

Collusion and Coercion with Naive Rivals*

Zach Y. Brown **Alexander MacKay**
University of Michigan[†] Harvard University[‡]

March 6, 2023

Abstract

Standard models of collusion require that all firms are forward-looking and strategic. When one firm displays naive behavior—i.e., when it is myopic, memoryless, or non-strategic—typical collusive strategies cannot be supported in equilibrium. Motivated by the increasing adoption of high-speed pricing algorithms that monitor rivals' prices, we consider instead a model in which a single firm can update prices faster than its rivals. We show that this expands the set of strategies that yield prices above the competitive equilibrium. With sufficiently fast pricing, a firm can unilaterally sustain price levels that maximize joint profits even when all of its rivals are naive. We also characterize a *coercive* equilibrium that maximizes the profits of the faster firm. Overall, faster pricing provides a single firm with the ability to induce market outcomes that yield higher profits and lower consumer surplus.

Keywords: Pricing Frequency, Pricing Algorithms, Online Competition, Collusion

JEL Classification: L40, D43, L81, L13, L86

*We thank Scott Kominers, Joe Harrington, and Nathan Miller for helpful comments. We are grateful for the research assistance of Harry Kleyer.

[†]University of Michigan, Department of Economics. Email: zachb@umich.edu.

[‡]Harvard University, Harvard Business School. Email: amackay@hbs.edu.

1 Introduction

In repeated games of price competition, strategies that maximize joint profits can be sustained in equilibrium when firms are sufficiently forward-looking. Typical models of collusion assume that all firms understand the dynamic strategies of rivals and value future profits. Several factors are viewed as important for facilitating collusion, including similar size and costs, predictability of demand, the observability of all rivals prices, and possibility for frequent direct communication.¹

Such conditions point to the fragility of collusion in standard models. In particular, collusive equilibria are not robust to the presence of short-sighted behavior by rivals. Myopic behavior by a single firm constrains the set of equilibria in repeated games, typically ruling out collusive strategies and joint profit maximization.

From this starting point, we consider the implications of pricing algorithms in environments where rival firms may act naively—i.e., they behave as if they are myopic, memoryless, or non-strategic. Firms are increasingly adopting algorithms that automate the monitoring of rivals' prices and generate price updates in response. However, these adoption decisions have not been uniform across firms. Even among firms that compete in the same market, some firms have a significant advantage in terms of their ability to observe rivals' prices and the speed in which prices can react. Building on empirical work that documents high-speed pricing algorithms in various markets,² we consider a model where a single firm's algorithm allows it to update prices more quickly than its rivals. The faster firm considers repeated-game dynamics, while its rivals may not. We then ask: what outcomes can a single, forward-looking firm achieve even when all of its rivals are naive?

We are motivated by three possibilities for naive behavior by rival firms. First, a naive rival may be *myopic*, in that it only considers current-period profits when setting prices. Second, a naive rival may be *memoryless*, in that it only conditions its actions on payoff-relevant variables and not the history of play. Third, firms may be *non-strategic*, in that they may be unaware that they are playing a game where rival prices influence their profits. For a motivating example, consider a firm that sets prices each quarter in order to maximize profits for the quarter, without considering future periods or conditioning on any information about competitor prices. This example displays all three of the above properties that we categorize as “naive.” Any one of these conditions would rule out standard collusive strategies in equilibrium.

In this paper, we show that the adoption of faster pricing by a single competitor expands the set of equilibrium strategies that allow for higher prices in a repeated game with differentiated

¹See, e.g., Scherer (1980), Tirole (1988), or Porter (2005).

²Brown and MacKay (2021) show that the pricing technology for large online retailers varies from once-per-week updates to updates that occur multiple times per hour. In addition, third party pricing software often highlight differences in speed. For instance, repricer.com offers a third-party pricing algorithm that prices hourly and a premium version that “reacts to changes your competitors make in 90 seconds.” The adoption of high-speed pricing algorithms has also been observed in other settings (e.g., Assad et al., 2022; Aparicio, Metzman and Rigobon, 2021).

products. First, we show how firm’s faster pricing affects the critical discount factor of its rival when playing a standard trigger strategy. We establish that a sufficiently fast pricing advantage can drive the rival’s critical discount factor to zero. In other words, the joint profit maximizing price levels can be achieved when the slower rival is myopic, when it is incapable of employing history-dependent strategies, or even when it is blind to presence of the fast firm. Second, we demonstrate that *any* pricing advantage enables the faster firm to sustain higher (partially collusive) prices with a naive rival. Third, we consider coercive trigger strategies that maximize profits for the faster firm, and we compare these outcomes to the outcomes that maximize joint profits.

The equilibrium strategies we develop are more robust than those in standard models that assume simultaneous price-setting behavior. In our model, the faster firm can punish the slower rival before the slower rival’s next opportunity to adjust prices. In this way, a firm with faster pricing can unilaterally coerce a slower rival into setting higher prices in the current period even if the slower rival does not internalize punishment in future periods.³ The results can be generalized to N -firm oligopolies where a single firm has a pricing advantage and the other $N - 1$ firms are naive.⁴

While we start by examining trigger strategies, the punishment need not be that drastic. Our analysis of pricing rules that are a linear function of the slower rival’s price demonstrates that such strategies can lead to collusive prices while potentially raising less antitrust scrutiny and being robust to the use of simple profit optimization by the slower firm. When the faster firm chooses the optimal strategy that maximizes its own profits, faster pricing can generate large asymmetries in prices and profits even if the firms are otherwise symmetric. In some environments, the fast firm can dictate that the slower firm set prices above the joint profit-maximizing price, leading to a greater deadweight loss and lower welfare than that obtained under joint profit maximization.

There is growing concern about collusion in online markets, especially when firms use algorithms (Harrington, 2018). The literature has largely focused on simultaneous move games, including a literature on collusion with artificial intelligence (Waltman and Kaymak, 2008; Calvano et al., 2020). It is well known that in simultaneous move games, high frequency pricing implies a smaller per period discount rate, making collusion easier to sustain when firms have perfect monitoring (e.g., Abreu, Milgrom and Pearce, 1991). However, asymmetries—typically in terms of costs or demand—are generally believed to make collusion more difficult to sustain (Scherer, 1980; Tirole, 1988). A smaller literature has focused on alternating-move games as in Maskin and Tirole (1988). In particular, Klein (2021) examines collusion with machine learning algorithms in a sequential game. Finally, the literature on price leadership and collusion has

³Our analysis differs from earlier work by Brown and MacKay (2021) in that we allow the fast firm to take into account repeated-game considerations when choosing its strategy.

⁴When the faster firm can observe each of the prices of its rivals, it can specify a punishment to each rival’s possible deviation that ensures no rival deviates in equilibrium.

examined the incentives to collude when one firm announces a price and a rival follows (e.g., Mouraviev and Rey, 2011). In previous work, Brown and MacKay (2021) examine competitive (Markov perfect) outcomes when firms may differ in the speed at which they set prices. We are not aware of research that examines collusion in settings where firms differ in their pricing frequency.⁵

The paper proceeds as follows. We introduce the model and present the non-cooperative Markov perfect equilibrium in Section 2. In Section 3, we provide the conditions under which collusion is incentive compatible, even in the presence of naive firms. Section 4 introduces the coercive equilibrium that maximize the profits of the sophisticated firm. In Section 5, we introduce the idea of learning for the naive firm, and we show how the sophisticated firm can obtain the collusive or coercive equilibria even with simple linear strategies. Section 6 concludes.

2 Model

2.1 Environment

Two firms, s and f , sell differentiated products in an infinitely repeated continuous-time game. We normalize the period length to 1. At the start of each period, both firms simultaneously set prices. Firm s , the slow firm, chooses p_s , which persists throughout the period. Firm f , the fast firm, has higher frequency pricing. The firm initially sets price p_f . Within each period, after the elapsed interval $\alpha \in [0, 1]$, the fast firm can observe p_s and adjust its price to r_f . We call α the *reaction time*.


Figure 1 shows the timing of the model. After f sets price r_f in the latter part of the period, s must wait fraction $1 - \alpha$ to adjust price in response. Thus, the model nests a standard simultaneous move game when $\alpha = 1$ and a sequential price-setting game when $\alpha = 0$. Throughout, we will implicitly assume that α represents the *subjective* reaction time that incorporates within-period discounting, and we will generally allow firms to use different short-term (within-period) and long-term (across-period) discount rates.⁶

In practice, differences in pricing frequency may provide the faster firm multiple opportuni-

⁵Our paper also relates to a previous literature examining games in which a single long-run player faces a succession of repeated short-run players. The classic application is one in which an incumbent faces a new (short-run) potential entrant in each period, as studied by Milgrom and Roberts (1982) and Kreps and Wilson (1982). Fudenberg and Levine (1989) and Fudenberg, Kreps and Maskin (1990) provide folk-theorem style analysis for feasible payoffs for a general class of games with a single long-run player. Our work contributes to this literature by considering repeated games of price competition and the advantage conferred to a single firm through the adoption of a pricing algorithm. In our setting, we consider repeated interactions with the same rival and interpret short-run behavior as arising from naive play.

⁶It is straightforward to explicitly account for discounting within the period. Consider any α and any instantaneous within-period discount rate ρ . Then there exists an objective (non-discounted) reaction time $\tilde{\alpha}$ such that $\frac{\int_0^{\tilde{\alpha}} e^{-\rho t} dt}{\int_0^{\alpha} e^{-\rho t} dt} = \frac{1 - e^{-\rho \tilde{\alpha}}}{1 - e^{-\rho \alpha}}$. For a given δ , the mapping of α to $\tilde{\alpha}$ is one-to-one.

Figure 1: Timing with Differences in Pricing Frequency



Notes: Figure shows timing of pricing when firm f , the fast firm, can react with lag α . Circle markers represent opportunities to adjust prices. The gray circle represents an opportunity to adjust price that is inconsequential in the model.

ties to adjust prices in each period.⁷ However, given the assumptions about our environment, only the first opportunity to adjust price within the period is consequential. In this way, α can be interpreted as the lag between s setting price p_s and the first opportunity for f to readjust its price in response. In the case of high-speed pricing algorithms, α is determined by the time required for software to observe a rival's price and update price in response.

Demand arrives in continuous time. The time-invariant instantaneous demand function for firm i is given by $D_i(p_i, p_{-i})$. We assume production costs are linear and equal to c_i . Thus, firm i receives instantaneous profits $\pi_i = (p_i - c_i)D_i(p_i, p_{-i})$. The stage game payoff for the fast and slow firm are then given by

$$\begin{aligned} v_f((p_f, r_f), p_s) &= \alpha\pi_f(p_f, p_s) + (1 - \alpha)\pi_f(r_f, p_s) \\ v_s(p_s, (p_f, r_f)) &= \alpha\pi_s(p_s, p_f) + (1 - \alpha)\pi_s(p_s, r_f). \end{aligned}$$

We make the following assumptions about demand:

Assumption A1. We assume that each $D_i(\mathbf{p})$ is continuous and twice differentiable, and that $\frac{\partial D_i(p_i, p_{-i})}{\partial p_i} < 0$ and $\frac{\partial^2 D_i(p_i, p_{-i})}{\partial p_i^2} \leq 0$. Further, we assume that $\frac{\partial D_i(p_i, p_{-i})}{\partial p_{-i}} > 0$ and $\frac{\partial^2 D_i(p_i, p_{-i})}{\partial p_i \partial p_{-i}} \geq 0$.

The first two conditions are to ensure that demand is downward-sloping and that solutions are well-behaved. The second two conditions ensure that the products are substitutes and that there weakly increasing differences in prices.

2.2 Noncooperative Markov Perfect Equilibrium

We begin by considering the noncooperative Markov perfect equilibrium. We follow closely Brown and MacKay (2021), and we consider this equilibrium to be *competitive*. It suffices

⁷Consider the differences in pricing frequency among major online retailers in the U.S. analyzed by Brown and MacKay (2021). A slower retailer updates prices once at the beginning of each week, whereas another retailer updates the price once each day. At the beginning of the week, the firms set prices simultaneously. Given that the fast firm can respond to the slow firm's price the following day, $\alpha = 1/7$ in this case.

to characterize the subgame perfect equilibrium (SPE) of the stage game. We use backward induction to solve for the SPE.

After α of the period has elapsed, f can respond to s 's price and chooses,

$$r_f^*(p_f, p_s) \in \operatorname{argmax}_{r_f} v_f((p_f, r_f(p_f, p_s)), p_s)$$

where r_f^* solves,

$$\frac{\partial \pi_f}{\partial p_f}(r_f^*, p_s) = (1 - \alpha) \left[D_f(r_f^*, p_s) + r_f^* \frac{\partial D_f}{\partial p_f}(r_f^*, p_s) \right] = 0.$$

Since r_f^* does not depend on p_f , let $r_f^*(p_s) = r_f^*(p_f, p_s)$.

Given that f will choose $r_f^*(p_s)$ in latter part of the stage game, f and s simultaneously choose p_f and p_s . In particular, f chooses,

$$p_f^*(p_s) \in \operatorname{argmax}_{p_f} v_f((p_f, r_f^*(p_s)), p_s)$$

where p_f^* solves,

$$\frac{\partial v_f}{\partial p_f}(p_f^*, p_s) = \alpha \left[D_f(p_f^*, p_s) + r_f^* \frac{\partial D_f}{\partial p_f}(p_f^*, p_s) \right] = 0. \quad (1)$$

This implies $p_f^* = r_f^*$ in equilibrium.

Turning to the slow firm, s chooses,

$$p_s^*(p_f) \in \operatorname{argmax}_{p_s} v_s(p_s, (p_f, r_f^*(p_s)))$$

where p_s^* solves,

$$\frac{\partial v_s}{\partial p_s}(p_s, p_f^*) = \alpha \underbrace{\frac{\partial \pi_s}{\partial p_s}(p_s, p_f^*)}_{\text{Bertrand Incentive}} + (1 - \alpha) \underbrace{\left[\frac{\partial \pi_s}{\partial p_f}(p_s^*, r_f^*(p_s^*)) \frac{\partial r_f^*}{\partial p_s}(p_s^*) + \frac{\partial \pi_s}{\partial p_s}(p_s^*, r_f^*(p_s^*)) \right]}_{\text{Sequential Incentive}} = 0 \quad (2)$$

The first-order condition in equation (2) highlights a key implication when a firm faces a rival with faster pricing. Since prices are initially set simultaneously, fraction α of the period creates incentives that are analogous to the standard Bertrand model with simultaneous price setting. However, the slow firm's profit in the latter part of the stage game, $\pi_s(r_f(p_f, p_s), p_s)$, is a function of the fast firm's updated price before the slow firm can react. This implies that for fraction $1 - \alpha$ of the stage game, there is a sequential pricing incentive in which the slow firm acts as the leader. In this part of the stage game, the slow firm internalizes the fast firm's reaction function $r_f^*(p_s)$. The relative speed that the fast firm can adjust prices, α , determines the weight that the slow firm places on simultaneous versus sequential incentives.

An SPE, $((p_f^N, r_f^N), p_s^N)$, is a solution to the system given by,

$$p_f^N = p_f^*(p_s^N), \quad p_s^N = p_s^*(p_f^N), \quad r_f^N = r_f^*(p_s^*(p_f^N))$$

As discussed in Brown and MacKay (2021), when prices are strategic complements and demand is symmetric, the faster firm has lower prices and higher profits than the slower firm.⁸ Both prices are higher than the Bertrand Nash equilibrium of the simultaneous pricing game. In our context, strategic complementary follows from assumption A1.

2.3 Joint Profit Maximization

The joint-profit maximizing prices satisfy

$$(p_f^C, p_s^C) = \arg \max_{(p_f, p_s)} \sum_{i \in \{f, s\}} \pi_i(p_i, p_{-i}).$$

For an interior solution, (p_f^C, p_s^C) solves,

$$D_i(p_i, p_{-i}) + p_i \cdot \frac{\partial D_i(p_i, p_{-i})}{\partial p_i} + p_{-i} \cdot \frac{\partial D_{-i}(p_i, p_{-i})}{\partial p_i} = 0, \quad \forall i.$$

Whenever the set of chosen strategies yields (p_f^C, p_s^C) , we say that the firms obtained the *collusive* outcome.

2.4 Example with Linear Demand

Throughout the paper, we use a simple symmetric linear demand system to illustrate key features of our model given by

$$D_i(p_i, p_{-i}) = \frac{1}{2} \left[1 - \left(1 + \frac{d}{2} \right) p_i + \frac{d}{2} p_{-i} \right] \quad (3)$$

where $d \geq 0$ is an inverse measure of product differentiation. This demand system can be derived from the quasilinear quadratic utility model (Singh and Vives, 1984). The goods do not compete when $d = 0$ and are perfect substitutes when $d = \infty$. Without loss of generality, marginal costs are normalized to zero.

The MPE of the stage game given this demand system is presented in Appendix A. For $\alpha \in (0, 1)$, then $p^B < p_f^N < p_s^N < p^C$ where p^B is the Bertrand price (under simultaneous price setting).

⁸Brown and MacKay (2021) refer to this as the pricing frequency equilibrium.

3 Collusive Prices

We consider the unilateral adoption of dynamic strategies by the faster firm that yield outcomes equivalent to those obtained by standard models of collusion. *Full collusion* prices are the prices that maximize joint profits, whereas *partial collusion* prices are those that are elevated above the Bertrand-Nash equilibria but not to the level of full collusion. In contrast to standard models of collusion, we show that full collusion prices may be obtained using a one-sided trigger strategy even when the slower firms are naive.

3.1 Full Collusion

We characterize a dynamic trigger strategy that the fast firm can employ so that it is incentive compatible for the slow firm to choose collusive prices in every period. As with standard models of tacit and explicit collusion, we require that the slow firm maximizes the present discounted value of profits given the faster firm's strategy. We defer a discussion of the fast firm's incentives until the end of this section.

We assume that the fast firm chooses a strategy that determines prices $p_{f,t}$ and $r_{f,t}$ based on the the history of previous prices. Specifically, the fast firm employs a modified grim trigger strategy. The fast firm initially chooses price p_f^C , the joint profit maximizing price. As long as the slow rival also chooses the joint profit maximizing price, p_s^C , the fast firm chooses p_f^C at the subsequent opportunity. If the rival defects and sets any other price p'_s , the fast firm chooses a within-period punishment price $r_f^P(p'_s)$. Importantly, the punishment occurs after lag α within the period, based on the advantage of the faster firm. Starting with the beginning of the subsequent period, the fast firm chooses p_f^P at every pricing opportunity.

We follow convention and assume that the punishment prices are chosen based on non-cooperative play. Specifically, we assume that the within-period punishment is given by $r_f^P(p'_s) = r_f^*(p'_s)$ and that $p_f^P = p_f^N$ following the SPE of the stage game.

We initially assume the slow firm is informed about the strategy of its rival. This is an important assumption that we relax in Section 5. Given the discount factor δ_s for the slow firm, the present discounted value of colluding for the slow firm is given by

$$\sum_{t=0}^{\infty} \delta_s^t v_{s,t}(p_s^C, (p_f^C, p_f^C)) = \frac{1}{1 - \delta_s} \pi_{s,t}(p_s^C, p_f^C)$$

The slow firm compares discounted profits of complying versus deviating. Define the *defection* and *defection reaction* prices as,

$$p_s^D \in \operatorname{argmax}_{p_s} \{ \alpha \pi_s(p_s, p_f^C) + (1 - \alpha) \pi_s(p_s, r_f^*(p_s)) \},$$

$$r_f^{DR} = r_f^*(p_s^D).$$

The defection price is the price that maximizes the slow firm's stage game profit given that f sets the collusive price at the beginning of the stage game and best responds when updating its price within the stage game. The defection reaction is the best response by f in the latter part of the stage game after s defects.

We impose additional assumptions on the solutions of the stage game that will allow analysis of the repeated game.

Assumption A2. *The following profit relations hold,*

1. $\pi_s(p_s^N, p_f^N) < \pi_s(p_s^C, p_f^C),$
2. $\pi_s(p_s^C, p_f^C) < \pi_s(p_s^C, p_f^D),$
3. $\pi_s(p_s^D, p_f^{DR}) < \pi_s(p_s^C, p_f^C).$

These assumptions eliminate trivial outcomes of the game. In particular, they guarantee that, for the slow firm, collusion yields higher profits than the SPE in the stage game, there is short-run incentive to defect when colluding, and that the fast firm can punish defection within the stage game.

Given the trigger strategy of the fast firm, the collusive price p_s^C is incentive compatible for the slow firm if

$$\frac{1}{1 - \delta_s} \pi_s(p_s^C, p_f^C) \geq \alpha \pi_s(p_s^D, p_f^C) + (1 - \alpha) \pi_s(p_s^D, p_f^{DR}) + \frac{\delta_s}{1 - \delta_s} \pi_s(p_s^N, p_f^N). \quad (4)$$

We define the critical discount factor as the minimum firm-specific discount factor above which collusive prices are incentive compatible. The critical discount factor for the slow firm can then be expressed as

$$\bar{\delta}_s(\alpha) = \frac{\alpha \pi_s(p_s^D, p_f^C) + (1 - \alpha) \pi_s(p_s^D, p_f^{DR}) - \pi_s(p_s^C, p_f^C)}{\alpha \pi_s(p_s^D, p_f^C) + (1 - \alpha) \pi_s(p_s^D, p_f^{DR}) - \pi_s(p_s^N, p_f^N)}. \quad (5)$$

When $\delta_s > \bar{\delta}_s(\alpha)$, it is not profitable for the slow firm to deviate from p_s^C . When $\alpha = 1$, the critical discount factor for standard collusion with simultaneous pricing is obtained.

Providing a firm with a faster pricing has three effects that shift the critical discount factor relative to simultaneous pricing. First, the benefit from deviating shrinks, as the slow firm receives the deviation profits for only a portion (α) of the period. Further, p_s^D will tend to be a less aggressive deviation with smaller α , as the slow firm puts more weight on the within-period punishment profits $\pi_s(p_s^D, p_f^{DR})$. Both of these forces reduce the benefits of a deviation and lower the critical discount factor. On the other hand, noncooperative profits $\pi_s(p_s^N, p_f^N)$ may be higher with a smaller value of α (Brown and MacKay, 2021). This limits the ability of the faster firm to punish.

Overall, the presence of a pricing advantage lowers the critical discount rate for the slower firm. In particular, for an arbitrary discount rate of the slower firm, there exists a pricing advantage that make collusion incentive compatible. We now state this formally:

Proposition 1. *For any $\delta_s \in [0, 1)$, there exists a threshold $\bar{\alpha} > 0$ such that, for any $\alpha < \bar{\alpha}$, collusive prices are incentive compatible for the slower firm, i.e., $\delta_s \geq \bar{\delta}_s(\bar{\alpha})$. The critical reaction time $\bar{\alpha}$ is given by*

$$\bar{\alpha} = \frac{\delta_s(\pi_s(p_s^D, p_f^{DR}) - \pi_s(p_s^N, p_f^N)) + \pi_s(p_s^C, p_f^C) - \pi_s(p_s^D, p_f^{DR})}{\delta_s(\pi_s(p_s^D, p_f^{DR}) - \pi_s(p_s^D, p_f^C)) + \pi_s(p_s^D, p_f^C) - \pi_s(p_s^D, p_f^{DR})}. \quad (6)$$

Proof. See Appendix.

Thus, any strategy that results in the slower firm playing p_s^C in every period is incentive-compatible for the slow firm. Example strategies include grim-trigger strategies or, in the case when $\delta_s = 0$, naive strategies that seek to maximize static profits. Further, as long as deviation provides any benefit to the slow firm ($\pi_s(p_s^D, p_f^C) > \pi_s(p_s^C, p_f^C)$), the fast firm *requires* a speed advantage to maintain collusive prices. The incentive compatibility of collusive prices with a zero discount factor is a key result of our model, and we discuss it below.

3.2 Considering Naive Rivals


We now discuss in more detail the behavior of the slower rival. We interpret the case in which $\delta_s = 0$ as a scenario in which the rival is *naive*. This case can be rationalized by different assumptions lead the slower firm to act as if it is maximizing its static period profits. We identify three behavioral assumptions that generate this naive behavior:

1. *Myopic*: The slow firm only values current-period profits.⁹
2. *Memoryless*: The slow firm is not able to condition its actions on the history of past play. More specifically, we make the Markovian assumption that the firm is only able to condition its actions on payoff-relevant variables.
3. *Non-strategic*: The slow firm may be unaware that it is playing a dynamic game with rival firms. For example, it may (falsely) believe that it is maximizing the profit function $\bar{\pi}_s(p_s)$ that does not depend on the strategic variable p_f .

In the latter two cases, the slower firm's solution to the repeated game is equivalent to the myopic solution, even though the slow firm values future profits. Thus, all three assumptions give rise to behavior that can be captured by the case when $\delta_s = 0$.

⁹Such "short-termism" may arise due to, e.g., misalignment between managerial incentives and shareholders.

Figure 2: Critical Discount Factor and Reaction Times



Notes: Panel (a) shows the slow firm’s critical discount factor that sustains collusive prices as a function of the reaction time, α , for different values of the (inverse) differentiation parameter d under linear demand given by Equation 3. Panel (b) shows the cutoff for the reaction time, $\bar{\alpha}$, such that full collusive prices can be sustained when the slow firm is myopic (i.e., $\bar{\delta}_s = 0$). The firms simultaneously set prices at the beginning of the period and then the fast firm can react with lag α .

The first two cases are distinguished from the third case in that the slow firm may be fully aware of the actions of the fast rival and be able to predict its actions. In these cases, the slow firm is limited in terms of how far forward or backward it looks in time when deciding its actions.

In the third case, the slow firm is unaware that it is playing a game with strategic rivals. The slow firm may be modeled as simply trying to maximize profits as a function of a single input (its own price). Consider, for example, a firm that sets prices at the beginning of the week, observes profits at the end of the week, and attempts to learn which price yields the highest profits. The firm specifies its profit function as $\pi_s(p_s)$. In truth, the profit function depends also on the price of the faster firm, i.e., $\pi_s(p_s, p_f)$, but the slow firm is not aware of this. The faster firm can vary p_f throughout the week to manipulate the slow firm into learning that the optimal p_s is a particular target price, for example, the fully collusive price.

Standard models of collusion do not hold up in equilibrium if any the above three conditions are satisfied. All firms (1) must have positive valuations of future periods, (2) must be able to condition on the history of play, and (3) must be aware of the strategies of rival firms.

By contrast, Proposition 1 shows that a fast firm can induce a slow firm to play collusive prices even when the slow firm is naive, i.e., $\delta_s = 0$. When α is small enough, the potential reaction by the fast firm within a current period is enough to incentivize the slow firm to maintain the collusive price. Thus, a pricing advantage enables “easy” collusion in ways that are robust to a broader set of behavioral assumptions about rival firms.

In Figure 2, we illustrate how a pricing advantage can incentivize a naive rival by solving for the critical discount factor using our example model with linear demand. Panel (a) plots the slow firm's critical discount factor as a function of α for three different levels of differentiation. For simultaneous pricing ($\alpha = 1$), greater differentiation (lower d) yields lower critical discount factors, implying collusion is easier to sustain. Regardless of the level of differentiation, faster relative pricing (lower α) leads to lower critical discount factors. Notably, for all three values of differentiation shown, the critical discount factor can be driven to zero with a short enough reaction time. For example, when $d = 1$, any $\alpha < 0.5$ yields a critical discount factor of 0. Panel (b) in Figure 2 plots the critical reaction time necessary to incentivize collusive prices for a naive rival as a function of the differentiation parameter, d .

3.3 Partial Collusion

Above, we consider the conditions under which it is possible for a single firm with a pricing advantage to support prices equivalent to full collusion. Even when joint profit maximization is not possible, a firm with a pricing advantage can employ a trigger strategy to support supracompetitive price levels. For partial collusion, we consider prices that are elevated proportionally above the noncooperative prices for both firm.

In particular, any reaction time less strictly than 1 provides an opportunity for the faster firm to increase prices above the MPE, even with naive rivals.

Proposition 2. *For any reaction time $\alpha \in [0, 1)$, there exists a trigger strategy for the fast firm that yields prices between the joint profit maximizing prices and the Bertrand-Nash prices, $p_i = \theta p_i^C + (1 - \theta)p_i^B$ for $i \in s, f$ and some $\theta \in (0, 1]$, for which p_s is incentive compatible for the slow firm when it is naive ($\delta_s = 0$).*

Proof. See Appendix.


The price and profits that can be sustained using a trigger strategy with a myopic rival are shown in Figure 3. In our numerical example, demand and marginal costs are symmetric, which yields symmetric prices. When $\alpha = 1$, the standard Bertrand-Nash prices are obtained. The profit-maximizing prices are plotted in panel (a). As α declines, these prices increase up until the point that a critical discount factor $\bar{\alpha}$ is obtained. For $\alpha < \bar{\alpha}$, the faster firm is able to support full collusion.

We discuss how these results may be generalized to an N -firm oligopoly in Appendix B.

3.4 Equilibrium

The above highlights that a pricing speed advantage can provide a fast firm with strategies that incentivize naive (myopic, memoryless, or non-strategic) rivals to set collusive prices. We

Figure 3: Maximum Symmetric Price and Profits by Reaction Time



Notes: Figure shows the symmetric prices and profits that can be sustained for different levels of product differentiation d under linear demand given by Equation 3. The firms simultaneously set prices at the beginning of the period and then the fast firm can react with lag α .

have not yet addressed whether such strategies would be played in equilibrium, as we have not addressed the incentives for the fast firm.

Formally, in our model, the incentive of the fast firm is similar to standard collusion. Consider the history of prices of the fast firm at a period τ , $h_{f,\tau} = \{p_{f,1}, r_{f,1}, \dots, p_{f,t}, r_{f,t}, \dots, p_{f,\tau}, r_{f,\tau}\}$. We assume that the slow firm employs a grim trigger strategy following

$$p_s(h_{f,\tau}) = \begin{cases} p_s^C, & \text{if } p_{f,t}, r_{f,t} = p_f^C \forall t < \tau \\ p_s^N, & \text{otherwise.} \end{cases}$$

Thus, after a deviation by the fast firm, we assume that the slow firm chooses the noncooperative MPE price. This yields the prices $\pi_f(p_f^N, p_s^N)$ after collusion has broken down.

The critical discount factor for the fast firm can be expressed as

$$\bar{\delta}_f = \frac{\pi_f(p_f^D, p_s^C) - \pi_f(p_f^C, p_s^C)}{\pi_f(p_f^D, p_s^C) - \pi_f(p_f^N, p_s^N)} \quad (7)$$

where $p_f^D \in \operatorname{argmax}_{p_f} \{\pi_f(p_f, p_s^C)\}$. Since p_f^D , p_f^C and p_s^C are not a function of α , these terms are the same as under standard collusion ($\alpha = 1$). However, the profits $\pi_f(p_f^N, p_s^N)$ depend on α , which may cause the critical discount factor for the fast firm to be higher than under simultaneous pricing.

Assume that $\delta_s > \bar{\delta}_s$ and $\delta_f > \bar{\delta}_f$. Then a subgame perfect Nash equilibrium is obtained when the fast firm employs the trigger strategy described in Section 3.1 and the slow firm

employs the above grim trigger strategy. This SPE yields the collusive prices.

There are two reasons why we have deferred the description of equilibrium to the end of this section. First, our view is that the equilibrium constructed above may place unnecessarily strong restrictions on the beliefs and behavior of the faster firm. The standard SPE conditions assume that each firm optimizes taking as given the strategies of its rivals. Considering the sophistication of the faster firm, we think it might be reasonable for it to internalize potential strategic reactions of its rivals, rather than taking them as fixed. For example, the fast firm may believe that its rivals choose to be naive precisely because they are earning collusive profits, but they may invest to become more sophisticated if they anticipate that the naive play would not be in their favor. Then, faced with a choice between, e.g., the MPE characterized in Brown and MacKay (2021) and the collusive outcome, we think it is reasonable that the fast firm may pick the collusive outcome if it yields higher profits.

Second, in the presence of naive rivals, there are alternative trigger strategies that can yield higher profits for the fast firm. We characterize these “coercive” strategies in the next section. However, as above, there may be features of the environment “outside of the model” that provide the fast firm incentives to maintain the joint-profit maximizing prices. For example, the fast firm may be compensated through transfers or through reciprocal behavior by its rivals in other (non-modeled) markets. In addition, one can consider environments where the slower rivals consider adopting higher-frequency pricing to match the technology of the fast firm. The fast firm may prefer to maintain the collusive prices to discourage its rivals from adopting the faster pricing. Such incentives may be more relevant in settings where collusive prices can be unilaterally initiated by a single firm.

4 Coercion with Faster Pricing

We now consider the adoption of trigger strategies by the faster firm that maximize its own profits. In contrast to collusion, coercive prices yield greater profits for the faster firm and lower profits for its rivals. We focus on the case where slower rivals are naive. If the slower firm can also take into account future periods, the scope for punishment is greater, which allows for a greater increase in profit for the faster firm.

4.1 Coercive Strategies

As in the previous section, we assume that the faster firm chooses a strategy that determines prices $p_{f,t}$ and $r_{f,t}$ based on the history of previous prices. Specifically, the faster firm chooses a modified grim trigger strategy s_f to maximize its net present value of profits, subject to two constraints. First, we maintain that slower firm’s short-run incentive-compatibility constraint is met and, second, that the fast firm punishes deviations with noncooperative best-responses.

Formally, given the history of prices of the slow rival $h_{s,\tau} = \{p_{s,1}, \dots, p_{s,t}, \dots, p_{s,\tau}\}$, the strategy s_f is determined by:

$$\begin{aligned}
p_f(h_{s,\tau}) &= \begin{cases} p_f^T, & \text{if } p_{s,t} = p_s^T \forall t < \tau \\ p_f^N, & \text{otherwise.} \end{cases} \\
r_f(h_{s,\tau}) &= \begin{cases} p_f^T, & \text{if } p_{s,t} = p_s^T \forall t \leq \tau \\ r_f^*(p_s), & \text{if } p_{s,t} = p_s^T \forall t < \tau \text{ and } p_{s,\tau} \neq p_s^T \\ p_f^N, & \text{otherwise.} \end{cases} \\
s.t. \quad (i) \quad p_f^T &= \operatorname{argmax}_{p_f^T} \sum_{t=0}^{\infty} \delta_f^t \pi_f(p_f^T, p_s^*(s_f)) \\
(ii) \quad p_s^*(s_f) &\in \operatorname{argmax}_{p_s | s_f} v_s(p_s, (p_f, r_f)) \\
(iii) \quad p_s^*(s_f) &= p_s^T
\end{aligned}$$

where (p_f^T, p_s^T) is the target set of prices for the fast firm. Note that the conditions do not depend on the discount factor of the slower firm, as it maximizes the current period profits v_s .

The trigger strategy of the faster firm s_f resembles the grim-trigger strategy in the previous section, where the fast firm responds to any deviation by the slower firm within the period with the best-response r_f^* , and then plays the noncooperative price p_f^N at every subsequent opportunity.

Conditional on the strategy of the fast firm, the slow firm's short-run incentive compatibility constraint can be expressed as

$$\pi_s(p_s^T, p_f^T) \geq \alpha \pi_s(p_s', p_f^T) + (1 - \alpha) \pi_s(p_s', r_f^*(p_s')) \quad \forall p_s'. \quad (8)$$

Denote the optimal deviation for the slow firm, conditional on p_f^T and r_f^* as \hat{p}_s^D and the deviation reaction prices as \hat{p}_f^{DR} . The fast firm chooses own price and rival's target price to maximize own profit such that the slow firm is indifferent between deviating and complying, i.e., such that equation (8) strictly binds. Note that the target prices (p_f^T, p_s^T) depend on the reaction time α of the fast firm. When $\alpha = 1$ (simultaneous price setting), the fast firm cannot coerce the slow firm from setting prices higher than the Bertrand-Nash equilibrium.

Proposition 3. *For any pricing speed advantage $\alpha \in [0, 1)$, there exists a coercive trigger strategy that yields higher price and greater profits for the faster firm than those obtained in (non-cooperative) Markov perfect equilibrium, even when the rival firm is naive. Profits for the fast firm are weakly greater than in the collusive equilibrium.*

Proof. See Appendix.

Note that Proposition 2 indicates that partial collusion (symmetric elevated prices) is incentive-compatible for naive rivals under identical conditions. Thus, in general, coercive strategies yield profits that are (weakly) greater for the fast firm than those received from partial or, when feasible, full collusion.

4.2 Characterizing the Optimal Coercive Strategy

We now solve for the optimal coercive trigger strategy of the fast firm. Given the punishment price and the optimal reaction by the slow firm, we can express the fast firm's constrained optimization problem as

$$\begin{aligned} \max_{p_s^T, p_f^T, \lambda} \quad & \pi_f(p_f^T, p_s^T) \\ & + \lambda (\pi_s(p_s^T, p_f^T) - \alpha \pi_s(p_s^D, p_f^T) - (1 - \alpha) \pi_s(\hat{p}_s^D, \hat{p}_f^{DR})) \end{aligned}$$

where λ is the Lagrange multiplier on the incentive compatibility constraint. This yields the first-order conditions:

$$[p_s^T] \quad 0 = \pi_f^2(p_f^T, p_s^T) + \lambda \pi_s^1(p_s^T, p_f^T) \quad (9)$$

$$[p_f^T] \quad 0 = \pi_f^1(p_f^T, p_s^T) + \lambda (\pi_s^2(p_s^T, p_f^T) - \alpha \pi_s^2(\hat{p}_s^D, p_f^T)) \quad (10)$$

$$[\lambda] \quad 0 = \pi_s(p_s^T, p_f^T) - \alpha \pi_s(\hat{p}_s^D, p_f^T) - (1 - \alpha) \pi_s(\hat{p}_s^D, \hat{p}_f^{DR}) \quad (11)$$

where we use superscripts to denote partial derivatives with respect to each argument. From the first relation, we obtain $\lambda = -\frac{\pi_f^2(p_f^T, p_s^T)}{\pi_s^1(p_s^T, p_f^T)}$. Thus, the optimal p_f^T solves

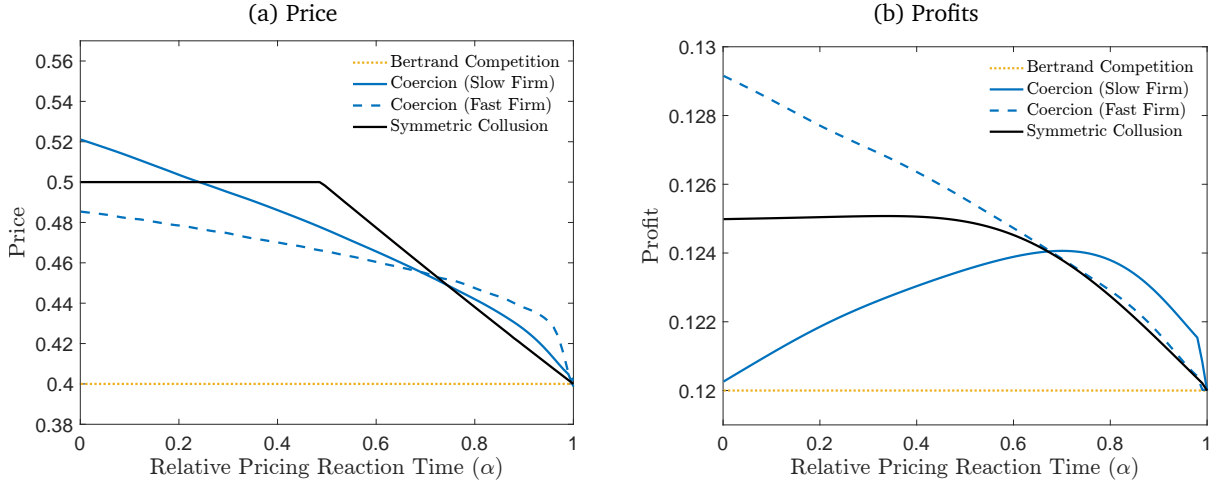
$$\pi_f^1(p_f^T, p_s^T) - \frac{\pi_f^2(p_f^T, p_s^T)}{\pi_s^1(p_s^T, p_f^T)} (\pi_s^2(p_s^T, p_f^T) - \alpha \pi_s^2(\hat{p}_s^D, p_f^T)) = 0 \quad (12)$$

where p_s^T is pinned down as a function of p_f^T from equation (11) and \hat{p}_s^D is an implicit function of p_s^T based on static profit maximization by the slower firm.

To illustrate, Figure 4 plots the prices (panel (a)) and profits (panel (b)) when the fast firm employs its optimal coercive trigger strategy as a function of its reaction time, α . The equilibria are generated from the linear demand system with a differentiation parameter of $d = 1$. The slow firm is assumed to be naive ($\delta_s = 0$). The solid blue line represents the slow firm, and the dashed blue line represents the fast firm. For comparison, we also plot the feasible full/partial collusion prices and profits (black line) and the prices and profits for Bertrand competition (yellow dotted line).

With a fast enough reaction $\alpha \gtrsim 0.7$, the fast firm is able to incentive its slower rival to set a higher price. The fast firm can undercut this price and earn higher profits. If the reaction

Figure 4: Prices and Profits with Coercive Trigger Strategies



Notes: Panel (a) shows prices and panel (b) shows profits under Bertrand competition with simultaneous pricing, coercion in which the faster firm uses a trigger strategy to maximize its own profits, and full/partial collusion with symmetric prices. The firms simultaneously set prices at the beginning of the period and then the fast firm can adjust price with reaction time α . Assumes $d = 1$ under linear demand given by Equation 3.

time is short enough $\alpha \lesssim 0.2$, the fast firm can coerce its slower rival into setting prices above the fully collusive price. Interestingly, if fast firm's timing advantage is small $0.7 \lesssim \alpha < 1$, it is optimal for the fast firm to target a higher price than the slow firm and to use its threat of punishment to prevent the slow firm from lowering its price further.

Panel (b) shows that the optimal coercive strategy always yields profits for the fast firm that are weakly greater than the full/partial collusion profits. When $0.7 \lesssim \alpha < 1$, the slower firm earns greater profits than the fast firm, and the joint profits are greater than the feasible partial collusion profits (yet, as one would expect, less than the full collusion profits). This shows that our definition of partial collusion, which has symmetric price increases across all firms, does not necessarily maximize joint profits when $\delta_s = 0$.

These equilibria illustrates the power of higher-frequency pricing. A single firm with superior pricing technology can unilaterally adopt a dynamic strategy that elevates prices above the competitive Bertrand equilibria. The faster firm obtain profits higher than the collusive profits without any coordination from its rival.¹⁰

4.3 Coercion in Equilibrium

Thus far, we have characterized the optimal modified trigger strategy for the fast firm subject to the incentive compatibility constraints of the slow firm. To support coercive strategies in equilibrium, the fast firm must have some incentive or constraint that prevents it from choosing

¹⁰See Appendix B for how these results may be extended to an N -firm oligopoly.

the short-run best response to the chosen price of the slow firm, p_s^T .

First, the slow firm may employ a trigger strategy following

$$p_s(h_{f,\tau}) = \begin{cases} p_s^T, & \text{if } p_{f,t}, r_{f,t} = p_f^T \forall t < \tau \\ p_s^N, & \text{otherwise.} \end{cases}$$

Thus, after a deviation by the fast firm, we assume that the slow firm chooses the noncooperative MPE price. This yields the prices $\pi_f(p_f^N, p_s^N)$ after collusion has broken down.

The critical discount factor for the fast firm can be expressed as

$$\hat{\delta}_f = \frac{\pi_f(\hat{p}_f^D, p_s^T) - \pi_f(p_f^T, p_s^T)}{\pi_f(\hat{p}_f^D, p_s^T) - \pi_f(p_f^N, p_s^N)} \quad (13)$$

where $\hat{p}_f^D \in \operatorname{argmax}_{p_f} \{\pi_f(p_f, p_s^T)\}$. If $\delta_f > \hat{\delta}_f$, then the two trigger strategies described above characterize a subgame perfect equilibrium yielding coercive prices. Note that, in general, $\hat{\delta}_f \neq \bar{\delta}_f$, so coercion may be sustained with lower or higher discount rates for the faster firm than collusion.

Another behavioral assumption that would maintain coercion in equilibrium is an environment in which the naive firm learns over time about its profit function. If the fast firm deviates from the price p_f^T with no deviation by the slow rival, the slow rival's incentive compatibility constraint is no longer satisfied. The slow firm may be able to learn this over time and adjust its behavior. In the limit, a repeated process of deviations by the fast firm and adjustments by the slow firm would converge to the noncooperative equilibrium. Thus, one interpretation of the trigger strategy for the slow firm described above is that it approximates this tâtonnement process. To the extent that the fast firm internalizes the effect that its deviations has on future behavior by the slow firm, the fast firm has an incentive to maintain its prices at p_f^T .

In the next section, we investigate learning in more detail.

5 Incorporating Learning

In the above analysis, we show how a faster firm may unilaterally implements supracompetitive prices when the slow firm understands (explicitly or implicitly) the potential punishment strategy and resulting profits. We now consider the case in which the slow firm is potentially uninformed about the strategy used by the fast firm or its own profit function. Instead, the firm learns over time using an optimization algorithm. We ask whether the fast firm can induce higher prices without announcing a strategy.

5.1 Simple Learning and Linear Strategies

We consider an environment in which the slower rival is unaware of the faster rival’s (punishment) strategy. We assume that the slower rival follows a simple learning approach with the following properties: The rival begins with a best guess of its optimal price (\hat{p}_s). In any period, it plays \hat{p}_s with positive probability or some other price on $[p, \bar{p}]$, where \bar{p} is the maximum price it considers. The slow firm uses the history of chosen prices and realized profits to predict a new optimal price, and it updates its best guess to that price. Examples of possible learning strategies include (i) choosing the price that yielded the highest profit in any single period in the past, (ii) using a more general reinforcement-learning method such as Q-learning, or (iii) using Newton’s method to predict the point that optimizes profits. The slow firm does not take into account the actions or strategies of the faster rival, nor does it form beliefs about the economic environment. These learning properties are standard for experimental A/B testing approaches used in practice.¹¹

In contrast to the previous sections, we assume here that the fast firm adopts a linear punishment strategy instead of a discontinuous trigger strategy. In particular, for a target price pair $(p_s^\dagger, p_f^\dagger)$, the faster firm chooses $p_f = p_f^\dagger$ at the beginning of the period and then updates its price according to the linear pricing rule:

$$r_f = \begin{cases} p_f^\dagger - \gamma(p_s^\dagger - p_s) & \text{if } p_f^\dagger - \gamma(p_s^\dagger - p_s) \geq 0 \\ 0 & \text{otherwise} \end{cases} \quad (14)$$

Thus, the punishment depends on how far from the target price the slow firm deviates, and the degree of punishment is captured by γ .

We focus on linear strategies because the linearity helps ensure that experimentation by the slower firm will converge to desired price of the faster firm, for many different learning approaches of the rival firm (e.g., using Newton’s method). In practice, pricing strategies that are linear in rivals price are common.

Throughout, we assume that the fast firm plays the long-run equilibrium strategy, i.e., the firm does not adjust its strategy to manipulate the rate of learning of the slower firm. For expositional clarity, we assume marginal costs are equal to zero in this section.

5.2 Joint Profit Maximization with Simple Learning

We now solve for the linear strategy that yields the fully collusive prices $(p_s^\dagger, p_f^\dagger) = (p_s^C, p_f^C)$. We maintain the (conservative) assumption that the slow firm maximizes static profits.

¹¹These properties also follow closely the simple “asynchronous” learning assumptions of Asker, Fershtman and Pakes (2021).

The optimal price for the slow firm is

$$\max_{p_s} \alpha p_s D_s(p_s, p_f^C) + (1 - \alpha) p_s D_s(p_s, p_f^C - \gamma(p_s^C - p_s)) \quad (15)$$

yielding the first-order condition

$$\alpha [D_s(p_s, p_f^C) + p_s D_s^1(p_s, p_f^C)] + (1 - \alpha) [D_s(p_s, p_f^C - \gamma(p_s^C - p_s)) + p_s \gamma D_s^2(p_s, p_f^C - \gamma(p_s^C - p_s)) + p_s \gamma D_s^1(p_s, p_f^C - \gamma(p_s^C - p_s))] = 0.$$

The slow firm sets the collusive price, p_s^C , when

$$\gamma^C = -\frac{1}{1 - \alpha} \frac{D_s(p_s^C, p_f^C) + p_s^C D_s^1(p_s^C, p_f^C)}{p_s^C D_s^2(p_s^C, p_f^C)} \quad (16)$$

Thus, the fast firm may employ a linear strategy to induce the collusive prices, and the slope of the punishment γ^C depends on its pricing advantage, α . When α is smaller (faster pricing technology), γ^C is smaller and the punishment need not be as drastic.

The fast firm may face a constraint on the aggressiveness of its punishment strategy, i.e., an upper bound on γ ($\bar{\gamma}$). One reason for this constraint is that a firm may believe that punishment that is too severe may raise suspicions from antitrust and competition authorities. Another reason is that a steeper slope (larger γ) can make learning more challenging for the rival firm, as it yields larger regions for which $r_f = 0$, where there is no (local) punishment and hence no incentive for the slow firm to marginally increase price.

The presence on of an upper bound on γ limits the ability of the firm to impose collusion with linear strategies. Specifically, an upper bound $\bar{\gamma}$ yields a critical pricing advantage α' . The faster firm must have sufficiently fast pricing $\alpha \leq \alpha'$ in order to support collusion with linear strategies. The critical pricing speed is characterized as:

$$\alpha' = 1 + \frac{1}{\bar{\gamma}} \frac{D_s(p_s^C, p_f^C) + p_s^C D_s^1(p_s^C, p_f^C)}{p_s^C D_s^2(p_s^C, p_f^C)} \quad (17)$$

To see how the bounds restrict the required pricing speed, consider the linear demand system characterized by equation (3). For this demand system, $\frac{D_s(p_s^C, p_f^C) + p_s^C D_s^1(p_s^C, p_f^C)}{p_s^C D_s^2(p_s^C, p_f^C)} = -1$, so that $\gamma^C = \frac{1}{1 - \alpha}$. Consider a restriction $\bar{\gamma} = 1$, so that the reaction by the fast firm may not be more aggressive than a price-matching strategy. This yields a critical $\alpha' = 0$, with a unique solution $\alpha = 0$ and $\gamma^C = 1$. In other words, price matching can sustain the collusive equilibrium when the fast firm can instantaneously respond to the price change by the slower rival. For $\alpha > 0$, the fast firm requires a more aggressive reaction than price matching for the particular demand system.

We summarize how a faster firm can obtain collusive prices even in the presence of a naive

rival with a proposition, which we prove for our example linear demand system:

Proposition 4. *For the the linear demand system characterized by equation (3), if firm f sets initial price $p_f = p_f^C$ at the beginning of the period and then uses pricing strategy $r_f = p_f^C - \gamma^C(p_s^C - p_s)$ that reacts to firm s after period α , where $\gamma^C < \bar{\gamma}$ and $\alpha < \alpha'$, firm s will set the collusive price. Moreover, if firm s is uninformed and uses an optimization routine that maximizes static profits, when the routine considers prices $p_s \geq \max\{0, \frac{\gamma^C-1}{\gamma^C}p_s^C\}$ it will converge to the collusive price because the firm's per-period profits are concave in p_s .*

Proof. See Appendix.

While the linear pricing rule provides a less severe punishment than a trigger strategy, it is still possible to sustain collusive prices. Unlike the trigger strategy analyzed in Section 3.1, a slow firm setting any price $p_s \in [\max\{0, \frac{\gamma^C-1}{\gamma^C}p_s^C\}, p_s^C)$ will be able to increase profits by marginally increasing price. Moreover, the fact that the slow firm's minimization problem is convex implies that the firm need not be informed about the strategy of their rival or demand since naive optimization will always converge to the collusive prices.

5.3 Coercive Linear Strategies with Simple Learning

With a naive rival with a simple learning strategy, the fast firm may instead choose to pursue a coercive strategy that maximizes its own profits. We again consider linear strategies to encourage convergence by the slower firm. The fast firm attempts to induce the target price vector $(p_s^\dagger, p_f^\dagger) = (p_s^T, p_f^T)$.

To constrain the slope of the reaction by the faster firm, we assume that the linear slope of the pricing rule passes through the point $(0, 0)$. Linear strategies of this form have the property that the faster firm's price changes in response to any non-negative price chosen by the slower firm. Using the pricing rule $r_f = p_f^T - \gamma^T(p_s^\dagger - p_s)$, this implies $\gamma^T = \frac{p_f^T}{p_s^T}$.

In each period, the fast firm chooses its initial price p_f and target price for slow firm to maximize its own profit, such that the slow firm is maximizing static profit

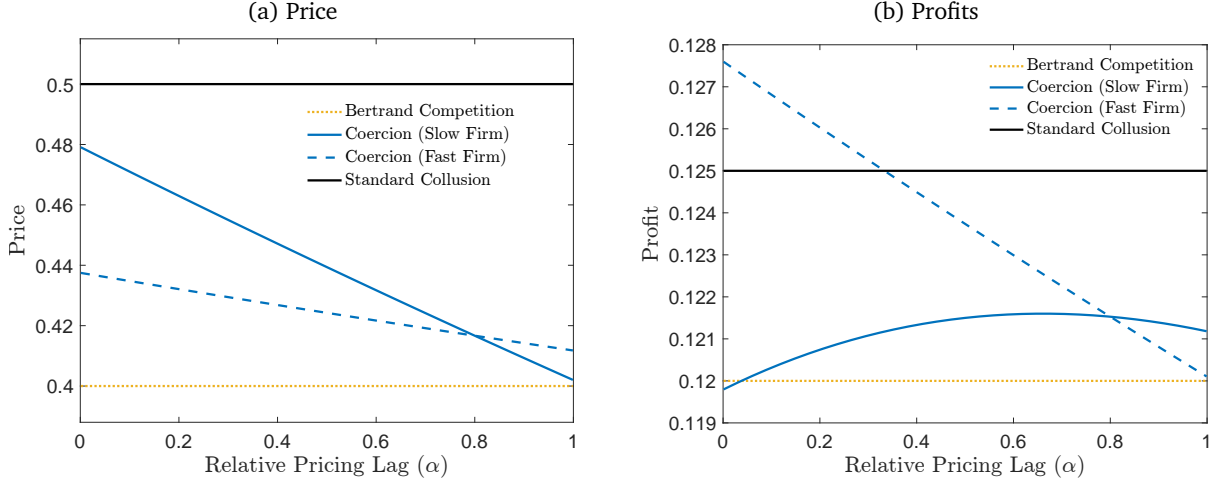
$$\begin{aligned} & \max_{p_f, p_s^T} \left[\alpha p_f D_f(p_f, p_s) + (1 - \alpha) \frac{p_f p_s}{p_s^T} D_f\left(\frac{p_f p_s}{p_s^T}, p_s\right) \right] \\ & s.t. \ p_s = \arg \max_{p_s'} \left[\alpha p_s' D_s(p_s', p_f) + (1 - \alpha) p_s' D_s\left(p_s', \frac{p_f p_s'}{p_s^T}\right) \right] \end{aligned} \quad (18)$$

The slow firm's first-order condition is given by

$$\alpha \left[D_s(p_s, p_f) + p_s D_s^1(p_s, p_f) \right] + (1 - \alpha) \left[D_s\left(p_s, \frac{p_f p_s}{p_s^T}\right) + p_s \frac{p_f}{p_s^T} D_s^2\left(p_s, \frac{p_f p_s}{p_s^T}\right) + p_s D_s^1\left(p_s, \frac{p_f p_s}{p_s^T}\right) \right] = 0 \quad (19)$$

The fast firm's pricing rule can be written in terms of the implicit function $p_s^*(p_f)$ that solves

Figure 5: Coercive Strategies with Linear Punishment



Notes: Panel (a) shows prices and Panel (b) shows profits under Bertrand competition with simultaneous pricing, joint profit maximization (standard collusion), and coercion in which the faster firm uses a linear punishment rule. In the coercion case, the firms simultaneously set prices at the beginning of the period and then the fast firm can react with lag α . Assumes $d = 1$ under linear demand given by Equation 3.

this first-order condition. The pricing-rule implies that

$$r_f = \frac{p_f p_s^*(p_f)}{p_s^T}. \quad (20)$$

In equilibrium, the fast firm sets the same price throughout the period so $p_f = r_f$. Let $p_f^*(p_s^T)$ be the fast firm price as a function of the target that solves $p_s^T = p_s^*(p_f)$ for p_f .

The fast firm then solves the following problem:

$$\max_{p_s^T} p_f^*(p_s^T) D_f(p_f^*(p_s^T), p_s^T). \quad (21)$$

The solution is given by the first-order condition

$$p_f^*(p_s^T) \left(\frac{\partial p_f^*(p_s^T)}{\partial p_s^T} D_f^1(p_f^*(p_s^T), p_s^T) + D_f^2(p_f^*(p_s^T), p_s^T) \right) + \frac{\partial p_f^*(p_s^T)}{\partial p_s^T} D_f(p_f^*(p_s^T), p_s^T) = 0 \quad (22)$$

This first-order condition provides the fast-firm optimal price for the slow firm, p_s^T . The optimal price for the fast firm is then $p_f^T = p_f^*(p_s^T)$. The solution reflect the fact that the fast firm chooses the target price for the slow firm knowing that the slow firm will maximize profit.

We plot an example solution for the example with linear demand ($d = 1$) in Figure 5. Panel (a) displays prices for different values of α . The fast firm prices are in the solid blue line. They are consistently higher than the Bertrand prices, and they increase with faster pricing (lower α). The fast firm prices are lower than the slow firm's when $\alpha < 0.8$. These patterns are similar

to the non-linear coercive strategies from the previous section (Figure 4). However, given the linear restriction on punishment, the fast firm cannot coerce its slow rival to set prices higher than the collusive price, indicating that the linearity of the strategies does limit the degree to which prices increase.

Panel (b) of Figure 5 displays the profits. Profits for the fast firm are declining in α , but profits for the slow firm are non-monotonic. With a fast enough reaction time, the fast firm can make higher profits than (its share of) the full collusion profits, even with the linear restriction. For a low enough α , the slow firm makes lower profits than in the Bertrand equilibrium.¹²

Proposition 5. *For the the linear demand system characterized by equation (3), coercive linear strategy profits are increasing in pricing speed for the fast firm. There exists cutoff $\bar{\alpha}'$ such that for $\alpha < \bar{\alpha}'$, profits for the fast firm are higher than under standard collusion.*

Proof. See Appendix.

6 Conclusion

Asymmetries in pricing technology expand the set of equilibrium strategies that yield supra-competitive profits. A firm with faster pricing may incentivize a slower firm to set prices that maximize joint profits even when the slower firm is myopic, memoryless, or non-strategic. Thus, a firm with a pricing advantage can unilaterally—without the cooperation of its rivals—obtain identical outcomes to those obtained from collusion.

Firms with faster pricing may not want to induce the collusive outcome. We characterize the set of coercive strategies that use punishment to maximize the faster firm’s profits. This provides a different equilibrium, with prices that typically differ across firms, that can be worse for consumers than collusion.

Overall, our results suggest a broader scope for firms to strategically increase prices. Sophisticated firms may be able to manipulate their rivals into setting prices above the competitive levels even when characteristics of the market would rule out traditional collusive strategies, such as short-termism or large differences in prices among similar firms. There is an opportunity for future research to examine the extent to which pricing strategies and algorithms used in practice may raise prices based on the features we identify here.

¹²This arises because we allow the fast firm to punish with a price of 0.

References

- Abreu, Dilip, Paul Milgrom, and David Pearce.** 1991. "Information and timing in repeated partnerships." *Econometrica*, 1713–1733.
- Aparicio, Diego, Zachary Metzman, and Roberto Rigobon.** 2021. "The Pricing Strategies of Online Grocery Retailers." National Bureau of Economic Research Working Paper.
- Asker, John, Chaim Fershtman, and Ariel Pakes.** 2021. "Artificial Intelligence and Pricing: The Impact of Algorithm Design." National Bureau of Economic Research Working Paper.
- Assad, Stephanie, Robert Clark, Daniel Ershov, and Lei Xu.** 2022. "Algorithmic Pricing and Competition: Empirical Evidence from the German Retail Gasoline Market." CESifo Working Paper 8521.
- Brown, Zach Y, and Alexander MacKay.** 2021. "Competition in Pricing Algorithms." *American Economic Journal: Microeconomics*. Forthcoming.
- Calvano, Emilio, Giacomo Calzolari, Vincenzo Denicolo, and Sergio Pastorello.** 2020. "Artificial Intelligence, Algorithmic Pricing, and Collusion." *American Economic Review*, 110(10): 3267–97.
- Fudenberg, Drew, and David K. Levine.** 1989. "Reputation and Equilibrium Selection in Games with a Patient Player." *Econometrica*, 57(4): 759–778.
- Fudenberg, Drew, David M. Kreps, and Eric S. Maskin.** 1990. "Repeated Games with Long-Run and Short-Run Players." *The Review of Economic Studies*, 57(4): 555.
- Harrington, Joseph E.** 2018. "Developing Competition Law for Collusion by Autonomous Artificial Agents." *Journal of Competition Law & Economics*, 14(3): 331–363.
- Klein, Timo.** 2021. "Autonomous Algorithmic Collusion: Q-learning under Sequential Pricing." *The RAND Journal of Economics*, 52(3): 538–558.
- Kreps, David M, and Robert Wilson.** 1982. "Reputation and Imperfect Information." *Journal of Economic Theory*, 27(2): 253–279.
- Maskin, Eric, and Jean Tirole.** 1988. "A Theory of Dynamic Oligopoly, I: Overview and Quantity Competition with Large Fixed Costs." *Econometrica*, 56(3): 549–569.
- Milgrom, Paul, and John Roberts.** 1982. "Predation, Reputation, and Entry Deterrence." *Journal of Economic Theory*, 27(2): 280–312.
- Mouraviev, Igor, and Patrick Rey.** 2011. "Collusion and Leadership." *International Journal of Industrial Organization*, 29(6): 705–717.

- Porter, Robert H.** 2005. "Detecting Collusion." *Review of Industrial Organization*, 26(2): 147–167.
- Scherer, Frederic M.** 1980. *Industrial Market Structure and Economic Performance: 2d Ed.* Rand McNally College.
- Singh, Nirvikar, and Xavier Vives.** 1984. "Price and Quantity Competition in a Differentiated Duopoly." *The RAND Journal of Economics*, 15(4): 546–554.
- Tirole, Jean.** 1988. *The Theory of Industrial Organization.* MIT Press, Cambridge, MA.
- Waltman, Ludo, and Uzay Kaymak.** 2008. "Q-learning agents in a Cournot oligopoly model." *Journal of Economic Dynamics and Control*, 32(10): 3275–3293.

Appendix

A Equilibrium of the Stage Game with Linear Demand

With linear demand given by Equation 3, solving FOC (1) and FOC (2) yields the SPE of the stage game. Prices are

$$p_s^N = \frac{6d + 8}{d((\alpha + 2)d + 16) + 16} \quad (\text{A-1})$$

$$p_f^N = \frac{\alpha d^2 + 5d^2 + 20d + 16}{(d + 2)(\alpha d^2 + 2d^2 + 16d + 16)} \quad (\text{A-2})$$

for the slow and fast firm, respectively. Profits are then

$$\pi_s^N = \frac{(3d + 4)^2(d(\alpha d + d + 8) + 8)}{2(d + 2)(d((\alpha + 2)d + 16) + 16)^2} \quad (\text{A-3})$$

$$\pi_f^N = \frac{(d((\alpha + 5)d + 20) + 16)^2}{4(d + 2)(d((\alpha + 2)d + 16) + 16)^2}. \quad (\text{A-4})$$

The above nests the Bertrand equilibrium when $\alpha = 1$. This equilibrium can be compared to the joint-profit maximizing profits given by $\pi_i^C = \frac{1}{8}$ which results when prices are $p_i^C = \frac{1}{2}$ for $i = f, s$.

B Extension to N -Firm Oligopoly

Throughout the paper, we consider the case where a fast firm faces a single, slower rival. It is straightforward to extend the results to a more general setting in which a single fast firm faces $N - 1$ naive slower rivals that all have the same pricing frequency. Thus, the reaction time α characterizes the reaction time that the faster firm has relative to each of the slower firms.

To extend the results, note that the profit function for any focal slow firm i can be written as

$$\pi_s(p_i, p_f, \{s_a\}) \quad (\text{A-5})$$

where $\{s_a\}$ is the set of strategies of all other slower rivals. Under the assumption that the slower rivals are naive, all of the slower firm pricing decisions will be made independently holding fixed the strategies of rivals.

For a given price vector, we can define $\tilde{\pi}_i(p_i, p_f) = \pi_s(p_i, p_f, \{s_a\})$ for each slow firm, construct the relevant incentive compatibility conditions, and the firm-specific critical reaction time $\bar{\alpha}_i$. The critical reaction time needed to support a target price vector for the market is given by

$\min\{\bar{\alpha}_i\}$. Thus, the fast firm must have a short enough reaction time in order to incentivize the naive rival that has the most to gain from a deviation. When this is satisfied, the fast firm can construct a trigger strategy that supports collusive prices by specifying an idiosyncratic reaction to individual deviations from each rival.

C Proofs

C.1 Proof of Proposition 1

The reaction time threshold for sustaining collusion when the slow firm has discount factor δ_s is given when Equation (4) strictly binds, implying

$$\bar{\alpha} = \frac{\frac{1}{1-\delta_s}(\pi_s(p_s^C, p_f^C) + \delta_s \pi_s(p_s^N, p_f^N)) - \pi_s(p_s^D, p_f^{DR})}{\pi_s(p_s^D, p_f^C) - \pi_s(p_s^D, p_f^{DR})}. \quad (\text{A-6})$$

From assumption A2, $\pi_s(p_s^C, p_f^C) > \pi_s(p_s^D, p_f^{DR})$. Therefore, for any $\delta_s \in [0, 1)$, $\frac{1}{1-\delta_s}(\pi_s(p_s^C, p_f^C) + \delta_s \pi_s(p_s^N, p_f^N)) - \pi_s(p_s^D, p_f^N) > 0$. Similarly, the denominator is positive by assumption A2, $\pi_s(p_s^D, p_f^C) > \pi_s(p_s^D, p_f^{DR})$. Therefore, $\bar{\alpha} > 0$ and any $\alpha < \bar{\alpha}$ can sustain collusion through a unilateral trigger strategy, even when the rival firm is myopic.

C.2 Proof of Proposition 2

We wish to show there exist prices $p_i^S = \theta p_i^C + (1 - \theta)p_i^B$ $i \in s, f$ such that

$$\pi_s(p_s^S, p_f^S) > \alpha \pi_s(p_s^D, p_f^S) + (1 - \alpha) \pi_s(p_s^D, p_f^{DR}) \quad (\text{A-7})$$

Because p_i^S is defined on the line segment between p_i^B and p_i^C , we can represent $dp_s^S = \beta dp_f^S$ for the appropriate $\beta > 0$.¹³ Evaluating derivatives with respect to p_i^S at $p_i^S = p_i^B$ yields

$$\pi_s^1(p_s^B, p_f^B) + \beta \pi_s^2(p_s^B, p_f^B) > \alpha \left(\frac{dp_s^D}{dp^S} \pi_s^1(p_s^B, p_f^B) + \beta \pi_s^2(p_s^B, p_f^B) \right) + (1 - \alpha) \left(\frac{dp_s^D}{dp^S} \pi_s^1(p_s^B, p_f^B) \right). \quad (\text{A-8})$$

where $p_s^D = p_s^B$ and $p_f^{DR} = p_f^B$ when $p^S = p^B$. Since $\pi_s^1(p_s^B, p_f^B) = 0$ and $\beta > 0$, we obtain

$$(1 - \alpha) \pi_s^2(p_s^B, p_f^B) > 0 \quad (\text{A-9})$$

which holds by assumption A2 (products are substitutes) for any $\alpha \in [0, 1)$.

¹³When demand is symmetric, $\beta = 1$.

C.3 Proof of Proposition 3

We proceed with a proof by contradiction. Suppose that, for some $\alpha < 1$, the optimal set of prices is equal to MPE prices, i.e., $(p_s^T, p_f^T) = (p_s^N, p_f^N)$. The MPE is characterized by $\pi_f^1(p_f^T, p_s^T) = 0$, since the fast firm is best-responding to the price of the slower firm.

Likewise, $\hat{p}_s^D = p_s^T$ in MPE, since the slow firm is choosing the optimal deviation subject to the within-period response by the faster firm. Because $\alpha < 1$ and $\pi_s^2(\cdot, \cdot) > 0$ (from assumption A1), we obtain $\pi_s^2(p_s^T, p_f^T) - \alpha\pi_s^2(\hat{p}_s^D, p_f^T) > 0$.

When $\alpha < 1$, $\pi_s^1(p_s^N, p_f^N) < 0$. From assumption A2, we also have that $\pi_f^2(p_f^T, p_s^T) > 0$, which yields $-\frac{\pi_f^2(p_s^T, p_f^T)}{\pi_s^1(p_s^T, p_f^T)} > 0$. Thus, for any $\alpha < 1$, the left hand side of equation (12) is positive when evaluated at (p_s^N, p_f^N) , and there exists a higher price $p_f^T > p_f^N$ that yields greater profits for the faster firm. Because of strategic complementarity, $p_s^T > p_s^N$ as well.

C.4 Proof of Proposition 4

When the fast firm uses pricing rule $r_f = p_f^C - \gamma(p_s^C - p_s)$ and demand is given by equation (3), the slow firm solves

$$\max_{p_s} \left[\alpha \frac{p_s}{2} \left(1 - \left(1 + \frac{d}{2} \right) p_s + \frac{dp^C}{2} \right) + (1 - \alpha) \frac{p_s}{2} \left(1 - \left(1 + \frac{d}{2} \right) p_s + \frac{dp^C - d\gamma(p^C - p_s)}{2} \right) \right]. \quad (\text{A-10})$$

This implies the slow firm sets price

$$p_s^* = \frac{(\alpha - 1)\gamma d + d + 4}{4((\alpha - 1)\gamma d + d + 2)}. \quad (\text{A-11})$$

Consistent with equation (16), the slow firm sets the collusive price, $p_s^* = p^C$, when the slope of the punishment is $\gamma^C = \frac{1}{1-\alpha}$. Therefore, given upper bound $\bar{\gamma}$, it must be that $\alpha \leq \alpha'$ where $\alpha' = \frac{\bar{\gamma}-1}{\bar{\gamma}}$. Substituting $\frac{1}{1-\alpha}$ for the slope of the punishment, the slow firm's profit is

$$\pi_s = \frac{1}{2} p_s (1 - p_s). \quad (\text{A-12})$$

This is defined for $p_s \geq \max\{0, \frac{\gamma^C-1}{\gamma^C} p_s^C\}$ since the fast firm's pricing rule returns a nonnegative price when $p_s \geq \frac{\gamma^C-1}{\gamma^C} p_s^C$. For $p_s \geq \max\{0, \frac{\gamma^C-1}{\gamma^C} p_s^C\}$, the slow firm's profit function is strictly concave in p_s so the minimization problem is convex and naive optimization will converge to the collusive price ($p_s^C = \frac{1}{2}$ for the specified demand system).

C.5 Proof of Proposition 5

Using our linear demand system from equation (3), the slow firm's problem is given by

$$\begin{aligned} \max_{p_s} \left[\alpha \frac{p_s}{2} \left(1 - \left(1 + \frac{d}{2} \right) p_s + \frac{dp_f}{2} \right) + (1 - \alpha) \frac{p_s}{2} \left(1 - \left(1 + \frac{d}{2} \right) p_s + \frac{p_f p_s}{2p_s^T} \right) \right] \\ \implies p_s = \frac{p_f(\alpha dp_f + 2)}{2d((\alpha - 1)p_f + p_s^T) + 4p_s^T}. \end{aligned}$$

Substituting this back into the pricing rule we have

$$r_f = \frac{p_f(\alpha dp_f + 2)}{2d((\alpha - 1)p_f + p_s^T) + 4p_s^T}.$$

If fast firm sets initial price $p_f = p_f$, the pricing rule must also return p_f at the equilibrium so

$$\begin{aligned} p_f &= \frac{p_f(\alpha dp_f + 2)}{2d((\alpha - 1)p_f + p_s^T) + 4p_s^T} \\ \implies p_f &= \frac{-2dp_s^T - 4p_s^T + 2}{\alpha d - 2d}, \end{aligned}$$

and the pricing rule is then

$$r_f = \frac{2 - 2(d + 2)p_s^T}{(\alpha - 2)d}.$$

Putting this all together, the fast firm chooses the target price p_s^T to hold the slow firm to in order to maximize profit

$$\begin{aligned} \max_{p_s^T} \left(\alpha \frac{1 - 2dp_s^T - 4p_s^T + 2}{\alpha d - 2d} \left(1 - \left(1 + \frac{d}{2} \right) \frac{-2dp_s^T - 4p_s^T + 2}{\alpha d - 2d} + \frac{dp_s^T}{2} \right) + \right. \\ \left. (1 - \alpha) \frac{1 - 2dp_s^T - 4p_s^T + 2}{\alpha d - 2d} \left(1 - \left(1 + \frac{d}{2} \right) \frac{-2dp_s^T - 4p_s^T + 2}{\alpha d - 2d} + \frac{dp_s^T}{2} \right) \right). \end{aligned}$$

Solving this yields the target price for the slow firm:

$$p_s^T = \frac{(6 - \alpha)d^2 + 4(6 - \alpha)d - 16}{2(d + 2)(\alpha d^2 + 8d + 8)}.$$

We can now solve for the fast firm's optimal target price

$$p_f^T = \frac{3d + 4}{d(\alpha d + 8) + 8}.$$

The profit for the slow firm is

$$\frac{((6 - \alpha)d^2 + 4(6 - \alpha)d - 16) (5\alpha d^2 + 4(\alpha + 4)d + 16)}{16(d + 2) (\alpha d^2 + 8d + 8)^2},$$

and the profit for the fast firm is

$$\frac{(3d + 4)^2}{8(d + 2) (\alpha d^2 + 8d + 8)}.$$

Note that the above is decreasing in α . The fast firm's profit from this strategy will be higher than under standard collusion for $\alpha < \bar{\alpha}''$ where

$$\bar{\alpha}'' = \frac{d^2}{d^3 + 2d^2},$$

which is decreasing in d . QED.