of price paths which may be observed in an equilibrium satisfying the XIC refinement. One possibility is that the vector of prices becomes constant after some finite number of periods. The only other possibility is that the prices exhibit a cyclic pattern which is a more version of the observed price cycles. Formally, this cyclic pattern consists of countably many blocks of finite length. Within each block, prices are non-increasing and every block ends with a period of low prices. Crucially, all price increases occur at the transition between blocks, when prices are sufficiently low. So any sequence of prices which fits this cyclic pattern will closely resemble gasoline price cycles.

Example 1 provides an explicit construction of strategies which satisfy the XIC refinement and generate the price cycles observed in retail gasoline markets. Although the resulting payoffs are stochastic, Theorem 2 shows that even in the worst-case outcome, the realized payoffs are not Pareto dominated by any stationary price equilibrium – defined as an equilibrium in which the prices do not change over time. Cyclical prices thus expand the set of equilibrium payoffs. One interpretation is that price cycles act as a compromise between alternative stationary equilibria: switching between periods with different relative prices allows the firms to achieve payoff vectors which lie between two extremes. We therefore say that price cycles are a mechanism for redistributing the gains from

cooperation without explicit side payments. Temporary phases of lower prices serve as implicit transfers, shifting demand and profits toward one firm.

Proposition 2 compares the equilibrium in Example 1 to all other equilibria which satisfy the XIC refinement. It shows that there is no equilibrium (stationary or otherwise) which can guarantee a higher payoff for both firms. That is because the only way to ensure a higher payoff for one firm is by changing the timing of price cuts, but such adjustments necessarily come at the expense of the other firm. Proposition 2 thus establishes that the equilibrium in Example 1 and, by extension the observed price cycles, lie on the frontier of feasible guarantees – every increase in the guaranteed payoff for one firm must come at the expense of the other.

When studying cooperation between price setting firms, it is natural to focus on equilibria in which both firms post the same price. However, when the firms are heterogeneous (as is generally the case in retail gasoline markets), one of the firms may demand a larger share of the gains from cooperation.¹ Existing papers typically sidestep this issue by allowing for cash transfers or assuming that the firms can split demand in any ratio they want. However, neither assumption is reasonable in our setting. Beyond their choice of posted price, gas stations have practically no control over where a consumer chooses to purchase gasoline. And since explicit collusion is illegal in most markets, direct transfers carry an increased risk of detection and prosecution. The ability to transfer surplus via price cycles is thus valuable and can help sustain cooperation between the firms.

The argument that price cycles are a mechanism for redistributing the gains from cooperation is supported by Clark and Houde (2013) who document court evidence from a price fixing case involving gas stations in four cities in Quebec, Canada. When the cartel was investigated by Canadian antitrust authorities, witness testimony revealed that the cartel members used adjustment delays during price changes to transfer surplus. Clark and Houde (2013) also used data on prices

¹This paper assumes that the two firms have different marginal costs. The firm with the lower cost therefore has a better outside option and more bargaining power in a Nash Bargaining setup. While a formal model of bargaining between the firms is beyond the scope of this paper, we claim that it is intuitive for the firm with the lower cost to demand a larger share of the surplus from cooperation.

and station characteristics to estimate the value of the resulting transfers. They show that this mechanism can be used to transfer substantial amounts and can significantly reduce the frequency of deviations.

This paper thus contributes to several distinct but related strands of the literature. First and foremost, we build on the rich empirical literature that has meticulously documented and analyzed price cycles across numerous international markets. Foundational work in this area has established the key stylized facts that any successful theory must explain: the “rockets and feathers” asymmetry; adjustment delays during the undercutting phase; and the puzzling observation that retail prices often change independently of wholesale costs. While this body of work provides compelling and robust evidence of the phenomenon, the underlying theoretical mechanism remains a subject of active debate. Our primary contribution here is to offer a novel micro-foundation for these cycles, demonstrating how they can emerge as a robust and Pareto efficient equilibrium outcome of strategic interaction between competing firms under a realistic form of imperfect monitoring.

Second, our work engages with the classic literature on tacit collusion and price wars in settings with imperfect monitoring. By introducing the XIC refinement, we select for equilibria which are sustained without recourse to costly on-path punishment phases. Under our refinement, a “bad” signal does not trigger a price war because the equilibrium is constructed to be robust to all possible on-path signal realizations. The XIC refinement thus offers a compelling explanation for the observed data while also being an intuitive restriction on human behaviour. In contrast, more general folk theorem constructions often rely on finely calibrated punishments and assume that the players are hyper rational agents who only wish to maximize expected utility.

Finally, we contribute to the literature on equilibrium refinements and behavioural game theory, specifically how anticipated regret influences strategic choice. The idea that decision-makers are motivated to avoid the ex-post feeling of having made a sub-optimal choice has a long history in economics. Our XIC refinement formalizes this concept within a dynamic game using a novel ex-post stability condition. While this paper applies the refinement to retail price cycles, its usefulness

extends to a much broader class of problems. The XIC condition provides a powerful selection criterion in any dynamic game where cooperation must be sustained under uncertainty and the players desire robust, stable agreements. The XIC refinement selects for agreements which do not require costly, hair-trigger punishments in response to ambiguous signals. The concept is also valuable for analyzing any strategic setting where trust is paramount and agents must justify their actions in hindsight. The refinement extends readily to games with imperfect private monitoring, making it a flexible tool for understanding cooperation in a wide range of economic and political environments.

2 Description of Price Cycles

Before describing our model, we use this section to establish some motivating facts about price cycles. We begin by emphasizing that retail gasoline is a largely homogeneous good. Consequently, individual gas stations face an extremely high own-price elasticity of demand. Even small differences in prices often cause very large differences in market share.

Figure 1 (from Byrne et. al. [2024]) provides an illustrative example of a price cycle from Melbourne in 2015. As one can see, price cycles are characterized by two distinct phases. In the undercutting phase, stations decrease their prices asynchronously and in small increments. Notably, some stations act as price leaders by lowering their prices earlier and consequently enjoy elevated market shares for a non-trivial amount of time. This is puzzling because, in most jurisdictions, competing stations face minimal restrictions (practical or regulatory) on how frequently they can update their own price. Stations can, and frequently do, adjust posted prices multiple times in a single day. Given these circumstances, we expect rapid price matching. Yet, in practice, we observe that competing stations often delay their responses, allowing the price leaders to temporarily benefit from their lower prices. This behaviour implies the existence of strategic considerations beyond simple short-run profit maximization.




Figure 1: Example of a Price Cycle

In stark contrast, the restoration phase occurs abruptly and simultaneously. After a sustained period of incremental price cuts, retail prices collectively jump back to a higher common level. Crucially, this coordinated price restoration occurs only after margins approach marginal cost. Intermediate levels of margins typically fail to trigger price increases. This distinctive timing suggests that restorations are not reactions to sudden changes in supply or demand conditions. Indeed, empirical analyses show that restorations frequently occur even during stable or declining wholesale prices, implying that external cost shocks or demand shifts cannot adequately explain these sudden price hikes. Moreover, given the frequency and regularity of these cycles, it seems highly unlikely that they are driven by *random* shocks.

2.1 Existing Theories

Given the high elasticity of demand, it is natural to expect that prices in a competitive market will be close to marginal cost. However, that is not the case in the retail gasoline market. As Figure 1 illustrates, retail prices are frequently much higher than the wholesale price and changes in the

two prices are often uncorrelated.

Maskin and Tirole (1988) showed that competitive behaviour can result in price cycles. But they assumed that firms can only update their price in alternating periods. This ensures that a price cut will result in at least one period of extra market share and motivates the firms to undercut their rival. However, as we have argued, gas stations are typically not constrained in how frequently they can update their price. In the extreme case, when firms can respond instantaneously, no firm has an incentive to undercut its rival because they know that their new price will be matched before they can enjoy any additional market share. It thus seems unlikely that price cycles are the result of competitive firms undercutting each other in an attempt to gain market share.

Additionally, *changes* in marginal cost tend to be highly correlated across gas stations. This correlation arises because the largest component of marginal cost is the wholesale price of gasoline, which is itself highly correlated with global oil prices. Thus, while different stations may have different marginal costs due to differences in supply agreements or ownership structures, we expect the relative costs to remain stable over short and medium time horizons. We can therefore exclude any theory which relies on unobserved changes in marginal cost. In particular, a decrease in the marginal cost may explain why a gas station lowers its price but it cannot explain why its rivals delay matching the new price.

Alternately, one might suspect that price cycles are the result of cooperative or collusive agreements between gas stations. Indeed, there are some documented examples of cartels yielding cyclical prices. However, if that is the case, it is natural to wonder why the cartel members choose to coordinate on a price cycle. Intuitively, we expect that a cartel could increase total profits by maintaining a stable price which is consistently higher than marginal cost. The extended periods of low prices, which are common in a price cycle, would appear to be sub-optimal if the objective is to maximize joint profits.

One possible explanation is that price cycles reflect recurring price wars caused by imperfect monitoring. However, in standard models with on-path price wars, the firms ability to maintain high

prices is restricted by the richness of the signal space.² Our model also features imperfect monitoring but because each firm's action (the posted price) is observed independently, we expect standard folk theorem arguments to hold and for more profitable equilibria to be available to the firms.³ Moreover, the observed price cycles appear too regular, consistent, and predictable to plausibly be attributed to random shocks or sporadic breakdowns in cooperation. Random disruptions or misunderstandings would naturally lead to irregular and unpredictable patterns, characterized by sudden and erratic price adjustments. Yet, the consistency and recurring nature of price cycles suggest deliberate, systemic behaviour rather than random fluctuations. Additionally, genuine price wars tend to manifest as aggressive and rapid price reductions, with firms actively trying to undercut rivals to protect or expand their market share. In such scenarios, one would expect firms to rapidly match or surpass rivals' price cuts, resulting in swift, cascading price declines. However, the observed delays and hesitation among firms in matching competitors' lower prices are inconsistent with such aggressive competition.

Taken together, these facts highlight the puzzle posed by gasoline price cycles. Any satisfactory explanation must account for the fact that prices cycle asymmetrically, with slow and incremental decreases followed by sudden and coordinated increases. Importantly, the explanation cannot hinge upon sticky prices or changes in relative costs. Moreover, the theoretical framework must accommodate markets where profits are highly sensitive to small price differences and where margins fluctuate substantially over relatively short cycles.

²For example in Green and Porter (1984) the firms do not observe their rival's action. They only observe their own demand which is a function of all firms' actions.

³See Fudenberg et al. (1994).

3 Model

3.1 Preliminaries

We study a game with two firms indexed by $i = \{1, 2\}$. Time is discrete and indexed by $t = \{1, 2, \dots\}$. The firms have a common discount rate $0 < \delta < 1$. The marginal cost of firm i is given by c_i . We assume throughout that $c_1 < c_2$.⁴ This assumption reflects the fact that a station's size and ownership structure affects the price it pays for wholesale gasoline.

We define $p_i^t \in P$ to be the actual price posted by firm i in period t . Assume that $P = \{0, \epsilon, 2\epsilon, \dots, \bar{p}\}$ where $\epsilon > 0$ represents the smallest possible change in price.⁵

Our main assumption is that firms do not observe the actual price posted by their rival. Instead, at the end of each period, the firms observe a pair of independent public signals. The signal generating process is given by:

$$y_1^t = \begin{cases} p_1^t & \text{with probability } 1 - \gamma \\ y_1^{t-1} & \text{with probability } \gamma \end{cases} \quad \text{and} \quad y_2^t = \begin{cases} p_2^t & \text{with probability } 1 - \gamma \\ y_2^{t-1} & \text{with probability } \gamma \end{cases}$$

We argue that this information structure reflects the monitoring constraints faced by gas station managers. Public data sources – like price comparison apps – do not update in real time. After a firm changes its posted price, there is typically some delay before the new price appears on the app. Consequently, when a firm attempts to verify its rival's price, there is some positive probability that the observed price is not the current price.

Note that only prices which were actually posted in some prior period can be observed. So when a firm observes a specific price, it is certain that the observed price was posted at some point in the past. What is uncertain is exactly when the price was posted. The monitoring technology thus

⁴Because of this asymmetry in cost, we will sometimes refer to firm 1 as the stronger player and firm 2 as the weaker player.

⁵In the United States, ϵ is typically one-tenth of a cent for gas stations.

obscures the timing of actions but not necessarily the actions themselves.

Also note that this monitoring technology is history dependent. Conditional on the action profile played in period t , the distribution of signals observed at the end of period t is a function of the signal profile which was observed at the end of period $t - 1$. The signals are therefore a state variable which is governed by an action dependent Markov process. This interpretation serves as part of the motivation for our equilibrium refinement which is introduced in the following section.

We assume that each firm's payoff depends on the firm's own action and the realization of the public signals. More specifically:

$$u_i(p_i^t, y_j^t) = D(p_i^t, y_j^t)(p_i^t - c_i)$$

In the retail gasoline market, $D(\cdot)$ represents the demand faced by firm i . Our key assumption is that demand is a function of the actual price posted by firm i and the noisy signal of firm j 's price. This is consistent with a world in which consumers always learn firm i 's actual price before they make a purchase from firm i . But they do not learn firm j 's price and must therefore rely on asymmetric information.

An alternate interpretation is that demand is a stochastic function of actual prices but firms do not observe their own payoffs in real time. So they must rely on the noisy signals to infer their rival's actions and choose their continuation strategy.

To maintain generality, we do not impose a specific functional form on demand. Instead, we only assume the following:

1. $D(p, y)$ is decreasing in p
2. $D(p, y)$ is increasing in y .
3. Fix any y . Then the absolute value of $\frac{D(p, y) - D(p', y)}{D(p, y)} \frac{p}{p - p'}$ is greater than $1 + \Delta$ where $\Delta > 0$

and $p' \neq p$.

4. Fix any $y > c_i$. Then $u_i(y + \epsilon, y) \leq u_i(y, y)$.
5. Fix any $y > c_i + \epsilon$. Then $u_i(y - \epsilon, y) \geq u_i(y, y)$.

Assumptions 1 and 2 are standard in any model with price-setting firms. They require that the demand faced by a firm is decreasing in its own price and increasing in the rival's price.

Assumption 3 requires that demand is always elastic. The additional Δ term is a technical assumption which is necessary only to rule out cases in which the own price elasticity of demand is approaching 1.

Assumptions 4 and 5 are simplifying assumptions which require that matching the rival's price is always better than posting a higher price and that undercutting the rival by ϵ is always better than matching.

The underlying assumption is that consumers are price sensitive but must pay a travel cost to purchase from a gas station which is further away. As a result, when a gas station undercuts its rival, it captures some but not all of its rival's demand. The elasticity of demand is determined by the distribution of consumers and the travel cost. We assume that the density is high (or the travel cost is low) so that the percentage change in demand is always greater than the percentage change in price.

Note that we make no *a priori* assumptions about the aggregate size of the market. Empirical studies indicate that the market demand for gasoline is inelastic in the short run. In such cases, a unilateral price adjustment by one firm primarily results in a reallocation of market share between the two rivals. Our model is sufficiently general to accommodate this scenario, as well as the broader case where price changes influence the total demand faced by the industry.

We define $h^t = \{y^1, y^2, y^3, \dots, y^{t-1}\}$ to be the public history at the beginning of period t . Then for each firm, a public strategy is a measurable mapping from the set of public histories to P .

Throughout this paper, we will restrict our analysis to public strategies.

3.2 Equilibria

Let $g_i(p^t, h^t) = \mathbb{E}_{y_j^t} [u_i(p_i^t, y_j^t)]$ be the expected payoff of firm i when the action profile p^t is played.

Let $v_i^t(s \mid h^t)$ be firm i 's continuation payoff at the beginning of period t when the public history is h^t and the players are using the public strategy profile $s = (s_i, s_j)$. Notice that:

$$\begin{aligned}
v_i^t(s \mid h^t) &= g_i(s(h^t), h^t) \\
&\quad + \delta \mathbb{E}_{h^{t+1}} [g_i(s(h^{t+1}), h^{t+1}) \mid h^t] \\
&\quad \vdots \\
&= \sum_{\tau=t}^{\infty} \delta^{\tau-t} \mathbb{E}_{h^\tau} [g_i(s(h^\tau), h^\tau) \mid h^t] \\
&= g_i(s(h^t), h^t) + \delta \mathbb{E}_{h^{t+1}} [v_i^{t+1}(s \mid h^{t+1})] \\
&= g_i(s(h^t), h^t) + \sum_{h^{t+1}} \delta \Pr[h^{t+1}] v_i^{t+1}(s \mid h^{t+1})
\end{aligned}$$

Note that s is a perfect public equilibrium (PPE) if for every $s'_i \neq s_i$:

$$v_i^t(s'_i, s_j \mid h^t) \leq v_i^t(s \mid h^t)$$

We say that a public history h^t is consistent with the strategy profile s if h^t is observed with positive probability when the firms are playing according to s .

We say that a signal y^t is consistent with s and h^t if y^t is observed with positive probability when the firms are playing according to s and the public history is h^t .

Definition 1. A PPE is ex-post incentive compatible if for every h^t consistent with s and every

y^t consistent with s and h^t :

$$u_i(p_i^t, y_j^t) + \delta v_i^{t+1}(s \mid h^{t+1}) \geq u_i(\hat{p}_i^t, y_j^t) + \sum_{\hat{h}^{t+1}} \delta \Pr[\hat{h}^{t+1}] v_i^{t+1}(s \mid \hat{h}^{t+1}) \quad (\text{XIC})$$

where \hat{h}^{t+1} is the public history at the beginning of period $t + 1$ after firm i deviates and plays $\hat{p}_i^t \neq p_i^t$ in period t . Note that s is a PPE so $v_i^t(s \mid \hat{h}^{t+1}) \geq v_i^t(s'_i, s_j \mid \hat{h}^{t+1})$ after any \hat{h}^{t+1} .

The left hand side of condition (XIC) is firm i 's equilibrium payoff when the public signal y^t is observed. The right hand side is the expected payoff after any deviation in period t . The XIC refinement therefore requires that the equilibrium action is a best response for every possible realization of the public signal. Compare this to a standard PPE which requires only that the equilibrium action is a best response in expectation.

The XIC refinement is an ex-post condition and is partially inspired by the ex-post equilibrium (XPE) defined by Carroll (2024). Both conditions require the period t action to remain optimal after all uncertainty has been resolved. However, there are some subtle but important differences. In particular, the uncertainty in Carroll's setting comes from the realization of future stage games, which are determined exogenously. On the other hand, the uncertainty in our setting comes from the realization of future signals, which depend on endogenous actions.

Our preferred interpretation is that the XIC refinement ensures no regret on the path of play. That is because equilibrium actions are optimal for any possible realization of the public signal. So regardless of which signal is observed, neither firm will wish to go back in time and change its action. Equilibria with this property are clearly desirable when the firms are trying to sustain cooperation in an environment with uncertainty. Intuitively, it helps reassure the firms that continued cooperation is optimal, even after "bad" signal realizations.

There is some evidence in the behavioural economics literature which supports our claim that the XIC refinement is a reasonable restriction on equilibrium behaviour. Several studies have found

that decision makers facing uncertainty often choose to minimize future regret.⁶ This is particularly salient when the decision maker is a manager, rather than the owner, of a gas station. Managers, who might have to justify their choices ex-post, will naturally prefer equilibria under which they won't be penalized for choosing the equilibrium action in the past.

The XIC refinement can also be interpreted as a condition on the worst-case payoffs. It requires that the worst possible payoff in equilibrium is no less than the expected payoff from a deviation. As a result, whenever a firm gives up current period payoff by not deviating, its future payoff is guaranteed to be at least enough to compensate it for this loss. In other words, the gap between the payoff from the equilibrium action and the payoff from the best deviation is always less than the worst-case continuation payoff in equilibrium.

We note that there is a strong asymmetry in the XIC refinement. It compares each individual signal which may be observed on-path to the expected signal realization after a deviation. This asymmetry stems from the fact that firms never deviate in equilibrium. So they never observe what would have happened if they had deviated.

4 Results

Theorem 1. A PPE satisfies the XIC refinement only if:

$$p_i^t \neq y_i^{t-1} \implies u_i(p_i^t, y_j^t) \geq u_i(y_i^{t-1}, y_j^t)$$

for all y_j^t which are consistent with s and h^t .

Proof. Assume that there exists some \tilde{y} such that $u_i(p_i^t, \tilde{y}) < u_i(y_i^{t-1}, \tilde{y})$ and \tilde{y} is consistent with s and h^t .

⁶See for example Filiz-Ozbay and Ozbay (2007); Strack and Viefers (2021); Jhunjhunwala (2021); Fioretti et al. (2022).

Define $h^* = \{h^t, (y_i^{t-1}, \tilde{y})\}$. Note that h^* is consistent with s . That is because $y_i^t = y_i^{t-1}$ with probability γ even when firm i plays the action given by s_i .

Next note that $y_i^t = y_i^{t-1}$ with certainty when firm i deviates and plays $p_i^t = y_i^{t-1}$. So $\hat{h}^{t+1} = h^*$ with certainty when $y_j^t = \tilde{y}$.

Then equation (XIC) requires that:

$$\begin{aligned} u_i(p_i^t, \tilde{y}) + \delta v_i^{t+1}(s \mid h^*) &\geq u_i(y_i^{t-1}, \tilde{y}) + \delta v_i^{t+1}(s \mid h^*) \\ u_i(p_i^t, \tilde{y}) &\geq u_i(y_i^{t-1}, \tilde{y}) \end{aligned}$$

This is a contradiction.

□

Corollary 1. The XIC refinement is satisfied by every PPE with $p_i^t = p_i$ and $p_j^t = p_j$ for all t .

Lemma 1. There exists $\underline{p} \in P$ such that if $y_i^t > \underline{p}$ then $p_i^{t+1} \leq y_i^t$ in every XIC equilibrium.

Proof. Fix any y , any $p > c$, and any $p' > p$.

Assume $y_j^{t+1} = y$, $y_i^t = p$, and $p_i^{t+1} = p'$.

Then the XIC refinement requires:

$$\begin{aligned}
D(p', y)(p' - c) &\geq D(p, y)(p - c) \\
\frac{p' - c}{p - c} &\geq \frac{D(p, y)}{D(p', y)} \\
\frac{p' - p}{p - c} &\geq \frac{D(p, y) - D(p', y)}{D(p', y)} \\
1 &\geq \frac{D(p, y) - D(p', y)}{D(p', y)} \frac{p - c}{p' - p}
\end{aligned} \tag{1}$$

Note that $D(p, y) > D(p', y)$. Then by Assumption 3:

$$\begin{aligned}
\frac{D(p, y) - D(p', y)}{D(p', y)} \frac{p - c}{p' - p} &\geq \frac{D(p, y) - D(p', y)}{D(p, y)} \frac{p - c}{p' - p} \\
&\geq (1 + \Delta) \frac{p' - p}{p} \frac{p - c}{p' - p} \\
&\geq (1 + \Delta) \frac{p - c}{p}
\end{aligned}$$

Therefore inequality (1) requires:

$$\begin{aligned}
1 &\geq (1 + \Delta) \frac{p - c}{p} \\
p &\leq \frac{1 + \Delta}{\Delta} c
\end{aligned}$$

□

Lemma 2. Let y^t be any on-path realization of the public signal profile at the end of period t .

Assume $y_i^t > \underline{p}$ for all i .

Then one of the following holds in any XIC equilibrium:

1. $y_i^\tau = y_i^t$ for all i and all $\tau \geq t + 1$.
2. There exists $\tau \geq t + 1$ such that $y_i^\tau < y_i^t$ for some i .

Proof. Let τ be the first period after period t such that $y_i^\tau \neq y_i^t$ for some i .

If $\tau = \infty$ then we are done.

So consider the case when $\tau < \infty$. By assumption $y^{\tau-1} = y^t$ and $y_i^t > \underline{p}$ for all i .

Then Lemma 1 tells us that $p_i^\tau \leq y_i^{\tau-1}$.

Recall that $y_i^\tau \in \{p_i^\tau, y_i^{\tau-1}\}$ and $y_i^\tau \neq y_i^{\tau-1}$. So $y_i^\tau < y_i^t$ for some i .

□

Theorem 2. Let $\{y^t\}$ be any sequence of signals which can be realized in an XIC equilibrium.

Then either there exists some T such that $y^t = y^T$ for every $t > T$.

Or there exists a strictly increasing sequence of natural numbers $\{\tau_n\}$ such that:

1. $\tau_1 = 1$.
2. y^t is non-increasing when $t \in [\tau_n, \tau_{n+1} - 1]$.
3. There exists i such that $y_i^t \leq \underline{p}$ when $t \in \{\tau_n - 1 : n \geq 2\}$.

Note that y^t is a vector. So y^t non-increasing when $t \in [\tau_n, \tau_{n+1} - 1]$ means $y_i^t \leq y_i^{t-1}$ for all i .

Theorem 2 shows that in any XIC equilibrium which is Pareto efficient, prices will either become constant or they will exhibit the cyclical pattern observed in the retail gasoline market. In the latter case, the sequence of prices can be separated into countably many blocks of finite length. Within each block, prices will be non-increasing and each block will end with a low price followed by a price increase.

5 Efficiency of Price Cycles

Corollary 1 tells us that any PPE with stationary prices satisfies the XIC refinement. Therefore standard results can be used to show that the following is an XIC equilibrium for any sufficiently high δ :

$$p_i^t = p_j^t = p = \arg \max_p u_1(p, p) + u_2(p, p)$$

Such uniform price equilibria are an easy and intuitive way for price setting firms to increase their joint payoffs by colluding. However, stationary price equilibria cannot achieve every vector of payoffs which is feasible in a XIC equilibrium. We have already shown that all non-stationary equilibria must display the cyclical dynamic observed in the retail gasoline market and the remainder of this section will show that at least one such equilibrium lies on the Pareto frontier.

To fix ideas, we will focus on the following strategy profile for the remainder of this section:

Example 1.

$$p_1^t = \begin{cases} p_1^{t-1} - \epsilon & \text{if } p_1 = y_2 \text{ for } k_1 \text{ periods} \\ p^{max} & \text{if } (p_1^{t-1}, y_2^{t-1}) = (p^{min}, p^{max}) \\ p_1^{t-1} & \text{otherwise} \end{cases}$$

$$p_2^t = \begin{cases} y_1^{t-1} & \text{if } p^{min} < y_1 < p_2 \text{ for } k_2 \text{ periods} \\ p^{max} & \text{if } y_1^{t-1} = p^{min} \\ p_2^{t-1} & \text{otherwise} \end{cases}$$

We fix p^{max} and pick p^{min} to be the highest price such that:

$$\begin{aligned}
u_2(p^{\min}, p^{\max}) &\geq u_2(p^{\min}, p^{\min} + \epsilon) \\
u_1(p^{\max}, p^{\max}) &\geq u_1(p^{\min}, p^{\max})
\end{aligned}$$

Remark 1. Example 1 specifies on-path strategies but not off-path punishments. So when we refer to the equilibrium in Example 1, we mean any PPE which supports the on-path strategies described above.

Remark 2. The XIC refinement only constrains actions on the path of play. It imposes no restriction on off-path strategies. It is therefore easy to verify that the equilibrium in Example 1 satisfies the XIC refinement.

Results from Hörner et al. (2011) can be used to show that there exists a PPE which supports the strategies in Example 1 whenever δ is sufficiently high and the min-max payoffs are sufficiently low.

Proof Sketch. First note that our setting satisfies the rank assumptions in Hörner et al. (2011) (Assumptions F1 and F2).

Let y^t be the realization of the public signal at the end of period t .

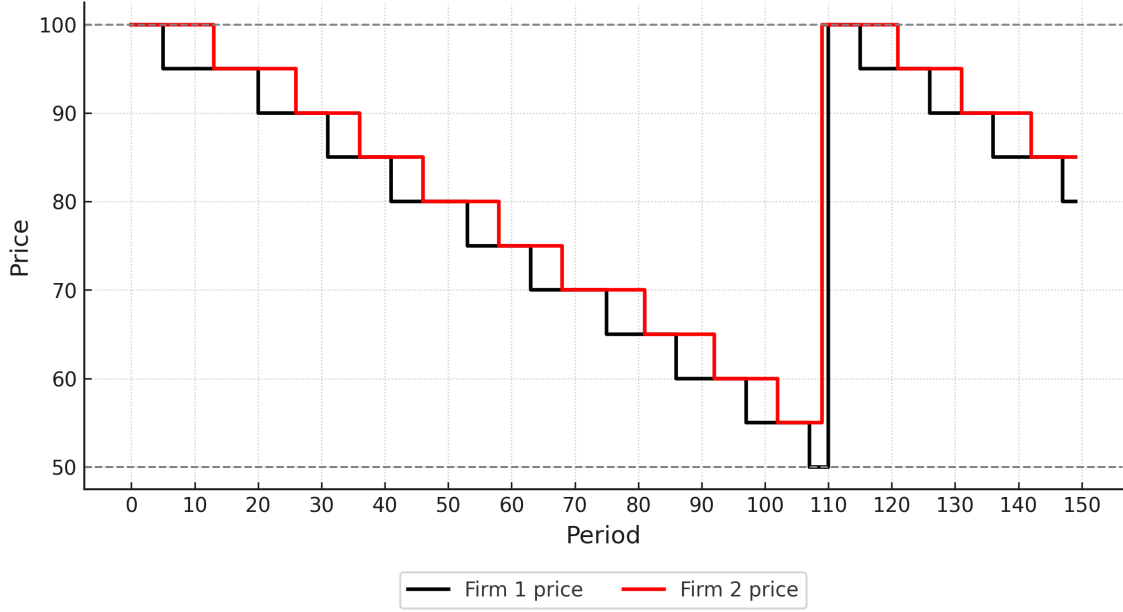
Let $F(y^t)$ be the set of feasible average payoffs in the sub-game starting in period $t + 1$.

It is clear that $F(y^t)$ is independent of y^t in the limit as $\delta \rightarrow 1$.

We can therefore apply Theorem 2 from Hörner et al. (2011). Specifically, for every v in the set of feasible and individually rational payoffs, there exists a PPE of the sub-game starting in period $t + 1$ such that the expected payoffs are v .

We can thus define punishment strategies such that the expected payoffs starting in period $t + 1$ are arbitrarily close to the min-max payoffs if a deviation is detected in period t .

Figure 2: Simulated Sequence of Prices



This allows us to incentivize any action in period t as long as the expected on-path payoffs are strictly greater than the min-max payoffs.⁷

□

Remark 3. An example of appropriate punishment strategies is that both firms respond to any off-path signal by switching to the min-max action profile for n periods. Then for a sufficiently large δ , we can pick n such that neither firm deviates.

The key assumption is that the min-max payoff for both firms is strictly less than the minimum of the on-path continuation payoffs across all periods.

Figure 3 depicts one possible sequence of prices which might be observed when the firms play the equilibrium in Example 1.

⁷The formal argument is essentially identical to the proof of the main result in Fudenberg et al. (1994).

We will now make two additional assumptions about the demand function:

Assumption 6. If $p > c_i$ then for every $p' \in P$:

$$u_i(p, p) > u_i(p', p' - \epsilon)$$

Assumption 7. $D(p, p)(p - c_i)$ is strictly concave and therefore single peaked for both i .

Assumption 6 states that both firms prefer to match their rival's price over posting a higher price. Assumption 7 states that there is a unique price which maximizes each firm's payoff when both firms are posting the same price. These additional assumptions allow us to prove the following results:

Proposition 1. The cycle equilibrium described in Example 1 is not Pareto dominated by any stationary equilibria.

Proof. Define $p^{max} = \arg \max_p D(p, p)(p - c_1)$.

Consider a stationary equilibrium with $p_i^t = p_j^t = p^{max}$. The payoff for firm 1 is given by:

$$\frac{1}{1 - \delta} u_1(p^{max}, p^{max})$$

Firm 1's payoff in the cycle equilibrium is *greater* than:

$$\frac{1 - \delta^{k_1}}{1 - \delta} u_1(p^{max}, p^{max}) + \delta^{k_1} \frac{1 - \delta^{k_2}}{1 - \delta} u_1(p^{max} - \epsilon, p^{max})$$

The cycle equilibrium is certain to give a higher payoff when:

$$(1 - \delta^{k_1})u_1(p^{max}, p^{max}) + \delta^{k_1}(1 - \delta^{k_2})u_1(p^{max} - \epsilon, p^{max}) > u_1(p^{max}, p^{max})$$

$$\delta^{k_1} \left[(1 - \delta^{k_2})u_1(p^{max} - \epsilon, p^{max}) - u_1(p^{max}, p^{max}) \right] > 0$$

The above inequality holds for any sufficiently large k_2 if $u_1(p^{max} - \epsilon, p^{max}) > u_1(p^{max}, p^{max})$.

Now define $p^* = \arg \max_p u_2(p - \epsilon, p)$. Firm 2 will strictly prefer the cycle equilibrium over a stationary equilibrium with prices $(p^* - \epsilon, p^*)$ if:

$$(1 - \delta^{k_1})u_2(p^{max}, p^{max}) > u_2(p^* - \epsilon, p^*)$$

$$u_2(p^{max}, p^{max}) - u_2(p^* - \epsilon, p^*) > \delta^{k_1}u_2(p^{max}, p^{max})$$

The left hand side of the above inequality is strictly positive. So we can always pick k_1 large enough to satisfy the inequality.

We have thus shown that firm 1 prefers the cycle equilibrium over any stationary price equilibrium in which both firms post the same price.⁸ We have also shown that firm 2 prefers the cycle equilibrium over any equilibrium in which firm 2 posts a higher price. These results combine to establish that at least one firm will be strictly worse off in any stationary price equilibrium. Therefore, the cycle equilibrium cannot be Pareto dominated by a stationary equilibrium.

□

Proposition 1 explains why the firms might prefer a cycle equilibrium over any stationary price equilibrium. It allows the firms to achieve an average payoff vector which is not feasible in the one-shot game.

Now define y^k to be the realization of the signal profile at the end of period $k_1 + k_2 = k$ and consider the sub-game starting in period $k_1 + k_2 + 1 = k + 1$. We know that $p_i^{k+1} \leq y_i^k$ for both i . So,

⁸Assumption 5 tells us that firm 1 also prefers the cycle equilibrium over any stationary price equilibrium in which firm 1 posts a higher price.

because of Assumption 6, we can use the argument made in the proof of Proposition 1 to show that the equilibrium in Example 1 is not Pareto dominated by any stationary price equilibrium of the sub-game starting in period $k + 1$. The same argument can also be repeated for any later sub-game.

Define $\Xi(s)$ to be the set of sequences of public signals which may be realized on-path when the firms play according to the strategy profile s . Then $V_i^s(\xi)$ is the payoff for firm i when the sequence $\xi \in \Xi(s)$ is realized and the firms are playing according to s .

Proposition 2. Let v^* be any vector of payoffs which may be realized when the firms play the equilibrium in Example 1.

Let s be the strategy profile in any alternate equilibrium which satisfies the XIC refinement.

If $\min_{\xi \in \Xi(s)} V_i^s(\xi) > v_i^*$ then $\max_{\xi \in \Xi(s)} V_j^s(\xi) < v_j^*$.

6 Conclusion

This paper introduces a novel equilibrium refinement to provide a theoretical foundation for the “rockets and feathers” price cycles commonly observed in retail gasoline markets. The model posits that these cycles are not a result of competitive undercutting or random shocks, but rather a sophisticated mechanism for sustaining tacit collusion between competing firms.

The model considers a repeated game between two heterogeneous firms with imperfect monitoring, a structure designed to reflect the realities of the retail gasoline market where firms observe rivals’ prices through randomly updated and potentially delayed public signals, such as online price-tracking applications. Under these conditions, the XIC refinement imposes a critical restriction on firm behaviour: a firm will only change its price if the new price is guaranteed to be at least as profitable as the old one, regardless of the signal observed by its rival.

This core restriction generates the key dynamics of price cycles. The paper demonstrates that while

price decreases are generally compatible with the XIC refinement, price increases are profitable only under specific conditions. Because demand is highly elastic, a firm will only raise its price if its current price is already very close to marginal cost. Small price increases from intermediate levels are ruled out. This finding explains both the sharp, coordinated price restorations from a low base (the “rockets”) and the subsequent phase of slow, asynchronous price decreases (the “feathers”). Any equilibrium satisfying the XIC refinement must therefore feature constant prices or exhibit this cyclical pattern.

Crucially, the paper establishes that price cycles are not just a feasible outcome but can be a Pareto efficient one. For heterogeneous firms, which differ in marginal costs and bargaining power, a single stationary cartel price may not be agreeable to all parties. Price cycles provide a dynamic mechanism to transfer surplus over time without resorting to illegal cash payments. By allowing for periods of price leadership where one firm enjoys a temporary market share advantage, the cycle facilitates a division of collusive profits that can be more desirable for all firms than any achievable stationary price. This theoretical conclusion is supported by empirical evidence from actual price-fixing cases, where adjustment delays were explicitly used to manage transfers within the cartel.

In conclusion, this paper makes a significant contribution by demonstrating that the puzzling and persistent phenomenon of gasoline price cycles can be understood as a rational, robust, and efficient collusive strategy. By integrating a behaviourally motivated equilibrium refinement with a realistic model of imperfect monitoring, the analysis provides a compelling micro-foundation for how competing firms can sustain cooperation and dynamically redistribute its gains in a complex and uncertain environment. The XIC refinement itself represents a valuable theoretical tool with broader applicability for analyzing strategic cooperation in any setting where agents must justify their actions and desire stable, punishment-free agreements.

References

- Gabriel Carroll. Dynamic incentives in incompletely specified environments. Working Paper, 2024.
- Robert Clark and Jean-François Houde. Collusion with asymmetric retailers: Evidence from a gasoline price-fixing case. *American Economic Journal: Microeconomics*, 5(3):97–123, 2013.
- Emel Filiz-Ozbay and Erkut Y. Ozbay. Auctions with anticipated regret: Theory and experiment. *American Economic Review*, 97(4):1407–1418, 2007.
- Michele Fioretti, Alexander Vostroknutov, and Giorgio Coricelli. Dynamic regret avoidance. *American Economic Journal: Microeconomics*, 14(1):70–93, 2022.
- Drew Fudenberg, David K. Levine, and Eric Maskin. The folk theorem with imperfect public information. *Econometrica*, 62(5):997–1039, 1994.
- Edward J. Green and Robert H. Porter. Noncooperative collusion under imperfect price information. *Econometrica*, 52(1):87–100, 1984.
- Johannes Hörner, Takuo Sugaya, Satoru Takahashi, and Nicolas Vieille. Recursive methods in discounted stochastic games: An algorithm for $\delta \rightarrow 1$ and a folk theorem. *Econometrica*, 79(4):1277–1318, 2011.
- Tanushree Jhunjhunwala. Searching to avoid regret: An experimental evidence. *Journal of Economic Behavior & Organization*, 189:298–319, 2021.
- Eric Maskin and Jean Tirole. A theory of dynamic oligopoly ii: Price competition, kinked demand curves, and edgeworth cycles. *Econometrica*, 56(3):571–599, 1988.
- Philipp Strack and Paul Viefers. Too proud to stop: Regret in dynamic decisions. *Journal of the European Economic Association*, 19(1):165–199, 2021.