

MARKET EXPERIMENTATION IN A DYNAMIC DIFFERENTIATED-GOODS DUOPOLY*

Godfrey Keller[†] Sven Rady[‡]

January 1998

This Version: July 1999

Abstract

We study the evolution of prices in a symmetric duopoly where firms are uncertain about the degree of product differentiation. Customers sometimes perceive the products as close substitutes, sometimes as highly differentiated. Firms learn about their competitive environment from the quantities sold and a background signal. As the informativeness of the market outcome increases with the price differential, there is scope for active learning. In a setting with linear demand curves, we derive firms' pricing strategies as payoff-symmetric mixed or correlated Markov perfect equilibria of a stochastic differential game where the common posterior belief is the natural state variable. When information has low value, firms charge the same price as would be set by myopic players, and there is no price dispersion. When firms value information more highly, on the other hand, they actively learn by creating price dispersion. This market experimentation is transient, and most likely to be observed when the firms' environment changes sufficiently often, but not too frequently.

Keywords: Duopoly Experimentation, Bayesian Learning, Stochastic Differential Game, Markov Perfect Equilibrium, Mixed Strategies, Correlated Equilibrium

JEL Classification: C73, D43, D83

*An earlier version of this paper was circulated as Research Paper No. 1529, Graduate School of Business, Stanford University. We are grateful for comments and suggestions from Kyle Bagwell, Patrick Bolton, Jeremy Bulow, Luís Cabral, Chris Harris, Mike Harrison, Sunil Kumar, George Mailath, Mike Riordan, John Roberts, Larry Samuelson, and seminar participants at Caltech, Columbia University, ECARE (Université Libre de Bruxelles), the London School of Economics, Northwestern University, Stern School of Business (New York University), University of Pennsylvania, Simon School of Business (University of Rochester), Yale University, the Stanford GSB Brown Bag Lunch, the 1998 European Summer Symposium in Economic Theory in Gerzensee, the 1998 European Meeting of the Econometric Society and the 1998 Annual Conference of the European Economic Association. We would like to thank the Financial Markets Group at the London School of Economics and the Studienzentrum Gerzensee for their hospitality, and Eric Stout for valuable research assistance.

[†]Department of Economics, London School of Economics, Houghton Street, London WC2A 2AE; tel.: (+44) 171 955 7524; e-mail: R.G.Keller@lse.ac.uk.

[‡]Graduate School of Business, Stanford University, Stanford, CA 94305-5015; tel.: (+1) 650 723 5512; e-mail: Rady_Sven@gsb.stanford.edu.

Introduction

This paper studies a situation where two firms, who are symmetric in every respect, face a conflict between their short-term strategic incentives and their long-term information-gathering incentives. In each period, the short-term strategic incentives derive from the expected payoffs in a stage game, calculated under a common posterior that summarizes the firms' assessment of their competitive environment. These incentives dictate that both firms charge the same price. The long-term information-gathering incentives derive from the fact that through their pricing decisions, the firms can generate information about their environment and thus potentially improve future expected payoffs. If the value of information is positive, these incentives dictate that the firms charge two different prices.

The firms can resolve this conflict in an equilibrium where, depending on their posterior belief, they either behave as if they were myopic (i.e. infinitely impatient), or randomize so as to generate some dispersion in their prices and, by the same token, share the costs of learning equally. We study this equilibrium behaviour in an infinite-horizon, continuous-time setting where the unknown state of the environment changes randomly. This allows us to give precise answers to the following questions: How much price dispersion do the firms generate at any given time? What are the dynamics of beliefs and prices? And what fraction of time do firms actually create price dispersion?

In a fixed environment, we would see price dispersion disappear for good after some finite time, once the firms feel sufficiently sure of the true state. When we let the environment change from time to time, by contrast, price dispersion emerges and disappears in a cyclical fashion. As the frequency of changes in the environment increases, the amount of time the firms spend creating price dispersion first increases, then decreases, and finally vanishes altogether. Thus, price dispersion is most likely to be observed if the environment changes sufficiently often, but not too frequently.

Specifically, we consider a symmetric differentiated-goods duopoly where price-setting firms are uncertain about the cross derivative of their own demand with respect to the other firm's price, i.e. about the degree to which consumers perceive their goods as substitutes. This parameter is allowed to change randomly over time. In each period, the quantities sold provide a noisy signal of this parameter. The information content of this signal increases with the difference between the firms' prices. If they charge the same price, the signal is completely uninformative. Additionally, firms have access to a background signal which reveals some information independent of their pricing. Given a common prior and full observability of actions and outcomes, firms have a common posterior belief at all times, which is the natural state variable of a stochastic differential game.

We characterize a class of (possibly mixed-strategy) Markov perfect equilibria (MPE) that are payoff-symmetric, i.e. that give the same expected payoff to the firms. First, we provide a necessary and sufficient condition for price matching to constitute an MPE in pure strategies. If a price matching MPE exists, it is unique, prescribing the same strategies as would be played by myopic firms. The reason for this is that, in our linear framework, the marginal informational benefit generated by a deviation from

price matching is zero. In a symmetric pure-strategy equilibrium, therefore, the prices chosen must coincide with the unique Nash equilibrium of the one-shot game where firms merely try to maximize expected current profits. The interpretation of the necessary and sufficient condition for a price matching MPE is that such an equilibrium arises if and only if the value of information to firms is small at all beliefs.

The price matching MPE breaks down if the value of information is sufficiently high at some beliefs. Then, any Markov perfect equilibrium must involve active learning through price dispersion. We consider equilibria where firms achieve price dispersion by playing mixed strategies. This allows us to maintain payoff symmetry, which simplifies the construction of equilibria considerably. We show that expected payoffs are identical across all these equilibria, and characterise the equilibrium value function as a solution to a nonlinear two-point boundary value problem. As a consequence of the continuous-time setup and the linear specification of demand curves in our model, only the first and second moments of the price distribution matter. These moments are uniquely determined in equilibrium. The average price is equal to the price prescribed by the myopic strategy. The price variance increases in the value of information, and is positive on an open interval of beliefs.

A mixed-strategy MPE induces the following dynamics. Suppose that the environment has been the same for some time, and the firms are close to subjective certainty about it (implying that they charge the same price). Now, if the environment changes, firms first learn about this passively through the background signal, while maintaining zero price dispersion. Once they believe with sufficiently high probability that a change has occurred, firms enter a phase of positive price dispersion, i.e. active learning. This phase ends once they have become fairly certain of the change and charge identical prices again.¹

The interplay between the frequency of price dispersion episodes and their average length leads to an interesting comparative statics result with respect to the instability of the environment. Given a positive value of information, the incentive to create price dispersion is highest in a fixed environment since none of the information acquired ever becomes obsolete. In the long run, however, we will not see any price dispersion in a fixed environment since beliefs will converge to the truth. For a low frequency of changes in the environment, we observe occasional price dispersion. However, the state of the environment tends to stay the same for quite long, so firms's beliefs are close to the true state most of the time, and the amount of time when firms create price dispersion is rather small. As the frequency of changes in the environment increases, there are two effects: first, it becomes harder for beliefs to track the true state, implying that they stay more often in the open interval where firms create price dispersion; second, firms value information less since it becomes obsolete faster, so the interval where they are willing to create price dispersion shrinks. The first effect dominates for low frequencies of a state change. Eventually the second effect takes over, until the firms' willingness to experiment dies off completely and the pure-strategy equilibrium arises. Overall, we find a hump-shaped relationship between the instability of the

¹These dynamics are reminiscent of Slade's (1992) findings in her analysis of the Vancouver gasoline market in the early Eighties.

environment and the fraction of time that price dispersion is observed.

In turn, the instability of the environment is closely related to the shape of the frequency distribution of observed prices in the long run. In rather stable environments, this price distribution is strongly bimodal with most of the probability mass concentrated at the two extremes of the equilibrium price range. As the environment becomes less stable, more prices lie in an intermediate range, and the price distribution eventually becomes unimodal with most of the probability mass concentrated in the middle of the equilibrium price range. Thus, while both a high and a low frequency of state changes are consistent with little observed price dispersion, these two cases are clearly distinguished by the likelihood of seeing extreme prices.

We can formulate explicit conditions on the model parameters that govern whether or not there will be price dispersion in equilibrium. Principally, price dispersion will arise if and only if the firms' products are perceived to be relatively close substitutes, firms are relatively patient, the signals not too noisy, the background information relatively imprecise, and the environment sufficiently stable. In fact, close substitutability ensures a positive value of information, as Harrington (1995) pointed out, and the remaining conditions then guarantee that the value of information is sufficiently high.

When firms randomize independently, they sometimes charge prices that are close to each other but far off the myopic best response. Such a price pair is not very informative, yet can be quite costly in terms of current revenues. Clearly, firms can improve their payoffs if they have access to a public randomization device that allows them to avoid costly miscoordination. We characterise a class of payoff-symmetric correlated MPE with maximal payoff. We show that whenever the firms randomize jointly at a given belief, they do so with perfect negative correlation: their prices are equidistant from, and on opposite sides of, the myopic price. As with independent randomization, firms engage in joint randomization only if the value of information is sufficiently high. Joint randomization has a lower threshold since it yields a given amount of information at a lower opportunity cost than independent randomization. In fact, there is a range of model parameters where firms are willing to engage in joint randomization, but not in independent randomization.

A large literature has derived equilibrium price dispersion as the result of producer or consumer heterogeneity such as differences in production costs (Reinganum 1979), store size and age (Burdett and Coles 1997), consumers' search costs (Salop and Stiglitz 1976), their willingness to pay for the good (Diamond 1987), or the number of price offers they receive (Burdett and Judd 1983). That learning can be a cause of price dispersion in a perfectly *symmetric* setup seems to have been noted first by Slade (1989) and Aghion, Espinoza and Jullien (1993). Harrington (1995) has investigated how price dispersion caused by producer heterogeneity is affected by firms' learning. We follow these three papers in that we also consider a symmetric differentiated-goods duopoly with linear demand and full observability of actions (prices) and outcomes (quantities).²

²Other papers in the literature on duopoly learning and experimentation use different frameworks. Mirman, Samuelson and Schlee (1994) investigate the case of homogeneous goods. Riordan (1985) and Mirman, Samuelson and Urbano (1993) study signal jamming when only outcomes, but not actions,

In Slade’s (1989) infinite-horizon model, demand is subject to periodic, but infrequent, shocks. For sufficiently low discount rates and probabilities of a shock, she constructs subgame ϵ -perfect equilibria where each firm’s price depends linearly on the rival’s price in the previous period, and where each shock triggers a learning phase of price dispersion. Slade emphasizes the role of price-matching behavior in explaining price wars as an equilibrium outcome; in fact, the “tit-for-tat” structure of intertemporal reaction functions implies that both firms price more aggressively during learning phases. In contrast, we wish to stress the experimentation trade-off that firms face. Imposing Markov perfection allows us to abstract from other aspects of infinitely repeated interaction.

Aghion *et al.* (1993) consider a finite-horizon Hotelling model where two firms learn about the (unknown but fixed) transportation cost for their good. They show existence of a unique symmetric subgame-perfect equilibrium, which, in contrast to our results, is always in mixed strategies. This difference is due to the specification of noise. In fact, the noise in their model is uniformly distributed over a compact interval, which means that firms either learn the truth in a given period, or do not learn anything at all. As a consequence, price dispersion turns out to have first-order gains in terms of future profits. Since firms set these gains against second order losses in current profits, they always want to create price dispersion. In our framework, the noise has full support, which means that firms always learn something when they create price dispersion, but never learn everything. This implies that both gains and losses from price dispersion are of second order, and allows the losses to dominate the gains sometimes. Our learning model can thus explain both the emergence and the disappearance of price dispersion over time.

Harrington (1995) studies a two-period model where symmetric firms are subject to i.i.d. cost shocks before they set prices. He proves existence of a pure-strategy subgame perfect equilibrium for sufficiently high discount rates. The main result is that, relative to myopic behavior, firms increase first-period price dispersion when goods are seen as highly substitutable, and reduce it when goods are perceived as highly differentiated. Since the underlying economic forces in our model are the same as in Harrington’s, his insight carries over to the equilibria that we construct. Removing within-period cost asymmetries and using a continuous-time setting, however, we are able to derive several new results. Most important, our framework allows us to address the dynamics and long-run properties of price dispersion.

Several authors have emphasized informational externalities and the resulting inefficiencies in experimentation and learning (Rob 1991, Caplin and Leahy 1993, Chamley and Gale 1994, Bergemann and Välimäki 1996, Bolton and Harris 1999, and others). While we also have an informational externality in our model (information is a public good, so one firm can free-ride on the information gathered by the other), equilibrium is inefficient not because of this externality, but because of imperfect competition. If firms had no market power, information would be of no value to them, there would be no need to experiment, and the informational externality would be completely irrele-

are observable. Fishman and Rob (1998) analyse the effects of competition on experimentation in a reservation-price model with search costs.

vant. Firms would set price equal to their common marginal cost, which is the efficient outcome. It is therefore of little interest to compare equilibrium experimentation with socially efficient experimentation in our model. In fact, this comparison is trivial: there is ‘excessive’ market experimentation whenever there is price dispersion.

Continuous-time models of multi-agent Bayesian learning go back to Bolton and Harris (1993), and have subsequently been used in Felli and Harris (1996) and Bergemann and Välimäki (1996, 1997). In these papers, experiments are indivisible in the sense that agents cannot control the information content of the signals they receive. In our model, the firms can finely tune the information content by adjusting their actions slightly.³ Our specification of background information is the same as in Bolton and Harris (1993), and we use the techniques of Keller and Rady (1999) to deal with a changing environment. Finally, our approach to continuous-time randomization follows Harris (1993). To the best of our knowledge, the present paper constitutes the first tractable application of continuous-time randomization to a game with continuous action space.

The next section sets up the model. Section 2 describes the evolution of beliefs. Section 3 introduces the benchmark of myopic behavior. Section 4 studies the circumstances under which price matching can be a Markov perfect equilibrium. Section 5 considers equilibria where firms achieve price dispersion by playing mixed strategies. Section 6 provides an existence and characterization result for payoff-symmetric MPE. Section 7 derives comparative statics for the value of information and price dispersion. Section 8 studies joint randomization. Section 9 contains some concluding remarks. Technical details and proofs are relegated to the Appendix.

1 The Model

We consider two firms, each producing a non-storable good in continuous time. There are two possible states of demand for these goods, $k = 0$ or 1 . If the current state is k , and the prices currently charged by the two firms are p_1 and p_2 , then the expected demand for the product of firm i in this period is

$$q_i = q(p_i, p_j; k) = \max\{a - bp_i - c_k(p_i - p_j), 0\}$$

where the intercept a , the slope b and the two possible substitutability parameters $c_0 < c_1$ are positive constants.⁴ Thus, the quantity that firm i expects to sell is decreasing in its own price p_i and the price *difference* with its rival, $p_i - p_j$. This price difference matters more in state 1, where consumers perceive the products as closer substitutes. The quantity actually sold by each firm is the expected quantity plus some noise term, specified below.

³Bolton and Harris (1999) give agents this capability by allowing them to allocate a *fraction* of each instant to experimentation, which is isomorphic to randomization.

⁴This structure of expected demand is the same as in Harrington (1995). The Hotelling model in Aghion, Espinoza and Jullien (1993) corresponds to the special case where $b = 0$.

The state changes according to a continuous time Markov process with the transition probabilities

$$\Pr(k_{t+\Delta t} = 0 \mid k_t = 0) = 1 - \lambda_0 \Delta t + o(\Delta t), \quad \Pr(k_{t+\Delta t} = 1 \mid k_t = 0) = \lambda_0 \Delta t + o(\Delta t),$$

$$\Pr(k_{t+\Delta t} = 0 \mid k_t = 1) = \lambda_1 \Delta t + o(\Delta t), \quad \Pr(k_{t+\Delta t} = 1 \mid k_t = 1) = 1 - \lambda_1 \Delta t + o(\Delta t)$$

where $\lambda_k \geq 0$ ($k = 0, 1$) are again known. In particular,

$$\Pr(k_s = k \forall s \in [t, t + \Delta t] \mid k_t = k) = \exp(-\lambda_k \Delta t).$$

During production, the firms know the parameters a, b and λ_k ($k = 0, 1$), but not the state of demand; furthermore, the noise in actual quantities sold prevents them from directly inferring the true state.

At each time t , the firms set non-negative prices $p_{1,t}$ and $p_{2,t}$ and observe the resulting increments in total quantities sold,

$$\begin{aligned} dQ_{1,t} &= q(p_{1,t}, p_{2,t}; k_t) dt + \sqrt{2} \sigma dZ_{1,t}, \\ dQ_{2,t} &= q(p_{2,t}, p_{1,t}; k_t) dt + \sqrt{2} \sigma dZ_{2,t}, \end{aligned}$$

where (Z_1, Z_2) is a standard two-dimensional Wiener process independent of the process k , and $\sigma > 0$ is a constant known to the firms.⁵ These equations capture the idea that consumers sometimes make mistakes or experience idiosyncratic taste shocks.

We assume that both firms observe each other's prices and sales as well as their own. Moreover, we assume that they both observe an additional background signal about the degree of product differentiation,

$$\sqrt{\xi} c_{k_t} dt + \sigma dW_t,$$

where W is another standard Wiener process independent of the processes k and (Z_1, Z_2) , and $\xi > 0$ is a constant known to the firms. This background information has the natural interpretation of what is freely available from consumer surveys, trade magazines, and the like.⁶ The ratio $\sqrt{\xi}/\sigma$ measures the strength of the public signal. Thus, using the same noise parameter σ as in the equations for total quantities sold is without loss of generality.

Each firm's objective is to maximize the expected present value of its profit flow. We assume that production has constant or linear marginal cost, normalized to zero without loss of generality, so revenue equals profit. Firm i 's total payoff, expressed in per-period terms, is then

$$\mathbb{E} \left[\int_0^\infty r e^{-rt} p_{i,t} dQ_{1,t} \right] = \mathbb{E} \left[\int_0^\infty r e^{-rt} p_{i,t} q(p_{i,t}, p_{j,t}; k_t) dt \right]$$

⁵This is the continuous-time limit of the equations $\Delta Q_{1,t} = q(p_{1,t}, p_{2,t}; k_t) \Delta t + \sqrt{2} \sigma \sqrt{\Delta t} \varepsilon_{1,t}$ and $\Delta Q_{2,t} = q(p_{2,t}, p_{1,t}; k_t) \Delta t + \sqrt{2} \sigma \sqrt{\Delta t} \varepsilon_{2,t}$ ($t = 0, \Delta t, 2\Delta t, \dots$) with $\varepsilon_{i,t} \sim \text{IIN}(0, 1)$. The factor $\sqrt{2}$ in the noise terms is a normalization that will simplify some expressions later on.

⁶Technically, it will play a dual role. First, it will guarantee a process of posterior beliefs that is regular on the open unit interval, i.e. that can reach every point with positive probability from any current position. Second, it will ensure that the differential equations characterizing the firms' value functions are non-degenerate.

where r is the firms' common discount rate. Implicit in this expectation is firms' common initial belief about the degree of product differentiation, summarized by π_0 , the subjective probability that $k_0 = 1$.

Our assumption of mutual and full observability of prices and quantities implies that the two firms' posterior beliefs will be identical at all times. Since there is no other intertemporal link in our model, the posterior belief is a natural state variable for Markov perfect equilibria. We therefore focus on stationary strategies where the price charged at time t is a time-invariant deterministic function of π_t , the subjective probability that $k_t = 1$. More precisely, a (*pure*) *strategy* is a Lebesgue measurable mapping from the interval $[0, 1]$ of possible beliefs to \mathbb{R}_+ .

For any belief π ,

$$R(\pi, p_i, p_j) = p_i [(1 - \pi)q(p_i, p_j; 0) + \pi q(p_i, p_j; 1)]$$

describes firm i 's expected instantaneous revenue if prices this period are p_i and p_j . Given two strategies $p_1(\cdot)$ and $p_2(\cdot)$ and a common prior belief π , therefore, we can rewrite firm i 's total payoff as

$$\mathbb{E} \left[\int_0^\infty r e^{-rt} R(\pi_t, p_i(\pi_t), p_j(\pi_t)) dt \mid \pi_0 = \pi \right]$$

using the law of iterated expectations. We will extend this framework to mixed strategies later.

2 The Evolution of Beliefs

For the above specification of the quantities sold and the background signal, a slight extension of the standard result in Liptser and Shiriyayev (1977, Theorem 9.1) leads to the following result. If the current belief is π and the firms charge prices p_1 and p_2 , respectively, then the change in beliefs $d\pi$ over the next time interval of infinitesimal length dt has mean $\lambda(\pi)dt$ and variance $\pi^2(1 - \pi)^2 S(p_1, p_2) \sigma^{-2} dt$ where

$$\lambda(\pi) = (1 - \pi)\lambda_0 - \pi\lambda_1$$

and

$$S(p_1, p_2) = \frac{1}{2}[q(p_1, p_2; 1) - q(p_1, p_2; 0)]^2 + \frac{1}{2}[q(p_2, p_1; 1) - q(p_2, p_1; 0)]^2 + (c_1 - c_0)^2 \xi.$$

The expression for the variance has two components. The first, $\pi^2(1 - \pi)^2$, is a measure of the diffuseness of the current belief. The second, $S(p_1, p_2)/\sigma^2$, is a signal-to-noise ratio that incorporates firm 1's sales, firm 2's sales, and the background signal. Firms control the informativeness of the quantity signals through their choice of prices: the further the expected quantity in state 1 is apart from the expected quantity in state 0, the more firms learn about the current state, and the stronger is the corresponding change of belief. If firms charge the same price, the state of demand k is irrelevant to the quantities firms expect to sell, and the quantity signals become completely

uninformative. In this event, the background signal still provides some information. The lower σ (the level of noise in the system), and the higher ξ (the quality of the background information), the stronger will be the revision of beliefs for a given pair of prices. The reaction of beliefs to new evidence also increases with the diffuseness of the current belief, and is strongest at $\pi = \frac{1}{2}$. For $\pi = 0$ or 1 , on the other hand, firms are subjectively certain of the current state and ignore the signals they receive.

Firms also update because of the possibility that the state may change over the next infinitesimal period of time. Given the current belief π , they assign probability $1 - \pi$ to state 0, hence probability $(1 - \pi)\lambda_0$ to a transition from state 0 to state 1 over the next instant dt ; in the same way, they assign probability $\pi\lambda_1$ to a transition from state 1 to state 0. The first possibility increases the probability of being in state 1 after the time dt has elapsed, the second reduces it, and the combined effect leads to the expected change $\lambda(\pi) dt$. If at least one of the transition intensities λ_0, λ_1 is nonzero, the linear function λ is downward sloping and vanishes at the *invariant belief*

$$\tilde{\pi} = \frac{\lambda_0}{\lambda_0 + \lambda_1}.$$

(This is simply the expected fraction of time spent in state 1.) We let $\Lambda = \lambda_0 + \lambda_1$ and write

$$\lambda(\pi) = -\Lambda (\pi - \tilde{\pi}).$$

This representation shows that state switching introduces mean reversion into the evolution of beliefs. Throughout the paper, we shall fix an invariant belief $\tilde{\pi}$ and use the parameter Λ to measure the intensity of demand curve switches, and hence the instability of the environment in which the firms operate.

3 Myopic Pricing

In preparation for the construction of Markov perfect equilibria, this section introduces the benchmark of myopic behavior.

In a symmetric one-shot game with payoffs $R(\pi, p_i, p_j)$ for fixed π , the unique Nash equilibrium is symmetric and has both firms set the price

$$p^m(\pi) = \frac{a}{2b + c(\pi)}$$

with $c(\pi) = (1 - \pi)c_0 + \pi c_1$. This yields each firm the payoff

$$m(\pi) = \frac{a^2[b + c(\pi)]}{[2b + c(\pi)]^2}.$$

The price $p^m(\pi)$ would also be chosen in our dynamic setting with repeated interactions and the possibility of learning if firms did not care about the future. We therefore call p^m the *myopic strategy* and m the *myopic payoff function*.

Given our convention that $c_0 < c_1$, the functions p^m and m are strictly decreasing; m is strictly concave on $\{\pi : c(\pi) < b\}$ and strictly convex on $\{\pi : c(\pi) > b\}$. The

monotonicity properties simply reflect the fact that competition becomes more intense as the two products are seen as closer substitutes. The curvature properties of m are the same as in Harrington (1995), so we briefly recap his discussion of the underlying economic forces. As profits fall to zero when we approach the limiting case of Bertrand competition ($c_1 \rightarrow \infty$), the myopic payoff function has to become convex for $c(\pi)$ sufficiently large. When firms see their goods as highly substitutable, therefore, they are risk-loving with respect to changes in beliefs, which means that the value of information is positive. Conversely, when products are seen as highly differentiated, firms behave like “local monopolists” and extract high profits from the buyers. Under these circumstances, additional *public* information might be bad for both firms. While each firm could improve its pricing decision holding the competitor’s price constant, the competitor could in fact react to new information so aggressively as to make the expected net effect on profits negative. This is exactly what happens in our linear setup. For $c(\pi)$ sufficiently small, the myopic payoff function is concave,⁷ the firms are risk-averse with respect to changes in beliefs, and the value of information is negative.

Harrington (1995) also spells out the implications of these findings for the firms’ behavior in a dynamic setup. If the duopolists perceive their goods as highly substitutable and thus place a positive value on information, they have an incentive to make market observations more informative. Therefore, they create at least as much price dispersion as predicted by myopic profit maximization. Conversely, if the products are perceived as highly differentiated, the negative value of information gives firms an incentive to make market outcomes less informative, so they create at most as much price dispersion as under myopic profit maximization.

In Harrington’s model, firms experience cost shocks in each period before producing. Thus, the myopic equilibrium generically exhibits price dispersion, and it is feasible for the firms to create strictly *less* dispersion than in the myopic benchmark. This is impossible in our model, where there are no cost differences and the myopic strategies imply price matching. A negative value of information will therefore always imply zero price dispersion.

4 Price Matching

Our next aim is to show that in a Markov perfect equilibrium, price matching can only occur at the myopic price, and only if the value of information is low.

Suppose we have an equilibrium where the firms employ the strategies $p_1^*, p_2^* : [0, 1] \rightarrow \mathbb{R}_+$. Let $u_1^*, u_2^* : [0, 1] \rightarrow \mathbb{R}_+$ be the corresponding equilibrium value functions, i.e.

$$u_i^*(\pi) = \mathbb{E} \left[\int_0^\infty r e^{-rt} R(\pi_t, p_i^*(\pi_t), p_j^*(\pi_t)) dt \mid \pi_0 = \pi \right].$$

⁷By the envelope theorem, the change in revenues for a small increase in π is the price $p^m(\pi)$ times the number of units lost to the competitor firm, $c(\pi)(p^m)'(\pi)$. Considering this revenue loss at a higher belief π , we see three changes: first, the price $p^m(\pi)$ is lower, so units lost are less valuable; second, the opponent’s price response $(p^m)'(\pi)$ is smaller; third, $c(\pi)$ is higher. The effect of the first two changes is to lower the loss in revenue, while the effect of the third is to increase it. For small $c(\pi)$, the latter effect dominates, implying that $m'(\pi)$ decreases.

Suppose the current belief is π and firm i considers deviating to a price p_i over the next instant dt . Equilibrium requires that

$$u_i^*(\pi) = \max_{p_i \geq 0} \left\{ r R(\pi, p_i, p_j^*(\pi)) dt + e^{-r dt} \mathbb{E} \left[u_i^*(\pi + d\pi) | \pi, p_i, p_j^*(\pi) \right] \right\} \quad (1)$$

with the maximum achieved at $p_i = p_i^*(\pi)$. The first term in the maximand on the right-hand side is the expected current payoff, and the second term is the discounted expected continuation value after setting the price p_i . The equilibrium price achieves the optimal trade-off between the two.

We approximate $e^{-r dt}$ by $1 - r dt$ and, granted sufficient regularity of u_i^* , use Itô's lemma to expand $\mathbb{E}[u_i^*(\pi + d\pi) | \pi, p_i, p_j^*(\pi)]$ as

$$u_i^*(\pi) + (u_i^*)'(\pi) \mathbb{E}[d\pi | \pi, p_i, p_j^*(\pi)] + \frac{1}{2} (u_i^*)''(\pi) \text{Var}[d\pi | \pi, p_i, p_j^*(\pi)].$$

As $\mathbb{E}[d\pi | \pi, p_i, p_j^*(\pi)] = \lambda(\pi) dt$ and $\text{Var}[d\pi | \pi, p_i, p_j^*(\pi)] = \pi^2(1 - \pi)^2 S(p_i, p_j^*(\pi)) \sigma^{-2} dt$, the discounted expected continuation value is

$$(1 - r dt) \left(u_i^*(\pi) + \lambda(\pi) (u_i^*)'(\pi) dt + \frac{1}{2} \pi^2(1 - \pi)^2 S(p_i, p_j^*(\pi)) (u_i^*)''(\pi) \sigma^{-2} dt \right).$$

Substituting this into (1), cancelling $u_i^*(\pi)$ on either side and considering remaining terms of order dt only, we obtain the following characterization of firm i 's value function given firm j 's strategy:

$$u_i^*(\pi) = \frac{\lambda(\pi)}{r} (u_i^*)'(\pi) + \max_{p_i \geq 0} \left\{ \pi^2(1 - \pi)^2 \frac{(u_i^*)''(\pi) S(p_i, p_j^*(\pi))}{2r \sigma^2} + R(\pi, p_i, p_j^*(\pi)) \right\}. \quad (2)$$

This equation shows that a firm's equilibrium value $u_i^*(\pi)$ has three components. The first component, $\lambda(\pi) (u_i^*)'(\pi)/r$, reflects changes in the environment. In fact, $\lambda(\pi)$ indicates the magnitude and direction of the likely change in belief due to a possible change in the underlying state of demand. This piece of information has the shadow price $(u_i^*)'(\pi)/r$. The resulting contribution to the value function is positive if the belief is expected to move in the direction which increases value. The next term, $(\pi^2(1 - \pi)^2 (u_i^*)''(\pi)/(2r)) S(p_i, p_j^*(\pi))/\sigma^2$, represents the value of information extracted from the noisy signals available to the firm. Indeed, $S(p_i, p_j^*(\pi))/\sigma^2$ provides a measure for the information content of the signals obtained after setting the price p_i . This information is valued at the shadow price $\pi^2(1 - \pi)^2 (u_i^*)''(\pi)/(2r)$. The last term, $R(\pi, p_i, p_j^*(\pi))$, represents the myopic payoff. The maximization on the right-hand side of (2) indicates that in choosing its price, the firm optimally balances the future benefits of information against short-term profits. Price p_i is a best response for firm i at belief π if it achieves the maximum in (2).

Now suppose that the two firms charge the same price at the belief π , i.e. $p_1^*(\pi) = p_2^*(\pi)$. Since the partial derivative of the signal-to-noise ratio with respect to either firm's price is zero in this situation, one firm's price must be a myopic best response to the other firm's price, hence $p_1^*(\pi) = p_2^*(\pi) = p^m(\pi)$.

Inserting these prices into (2), we see that at any belief where both firms charge the myopic price, their value functions u_1^* and u_2^* must satisfy the ODE

$$\tau(\pi) \frac{u''(\pi)}{2r} \xi = u(\pi) - \lambda(\pi) u'(\pi)/r - m(\pi) \quad (3)$$

with

$$\tau(\pi) = \pi^2(1 - \pi)^2 (c_1 - c_0)^2 / \sigma^2.$$

If they adopt the myopic pricing strategy at *all* beliefs, then, by a standard result, their common payoff function

$$u^m(\pi) = \mathbb{E} \left[\int_0^\infty r e^{-rt} m(\pi_t) dt \mid \pi_0 = \pi \right]$$

is the unique solution of (3) subject to the boundary conditions

$$u(0) - \lambda(0) u'(0)/r = m(0), \tag{4}$$

$$u(1) - \lambda(1) u'(1)/r = m(1). \tag{5}$$

These conditions are obtained on letting π tend to 0 or 1 in (2) and using boundedness of the second derivative together with the fact that $\tau(0) = \tau(1) = 0$. They have the natural interpretation that under subjective certainty, optimal and myopic payoffs are the same (up to the possibility of a state change).

The second-order condition for (2) is

$$\tau(\pi) \frac{(u_i^*)''(\pi)}{2r} \leq b + c(\pi). \tag{6}$$

This means that the shadow price of information must be sufficiently low so that no firm has an incentive to deviate from symmetric pricing and thereby generate additional information. Multiplying through by ξ , the *amount* of information under myopic behavior, we can reinterpret condition (6) as stating that the *value* of information under myopic behavior must be sufficiently low. By the ODE (3), this value of information is $u_i^*(\pi) - \lambda(\pi) (u_i^*)'(\pi)/r - m(\pi)$, so (6) is equivalent to

$$u_i^*(\pi) - \lambda(\pi) (u_i^*)'(\pi)/r - m(\pi) \leq [b + c(\pi)] \xi. \tag{7}$$

This formulation will be more convenient to work with in the remainder of the paper.

We have the following result.

Theorem 4.1 *If there exists a pure-strategy Markov perfect equilibrium with price matching at all beliefs, it is unique, with both firms using the myopic pricing strategy p^m and achieving the equilibrium value $u^* = u^m$.*

Such an equilibrium exists if and only if the value of information under myopic behavior is sufficiently low, i.e. if and only if

$$u^m(\pi) - \lambda(\pi) (u^m)'(\pi)/r - m(\pi) \leq [b + c(\pi)] \xi \quad \text{for all } \pi. \tag{8}$$

We have already shown necessity of the stated condition; the proof of sufficiency is straightforward.

Condition (8) is formulated in terms of the payoff function u^m , which is not available in closed form. It is therefore desirable to have conditions that are expressed directly in terms of model parameters. Our next results provide such conditions.

Proposition 4.1 *Price matching at the myopic price is an equilibrium if the myopic payoff function m is concave throughout ($c_1 \leq b$); or (r, Λ, σ) is sufficiently far from the origin in the Euclidean space \mathbb{R}^3 ; or ξ is sufficiently large.*

The first condition covers the case where goods are highly differentiated in either state of demand, so the myopic payoff function is concave at all beliefs, implying a negative value of information throughout. Under these circumstances, it is clear that neither firm has an incentive to deviate from the myopic strategy. For price dispersion to arise at some beliefs, therefore, the duopolists' goods must be sufficiently close substitutes in at least one state of demand. But even when the value of information is positive at some beliefs, it may still be too small to warrant a deviation from myopic pricing. This is the content of the second condition. In fact, if the firms are too impatient, if their competitive environment changes too frequently, or if the quantity signals they receive are too noisy, then the firms will not gather information actively in equilibrium. The third condition confirms the intuition that firms will not experiment if they can rely on a very precise background signal.

The words “sufficiently far” in the second condition can be made more precise. Given a myopic payoff function with a convex part, there are positive constants ρ_0 and ℓ_0 such that price matching is an equilibrium for all parameter combinations (r, Λ, σ) that satisfy the inequality $\sigma^{-2} \leq r/\rho_0 + \Lambda/\ell_0$. The proof of Proposition 4.1 shows how such constants can be calculated explicitly in terms of the parameters of the demand structure.

Conversely, price matching will fail to be an equilibrium if products are perceived as fairly close substitutes, firms are sufficiently patient, their competitive environment is relatively stable, there is little noise, and the background information is of low quality. We formalize this in the special case where the value of information is positive at all non-degenerate beliefs.

Proposition 4.2 *Suppose that m is convex throughout ($c_0 \geq b$). Then, for any parameter combination (r, Λ, σ) sufficiently close to the origin in \mathbb{R}^3 and ξ sufficiently small, all Markov perfect equilibria involve experimentation, i.e. price dispersion.*

More precisely, the proof in Appendix B constructs positive constants ρ_1 , ℓ_1 and $\bar{\xi}$ such that (8) is violated whenever $\sigma^{-2} \geq r/\rho_1 + \Lambda/\ell_1$ and $\xi < \bar{\xi}$.

5 Price Dispersion

When condition (8) does not hold, any Markov perfect equilibrium must involve price dispersion at some beliefs. We consider equilibria where firms achieve price dispersion by playing mixed strategies. This allows us to maintain payoff symmetry ($u_1^* = u_2^*$), which simplifies the construction of equilibria considerably.

A *mixed strategy* is a Lebesgue measurable function from the state space $[0, 1]$ to the set of Borel probability measures on \mathbb{R}_+ . In the usual way, we can identify each pure strategy with a mixed strategy that, at each belief, puts full probability mass on the

action prescribed by the pure strategy. Note also that a mixed strategy can prescribe a pure action at some beliefs while it requires randomization at others.

We interpret randomization in continuous time as the limit of more and more frequent switching between pure strategies. In this limit, firms' payoffs and the law of motion of their beliefs can be calculated just by taking expectations over the distribution of actions.⁸ By the same arguments as in Section 4, therefore, the equilibrium value functions u_1^* and u_2^* in a mixed-strategy MPE $(\tilde{p}_1, \tilde{p}_2)$ must satisfy

$$\begin{aligned} u_i^*(\pi) - \frac{\lambda(\pi)}{r} (u_i^*)'(\pi) \\ = \pi^2(1-\pi)^2 \frac{(u_i^*)''(\pi)}{2r} \frac{\mathbb{E}[S(\tilde{p}_i(\pi), \tilde{p}_j(\pi))]}{\sigma^2} + \mathbb{E}[R(\pi, \tilde{p}_i(\pi), \tilde{p}_j(\pi))] \end{aligned} \quad (9)$$

$$= \max_{p_i \geq 0} \left\{ \pi^2(1-\pi)^2 \frac{(u_i^*)''(\pi)}{2r} \frac{\mathbb{E}[S(p_i, \tilde{p}_j(\pi))]}{\sigma^2} + \mathbb{E}[R(\pi, p_i, \tilde{p}_j(\pi))] \right\}, \quad (10)$$

with the maximum being achieved at any price p_i in the support of $\tilde{p}_i(\pi)$.

To keep matters as simple as possible, we focus on *interior* equilibria where prices and expected quantities are always positive, so that we can work in the truly linear part of demand where expected sales in state k are $q(p_i, p_j; k) = a - bp_i - c_k(p_i - p_j) > 0$. Thus, expected revenues at belief π are

$$\begin{aligned} R(\pi, p_i, p_j) &= p_i [a - bp_i - c(\pi)(p_i - p_j)] \\ &= m(\pi) - [b + c(\pi)](p_i - p^m(\pi))^2 + c(\pi)p_i(p_j - p^m(\pi)) \end{aligned} \quad (11)$$

and the information content of the firms' quantity and background signals simplifies to

$$S(p_i, p_j) = (c_1 - c_0)^2 [(p_1 - p_2)^2 + \xi]. \quad (12)$$

Using (11)–(12) and the notation $\tau(\pi) = \pi^2(1-\pi)^2(c_1 - c_0)^2/\sigma^2$, we see that the maximization problem (10) is equivalent to maximizing

$$\tau(\pi) \frac{(u_i^*)''(\pi)}{2r} (p_i - \mathbb{E}[\tilde{p}_j(\pi)])^2 - [b + c(\pi)](p_i - p^m(\pi))^2 + c(\pi)p_i(\mathbb{E}[\tilde{p}_j(\pi)] - p^m(\pi)). \quad (13)$$

Now suppose that firm 1 randomizes, that is, $\text{Var}[\tilde{p}_1(\pi)] > 0$. Then firm 1's maximand (13), which is quadratic in the choice variable over the range compatible with an interior equilibrium, achieves its global maximum in at least two different points, hence must be constant over this range. This implies

$$\tau(\pi) \frac{(u_1^*)''(\pi)}{2r} = b + c(\pi) \quad (14)$$

and $\mathbb{E}[\tilde{p}_2(\pi)] = p^m(\pi)$. Thus, for firm 1 to be willing to randomize over the range where prices and expected quantities are positive, the shadow price of information must be at the maximum level compatible with the existence of an interior equilibrium, and firm 2 must charge the myopic price on average.

⁸See Appendix A for details.

We do not know yet whether firm 2 randomizes as well, so we cannot conclude immediately that firm 1 is also charging the myopic price on average. However, imposing payoff symmetry ($u_1^* = u_2^*$), using (14) and simplifying, we see that firm 2 maximizes

$$(E[\tilde{p}_1(\pi)] - p^m(\pi)) \{[b + c(\pi)] (E[\tilde{p}_1(\pi)] + p^m(\pi)) - [2b + c(\pi)] p_2\}$$

when choosing its best response to firm 1's price distribution. So $E[\tilde{p}_1(\pi)] = p^m(\pi)$, since otherwise firm 2 would have no best response in the range compatible with an interior equilibrium.

By (11)–(12) and the fact that the average price charged by either firm equals $p^m(\pi)$, equation (9) becomes

$$\begin{aligned} u_i^*(\pi) - \frac{\lambda(\pi)}{r} (u_i^*)'(\pi) - m(\pi) \\ = \tau(\pi) \frac{(u_i^*)''(\pi)}{2r} (\text{Var}[\tilde{p}_i(\pi)] + \text{Var}[\tilde{p}_j(\pi)] + \xi) - [b + c(\pi)] \text{Var}[\tilde{p}_i(\pi)]. \end{aligned} \quad (15)$$

Payoff symmetry thus implies $\text{Var}[\tilde{p}_2(\pi)] = \text{Var}[\tilde{p}_1(\pi)]$, so firm 2 must randomize whenever firm 1 does. Moreover, (15) shows that the common price variance is

$$\frac{u^*(\pi) - \lambda(\pi) (u^*)'(\pi)/r - m(\pi)}{b + c(\pi)} - \xi$$

where u^* is the common value function.

In a payoff-symmetric interior MPE, therefore, necessary conditions for randomization at a belief π are

$$\tau(\pi) \frac{(u^*)''(\pi)}{2r} = b + c(\pi) \quad (16)$$

and

$$u^*(\pi) - \lambda(\pi) (u^*)'(\pi)/r - m(\pi) > [b + c(\pi)] \xi. \quad (17)$$

We can again interpret the left-hand side of (17) as the equilibrium value of information. Indeed, we can rewrite (15) as

$$\begin{aligned} u^*(\pi) - \frac{\lambda(\pi)}{r} (u^*)'(\pi) - m(\pi) \\ = \tau(\pi) \frac{(u^*)''(\pi)}{2r} (\text{Var}[\tilde{p}_j(\pi)] + \xi) - \left([b + c(\pi)] - \tau(\pi) \frac{(u^*)''(\pi)}{2r} \right) \text{Var}[\tilde{p}_i(\pi)]. \end{aligned} \quad (18)$$

The first term on the right-hand side is the value of the information that firm i acquires from the background signal and firm j 's randomization (which is perfect substitute for the background signal from firm i 's point of view). The second term represents firm i 's *net* value of its own experimentation, which is zero by (16); indeed, the benefit of deviating from the myopic price and thereby generating more information just equals the corresponding loss in revenue. In other words, a firm never gains directly from its own randomization; it only randomizes because this makes the other firm willing

to randomize as well, and it is *that* randomization which provides information with a positive net value. In summary,

$$u^*(\pi) - \frac{\lambda(\pi)}{r} (u^*)'(\pi) - m(\pi) = [b + c(\pi)] (\text{Var}[\tilde{p}(\pi)] + \xi)$$

measures firms' equilibrium value of information at a belief π where they randomize. Condition (17) can now be interpreted as stating that firms create price dispersion only at beliefs where the value of information is sufficiently high, i.e. higher than the value of background information, calculated at the highest shadow price compatible with interior equilibrium.

6 Characterization of Equilibrium

We found that in a payoff-symmetric interior MPE, the value function u^* satisfies

$$\frac{\tau(\pi)}{2r} (u^*)''(\pi) \xi = u^*(\pi) - \frac{\lambda(\pi)}{r} (u^*)'(\pi) - m(\pi) \leq [b + c(\pi)] \xi$$

at beliefs where firms charge identical (deterministic) prices, and

$$\frac{\tau(\pi)}{2r} (u^*)''(\pi) \xi = [b + c(\pi)] \xi < u^*(\pi) - \frac{\lambda(\pi)}{r} (u^*)'(\pi) - m(\pi)$$

at beliefs where they create (random) price dispersion. The equilibrium value function is thus a solution to

$$\frac{\tau(\pi)}{2r} u''(\pi) \xi = G\left(\pi, u(\pi) - \frac{\lambda(\pi)}{r} u'(\pi)\right) \quad (19)$$

where the right-hand side is given by the function

$$G(\pi, v) = \min \{v - m(\pi), [b + c(\pi)] \xi\}. \quad (20)$$

As before, the boundary conditions are (4)–(5). The following result is proved in Appendix B.

Proposition 6.1 *The boundary value problem consisting of the ODE (19) and the boundary conditions (4)–(5) has a unique solution $u^\dagger : [0, 1] \rightarrow \mathbb{R}_+$.*

We write \mathcal{P}^\dagger for the set of all positive mixed strategies \tilde{p} with mean $\text{E}[\tilde{p}(\pi)] = p^m(\pi)$ and variance

$$\text{Var}[\tilde{p}(\pi)] = \max \left\{ 0, \frac{u^\dagger(\pi) - \lambda(\pi)(u^\dagger)'(\pi)/r - m(\pi)}{b + c(\pi)} - \xi \right\} \quad (21)$$

for $\pi \in [0, 1]$. We define \mathcal{I} as the set of all pairs of mixed strategies $(\tilde{p}_1, \tilde{p}_2)$ in $\mathcal{P}^\dagger \times \mathcal{P}^\dagger$ with the property that the probability of the event $\{q(\tilde{p}_i(\pi), \tilde{p}_j(\pi); k) = 0 \text{ at some } \pi\}$ is zero for all $i = 1, 2$, $j \neq i$ and $k = 0, 1$.

We have already shown that any payoff-symmetric interior MPE lies in the set \mathcal{I} . Conversely, any strategy profile in \mathcal{I} gives both firms the payoff function u^\dagger and has the mutual best response property, hence constitutes an interior payoff-symmetric MPE. To decide whether the set \mathcal{I} is empty or not, we have to ask whether a price distribution with mean equal to the myopic price and variance prescribed by (21) can be implemented on the open half axis \mathbb{R}_{++} in a way that maintains positive expected quantities.

Given a belief π and the required mean $p^m(\pi)$, the maximal variance that can be achieved on \mathbb{R}_+ is $p^m(\pi)^2$, obtained from a binomial distribution which assigns probability $\frac{1}{2}$ to each of $p = 0$ and $p = 2p^m(\pi)$. By (21), therefore, the set \mathcal{P}^\dagger is non-empty if and only if

$$u^\dagger(\pi) - \lambda(\pi) (u^\dagger)'(\pi)/r - m(\pi) < [b + c(\pi)] (\xi + p^m(\pi)^2) \quad \text{for all } \pi. \quad (22)$$

Next, writing $s(\pi)$ for the standard deviation corresponding to (21), we see that the set \mathcal{I} is non-empty if and only if $q(p^m(\pi) + s(\pi), p^m(\pi) - s(\pi); k) > 0$, or $a - bp^m(\pi) - [b + 2c_k]s(\pi) > 0$, for all π and both $k = 0$ and 1 . Since $c_1 > c_0$, this is equivalent to the condition that

$$s(\pi) < \frac{a - bp^m(\pi)}{b + 2c_1} = \frac{b + c(\pi)}{b + 2c_1} p^m(\pi) =: \kappa(\pi) p^m(\pi) \quad (23)$$

at all beliefs. By (21), this translates into the requirement that

$$u^\dagger(\pi) - \lambda(\pi) (u^\dagger)'(\pi)/r - m(\pi) < [b + c(\pi)] (\xi + \kappa(\pi)^2 p^m(\pi)^2) \quad \text{for all } \pi. \quad (24)$$

As $\kappa(\pi) < 1$ everywhere, this condition is stronger than (22).

We summarise our findings in

Theorem 6.1 *The set of all payoff-symmetric interior Markov perfect equilibria, \mathcal{I} , is non-empty if and only if condition (24) holds. The firms' value function in any such equilibrium is $u^* = u^\dagger$.*

Equilibrium as in Theorem 6.1 does not impose any restrictions on price distributions beyond support, mean and variance. This is the combined effect of two modeling choices: the linearity of expected demand, and the continuous-time framework with Brownian noise. In the simplest possible mixed-strategy MPE, each firm flips a coin so as to decide whether it will charge a price above or below the myopic price, and then uses the expression (21) for the variance to calculate by how much it will deviate from that price.

We have interior equilibria with price dispersion when the value of information is high enough to violate condition (8), but small enough to respect condition (24).⁹ Given a myopic payoff function m that is not concave throughout, one can proceed as in the proof of Proposition 4.1 and construct constants ρ_2 and ℓ_2 such that (24) holds

⁹It is possible to construct non-interior equilibria when condition (24) is violated. Such a construction is cumbersome without generating additional insights. We therefore do not present it here.

whenever $\sigma^{-2} < r/\rho_2 + \Lambda/\ell_2$. In the case where the myopic payoff function m is convex throughout, we see interior price-dispersion equilibria for all parameter combinations $(r, \Lambda, \sigma, \xi)$ satisfying $r/\rho_1 + \Lambda/\ell_1 \leq \sigma^{-2} < r/\rho_2 + \Lambda/\ell_2$ and $\xi < \bar{\xi}$, where ρ_1 , ℓ_1 and $\bar{\xi}$ are the constants constructed in the proof of Proposition 4.2.

At beliefs close to subjective certainty, firms do not value information very highly, hence will not create price dispersion. So, if we are to see price dispersion at all, it will have to arise at ‘diffuse’ beliefs. This is made precise in the following result, which we prove in Appendix B.

Proposition 6.2 *In any equilibrium in \mathcal{I} , the set of beliefs where firms create price dispersion is either empty or an open interval $]\underline{\pi}, \bar{\pi}[$ with $0 < \underline{\pi} < \bar{\pi} < 1$.*

As an illustration of an interior equilibrium with price dispersion, we have the sample paths in Figure 1.¹⁰ The upper panel shows beliefs evolving over time¹¹ and the horizontal band between the dot-dash line is the interval of ‘diffuse’ beliefs from the above proposition. Whenever beliefs traverse this band, the firms’ prices have a positive standard deviation, shown in the lower panel. Thus, experimentation cycles arise where periods of price dispersion, triggered by a change in the environment, alternate with periods of price matching, once the firms learn enough about the new state. (Of course, we see price dispersion occurring occasionally at other times, as the firms are hit by short-term demand shocks or are pulled towards the invariant belief.)

7 Comparative Statics of Price Dispersion

Our next goal is to study how the occurrence of price dispersion varies with the basic model parameters. We start with an observation that follows immediately from expression (21) for the price variance.

Proposition 7.1 *Equilibrium price dispersion is weakly increasing in the value of information, and strictly increasing at all beliefs where price dispersion is positive.*

It is therefore enough to consider the comparative statics of the value of information.

Proposition 7.2 *If m is convex throughout ($c_0 \geq b$), then the value of information in all equilibria in \mathcal{I} decreases strictly in the firms’ discount rate, the noise intensity and, for sufficiently small discount rates, the frequency of state switches.*

These are the same comparative statics as in the single-agent problem of Keller and Rady (1999), and they follow from a (somewhat simpler) version of the arguments given there. We therefore omit the proof.

¹⁰In this and the subsequent figures at the end of the paper, the demand curve parameters are $a = 5.0$, $b = 0.3$, $c_0 = 0.5$ and $c_1 = 1.0$; the discount rate is $r = 0.1$, the invariant belief is $\bar{\pi} = 0.5$, and the noise parameter is $\sigma = 0.1$. Only Λ , ξ and the type of randomization vary across figures.

¹¹The firms’ common prior is $\pi_0 = 0.75$. The bold dashed line in either panel represents the true state, the initial state being $k_0 = 0$.

If the myopic payoff function is concave throughout, then the comparative statics change direction: the equilibrium value of information (which is now negative) *increases* in the discount rate, etc. Finally, if the myopic payoff function has a concave as well as a convex part, we see comparative statics as in the proposition for beliefs close to 1, and the reverse for beliefs close to 0.

Numerical solutions suggest further comparative statics properties when the value of information is positive throughout. For example, it appears that with regard to changes in Λ , the restriction to small discount rates in Proposition 7.2 can be dropped. Figure 2 illustrates how the equilibrium value of information decreases as Λ increases. Firms' price variance increases with the vertical distance between the value of information curve and the threshold line $[b + c(\pi)]\xi$, and is zero throughout for higher values of Λ .

An increase in the quality of background information seems to increase the value of information, but since the threshold that triggers randomization shifts upward at the same time, we find that price dispersion diminishes overall: firms randomize over a smaller set of beliefs, and with lower variances.¹² An extreme case of this is illustrated in Figure 3, in which the higher value of ξ leads to no dispersion (the higher value of information curve is never above the higher threshold line for the higher value of ξ), whereas the lower value of ξ leads to some dispersion (the lower value of information curve is sometimes above the lower threshold line for the lower value of ξ).

Our next comparative statics result deals with the likelihood of observing price dispersion at any given *time* (rather than at any given *belief*, as in the above propositions). The changing environment and the availability of background information ensure that the process of firms' posterior beliefs is ergodic, so the long-run fraction of time that price dispersion is positive along any sample path can be calculated simply as the probability of the event $\{\pi : \text{Var}[\tilde{p}(\pi)] > 0\}$ under the invariant distribution of the belief process. Given the firms' strategies, we can characterize this invariant distribution using standard results from the theory of diffusion processes; see Appendix C for details and Figure 4 for an illustration. We find that price dispersion is most likely to be observed when the competitive environment is neither too stable nor too variable. The proof of this result is also in Appendix C.

Proposition 7.3 *Suppose that m is convex throughout ($c_0 \geq b$). For a sufficiently small interest rate and precision of the background signal, the likelihood of price dispersion (under the stationary distribution induced by any equilibrium in \mathcal{I}) first increases, but eventually decreases with the frequency of state switches.*

In fact, as changes in the environment become more frequent, there are two effects: first, it becomes harder for beliefs to track the true state, implying that they stay more often in the open interval where firms create price dispersion; second, information

¹²Formally, the increase in the value of information corresponds to the encouragement effect described by Bolton and Harris (1999), and the upward shift of the threshold to their free-rider effect. The latter is easily explained: as background information becomes more reliable, the shadow price of information falls, and there is less of a reason to incur the opportunity costs of experimentation – it is cheaper to “free-ride” on the public signal. The first effect does not seem to have a natural explanation in our model. In any case, the free-rider effect appears to dominate in equilibrium.

decays faster, so the interval where firms are willing to create price dispersion shrinks. The first effect dominates for low frequencies of a state change. Eventually the second effect takes over, until the firms' willingness to experiment dies off completely and the pure-strategy equilibrium arises. An illustration of this result is given in Figure 5.

The shape of the long-run frequency distribution of observed prices also depends on the stability of the environment. In rather stable environments, this price distribution is strongly bimodal with most of the probability mass concentrated at the two extremes of the equilibrium price range. As the environment becomes less stable, more prices lie in an intermediate range, and the price distribution eventually becomes unimodal with most of the probability mass concentrated in the middle of the equilibrium price range. Thus, while both a high and a low frequency of state changes are consistent with little observed price dispersion, these two cases are clearly distinguished by the likelihood of seeing extreme prices.

8 Joint Randomization

So far, we have studied independent randomization by the two firms. Clearly, we expect them to achieve better expected payoffs if they are able to coordinate their pricing decisions via some public randomization device. In fact, such coordination will allow them to avoid the main drawback of independent mixing, i.e. prices that are close to each other, but away from the myopic price; such prices do not provide much information, yet can be quite costly in terms of current revenue forgone.

A *joint strategy* is a Lebesgue measurable function from the state space $[0, 1]$ to the set of Borel probability measures on $\mathbb{R}_+ \times \mathbb{R}_+$, assigning a distribution over pairs of prices to each belief. A joint strategy constitutes a *correlated MPE* if all price pairs that are assigned by the public randomization device are mutually best responses.

Similar to our treatment of private randomization, we interpret joint randomization as the limit of more and more frequent switching between pairs of pure strategies. Again, firms' payoffs and the law of motion of their beliefs can be calculated by taking expectations over the distribution of actions.¹³ As in Section 5, therefore, the equilibrium functions u_1^* and u_2^* in a correlated MPE $(\tilde{p}_1, \tilde{p}_2)$ must satisfy (9). In addition, the best response property requires any pair of prices (p_1, p_2) in the support of $(\tilde{p}_1(\pi), \tilde{p}_2(\pi))$ to satisfy

$$p_i \in \arg \max_{p \geq 0} \left\{ \pi^2 (1 - \pi)^2 \frac{(u_i^*)''(\pi)}{2r} \frac{S(p, p_j)}{\sigma^2} + R(\pi, p, p_j) \right\} \quad (25)$$

($i = 1, 2$), striking the optimal balance between value of information and current revenues.

We restrict our attention again to payoff-symmetric interior equilibria. By (11)–(12), firm i 's first-order condition for (25) thus reads

$$\tau(\pi) \frac{(u^*)''(\pi)}{2r} (p_i - p_j) - [b + c(\pi)] (p_i - p^m(\pi)) + \frac{1}{2} c(\pi) (p_j - p^m(\pi)) = 0$$

¹³See Appendix A for details.

where u^* is the common value function. Summing the two firms' first-order conditions leads immediately to $p_1 + p_2 = 2p^m(\pi)$, so if $p_1 \neq p_2$, the two prices must be equidistant from, and on opposite sides of, the myopic price. Inserting this in the first-order condition of either firm shows that the shadow price of information must be

$$\frac{\tau(\pi)}{2r} (u^*)''(\pi) = \frac{2b + 3c(\pi)}{4} \quad (26)$$

whenever the support of $(\tilde{p}_1(\pi), \tilde{p}_2(\pi))$ contains a pair of prices (p_1, p_2) with $p_1 \neq p_2$. Since the right-hand side of (26) is strictly less than $b + c(\pi)$, we see in particular that the second-order condition for (25) holds.

Given condition (26), a pair of non-negative prices (p_1, p_2) has the mutual best response property if and only if $p_1 + p_2 = 2p^m(\pi)$. This means that $\tilde{p}_1(\pi)$ and $\tilde{p}_2(\pi)$ must be perfectly negatively correlated: $\text{Var}[\tilde{p}_1(\pi)] = \text{Var}[\tilde{p}_2(\pi)] = -\text{Cov}[\tilde{p}_1(\pi), \tilde{p}_2(\pi)]$.

By (9) and the symmetry of $S(\cdot, \cdot)$ in its two arguments, payoff symmetry requires that the firms earn the same expected current revenues: $\text{E}[R(\pi, \tilde{p}_1(\pi), \tilde{p}_2(\pi))] = \text{E}[R(\pi, \tilde{p}_2(\pi), \tilde{p}_1(\pi))]$. Given that $\text{Var}[\tilde{p}_1(\pi)] = \text{Var}[\tilde{p}_2(\pi)]$ and $\text{E}[\tilde{p}_1(\pi)] + \text{E}[\tilde{p}_2(\pi)] = 2p^m(\pi)$, it is straightforward to show that this implies $\text{E}[\tilde{p}_1(\pi)] = \text{E}[\tilde{p}_2(\pi)] = p^m(\pi)$. We can thus cast public randomization as the observation of a real random variable $\tilde{\delta}(\pi)$ with mean zero, the assigned prices being $\tilde{p}_1(\pi) = p^m(\pi) + \tilde{\delta}(\pi)$ and $\tilde{p}_2(\pi) = p^m(\pi) - \tilde{\delta}(\pi)$.

Note that a lower shadow price of information is necessary for joint randomization than for independent randomization. This is due to the fact that joint randomization yields a higher net benefit. Indeed, using (11)–(12), we can rewrite equation (9) for firm 1 in the following way:

$$\begin{aligned} & u^*(\pi) - \frac{\lambda(\pi)}{r} (u^*)'(\pi) - m(\pi) \\ &= \frac{\tau(\pi)}{2r} (u^*)''(\pi) (\text{Var}[\tilde{p}_2(\pi)] + \xi) - \left([b + c(\pi)] - \frac{\tau(\pi)}{2r} (u^*)''(\pi) \right) \text{Var}[\tilde{p}_1(\pi)] \\ & \quad - 2 \left(\frac{\tau(\pi)}{2r} (u^*)''(\pi) - \frac{c(\pi)}{2} \right) \text{Cov}[\tilde{p}_1(\pi), \tilde{p}_2(\pi)]. \end{aligned} \quad (27)$$

The first two terms on the right-hand side of (27) are formally the same as in equation (18). The first term is again the value of the information that firm 1 acquires from the background signal and directly from firm 2's randomization. The second term represents firm 1's net value of its own experimentation. By (26), this is negative; using the fact that the assigned prices are $p^m(\pi) \pm \tilde{\delta}(\pi)$, we see that it equals $-[2b + c(\pi)] \text{Var}[\tilde{\delta}(\pi)]/4$. The third term is specific to joint randomization and captures firm 1's net value of coordinating with firm 2. Since the assigned prices are perfectly negatively correlated, this term equals $[2b + c(\pi)] \text{Var}[\tilde{\delta}(\pi)]/2$. Thus, the net loss incurred by deviating from the myopic price is doubly offset by the net benefit derived from coordinating with the other firm.

Combined with (26), equation (27) shows that the common variance of firms' price distributions is

$$\frac{u^*(\pi) - \lambda(\pi) (u^*)'(\pi)/r - m(\pi)}{b + c(\pi)} - \frac{2b + 3c(\pi)}{4[b + c(\pi)]} \xi.$$

Therefore, a necessary condition for joint randomization at a belief π is

$$u^*(\pi) - \lambda(\pi) (u^*)'(\pi)/r - m(\pi) > [2b + 3c(\pi)] \xi/4. \quad (28)$$

As before, we can interpret the left-hand side as the equilibrium value of information, so condition (28) states again that firms create price dispersion only at beliefs where the value of information is sufficiently high.

The construction of payoff-symmetric correlated Markov perfect equilibria now parallels the development in Section 6. The equilibrium value function u^* will solve the ODE

$$\frac{\tau(\pi)}{2r} u''(\pi) \xi = H\left(\pi, u(\pi) - \frac{\lambda(\pi)}{r} u'(\pi)\right) \quad (29)$$

where the right-hand side is given by the function

$$H(\pi, v) = \min \{v - m(\pi), [2b + 3c(\pi)] \xi/4\}. \quad (30)$$

As before, the boundary conditions are (4)–(5).

Proposition 8.1 *The boundary value problem consisting of the ODE (29) and the boundary conditions (4)–(5) has a unique solution $u^\dagger : [0, 1] \rightarrow \mathbb{R}_+$, and $u^\dagger \geq u^*$.*

The existence and uniqueness result is proved exactly as Proposition 6.1. The comparison result follows easily from the fact that the right-hand sides of the relevant ODEs satisfy $H \leq G$.

Let \mathcal{J} denote the set of all joint strategies $(\tilde{p}_1, \tilde{p}_2)$ with the following properties: (i) \tilde{p}_1 and \tilde{p}_2 have mean $p^m(\pi)$ and variance

$$\max \left\{ 0, \frac{u^\dagger(\pi) - \lambda(\pi)(u^\dagger)'(\pi)/r - m(\pi)}{b + c(\pi)} - \frac{2b + 3c(\pi)}{4[b + c(\pi)]} \xi \right\}$$

at all π ; (ii) $\tilde{p}_1(\pi) + \tilde{p}_2(\pi) = 2p^m(\pi)$ at all π ; (iii) \tilde{p}_1 and \tilde{p}_2 are positive; (iv) the probability of the event $\{q(\tilde{p}_i(\pi), \tilde{p}_j(\pi); k) = 0 \text{ at some } \pi\}$ is zero for all $i = 1, 2, j \neq i$ and $k = 0, 1$. Arguing as in Section 6, we see that \mathcal{J} is non-empty if and only if

$$u^\dagger(\pi) - \lambda(\pi) (u^\dagger)'(\pi)/r - m(\pi) < [2b + 3c(\pi)] \xi/4 + [b + c(\pi)] \kappa(\pi)^2 p^m(\pi)^2 \quad \text{for all } \pi, \quad (31)$$

where the function κ is defined in equation (23).

We have the following result.

Theorem 8.1 *All joint strategies in the set \mathcal{J} are correlated Markov perfect equilibria of the dynamic duopoly game. The firms' value function in all these equilibria is $u^* = u^\dagger$, and this is the maximum expected payoff that firms can achieve in any payoff-symmetric interior correlated MPE.*

That these joint strategies are equilibria is clear from the arguments given above. In any payoff-symmetric interior correlated MPE outside \mathcal{J} , firms must use independent randomization or even the myopic strategy at some beliefs where an equilibrium in

\mathcal{J} would prescribe joint randomization. The corresponding value function \tilde{u} is thus a solution of an ODE whose right-hand side is given by a ‘mixture’ of the functions G and H . Since $H \leq G$, it must be the case that H is less than or equal to this mixture, so the argument used to derive the comparison result of Proposition 8.1 also shows that u^\ddagger dominates \tilde{u} . Of course, this is intuitively clear since any element of independent randomization brings with it harmful miscoordination.

Similar to independent randomization, only the first and second moments of a joint strategy are pinned down in equilibrium. In particular, the observation of a public coin flip is sufficient for the firms to coordinate: they can use the outcome to determine which firm will charge above the myopic price and which below, and then deviate from the myopic price by a distance that generates the required price variance.

We have $u^\ddagger = u^m$ and $\mathcal{J} = \{(p^m, p^m)\}$ if and only if

$$u^m(\pi) - \lambda(\pi) (u^m)'(\pi)/r - m(\pi) \leq [2b + 3c(\pi)] \xi/4 \quad \text{for all } \pi, \quad (32)$$

so firms will not engage in joint randomization when the value of information resulting from the myopic strategies is sufficiently low. Condition (32) is tighter than its counterpart (8), reflecting once more the fact that joint randomization is more beneficial than independent randomization.

Exploiting the similarity of the relevant differential equations, we can derive counterparts of Propositions 4.1–4.2 (sufficient conditions for price matching and price dispersion), Proposition 6.2 (if firms create price dispersion, they do so over an open interval of beliefs) and the comparative statics of Section 7 for the correlated equilibria in \mathcal{J} .¹⁴ The same intuitions apply. One also shows easily that there is a range of model parameters where firms engage in joint randomization at some beliefs, but not in independent randomization. This is illustrated in Figure 6, in which independent randomization leads to no dispersion (the lower value of information curve is never above the higher threshold line $[b + c(\pi)] \xi$), whereas joint randomization leads to some dispersion (the higher value of information curve is sometimes above the lower threshold line $[2b + 3c(\pi)] \xi/4$).

9 Concluding Remarks

We considered a symmetric differentiated-goods duopoly which evolves over an infinite horizon in continuous time. Firms are uncertain about their competitive environment and can learn actively about it by creating price dispersion. We characterized all payoff-symmetric interior MPE with the firms’ common posterior belief as the state variable. We showed that these equilibria can be of two types, and provided sufficient conditions for either type. If the value of information is low, a pure-strategy no-experimentation equilibrium arises: firms behave as if they were myopic, and charge the static duopoly price corresponding to their current assessment of the competitive environment. If the value of information is high, firms use mixed strategies (randomizing either independently or jointly) to create price dispersion and thus increase the information content

¹⁴We just have to replace u^\ddagger by u^\ddagger and $b + c(\pi)$ by $[2b + 3c(\pi)]/4$ at all occurrences.

of the signals which they observe. In such an equilibrium, the duopoly goes through experimentation cycles where periods of price dispersion (which are triggered by a change in the environment) are alternating with periods of price symmetry (which begin once the firms have learnt enough about the new state). The likelihood of observing price dispersion at any given time is highest when the environment changes sufficiently often, but not too frequently.

The tractability of our model is due to a combination of several modeling choices: the assumption that there are only two possible states of demand; the linearity of expected demand curves; and the continuous-time framework with Brownian noise. While each of these assumptions simplifies the analysis of the model considerably, none of them affects our main results on the dynamics of the value of information and price dispersion.

The presence of background information in our model gives rise to an equilibrium process of posterior beliefs that is regular on the open unit interval, i.e. that can reach any point in this interval with positive probability irrespective of the current position. Without background information, this is no longer true: either the invariant belief $\tilde{\pi}$ lies outside the range where the price variance is positive, and the belief converges to $\tilde{\pi}$; or $\tilde{\pi}$ lies inside the range of positive variance, and the process of beliefs is eventually confined to this range. In the former case, price dispersion disappears for good after some time. In the latter case, price dispersion will continue uninterrupted forever. In either case, the firms track the true state poorly in the sense that their belief will be bounded away from the true state at all times. If we also let the frequency of state switches go to zero, our Markov perfect equilibria produce an extreme form of error persistence as identified originally by Rothschild (1974) and McLennan (1984). In fact, a non-degenerate prior belief at which the price variance is zero will not be updated at all, while a belief inside the range of positive variance will converge to one of the boundaries of this range. In either case, beliefs cannot converge to the truth.

We have focused on payoff-symmetric equilibria, i.e. equilibria where price dispersion is created through randomization. Asymmetric pure-strategy MPE would imply the same experimentation cycles as we have described here. Such equilibria, however, would be more complicated to analyse because of the need to solve a pair of coupled second-order differential equations for the firms' value functions.

The restriction to payoff-symmetric equilibria allows us in particular to derive analytical comparative statics. While our main focus in this paper has been on price dispersion, our methods also yield comparative statics results for equilibrium payoffs. We can show that each firm's payoff is strictly decreasing in the discount rate and the noise intensity, and strictly increasing in the quality of background information. Numerically, we find that the gains from a given increase in background information are non-monotonic in the parameter Λ : they are small for very low frequencies of state switching, increase with Λ over some range, and become smaller again as the environment becomes very unstable. This suggests that, at the margin, background information is most valuable when the environment is variable (so the information has a role in detecting a change of demand conditions more rapidly), but not too variable (so there is enough time to exploit improved knowledge of demand conditions between state switches).

The firms in our model are assumed to compete in prices, which are strategic complements. Again, our findings are not affected by this assumption. One can easily rewrite the expected demand structure with quantities as the decision variables, and then let the Brownian noise operate on realized prices. This situation with strategic substitutes implies the same dynamics of dispersion in the firms' actions as the scenario that we analyse in the paper.

It is straightforward to extend the model to an arbitrary number of firms each of which is learning about the sensitivity of its own demand to the *average* price charged by its competitors. In the absence of a fully-fledged model of (stochastic) consumer behavior, however, one has to rely on *ad hoc* assumptions when it comes to specifying how the demand structure and the noise that affects it vary with the number of firms. As the "reduced-form" approach adopted in this paper does not seem to permit any robust comparisons between oligopolies with different numbers of firms, we have only presented the two-firm case.

Appendix

A Randomization

While defining mixed strategies is straightforward, it is not obvious how randomization should be interpreted in continuous time. We use a time-alternation approach, viewing randomization as the limit of more and more frequent switching between pure strategies over time.¹⁵

By Skorohod's theorem, we can associate with any mixed strategy \tilde{p} a Lebesgue \times Borel measurable function $P : [0, 1] \times [0, 1] \rightarrow \mathbb{R}_+$ such that for any π , the distribution of $\tilde{p}(\pi)$ is the same as that of $P(\pi, \tilde{y})$ where \tilde{y} is uniformly distributed on the unit interval. That is, we can associate with \tilde{p} a parameterized family of pure strategies such that a uniform distribution over these strategies induces the same distribution over a player's actions at each belief as the mixed strategy \tilde{p} .

Time alternation can now be introduced as follows. Players start with a common prior π_0 and initial values $y_i \in [0, 1]$ ($i = 1, 2$). Player i employs the pure strategy $P(\cdot, y_i)$ until he is given the opportunity to revise the strategy by drawing a new y_i from the uniform distribution on $[0, 1]$. He uses the corresponding pure strategy until he is again called to revise it, and so on. Opportunities to revise the strategy arrive according to a Poisson process with intensity κ_i ; these processes are independent across players. On letting κ_1 and κ_2 tend to ∞ , one can show that the game converges to a limit game with the following properties: if the firms employ the mixed strategies \tilde{p}_1 and \tilde{p}_2 , then (i) firm i obtains expected flow revenues $\mathbb{E}[R(\pi_t, \tilde{p}_i(\pi_t), \tilde{p}_j(\pi_t))]$ where \mathbb{E} denotes the expectation with respect to the distribution of $\tilde{p}_1(\pi_t)$ and $\tilde{p}_2(\pi_t)$; (ii) the change in belief $d\pi_t$ has mean $\lambda(\pi_t) dt$ and variance $\tau(\pi_t) \mathbb{E}[(\tilde{p}_1(\pi_t) - \tilde{p}_2(\pi_t))^2 + \xi] dt$. In the time-alternation limit, therefore, the firms' payoffs and the law of motion of their beliefs are obtained simply by taking expectations over the mixed strategies.

This defines a limiting game in the sense of a mapping from mixed strategies into payoffs, but not a limiting *outcome*, i.e. process of prices. In fact, the attempt to construct a limiting price process runs into the usual problems of dealing with a continuum of random variables. For our purposes, however, this is of no consequence. First, the time-alternation limit still gives us a well-defined process for the state variable, which will be enough to formulate results on the emergence and long-run frequency of price dispersion. Second, there always exists a well-defined price process for any pair of *finite* arrival rates (κ_1, κ_2) , so results based on the properties of the limiting belief process can easily be interpreted in terms of price processes for sufficiently large κ_i .

¹⁵Harris (1993) contrasts this with a state-alternation approach where players use the state itself as a randomizing device. The rest of this appendix follows Harris (1993, Section 10.1).

For joint randomization, we associate with any joint strategy $(\tilde{p}_1, \tilde{p}_2)$ a Lebesgue \times Borel measurable function $P : [0, 1] \times [0, 1] \rightarrow \mathbb{R}_+ \times \mathbb{R}_+$ such that for any π , the distribution of price pairs $(\tilde{p}_1(\pi), \tilde{p}_2(\pi))$ is the same as that of $P(\pi, \tilde{y})$ where \tilde{y} is again uniformly distributed on the unit interval. The corresponding time-alternation limit is obtained by drawing values y at random times and letting the arrival rate of these times tend to infinity.

B Various Proofs

In preparation for the proofs of Propositions 4.1 and 4.2, we define

$$v^m(\pi) = u^m(\pi) - \lambda(\pi) (u^m)'(\pi)/r \quad (\text{B.1})$$

on $]0, 1[$ and

$$g(\pi) = \frac{\tilde{\pi}(1-\pi)}{\pi} + \frac{(1-\tilde{\pi})\pi}{1-\pi}$$

on $]0, 1[$. Arguing as in Keller and Rady (1999, Section 4.4), one shows easily that v^m is the unique solution of the ODE

$$\frac{\tau(\pi)}{2} v''(\pi) \xi = r [v(\pi) - m(\pi)] + \Lambda \left\{ 2g(\pi) [v(\pi) - m(\pi)] + (\pi - \tilde{\pi}) [v'(\pi) - m'(\pi)] \right\} \quad (\text{B.2})$$

subject to the boundary conditions $v^m(0) = m(0)$ and $v^m(1) = m(1)$.

Define $\check{m}(\pi) = m(\pi) + [b + c(\pi)] \xi$ on $[0, 1]$, so that (8) is equivalent to $v^m \leq \check{m}$.

PROOF OF PROPOSITION 4.1: If m is concave, then the left-hand side of (B.2) with $v = m$ is negative, while the right-hand side is zero. This means that m is a supersolution of (B.2) on $]0, 1[$. A standard argument shows that the difference $v^m - m$ cannot attain a positive maximum in the open unit interval. Since $v^m - m$ vanishes at $\pi = 0$ and 1 , we have $v^m \leq m$ on $[0, 1]$, implying (8).

Next, note that the function \check{m} is a supersolution of (B.2) if

$$\frac{\tau(\pi)}{2} m''(\pi) \leq r [b + c(\pi)] + \Lambda \left\{ 2g(\pi) [b + c(\pi)] + (\pi - \tilde{\pi})(c_1 - c_0) \right\} \quad \text{for } 0 < \pi < 1. \quad (\text{B.3})$$

The coefficient of r on the right-hand side of this inequality is clearly positive, and the same can be shown for the coefficient of Λ . Now, suppose that m is not concave throughout, and let $\hat{\pi} = \min\{\pi : c(\pi) \geq b\}$. Define the positive constants

$$\begin{aligned} \rho_0 &= \frac{1}{2} \max_{\pi \geq \hat{\pi}} \frac{(c_1 - c_0)^2 \pi^2 (1 - \pi)^2 m''(\pi)}{b + c(\pi)}, \\ \ell_0 &= \frac{1}{2} \max_{\pi \geq \hat{\pi}} \frac{(c_1 - c_0)^2 \pi^2 (1 - \pi)^2 m''(\pi)}{2g(\pi) [b + c(\pi)] + (\pi - \tilde{\pi})(c_1 - c_0)}. \end{aligned}$$

Whenever $\sigma^{-2} \leq r/\rho_0 + \Lambda/\ell_0$, we then have (B.3), hence $v^m \leq \check{m}$ by the same argument as in the previous paragraph.

Finally, we can find $K > 0$ such that the constant function $v = K$ is another supersolution of (B.2), implying $v^m \leq K$. For sufficiently large ξ , we then have $K \leq \check{m}$, hence $v^m \leq \check{m}$. \blacksquare

PROOF OF PROPOSITION 4.2: Define $\overline{m}(\pi) = (1 - \pi) m(0) + \pi m(1)$ on $[0, 1]$, and set $\bar{\xi} = \min\{\xi > 0 : \check{m}(\pi) \geq \overline{m}(\pi) \text{ for all } \pi\}$. Fix $\xi < \bar{\xi}$, so $\overline{m}(\pi) > \check{m}(\pi)$ for some π . Choose a strictly convex function $\underline{m} : [0, 1] \rightarrow \mathbb{R}$ with the following properties: $\underline{m} = m$ on $[0, \pi_\ell]$ and $[\pi_r, 1]$ with $0 < \pi_\ell < \pi_r < 1$; $\underline{m} > m$ on $]\pi_\ell, \pi_r[$; $\underline{m}(\pi) > \check{m}(\pi)$ for some π ; \underline{m} has a continuous first derivative on $[0, 1]$ and a continuous second derivative on $[1, 0] - \{\pi_\ell, \pi_r\}$. (Such a function exists because of the assumption made on ξ ; one has to choose \underline{m} close enough to \overline{m} .) Note that \underline{m} is a subsolution of (B.2) if

$$\frac{\tau(\pi)}{2} \underline{m}''(\pi) \xi \geq r [\underline{m}(\pi) - m(\pi)] + \Lambda \left\{ 2g(\pi) [\underline{m}(\pi) - m(\pi)] + (\pi - \tilde{\pi}) [\underline{m}'(\pi) - m'(\pi)] \right\} \quad \text{for } 0 < \pi < 1. \quad (\text{B.4})$$

Define the positive constants

$$\begin{aligned}\rho_1 &= \frac{1}{2}\xi \min_{\pi_\ell \geq \pi \geq \pi_r} \frac{(c_1 - c_0)^2 \pi^2 (1 - \pi)^2 \underline{m}''(\pi)}{\underline{m}(\pi) - m(\pi)}, \\ \ell_1 &= \frac{1}{2}\xi \frac{\min_{\pi_\ell \geq \pi \geq \pi_r} (c_1 - c_0)^2 \pi^2 (1 - \pi)^2 \underline{m}''(\pi)}{\max_{\pi_\ell \geq \pi \geq \pi_r} \{2g(\pi) [\underline{m}(\pi) - m(\pi)] + (\pi - \tilde{\pi}) [\underline{m}'(\pi) - m'(\pi)]\}}.\end{aligned}$$

Whenever $\sigma^{-2} \geq r/\rho_1 + \Lambda/\ell_1$, we then have (B.4), hence $v^m \geq \underline{m}$; in particular, $v^m(\pi) > \check{m}(\pi)$ for some π , so (8) does not hold. \blacksquare

PROOF OF PROPOSITION 6.1: If condition (8) holds, existence is trivial: we can simply take $u^\dagger = u^m$. Suppose therefore that (8) does not hold. We can then find positive constants K and L such that the functions $\bar{u} = K$ and $\underline{u} = u^m - L$ satisfy

$$\underline{u}(\pi) - \lambda(\pi) \underline{u}'(\pi)/r \leq \check{m}(\pi) \leq \bar{u}(\pi) - \lambda(\pi) \bar{u}'(\pi)/r$$

on $[0, 1]$. It is easy to see that \bar{u} is a supersolution of the ODE (19), and \underline{u} a subsolution, i.e.

$$\frac{\tau(\pi)}{2r} \bar{u}''(\pi) \xi \leq G(\pi, \bar{u}(\pi) - \lambda(\pi) \bar{u}'(\pi)/r)$$

and

$$\frac{\tau(\pi)}{2r} \underline{u}''(\pi) \xi \geq G(\pi, \underline{u}(\pi) - \lambda(\pi) \underline{u}'(\pi)/r).$$

Moreover, we have

$$\begin{aligned}\bar{u}(0) - \lambda(0) \bar{u}'(0)/r &> m(0), & \bar{u}(1) - \lambda(1) \bar{u}'(1)/r &> m(1), \\ \underline{u}(0) - \lambda(0) \underline{u}'(0)/r &< m(0), & \underline{u}(1) - \lambda(1) \underline{u}'(1)/r &< m(1).\end{aligned}$$

Now fix $\epsilon > 0$ such that $\lambda(\epsilon) > 0 > \lambda(1 - \epsilon)$ and

$$\begin{aligned}\bar{u}(\epsilon) - \lambda(\epsilon) \bar{u}'(\epsilon)/r &> m(\epsilon), & \bar{u}(1 - \epsilon) - \lambda(1 - \epsilon) \bar{u}'(1 - \epsilon)/r &> m(1 - \epsilon), \\ \underline{u}(\epsilon) - \lambda(\epsilon) \underline{u}'(\epsilon)/r &< m(\epsilon), & \underline{u}(1 - \epsilon) - \lambda(1 - \epsilon) \underline{u}'(1 - \epsilon)/r &< m(1 - \epsilon).\end{aligned}$$

We can find a positive constant C such that

$$|G(\pi, u_0 - \lambda(\pi)u_1/r)| \leq C(1 + |u_1|)$$

for all $\pi \in [\epsilon, 1 - \epsilon]$, $u_0 \in [\underline{u}(\pi), \bar{u}(\pi)]$ and $u_1 \in \mathbb{R}$. Since

$$\int_x^\infty \frac{s}{1+s} ds = \infty$$

for all $x \geq 0$, this implies that the ODE (19) satisfies Nagumo's condition¹⁶ on $[\epsilon, 1 - \epsilon]$ relative to \underline{u} , \bar{u} . By Bernfeld and Lakshmikantham (1974, Theorem 1.5.4), there is a solution $u : [\epsilon, 1 - \epsilon] \rightarrow \mathbb{R}$ of (19) such that $\underline{u} \leq u \leq \bar{u}$ and

$$u(\epsilon) - \lambda(\epsilon) u'(\epsilon)/r = m(\epsilon), \quad u(1 - \epsilon) - \lambda(1 - \epsilon) u'(1 - \epsilon)/r = m(1 - \epsilon).$$

Finally, consider a sequence $u_k : [\epsilon_k, 1 - \epsilon_k] \rightarrow \mathbb{R}$ of such solutions for small positive numbers $(\epsilon_k)_{k=1,2,\dots}$ converging monotonically to 0. By Bernfeld and Lakshmikantham (1974, Theorem 1.4.1), there is an $N_k > 0$ such that $|u'| \leq N_k$ on $[\epsilon_k, 1 - \epsilon_k]$ for *any* solution of (19) lying between \underline{u} and \bar{u} on this interval. Thus for any fixed integer $K \geq 1$ and all $k \geq K$, u_k is a solution satisfying $\underline{u} \leq u_k \leq \bar{u}$ and $|u'_k| \leq N_K$ on $[\epsilon_K, 1 - \epsilon_K]$, so the sequences $(u_k)_{k \geq K}$ and $(u'_k)_{k \geq K}$ are both uniformly bounded and equicontinuous on that interval. Employing the standard diagonalization argument, we obtain a

¹⁶See Bernfeld and Lakshmikantham (1974, page 25).

subsequence which converges uniformly on all compact subintervals of $]0, 1[$ to a solution u^\dagger with the desired properties.

As to uniqueness, suppose that there are two such solutions, u^\dagger and u^\ddagger , and that $u^\dagger - u^\ddagger$ has a positive maximum. As $u^\dagger(0) - \lambda(0)(u^\dagger)'(0)/r = u^\ddagger(0) - \lambda(0)(u^\ddagger)'(0)/r$ and $\lambda(0) > 0$, this maximum cannot be attained at $\pi = 0$. In fact, $u^\dagger(0) > u^\ddagger(0)$ implies $(u^\dagger)'(0) > (u^\ddagger)'(0)$ while a maximum would require $(u^\dagger)'(0) \leq (u^\ddagger)'(0)$. A similar argument rules out a maximum of $u^\dagger - u^\ddagger$ at $\pi = 1$. The maximum must therefore be attained in the interior of the unit interval, say at π_0 . This requires $u^\dagger(\pi_0) > u^\ddagger(\pi_0)$, $(u^\dagger)'(\pi_0) = (u^\ddagger)'(\pi_0)$ and $(u^\dagger)''(\pi_0) \leq (u^\ddagger)''(\pi_0)$, hence, by weak monotonicity of G in its second argument,

$$\begin{aligned} 0 &\geq \frac{\tau(\pi_0)}{2r} [(u^\dagger)''(\pi_0) - (u^\ddagger)''(\pi_0)] \xi \\ &= G(\pi_0, u^\dagger(\pi_0) - \lambda(\pi_0)(u^\dagger)'(\pi_0)/r) - G(\pi_0, u^\ddagger(\pi_0) - \lambda(\pi_0)(u^\ddagger)'(\pi_0)/r) \\ &\geq 0. \end{aligned}$$

So the maximum is in fact attained at a point where G is constant in its second argument, i.e.

$$\frac{\tau(\pi_0)}{2r} (u^\dagger)''(\pi_0) = \frac{\tau(\pi_0)}{2r} (u^\ddagger)''(\pi_0) = b + c(\pi_0).$$

Now let $\check{\pi} = \inf \{ \pi \geq \pi_0 : u^\dagger(\pi) - \lambda(\pi)(u^\dagger)'(\pi)/r < \check{m}(\pi) \}$; this is well-defined and smaller than 1 because of the boundary condition at the right end of the unit interval. The difference $u^\dagger - u^\ddagger$ is constant on $[\pi_0, \check{\pi}]$, and we have in particular $u^\dagger(\check{\pi}) > u^\ddagger(\check{\pi})$, $(u^\dagger)'(\check{\pi}) = (u^\ddagger)'(\check{\pi})$ and $(u^\dagger)''(\check{\pi}) = (u^\ddagger)''(\check{\pi})$. Immediately to the right of $\check{\pi}$, we have

$$u^\dagger(\pi) - \lambda(\pi)(u^\dagger)'(\pi)/r < \check{m}(\pi) < u^\ddagger(\pi) - \lambda(\pi)(u^\ddagger)'(\pi)/r$$

and therefore

$$\frac{\tau(\pi)}{2r} (u^\dagger)''(\pi) < b + c(\pi) = \frac{\tau(\pi)}{2r} (u^\ddagger)''(\pi).$$

This implies that $u^\dagger - u^\ddagger$ is strictly increasing immediately to the right of $\check{\pi}$ – which contradicts our assumption that this difference is maximal in π_0 . We conclude that $u^\dagger \leq u^\ddagger$ on the entire unit interval. The converse inequality, $u^\dagger \geq u^\ddagger$, follows by symmetry. \blacksquare

For the proof of Proposition 6.2, we define

$$v^\dagger(\pi) = u^\dagger(\pi) - \lambda(\pi)(u^\dagger)'(\pi)/r. \tag{B.5}$$

It is straightforward to see that $v^\dagger = v^m$ if and only if $u^\dagger = u^m$. If $v^\dagger \neq v^m$, then

$$\frac{\tau(\pi)}{2} (v^\dagger)''(\pi) = r [b + c(\pi)] + \Lambda \left\{ 2g(\pi) [b + c(\pi)] + (\pi - \check{\pi})(c_1 - c_0) \right\} \tag{B.6}$$

on the set $\{ \pi : v^\dagger(\pi) > \check{m}(\pi) \}$. In view of this, it will be convenient to use the following notation:

$$F(\pi) = r [b + c(\pi)] + \Lambda \left\{ 2g(\pi) [b + c(\pi)] + (\pi - \check{\pi})(c_1 - c_0) \right\}.$$

PROOF OF PROPOSITION 6.2: Let $\underline{\pi} = \inf \{ \pi : v^\dagger(\pi) > \check{m}(\pi) \}$ and $\bar{\pi} = \sup \{ \pi : v^\dagger(\pi) > \check{m}(\pi) \}$. Suppose that there is a belief π_0 strictly between $\underline{\pi}$ and $\bar{\pi}$ such that $v^\dagger(\pi_0) \leq \check{m}(\pi_0)$. Then there exist beliefs $\pi_1 < \pi_2 < \pi_3$ in $]\underline{\pi}, \bar{\pi}[$ such that the function $v^\dagger - \check{m}$ has positive local maxima at π_1 and π_3 , and a non-positive local minimum at π_2 . In particular, we have $v^\dagger(\pi_1) > \check{m}(\pi_1)$ and $(v^\dagger)''(\pi_1) \leq \check{m}''(\pi_1) = m''(\pi_1)$, hence

$$\tau(\pi_1)m''(\pi_1)/2 \geq \tau(\pi_1)(v^\dagger)''(\pi_1)/2 = F(\pi_1).$$

The same inequality holds at π_3 . On the other hand, we have $v^\dagger(\pi_2) \leq \check{m}(\pi_2)$, $(v^\dagger)'(\pi_2) = \check{m}'(\pi_2) = m'(\pi_2) + (c_1 - c_0)\xi$ and $(v^\dagger)''(\pi_2) \geq \check{m}''(\pi_2) = m''(\pi_2)$, hence

$$\begin{aligned} \tau(\pi_2)m''(\pi_2)/2 &\leq \tau(\pi_2)(v^\dagger)''(\pi_2)/2 \\ &= (r[v^\dagger(\pi_2) - m(\pi_2)] + \Lambda \{2g(\pi_2)[v^\dagger(\pi_2) - m(\pi_2)] + (\pi_2 - \tilde{\pi})[(v^\dagger)'(\pi_2) - m'(\pi_2)]\})/\xi \\ &\leq F(\pi_2). \end{aligned}$$

This is the desired contradiction since $\{\pi : \tau(\pi)m''(\pi)/2 \geq F(\pi)\}$ and $\{\pi : \tau(\pi)m''(\pi)/2 > F(\pi)\}$ are intervals, the latter being the interior of the former, hence containing π_2 of necessity. In fact, as the function g is strictly convex, so is F , and tedious but straightforward algebra reveals that the function $\pi \mapsto \tau(\pi)m''(\pi)/2$ is quasiconcave on the set $\{\pi : m''(\pi) \geq 0\}$. \blacksquare

C Stationary Distribution

In all the Markov perfect equilibria of Theorem 6.1, the process of posterior beliefs is a diffusion whose increment $d\pi$ has mean $\lambda(\pi)dt$ and variance $\tau(\pi)(2\text{Var}[\tilde{p}(\pi)] + \xi)dt$ with $\text{Var}[\tilde{p}(\pi)]$ given by equation (21). According to Karlin and Taylor (1981, p. 221), the stationary density of this process, if it exists, has the form

$$\psi(\pi) = \mu(\pi)[C_1S(\pi) + C_2]$$

where

$$\mu(\pi) = \frac{1}{s(\pi)\tau(\pi)(2\text{Var}[\tilde{p}(\pi)] + \xi)}$$

and

$$S(\pi) = \int_{1/2}^{\pi} s(x) dx$$

with

$$s(\pi) = \exp\left\{-\int_{1/2}^{\pi} \frac{2\lambda(x)}{\tau(x)(2\text{Var}[\tilde{p}(x)] + \xi)} dx\right\}.$$

A stationary density exists if it is possible to find constants C_1 and C_2 such that $\psi(\pi) \geq 0$ on $]0, 1[$ and $\int_0^1 \psi(\pi) d\pi = 1$.

Near the boundaries of the unit interval, $\text{Var}[\tilde{p}(\pi)]$ vanishes, and we have

$$s(\pi) \propto \exp\left\{\frac{2\sigma^2\Lambda}{(c_1 - c_0)^2\xi} \left[\frac{\tilde{\pi}}{\pi} + \frac{1 - \tilde{\pi}}{1 - \pi} + (1 - 2\tilde{\pi}) \ln \frac{\pi}{1 - \pi}\right]\right\}.$$

As we let π tend to 0, $\tau(\pi)$ decreases like π^2 , $s(\pi)$ grows like $\pi^K e^{L/\pi}$ with positive constants K and L , so $\tau(\pi)s(\pi) \rightarrow \infty$ and $\mu(\pi) \rightarrow 0$. The product $\mu(\pi)S(\pi)$ behaves like the function $f(\pi) = -\pi^{-2-K} e^{-L/\pi} \int_{\pi}^{1/2} x^K e^{L/x} dx$. The integral in this expression satisfies

$$\int_{\pi}^{1/2} x^K e^{L/x} dx > \int_{\pi}^{\pi+\pi^2} x^K e^{L/x} dx > \pi^2 \min\{x^K e^{L/x} : \pi \leq x \leq \pi + \pi^2\} = \pi^2(\pi + \pi^2)^K e^{L/(\pi+\pi^2)},$$

so $f(\pi) < -(1 + \pi)^K e^{-L/(1+\pi)}$ for π close to 0. This shows that $\limsup_{\pi \rightarrow 0} \mu(\pi)S(\pi) < 0$. To insure that $\psi(\pi)$ remains non-negative, therefore, we must have $C_1 \leq 0$. Conducting the same analysis at the other end of the unit interval, one shows $\liminf_{\pi \rightarrow 1} \mu(\pi)S(\pi) > 0$, and hence $C_1 \geq 0$. Thus, the stationary density is $\psi(\pi) = C_2\mu(\pi)$ with $C_2 = 1/\int_0^1 \mu(\pi) d\pi$.

PROOF OF PROPOSITION 7.3: Given sufficiently small r and ξ , Propositions 4.2 and 7.2 imply the existence of a cutoff level $\bar{\Lambda}$ such that the interval of beliefs where $\text{Var}[\tilde{p}(\pi)] > 0$ is non-empty if and only if $0 \leq \Lambda < \bar{\Lambda}$. Let $]\pi_\ell, \pi_r[$ be this interval when $\Lambda = 0$. As Λ increases from 0 to $\bar{\Lambda}$, the interval shrinks monotonically to the empty set, while the probability mass assigned to $]\pi_\ell, \pi_r[$ by the stationary distribution increases monotonically from its lower bound 0. \blacksquare

References

- AGHION, P., ESPINOZA, M.P. and JULLIEN, B. (1993): “Dynamic Duopoly with Learning through Market Experimentation”, *Economic Theory*, **3**, 517–539.
- BERGEMANN, D. and VÄLIMÄKI, J. (1996): “Market Experimentation and Pricing” (Cowles Foundation Discussion Paper, Yale University).
- BERGEMANN, D. and VÄLIMÄKI, J. (1997): “Market Diffusion with Two-Sided Learning”, *RAND Journal of Economics*, **28**, 773–795.
- BERNFELD, S.R. and LAKSHMIKANTHAM, V. (1974): *An Introduction to Nonlinear Boundary Value Problems* (New York and London: Academic Press).
- BOLTON, P. and HARRIS, C. (1993): “Strategic Experimentation” (STICERD Discussion Paper No. TE/93/261, London School of Economics).
- BOLTON, P. and HARRIS, C. (1999): “Strategic Experimentation”, *Econometrica*, **67**, 349–374.
- BURDETT, K. and COLES, M.G. (1997): “Steady State Price Distributions in a Noisy Search Equilibrium”, *Journal of Economic Theory*, **72**, 1–32.
- BURDETT, K. and JUDD, K. (1983): “Equilibrium Price Dispersion”, *Econometrica*, **51**, 955–969.
- CAPLIN, A. and LEAHY, J. (1993): “Sectoral Shocks, Learning, and Aggregate Fluctuations”, *Review of Economic Studies*, **60**, 777–794.
- CHAMLEY, C. and GALE, D. (1994): “Information Revelation and Strategic Delay in a Model of Investment”, *Econometrica*, **62**, 1065–1085.
- DIAMOND, P. (1987): “Consumer Differences and Prices in a Search Model”, *Quarterly Journal of Economics*, **102**, 429–436.
- FELLI, L. and HARRIS, C. (1996): “Learning, Wage Dynamics and Firm-Specific Human Capital”, *Journal of Political Economy*, **104**, 838–868.
- FISHMAN, A. and R. ROB (1998): “Experimentation and Competition”, *Journal of Economic Theory*, **78**, 299–320.
- HARRINGTON, J.E. JR. (1995): “Experimentation and Learning in a Differentiated-Products Duopoly”, *Journal of Economic Theory*, **66**, 175–288.
- HARRIS, C. (1993): “Generalized Solutions to Stochastic Differential Games in One Dimension” (working paper, Nuffield College, Oxford).
- KARLIN, S. and TAYLOR, H.M. (1981): *A Second Course in Stochastic Processes*, 2nd edition (New York: Academic Press).
- KELLER, G. and RADY, S. (1999): “Optimal Experimentation in a Changing Environment”, *Review of Economic Studies*, **66**, in print.
- LIPTSER, R.S. and SHIRYAYEV, A.N. (1977): *Statistics of Random Processes I* (New York: Springer-Verlag).
- MCLENNAN, A. (1984): “Price Dispersion and Incomplete Learning in the Long Run”, *Journal of Economic Dynamics and Control*, **7**, 331–347.
- MIRMAN, L.J., SAMUELSON, L. and SCHLEE, E. (1994): “Strategic Information Manipulation in Duopolies”, *Journal of Economic Theory*, **62**, 363–384.
- MIRMAN, L.J., SAMUELSON, L. and URBANO, A. (1993): “Duopoly Signal Jamming”, *Economic Theory*, **3**, 129–149.
- RIORDAN, M.H. (1985): “Imperfect Information and Dynamic Conjectural Variations”, *RAND Journal of Economics*, **16**, 41–50.
- REINGANUM, J. (1979): “A Simple Model of Equilibrium Price Dispersion”, *Journal of Political Economy*, **87**, 851–858.

- ROB, R. (1991): “Learning and Capacity Expansion under Demand Uncertainty”, *Review of Economic Studies*, **58**, 655–675.
- ROTHSCHILD, M. (1974): “A Two-Armed Bandit Theory of Market Pricing”, *Journal of Economic Theory*, **9**, 185–202.
- SALOP, S. and J. STIGLITZ (1976): “Bargains and Ripoffs: A Model of Monopolistically Competitive Prices”, *Review of Economic Studies*, **44**, 493–510.
- SLADE, M. (1989): “Price Wars in Price-Setting Supergames”, *Economica*, **56**, 295–310.
- SLADE, M. (1992): “Vancouver’s Gasoline Price Wars: An Empirical Exercise in Uncovering Supergame Strategies”, *Review of Economic Studies*, **59**, 257–276.

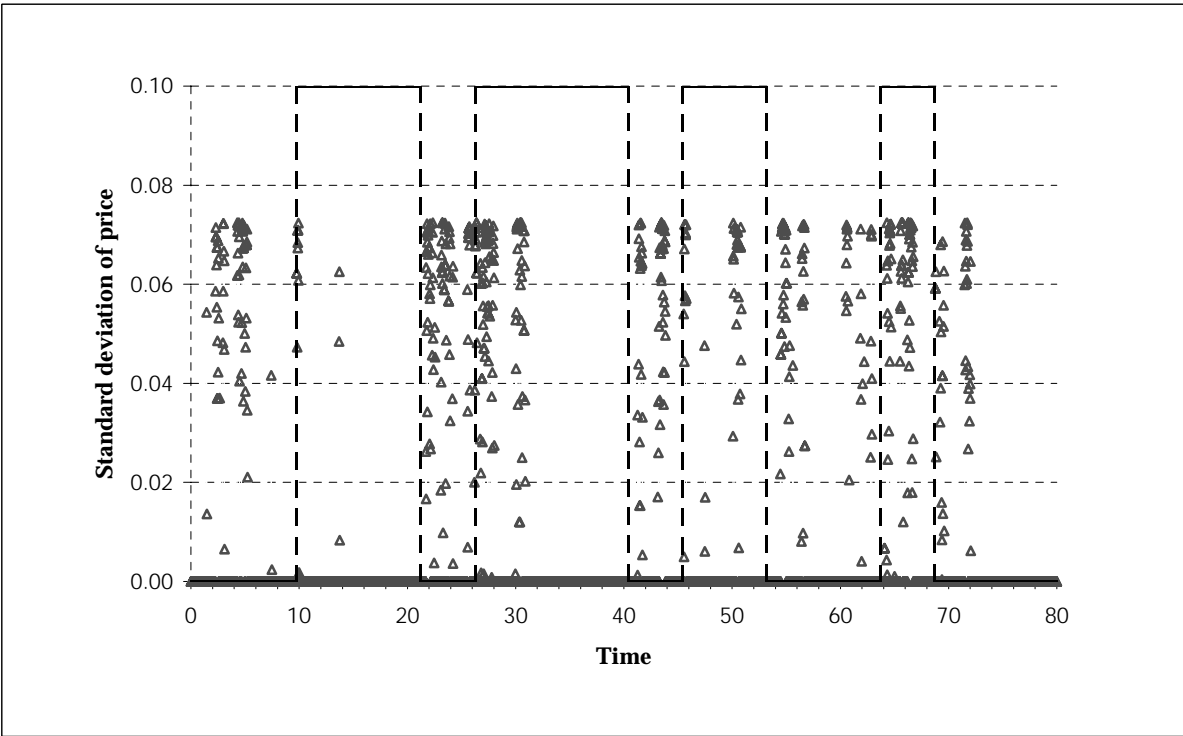



Figure 1: Sample paths for $\xi = 0.1$ and $\Lambda = 0.2$; independent randomization

The dashed line indicates the state switches. In the upper panel, the horizontal band between the dot-dash lines indicates the dispersion region. In the lower panel, the markers indicate the standard deviation of the firms' prices.




Figure 2: Value of information for $\xi = 0.1$ and $\Lambda = 0, 0.2, 0.4$;
independent randomization

The straight line is the threshold $[b + c(\pi)] \xi$.

The curves are the graphs of $u^{\dagger}(\pi) - \lambda(\pi)(u^{\dagger})'(\pi) / r - m(\pi)$ for the different values of Λ .




Figure 3: Value of information for $\xi = 0.1, 0.2$ and $\Lambda = 0.2$;
independent randomization

The straight lines are the thresholds $[b + c(\pi)] \xi$ for the two values of ξ .
The curves are the graphs of $u^\dagger(\pi) - \lambda(\pi)(u^\dagger)'(\pi) / r - m(\pi)$ for the two values of ξ .




Figure 4: Density functions of the invariant distribution of posterior beliefs for $\Lambda = 0.1, 0.2, 0.4$; $\xi = 0.1$; independent randomization




Figure 5: Likelihood of price dispersion (under the invariant distribution of posterior beliefs) as a function of Λ ; $\xi = 0.1$; independent randomization




Figure 6: Value of information for $\xi = 0.1$ and $\Lambda = 0.4$;
independent versus joint randomization

The two straight lines are the thresholds $[b + c(\pi)] \xi$ and $[2b + 3c(\pi)] \xi/4$.
The lower curve is the graph of $u^i(\pi) - \lambda(\pi) (u^i)'(\pi) / r - m(\pi)$ [independent randomization],
the upper curve that of $u^i(\pi) - \lambda(\pi) (u^i)'(\pi) / r - m(\pi)$ [joint randomization].