

DISCUSSION PAPER SERIES

No. 2919

**PRICE DISPERSION AND LEARNING
IN A DYNAMIC DIFFERENTIATED-
GOODS DUOPOLY**

R Godfrey Keller and Sven Rady

INDUSTRIAL ORGANIZATION



Centre for Economic Policy Research

www.cepr.org

Available online at:

www.cepr.org/pubs/dps/DP2919.asp

PRICE DISPERSION AND LEARNING IN A DYNAMIC DIFFERENTIATED- GOODS DUOPOLY

Godfrey Keller, University of Oxford
Sven Rady, Universität München and CEPR

Discussion Paper No. 2919
August 2001

Centre for Economic Policy Research
90–98 Goswell Rd, London EC1V 7RR, UK
Tel: (44 20) 7878 2900, Fax: (44 20) 7878 2999
Email: cepr@cepr.org, Website: www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programme in **Industrial Organization**. Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as a private educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions. Institutional (core) finance for the Centre has been provided through major grants from the Economic and Social Research Council, under which an ESRC Resource Centre operates within CEPR; the Esmée Fairbairn Charitable Trust; and the Bank of England. These organizations do not give prior review to the Centre's publications, nor do they necessarily endorse the views expressed therein.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Godfrey Keller and Sven Rady

ABSTRACT

Price Dispersion and Learning in a Dynamic Differentiated-Goods Duopoly*

We study the evolution of prices set by duopolists who are uncertain about the perceived degree of product differentiation. Customers sometimes view the products as close substitutes, sometimes as highly differentiated. As the informativeness of the quantities sold increases with the price differential, there is scope for active learning by firms. When information has low value to the firms, they 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. Such price dispersion arises in a cyclical fashion, and is most likely to be observed when the firms' environment changes sufficiently often, but not too frequently. Firms' pay-offs are higher when they use correlated pricing strategies. Contrary to what one might expect, such coordination need not hurt consumers, provided they are sufficiently impatient.

JEL Classification: C73, D43 and D83

Keywords: Bayesian learning, correlated equilibrium, duopoly experimentation, Markov perfect equilibrium and stochastic differential game

R Godfrey Keller
Department of Economics
University of Oxford
Manor Road Building
Oxford
OX1 3DQ
UK
Tel: (44 1865) 281173
Email: godfrey.keller@economics.ox.ac.uk

Sven Rady
Department of Economics
Universität München
Kaulbachstrasse 45
D-80539 Munich
GERMANY
Tel: (49 89) 2180 5642
Fax: (49 89) 2180 5650
Email: sven.rady@lrz.uni-muenchen.de

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=139115

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=126962

* Earlier versions of this work were circulated as discussion papers entitled 'Market Experimentation in a Dynamic Differentiated-Goods Duopoly'. We are grateful for comments and suggestions from an anonymous referee, Kyle Bagwell, Patrick Bolton, Jeremy Bulow, Luis Cabral, Chris Harris, Mike Harrison, Sunil Kumar, George Mailath, John Hardman Moore, George Papanicolaou, Mike Riordan, John Roberts, Larry Samuelson, Juuso Valimaki, and seminar participants at Caltech, Columbia, ECARE (Brussels), LSE, Northwestern, NYU, Penn, Rochester, Southampton, Stanford, Yale, the European Summer Symposium in Economic Theory in Gerzensee, 1998, the European Meeting of the Econometric Society 1998, and the Annual Conference of the European Economic Association 1998. We would like to thank the Financial Markets Group at the London School of Economics and the Studienzentrums Gerzensee for their hospitality, and Eric Stout for valuable research assistance.

Submitted 20 July 2001

1 Introduction

When firms are uncertain about their competitive environment, the choice of an action influences their overall profits in two ways. On the one hand, the action determines *current* profits. On the other hand, it may determine the information content of the market outcome, and thus *future* profits. This gives rise to a trade-off between short-term strategic incentives and long-term information-gathering incentives. In the present paper, we study this trade-off in the context of dynamic price competition. We show that firms' desire to learn about changing market conditions can generate distinct cyclical pricing patterns, where transient bouts of price dispersion alternate with phases of price matching.

Specifically, we consider a symmetric differentiated-goods duopoly in an infinite-horizon, continuous-time setting. 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 changes randomly over time. In each period, the quantities sold provide a noisy signal of the current state of demand, and the information content of the market outcome increases with the difference between the firms' prices. If they match each other's price, the market outcome is completely uninformative, but additionally, firms have access to some background information which is revealed independently of their pricing. There is a conflict between short-term and long-term incentives because short-run profit maximization requires both firms to charge the same price, while they must set different prices if they want to generate additional information about the state of demand. Equilibrium behavior must resolve this conflict of incentives.

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 Markov perfect equilibria (MPE) that are payoff-symmetric, i.e. that give the firms symmetric expected continuation payoffs after each history. At any given belief, such an equilibrium involves either price matching or random price dispersion. If firms match each other's price, this price must necessarily be the one that myopic firms would charge. This is because the marginal informational benefit generated by a deviation from price matching is zero in our model. Price matching arises in equilibrium at beliefs where the value of information to firms is sufficiently small. Whenever the value of information reaches a certain threshold, firms have an incentive to deviate from price matching; any Markov perfect equilibrium must then involve active learning through price dispersion. In this case, payoff symmetry requires the use of mixed strategies. 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, and these moments are uniquely determined in equilibrium. The average price is equal to the price prescribed by the myopic strategy, and the price variance is positive on an open interval of "diffuse" beliefs.

We provide explicit conditions on the model parameters that govern whether or not there is price dispersion in equilibrium. Principally, price dispersion arises 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.

Any MPE with price dispersion 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 an episode of positive price dispersion, i.e. active learning. This ends once they have become fairly certain of the change and charge identical prices again.¹ Thus, we obtain experimentation cycles where bouts of price dispersion alternate with phases of price matching, the latter occurring towards the extremes of the equilibrium price range.

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' 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 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 price-matching equilibrium arises. Overall, we find a hump-shaped relationship between the instability of the environment and the fraction of time that price dispersion is observed.

In turn, the shape of the frequency distribution of observed prices is closely related to the instability 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. If the market is quite stable, therefore, the model predicts high amplitude price cycles with price dispersion occurring infrequently but over a wide range of prices; in a more volatile market, the amplitude of the price cycles is smaller, while price dispersion arises more frequently but over a narrower range of prices. In a very unstable market, finally, the amplitude of the price cycles drops even further and price dispersion disappears

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

completely.

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 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.

The demand structure in our model implies that consumer surplus in a given period is convex in prices. As a consequence, joint randomization, with its higher price dispersion, leaves consumers better off in the short run than independent mixing. This leads to the surprising result that the more “collusive” type of equilibrium is actually beneficial to those consumers who discount the future sufficiently strongly.

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

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,

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

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.

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. We characterize equilibria over the full range of model parameters (including low discount rates), provide explicit sufficient conditions on the model parameters for each type of equilibrium, and analyse the comparative statics of equilibrium payoffs and price dispersion. Most important, we are able to address the dynamics and long-run properties of price dispersion.

The resulting interpretation of cyclical pricing patterns is distinct from that offered by theories of collusion over the business cycle (Green and Porter, 1984; Rotemberg and Saloner, 1986; Bagwell and Staiger, 1997) or sales (Sobel, 1984). Assuming imperfectly observable demand, Green and Porter explain cyclical movements in the price level as the alternation between collusive phases (when prices are relatively high) and punishment phases (when prices are relatively low). Rotemberg and Saloner as well as Bagwell and Staiger study setups with observable demand and explain price cycles by changes in the incentive to deviate from the collusive price. Sobel explains occasional price cuts as firms' attempt to sell a durable good to low-value customers that have not been served before. Price dispersion is incidental to all these theories, which above all are concerned with fluctuations in the price *level*.³ If a firm ever deviated from the "predominant" price, it would do so in the hope of reaping higher short-term profits. In contrast, we stress price *differences* across firms, and these differences arise not because firms hope to gain a short-term advantage, but because they make an investment with long-term payoffs.

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, 1997; 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

³Sobel's (1984) theory, for example, is compatible with equilibria where all firms lower their price at the same time and to the same level.

to experiment, and the informational externality would be completely irrelevant. 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 (1997, 2000). 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.⁴ 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 3 introduces the benchmark of myopic behavior. Section 4 describes the pay-off symmetric MPE of the dynamic pricing game. Section 5 provides necessary and sufficient conditions for equilibria with and without price dispersion. Section 6 derives comparative statics for price dispersion. Section 7 studies joint randomization and its effect on consumer welfare. Section 8 contains some concluding remarks. Technical details and proofs are relegated to the Appendix.

2 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

$$q_i = q(p_i, p_j; k) = a - bp_i - c_k(p_i - p_j)$$

where the intercept a , the slope b and the two possible substitutability parameters $c_0 < c_1$ are positive constants.⁵ 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.

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

⁴Bolton and Harris (1999) give agents this capability by allowing them to allocate a *fraction* of each instant to experimentation, which is isomorphic to randomization.

⁵This 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 $b = 0$. For reasons that will become obvious in a while, we will not impose non-negativity constraints on prices or expected quantities.

probabilities

$$\begin{aligned}\Pr(k_{t+\Delta t} = 0 \mid k_t = 0) &= 1 - \lambda_0 \Delta t + o(\Delta t), & \Pr(k_{t+\Delta t} = 1 \mid k_t = 0) &= \lambda_0 \Delta t + o(\Delta t), \\ \Pr(k_{t+\Delta t} = 0 \mid k_t = 1) &= \lambda_1 \Delta t + o(\Delta t), & \Pr(k_{t+\Delta t} = 1 \mid k_t = 1) &= 1 - \lambda_1 \Delta t + o(\Delta t)\end{aligned}$$

where $\lambda_k > 0$ ($k = 0, 1$) 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 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.⁶ 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.⁷ 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

⁶This 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. Note that the full support of the normal distribution implies a positive probability that quantities actually demanded are negative, even if the expected demand is positive. This is why there is no point imposing a non-negativity constraint on expected quantities in the first place. Of course, normal noise distributions are chosen for the sake of tractability of the updating problem.

⁷Technically, 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' payoff functions are non-degenerate.

loss of generality, so revenue equals profit. Firm i 's total payoff, expressed in per-period terms, is then

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

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 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 firms' prices depend only on π_t , the subjective probability that $k_t = 1$, and possibly the outcome of some randomization device. More precisely, a *pure strategy* is a Lebesgue measurable mapping from the interval $[0, 1]$ of possible beliefs to the real line. A *mixed strategy* is a Lebesgue measurable function from the state space $[0, 1]$ to the set of Borel probability measures on the real line. 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.⁸ In this limit, the stochastic process of beliefs is a diffusion with the following characteristics. If the current belief is π and the firms use mixed strategies \tilde{p}_1 and \tilde{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 $\pi^2(1 - \pi)^2 \mathbb{E}[S(\tilde{p}_1(\pi), \tilde{p}_2(\pi))] dt$ where

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

and

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

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, $\mathbb{E}[S(\tilde{p}_1(\pi), \tilde{p}_2(\pi))]$, is an expected 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.⁹ If firms charge the same price, the state

⁸See Appendix A for details.

⁹In particular, a firm may want to price below marginal cost in order to make the quantity signal more informative. As we have normalized marginal cost to zero, it would therefore not be appropriate to impose a non-negativity constraint on prices.

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 switch over the next infinitesimal period of time. Given the current belief π , they assign probability $1 - \pi$ to state 0, hence probability $(1 - \pi)\lambda_0 dt$ to a transition from state 0 to state 1 over the next instant dt ; in the same way, they assign probability $\pi\lambda_1 dt$ 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 $\lambda(\pi) dt$. The linear function λ 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.

With a belief π , the firms' best assessment of their competitive environment is

$$c(\pi) = (1 - \pi)c_0 + \pi c_1.$$

Now, given two mixed strategies \tilde{p}_1 and \tilde{p}_2 and a common prior belief π , firm i 's total payoff can be written as

$$u_i(\pi) = \mathbb{E} \left[\int_0^\infty r e^{-rt} \mathbb{E}[R(\pi_t, \tilde{p}_i(\pi_t), \tilde{p}_j(\pi_t))] dt \mid \pi_0 = \pi \right]$$

where

$$\begin{aligned} R(\pi, p_i, p_j) &= p_i [(1 - \pi)q(p_i, p_j; 0) + \pi q(p_i, p_j; 1)] \\ &= p_i [a - bp_i - c(\pi)(p_i - p_j)] \end{aligned}$$

is firm i 's expected instantaneous revenue given belief π and prices p_i and p_j . Note that the first expectations operator in the above expression is with respect to the law of motion of beliefs as described above, and the second with respect to the distribution of prices given a belief.

More useful for our analysis of payoff functions and equilibrium strategies is the characterization of u_i as a solution to the ODE

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

For pure strategies, this is a standard result; a heuristic derivation can be found in Keller and Rady (1999, Section 4). The extension to mixed strategies is then very intuitive: what matters is the average signal-to-noise ratio and the average current revenue given the firms' price distributions. Equation (1) shows that a firm's payoff has three components. The first component 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 second term represents the value of information extracted from the noisy signals available to the firm. Indeed, $\mathbb{E}[S(\tilde{p}_i(\pi), \tilde{p}_j(\pi))]$ measures the expected information content of the signals obtained after setting prices. This information is valued at the shadow price $\pi^2(1-\pi)^2 u_i''(\pi)/(2r)$. The third term is simply the expected current payoff.

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

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

which 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 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,¹⁰ 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 Equilibria of the Dynamic Pricing Game

Our next aim is to characterize Markov perfect equilibria when firms are forward-looking and take into account that their actions not only provide current rewards, but also information about the underlying state of demand. First, we will show that firms’ expected payoffs depend on the mean and variance of prices only. In a second step, we will see that the average price charged by each firm must equal the myopic price. Finally, we will pin down price variances by imposing payoff symmetry.

Suppose the firms employ the mixed strategies \tilde{p}_1 and \tilde{p}_2 . Write $\bar{p}_i(\pi)$ for the mean and $V_i(\pi)$ for the variance of firm i ’s price given belief π . Straightforward computations show that

$$\begin{aligned} \mathbb{E}[R(\pi, \tilde{p}_i(\pi), \tilde{p}_j(\pi))] &= a \bar{p}_i(\pi) - [b + c(\pi)] [\bar{p}_i(\pi)^2 + V_i(\pi)] + c(\pi) \bar{p}_i(\pi) \bar{p}_j(\pi) \\ &= R(\pi, \bar{p}_i(\pi), \bar{p}_j(\pi)) - [b + c(\pi)] V_i(\pi) \end{aligned} \quad (2)$$

¹⁰By 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.

and

$$E[S(\tilde{p}_i(\pi), \tilde{p}_j(\pi))] = \frac{(c_1 - c_0)^2}{\sigma^2} \left[[\bar{p}_1(\pi) - \bar{p}_2(\pi)]^2 + V_i(\pi) + V_j(\pi) + \xi \right]. \quad (3)$$

Thus, expected current rewards and the information content of observed quantities depend on firms' strategies only through the means and variances of prices (there are no covariance terms because firms randomize independently). We also see the fundamental trade-off that firms face: creating random price dispersion reduces expected current revenues, but increases the information content of observed quantities.

Using (2)–(3), we can reformulate the dynamic pricing game as a game where firms just choose functions \bar{p}_i and V_i that determine average prices and price variances, respectively, for all possible beliefs. Current reward and signal-to-noise ratio in this mean-variance game are given by the right-hand side of (2) and (3), respectively, so the ODE (1) for a firm's payoff function becomes

$$\begin{aligned} u_i(\pi) &= \lambda(\pi) \frac{u_i'(\pi)}{r} \\ &+ \pi^2(1 - \pi)^2 \frac{u_i''(\pi)}{2r} \frac{(c_1 - c_0)^2}{\sigma^2} \left[[\bar{p}_1(\pi) - \bar{p}_2(\pi)]^2 + V_i(\pi) + V_j(\pi) + \xi \right] \\ &+ R(\pi, \bar{p}_i(\pi), \bar{p}_j(\pi)) - [b + c(\pi)] V_i(\pi). \end{aligned} \quad (4)$$

Moreover, a standard argument implies that (\bar{p}_i, V_i) is a best response to (\bar{p}_j, V_j) if and only if $(\bar{p}_i(\pi), V_i(\pi))$ maximizes the right-hand side of (4) at almost all π , optimally balancing the future benefits of information against short-term profits.

It is easy to see that this maximization splits in two separate problems: to choose a mean price $\bar{p}_i(\pi)$ that maximizes

$$\pi^2(1 - \pi)^2 \frac{u_i''(\pi)}{2r} \frac{(c_1 - c_0)^2}{\sigma^2} [\bar{p}_i(\pi) - \bar{p}_j(\pi)]^2 + R(\pi, \bar{p}_i(\pi), \bar{p}_j(\pi)) \quad (5)$$

given the opponent's mean price $\bar{p}_j(\pi)$, and to choose a price variance $V_i(\pi)$ that maximizes

$$\left(\pi^2(1 - \pi)^2 \frac{u_i''(\pi)}{2r} \frac{(c_1 - c_0)^2}{\sigma^2} - [b + c(\pi)] \right) V_i(\pi). \quad (6)$$

Existence of a best-response variance for firm i clearly requires

$$\pi^2(1 - \pi)^2 \frac{u_i''(\pi)}{2r} \frac{(c_1 - c_0)^2}{\sigma^2} \leq b + c(\pi), \quad (7)$$

and the optimal variance $V_i(\pi) = 0$ whenever the inequality in (7) is strict. This implies that (6) is always zero in equilibrium. If a firm creates random price dispersion and thus generates additional information, the benefit that it derives from this information is just offset by the corresponding expected loss in revenue.

For the maximization of (5), condition (7) implies that the objective function is a concave quadratic, so the first-order condition

$$\pi^2(1 - \pi)^2 \frac{u_i''(\pi)}{r} \frac{(c_1 - c_0)^2}{\sigma^2} [\bar{p}_i(\pi) - \bar{p}_j(\pi)] + a - 2[b + c(\pi)] \bar{p}_i(\pi) + c(\pi) \bar{p}_j(\pi) = 0 \quad (8)$$

is necessary and sufficient for $\bar{p}_i(\pi)$ to be a best response to $\bar{p}_j(\pi)$. In particular, it can be seen that setting a mean price equal to the myopic price, i.e. $\bar{p}_i(\pi) = p^m(\pi)$, is a best response to the other firm doing the same, i.e. to $\bar{p}_j(\pi) = p^m(\pi)$. Generically, $\bar{p}_i(\pi) = \bar{p}_j(\pi) = p^m(\pi)$ is in fact the unique pair of mean prices that solves both firms' first-order conditions.¹¹

Inserting these mean prices in (4) and using the fact that (6) is zero, we obtain the following differential equation for the equilibrium payoff function of firm i :

$$u_i(\pi) - \lambda(\pi) \frac{u'_i(\pi)}{r} - m(\pi) = \pi^2(1 - \pi)^2 \frac{u''_i(\pi)}{2r} \frac{(c_1 - c_0)^2}{\sigma^2} [V_j(\pi) + \xi]. \quad (9)$$

This shows that in equilibrium, the value to firm i of the dynamic game, adjusted for the possibility of a state switch and net of the myopic payoff, equals the net value to firm i of information extracted from the noisy signals. Since the net benefit that firm i derives from its own randomization is zero, only the price variance of the other firm contributes to firm i 's payoff, and it does so through its impact on the speed of learning.

Because of the perfect symmetry of the model, it is natural to consider equilibria where the firms' payoff functions are identical, i.e. $u_1 = u_2 = u$. Equation (9) shows that in such a payoff-symmetric equilibrium, the price variances must be identical as well, i.e. $V_1 = V_2 = V$.

There are then two cases at any given belief π , depending on whether $V(\pi)$ is zero or positive. If $V(\pi) = 0$, then

$$u(\pi) - \lambda(\pi) \frac{u'(\pi)}{r} - m(\pi) = \pi^2(1 - \pi)^2 \frac{u''(\pi)}{2r} \frac{(c_1 - c_0)^2}{\sigma^2} \xi \leq [b + c(\pi)] \xi$$

by (9) and the inequality (7). If $V(\pi) > 0$, then

$$u(\pi) - \lambda(\pi) \frac{u'(\pi)}{r} - m(\pi) > \pi^2(1 - \pi)^2 \frac{u''(\pi)}{2r} \frac{(c_1 - c_0)^2}{\sigma^2} \xi = [b + c(\pi)] \xi$$

by (9) and the fact that (7) holds as an equality; moreover, (9) implies

$$V(\pi) = \frac{u(\pi) - \lambda(\pi) u'(\pi)/r - m(\pi)}{b + c(\pi)} - \xi.$$

Overall, therefore, the common equilibrium payoff function u must be a solution to the ODE

$$\pi^2(1 - \pi)^2 \frac{u''(\pi)}{2r} \frac{(c_1 - c_0)^2}{\sigma^2} \xi = \min \left\{ u(\pi) - \lambda(\pi) \frac{u'(\pi)}{r} - m(\pi), [b + c(\pi)] \xi \right\}. \quad (10)$$

¹¹The corresponding linear system has full rank unless

$$\pi^2(1 - \pi)^2 \frac{u''_1(\pi) + u''_2(\pi)}{r} \frac{(c_1 - c_0)^2}{\sigma^2} = 2b + 3c(\pi).$$

The relevant boundary conditions are

$$u(0) - \lambda(0) u'(0)/r = m(0), \quad (11)$$

$$u(1) - \lambda(1) u'(1)/r = m(1). \quad (12)$$

They have the natural interpretation that under subjective certainty, the net value of information extracted from the noisy signals is zero, so optimal and myopic payoffs coincide (up to the possibility of a state switch).

We have the following result.

Proposition 4.1 *The boundary value problem consisting of the ODE (10) and the boundary conditions (11)–(12) has a unique solution u^* . A pair of mixed strategies constitutes a payoff-symmetric Markov perfect equilibrium of the dynamic pricing game if and only if both strategies prescribe a mean price equal to the myopic price $p^m(\pi)$ and a price variance equal to*

$$V^*(\pi) = \max \left\{ 0, \frac{u^*(\pi) - \lambda(\pi)(u^*)'(\pi)/r - m(\pi)}{b + c(\pi)} - \xi \right\} \quad (13)$$

at almost all beliefs $\pi \in [0, 1]$. The firms' common payoff function in any such equilibrium is u^ , and the set of beliefs where firms create price dispersion is either empty or an open interval strictly smaller than the open unit interval.*

Existence and uniqueness of u^* are proved in Appendix B. As to the remainder of the proposition, we have already shown that each payoff-symmetric MPE must have the stated properties. Conversely, each pair of strategies with these properties gives both firms the payoff function u^* and has the mutual best-response property, hence constitutes a payoff-symmetric MPE. Finally, we prove in Appendix B that if there are beliefs at which V^* is positive, these beliefs form an open interval with left boundary strictly bigger than 0, and right boundary strictly smaller than 1.

This last part of Proposition 4.1 is very intuitive. At beliefs close to subjective certainty, firms do not value information very highly, hence will not create price dispersion. If we are to see price dispersion at all, it will have to occur at diffuse beliefs where firms are fairly uncertain about the state of demand.

5 Price Matching vs. Price Dispersion

As it stands, the characterization of equilibria in Proposition 4.1 does not tell us whether there will actually be price dispersion. The aim of this section is to provide sufficient conditions for equilibria with and without price dispersion.

Proposition 4.1 implies that the only candidate for a Markov perfect equilibrium without price dispersion is price matching at the myopic price. If both firms adopt the myopic pricing strategy at all beliefs, their common payoff function is

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

and solves the ODE

$$\pi^2(1-\pi)^2 \frac{u''(\pi)}{2r} \frac{(c_1 - c_0)^2}{\sigma^2} \xi = u(\pi) - \lambda(\pi) \frac{u'(\pi)}{r} - m(\pi), \quad (14)$$

which is a special case of (9), subject to the boundary conditions (11)–(12).

Corollary 5.1 *Price matching at the myopic price is an equilibrium if*

$$u^m(\pi) - \lambda(\pi) \frac{(u^m)'(\pi)}{r} - m(\pi) \leq [b + c(\pi)] \xi \quad \text{for all } \pi. \quad (15)$$

If (15) does not hold, then any Markov perfect equilibrium involves price dispersion.

This is an immediate consequence of Proposition 4.1 since u^m solves the ODE (10), and hence is identical with u^* , if and only if (15) holds.

The role of condition (15) is easy to understand: it ensures that no firm has an incentive to deviate from myopic pricing. Combined with the ODE (14), in fact, (15) implies that the shadow price of information under myopic pricing is low enough for myopic pricing to be an equilibrium. If condition (15) does not hold, on the other hand, the shadow price of information under myopic pricing exceeds the critical value $b + c(\pi)$ at some beliefs, and deviations from the myopic strategy become profitable.

Condition (15) 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.

The proposition below gives three conditions, each of which is sufficient for price matching at the myopic price to be an equilibrium. 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. 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 – this is the content of the second condition. The third condition confirms the intuition that firms will not experiment if they can rely on a very precise background signal.

Proposition 5.1 *Price matching at the myopic price is an equilibrium (i) if the myopic payoff function m is concave throughout ($c_1 \leq b$). When m has a convex part ($c_1 > b$), there are positive constants ρ_m , ℓ_m and ξ_m such that price matching at the myopic price is an equilibrium (ii) if $r/\rho_m + \Lambda/\ell_m \geq \sigma^{-2}$, or (iii) if $\xi \geq \xi_m$.*

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 5.2 *Suppose that the myopic payoff function m is convex throughout ($c_0 \geq b$). Then there are positive constants ρ_d , ℓ_d and ξ_d such that all Markov perfect equilibria involve price dispersion whenever $r/\rho_d + \Lambda/\ell_d \leq \sigma^{-2}$ and $\xi < \xi_d$.*

An equilibrium with price dispersion is illustrated in Figures 1 and 3.¹² Figure 1, with $\xi = 0.1$ and $\Lambda = 0.1$, plots the standard deviation of each firm's price as a function of beliefs. Price dispersion arises whenever beliefs lie in the open interval ranging from 0.2 to 0.6, approximately, and the standard deviation reaches a maximum of about 0.17 near the belief 0.4. The upper panel of Figure 3 shows beliefs evolving over time, starting from the common prior $\pi_0 = 0.75$. (The bold dashed line in either panel of Figure 3 represents the true state, the initial state being $k_0 = 0$.) The horizontal band between the dot-dash lines indicates the interval of diffuse beliefs where the firms create price dispersion. Whenever beliefs traverse this band, the firms' prices have a positive standard deviation, shown in the lower panel. As the prior belief and the initial state are on opposite sides of the band, the background signal pulls the belief towards the truth. After a short while, the firms become fairly unsure of the true state and engage in a bout of price dispersion. This moves them quite rapidly towards the truth, at which point they subside into a phase of price matching. However, even before the state changes for the first time, the firms are hit by short-term demand shocks and are also pulled back towards the invariant belief, and we see some further brief bouts of price dispersion. Later, a change in the environment triggers another period of active learning, once again followed by lengthy phases of no price dispersion when the firms have learnt enough about the new state. Thus, a cyclical pattern emerges where short bursts of price dispersion, triggered by a change in the environment, alternate with longer periods of price matching. The latter are occasionally interrupted when price dispersion is triggered by demand shocks or mean reversion of beliefs. For the parameter values we use, this is more likely to occur in state 0 rather than state 1 for two reasons: first, the interval of beliefs where dispersion arises is asymmetric and closer to 0 than 1; secondly, the invariant belief ($\tilde{\pi} = 0.5$) is much farther from the lower limit than the upper limit of this interval.

A second example is given in Figures 2 and 4, for the same level of background signal ($\xi = 0.1$) but a higher variability of the environment ($\Lambda = 0.2$). Figure 2 shows us that the range of beliefs where price dispersion arises is narrower, and that, where it does occur, it is smaller. Nevertheless, the cyclical pattern evident in the sample paths in Figure 4 is qualitatively the same as in the first example. Of course, the fact that the price variance at any belief is always lower manifests itself in the lower panel; yet, given that, the amount of price dispersion over time seems to be broadly the same.

We shall further explore these and other comparative statics in the next section.

6 Comparative Statics of Price Dispersion

Our next goal is to study how the occurrence of price dispersion varies with the basic model parameters. We shall say that price dispersion decreases for a given parameter

¹²In these and all subsequent figures, 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 $\tilde{\pi} = 0.5$, and the noise parameter is $\sigma = 0.1$. Only Λ , ξ and the type of randomization vary across figures.

change if, starting from a situation where there is price dispersion, (i) the variance of prices decreases at those beliefs where it is positive, and (ii) the interval of beliefs where the variance is positive shrinks.

Proposition 6.1 *Suppose that the myopic payoff function m is convex throughout ($c_0 \geq b$). Then price dispersion decreases in the firms' discount rate, the frequency of state switches (for sufficiently small discount rates), the noise level, and the quality of background information.*

The intuition behind these comparative statics is straightforward. As the discount rate increases, the future becomes less important to the firms. This reduces their willingness to sacrifice current revenue for potential future gains from experimentation, and so they create less price dispersion. A higher frequency of state switches increases the risk of information becoming obsolete, so the trade-off between current revenues and potential gains from experimentation again shifts in favor of the former. A higher level of noise renders the quantity signals less informative, which also decreases the expected benefit of any given amount of price dispersion. Finally, if better background information is available for free, the shadow price of information falls, and there is less of a reason to incur the opportunity cost of experimentation; on the margin, it is cheaper to rely on the public signal.¹³

Figure 5 illustrates how price dispersion decreases as the environment becomes less stable, i.e. as Λ increases. The four curves shown there are the graphs of the function $[u^*(\pi) - \lambda(\pi)(u^*)'(\pi)/r - m(\pi)]/[b + c(\pi)]$ for $\Lambda = 0, 0.1, 0.2$ and 0.4 ; a horizontal line at the threshold level ξ is also plotted. By equation (13), firms' price variance is positive whenever the curve is above the threshold; the variance is then the vertical distance between the curve and the line. As Λ increases from 0 to 0.1 (the benchmark level underlying Figure 1), the price variance decreases at all beliefs where it was positive, and the range of beliefs where firms create price dispersion shrinks. A further increase in Λ to 0.2 leads to a further decrease in price dispersion. (The comparison between the middle two curves here echoes the comparison between Figures 1 and Figure 2, which plot standard deviations.) For $\Lambda = 0.4$, they create no price dispersion at all.

For increases in the firms' discount rate, r , or the noise level, σ , we obtain graphs that look exactly like Figure 5, and we therefore do not present them here.

The effect of changes in the quality of background information is illustrated in Figure 6, which plots $[u^*(\pi) - \lambda(\pi)(u^*)'(\pi)/r - m(\pi)]/[b + c(\pi)]$ for $\xi = 0.1$ (the benchmark level underlying Figure 1) and 0.2 together with the corresponding threshold lines. Note that the higher ξ is associated with the higher curve; this is because firms are better off when more information is available for free. But since the threshold that triggers randomization shifts upward by more, price dispersion diminishes overall: while the lower curve is above the lower threshold line over some range of beliefs, the higher curve never reaches the higher threshold.

¹³Numerical solutions suggest that Proposition 6.1 continues to hold when the myopic payoff function is not convex throughout, but merely has a convex part ($c_1 > b$). It also appears that with regard to changes in Λ , the restriction to small discount rates in the proposition can be dropped.

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 proposition). 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 $\{\pi : V^*(\pi) > 0\}$ under the invariant distribution of the belief process. Given the firms' strategies, we can characterize this distribution using standard results from the theory of diffusion processes; see Appendix C for details.

Figure 7 illustrates how the shape of the long-run frequency distribution of beliefs depends on the stability of the environment. In rather stable environments, this distribution is strongly bimodal with most of the probability mass concentrated at the two extremes of the range. As the environment becomes less stable, more beliefs lie in an intermediate range, and the belief distribution eventually becomes unimodal with most of the probability mass concentrated in the middle of the equilibrium price range. However, as we showed above, for very stable environments the range of beliefs where firms engage in price dispersion is quite wide, whereas for less stable environments it is much narrower (about 0.2 to 0.6 for $\Lambda = 0.1$ compared with about 0.3 to 0.5 for $\Lambda = 0.2$; see Figure 5). So, on the one hand we have low probabilities over a wide range, and on the other we have higher probabilities over a narrower range. The trade-off between these two effects is the content of the next proposition, where 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 6.2 *Suppose that the myopic payoff function m is convex throughout ($c_0 \geq b$). For a sufficiently small interest rate and precision of the background signal, the likelihood of observing price dispersion first increases, but eventually decreases with the frequency of state switches.*

As mentioned in the preamble to the proposition, an increase in the frequency of changes in the environment has two effects: first, it becomes harder for beliefs to track the true state, implying that they stay more often in the interval where firms create price dispersion; second, information 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 8, which plots the likelihood of observing price dispersion as a function of Λ . The sample paths shown in Figures 3 and 4 are consistent with these findings. The former has few transits through a broad band of beliefs, the latter nearly twice as many transits through a band about half the width. The likelihood of price dispersion will thus be determined by the speed of these transits and the figures read off from Figure 8 indicate that this is higher the greater the intensity of demand switches – 8% likelihood for $\Lambda = 0.2$ as opposed to 10% for $\Lambda = 0.1$ indicates more rapid transits for the higher of the two intensities, and this is borne out by the evidence from the lower panels of the two sample paths.

The shape of the long-run frequency distribution of observed prices is clearly directly linked to that of beliefs in Figure 7, skewed higher when firms believe they have some

monopoly power and lower when competition is keener. Figure 8 predicts that observed price dispersion should be small for both a low and a high frequency of state changes, but in the former case we are highly likely to observe extreme prices and in the latter case prices are more likely to reside in the middle range, and so these two cases are easily distinguished.

7 Joint Randomization

So far, we have studied independent randomization by the two firms. Clearly, we expect them to achieve higher 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. Interestingly, consumers can also benefit from coordinated randomization, provided they are impatient.

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. Following the same approach as for independent randomization, we interpret joint randomization as the limit of more and more frequent switching between pairs of pure strategies. This means that firms' payoff functions given a joint strategy $(\tilde{p}_1, \tilde{p}_2)$ solve the ODE (1) exactly as in the case of independent mixing. The only difference is that with a joint strategy, the expected reward $E[R(\pi, \tilde{p}_i(\pi), \tilde{p}_j(\pi))]$ and signal-to-noise ratio $E[S(\tilde{p}_i(\pi), \tilde{p}_j(\pi))]$ at a given belief depend not just on means and variances, but also on the covariance of prices.

A joint strategy constitutes a *joint-action* Markov perfect equilibrium if all price pairs assigned by the public randomization device are mutually best responses.¹⁴ For a payoff-symmetric joint-action equilibrium $(\tilde{p}_1, \tilde{p}_2)$ with common payoff function u , this means that given a belief π , any price pair (p_1, p_2) in the support of $(\tilde{p}_1(\pi), \tilde{p}_2(\pi))$ must have the following property: p_i , the price assigned to firm i , maximizes the sum

$$\pi^2(1 - \pi)^2 \frac{u''(\pi)}{2r} S(p_i, p_j) + R(\pi, p_i, p_j),$$

striking the optimal trade-off between value of information and current revenue given p_j , the price assigned to firm j .

Firm i 's first-order condition is

$$\pi^2(1 - \pi)^2 \frac{u''(\pi)}{r} \frac{(c_1 - c_0)^2}{\sigma^2} [p_i - p_j] + a - 2[b + c(\pi)] p_i + c(\pi) p_j = 0.$$

The linear system consisting of the two firms' first-order conditions has full rank and the unique solution $p_1 = p_2 = p^m(\pi)$ unless

$$\pi^2(1 - \pi)^2 \frac{u''(\pi)}{2r} \frac{(c_1 - c_0)^2}{\sigma^2} = \frac{2b + 3c(\pi)}{4}, \quad (16)$$

¹⁴This can be seen as the special case of a *correlated* equilibrium where the conditional distribution of one firm's prices given the other firm's assigned price is degenerate.

in which case all price pairs with $p_1 + p_2 = 2p^m(\pi)$ are mutually best responses.¹⁵ Whenever the public randomization device assigns different prices to the firms, therefore, these prices must be equidistant from, and on opposite sides of, the myopic price. As a consequence, $\tilde{p}_1(\pi)$ and $\tilde{p}_2(\pi)$ must be perfectly negatively correlated.

Writing $V(\pi)$ for each firm's price variance, it is then straightforward to calculate

$$\mathbb{E}[R(\pi, \tilde{p}_i(\pi), \tilde{p}_j(\pi))] = R(\pi, \bar{p}_i(\pi), \bar{p}_j(\pi)) - [b + 2c(\pi)]V(\pi) \quad (17)$$

and

$$\mathbb{E}[S(\tilde{p}_i(\pi), \tilde{p}_j(\pi))] = \frac{(c_1 - c_0)^2}{\sigma^2} [\bar{p}_1(\pi) - \bar{p}_2(\pi)]^2 + 4V(\pi) + \xi \quad (18)$$

where $\bar{p}_i(\pi)$ denotes again firm i 's average price. A simple argument establishes that firms set the myopic price on average.¹⁶ Inserting all these results into (1) and arguing exactly as in Section 4, we see that the payoff function u must be a solution to the ODE

$$\pi^2(1 - \pi)^2 \frac{u''(\pi)}{2r} \frac{(c_1 - c_0)^2}{\sigma^2} \xi = \min \left\{ u(\pi) - \lambda(\pi) \frac{u'(\pi)}{r} - m(\pi), \frac{2b + 3c(\pi)}{4} \xi \right\}. \quad (19)$$

The relevant boundary conditions are the same as before.

The following result is the direct analogue of Proposition 4.1 and hence stated without proof.

Proposition 7.1 *The boundary value problem consisting of the ODE (19) and the boundary conditions (11)–(12) has a unique solution u^\dagger . A joint strategy constitutes a payoff-symmetric joint-action Markov perfect equilibrium of the dynamic pricing game if and only if prices are perfectly negatively correlated with mean equal to the myopic price $p^m(\pi)$ and variance equal to*

$$V^\dagger(\pi) = \max \left\{ 0, \frac{u^\dagger(\pi) - \lambda(\pi)(u^\dagger)'(\pi)/r - m(\pi)}{b + c(\pi)} - \frac{2b + 3c(\pi)}{4[b + c(\pi)]} \xi \right\} \quad (20)$$

at almost all beliefs $\pi \in [0, 1]$. The firms' common payoff function in any such equilibrium is u^\dagger , and the set of beliefs where firms create price dispersion is either empty or an open interval strictly smaller than the open unit interval.

In the remainder of this section, we want to compare payoff-symmetric Markov perfect equilibria under joint and independent randomization. We note first that the possibility of coordinating their prices through joint randomization is of no value to firms if

$$u^m(\pi) - \lambda(\pi) \frac{(u^m)'(\pi)}{r} - m(\pi) \leq \frac{2b + 3c(\pi)}{4} \xi \quad \text{for all } \pi, \quad (21)$$

¹⁵Since the right-hand side of (16) is strictly less than $b + c(\pi)$, each firm's maximand is strictly concave, so the first-order condition is sufficient for a maximum.

¹⁶Equation (1), payoff symmetry and the symmetry of the expected signal-to-noise ratio in the two average prices imply that firms must earn identical expected current revenues at any given belief. Since $\bar{p}_1 + \bar{p}_2 = 2p^m$, this in turn can be used to show that $\bar{p}_1 = \bar{p}_2 = p^m$.

because in this case price matching at the myopic price is the unique payoff-symmetric Markov perfect equilibrium under both modes of randomization.¹⁷ The interesting case is therefore when (21) does not hold.¹⁸

Proposition 7.2 *If (21) does not hold, firms are better off under joint randomization ($u^\dagger > u^*$), the range of beliefs where they create price dispersion is larger, and the price variance at each belief in this range is higher.*

This result confirms the intuition with which we started this section, and is illustrated in Figure 9. The curves are the graphs of the functions $[u^*(\pi) - \lambda(\pi)(u^*)'(\pi)/r - m(\pi)]/[b + c(\pi)]$ and $[u^\dagger(\pi) - \lambda(\pi)(u^\dagger)'(\pi)/r - m(\pi)]/[b + c(\pi)]$ and the straight lines are the thresholds ξ and $[2b + 3c(\pi)]\xi / 4[b + c(\pi)]$ for $\xi = 0.1$, $\Lambda = 0.2$, the parameters used in Figure 2. The lower curve is above the higher threshold for a narrow range of beliefs as we have already seen, whereas, for joint randomization, the higher curve is above the lower threshold for a much wider range. Interestingly, the difference in equilibrium payoffs is driven entirely by the dynamics of the problem, that is, by a difference in the speed of learning. In fact, comparing (2) and (17), we find that expected current revenues under joint randomization, $m(\pi) - [b + 2c(\pi)]V^\dagger(\pi)$, are *smaller* than those under independent randomization, $m(\pi) - [b + c(\pi)]V^*(\pi)$, at any belief where $V^\dagger(\pi) > 0$ and hence $V^\dagger(\pi) > V^*(\pi)$. For joint randomization to be more profitable overall, therefore, it must involve a higher speed of learning. A comparison of (3) and (18) supports this conclusion: at each belief where $V^\dagger(\pi) > 0$, the expected signal-to-noise ratio is higher under joint randomization.

Given that firms prefer joint randomization, one would naturally expect consumers to be better off under independent randomization; after all, this is the less “collusive” arrangement. However, this intuition is not necessarily correct. We just showed that firms incur higher short-run revenue losses when they randomize jointly. This suggests that joint randomization may well leave consumers with a higher surplus at least in the short run. To verify this, note that the per-period demand structure in state k is generated by the quasilinear utility function

$$U(q_1, q_2, y) = \frac{a(q_1 + q_2)}{b} - \frac{(q_1 + q_2)^2}{4b} - \frac{(q_1 - q_2)^2}{4[b + 2c_k]} + y$$

where q_i denotes consumption of the good produced by firm i and y stands for money left over. A straightforward calculation reveals that with prices p_1 and p_2 (and sufficiently high income), per-period consumer surplus in state k is

$$a^2/b - a(p_1 + p_2) + \frac{1}{2} [b(p_1^2 + p_2^2) + c_k(p_1 - p_2)^2],$$

a convex function of prices. If firms randomize independently with mean $p^m(\pi)$ and variance $V^*(\pi)$, expected consumer surplus in state k is

$$a^2/b - 2ap^m(\pi) + bp^m(\pi)^2 + (b + c_k)V^*(\pi);$$

¹⁷Condition (21) is the direct analogue of condition (15) from Corollary 5.1. Clearly, (21) implies (15). This reflects the fact that joint randomization is triggered at a lower shadow price of information, $[2b + 3c(\pi)]/4$ rather than $b + c(\pi)$.

¹⁸A sufficient condition for this case can easily be provided along the lines of Proposition 5.2.

if they randomize jointly with the same mean, variance $V^\dagger(\pi)$ and perfectly negative correlation, it is

$$a^2/b - 2ap^m(\pi) + bp^m(\pi)^2 + (b + 2c_k)V^\dagger(\pi).$$

Whenever $V^\dagger(\pi) > 0$, therefore, expected consumer surplus in a given period (and for a given belief held by firms) is higher under joint randomization, independent of the true state k . This has a simple explanation. Consumers are risk-loving with respect to price changes, so at any given time, the mode of randomization that creates more price variance is better for them from a purely static point of view. Moreover, joint randomization leads to higher price differences as firms randomize with perfect negative correlation, and this also increases short-run surplus.

When it comes to total consumer surplus in the infinite-horizon game, however, there is a second effect, reflecting the dynamics of information accrual to the firms. Under joint randomization, firms learn faster, so they can track the true state more closely. As better assessments of the true state translate into a stronger capability to extract surplus from consumers, joint randomization could well mean lower consumer surplus in the longer term.

Which of the two effects dominates obviously depends on a consumer's discount rate. For very impatient consumers, the dynamic effect is negligible – these consumers are clearly better off under joint randomization.

Proposition 7.3 *Suppose that (21) does not hold, so that joint randomization involves price dispersion over some range of beliefs. Whenever firms hold a belief in this range, consumers who discount future utility sufficiently strongly are better off if firms randomize jointly.*

This result follows directly from the above comparison of static surpluses and the fact that as the discount rate tends to infinity, overall surplus converges to the static one.¹⁹

8 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 payoff-symmetric 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 of the signals they observe. In such an equilibrium, the duopoly goes through experimentation cycles

¹⁹Showing that the dynamic effect may dominate for sufficiently patient consumers turns out to be more cumbersome. This is because one has to deal with the joint evolution of *two* state variables: the firms' belief and the true state of demand. We therefore do not pursue this issue any further here.

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. Finally, joint randomization by the firms does not necessarily hurt consumers.

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 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 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' payoff 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. 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 sit-

uation 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. We have therefore presented only 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. This approach is inspired by Harris (1993, Section 10.1), who contrasts it with a state-alternation approach where players use the state itself as a randomizing device. To ease the notational burden in what follows, we write $\tau(\pi) = \pi^2(1 - \pi)^2(c_1 - c_0)^2/\sigma^2$.

By Skorohod's theorem, we can associate with any mixed strategy \tilde{p}_i a Lebesgue×Borel measurable function $P_i : [0, 1] \times [0, 1] \rightarrow \mathbb{R}$ such that for any π , the distribution of $\tilde{p}_i(\pi)$ is the same as that of $P_i(\pi, \tilde{y})$ where \tilde{y} is uniformly distributed on the unit interval. That is, we can associate with \tilde{p}_i 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}_i .

Time alternation can now be introduced as follows. Firms start with a common prior π_0 and initial draws $y_i \in [0, 1]$ ($i = 1, 2$). Firm i employs the pure strategy $P(\cdot, y_i)$ until it is given the opportunity to revise the strategy by drawing a new y_i from the uniform distribution on $[0, 1]$. The firm uses the corresponding pure strategy until the next opportunity to revise it, and so on. Opportunities to revise the strategy arrive according to a Poisson process with intensity κ ; these processes are independent across players.

Given a current belief π and draws y_1 and y_2 , firm i expects flow revenues $R(\pi, P_i(\pi, y_i), P_j(\pi, y_j))$, and a slight extension of the standard result in Liptser and Shirayev (1997, Theorem 9.1) implies that the change in belief $d\pi$ has mean $\lambda(\pi) dt$ and variance $\tau(\pi) [(P_1(\pi, y_1) - P_2(\pi, y_2))^2 + \xi] dt$. For given κ , finally, firm i 's total payoff function (with arguments π, y_i, y_j) is characterized as the solution to a integro-differential equation that reflects not just the infinitesimal changes in firms' beliefs, but also the discrete jumps in the draws y_1 and y_2 .

For $\kappa \rightarrow \infty$ we obtain a limit game with the following properties. If the firms employ the mixed strategies \tilde{p}_1 and \tilde{p}_2 , and their current belief is π , then (i) firm i obtains expected flow revenues $E[R(\pi, \tilde{p}_i(\pi), \tilde{p}_j(\pi))]$ where E denotes the expectation with respect to the distribution of $\tilde{p}_1(\pi)$ and $\tilde{p}_2(\pi)$; (ii) the change in belief $d\pi$ has mean $\lambda(\pi) dt$ and variance $\tau(\pi) E[(\tilde{p}_1(\pi) - \tilde{p}_2(\pi))^2 + \xi] dt$; (iii) firm i 's total payoff function depends on π only, and is characterized as a solution to the ODE (1) subject to the boundary conditions (11)–(12). In the limit, therefore, the firms' flow payoffs and the law of motion of their beliefs are obtained simply by taking expectations over the mixed strategies, and payoff functions solve the standard differential equations associated with these flow payoffs and law of motion.

A mathematically rigorous treatment of this transition to the limit is beyond the scope of the present paper, so we merely indicate the main steps. The results of Harris (1993, Section 10.1) do not cover the case at hand; instead, one can rely on an asymptotic expansion technique going back to Papanicolaou, Stroock and Varadhan (1977), Blankenship and Papanicolaou (1978) or Papanicolaou (1978); a more accessible presentation in an option pricing context can be found in Fouque, Papanicolaou and Sircar

(2000). Let u_i^ϵ denote firm i 's payoff function under time-alternation with intensity $\kappa = 1/\epsilon$. The idea is to expand u_i^ϵ as a sum $u_i^0 + \epsilon u_i^1 + \epsilon^2 u_i^2 + \dots$ with functions $u_i^\ell: [0, 1] \times [0, 1] \times [0, 1] \rightarrow \mathbb{R}$ ($\ell = 0, 1, 2, \dots$) that do not depend on ϵ . A straightforward powers-of- ϵ comparison in the integro-differential equation for u_i^ϵ yields candidate functions u_i^0, u_i^1, u_i^2 ; these functions are characterized as solutions to certain boundary value problems. In particular, u_i^0 solves the ODE (1) subject to the boundary conditions (11)–(12). Examining the boundary value problems for the other functions, one finds that u_i^1 and u_i^2 are bounded, while $u_i^\epsilon - u_i^0 - \epsilon u_i^1 - \epsilon^2 u_i^2$ is $O(\epsilon)$ uniformly in (π, y_i, y_j) . This implies that $u_i^\epsilon - u_i^0$ is also uniformly $O(\epsilon)$. So, we have uniform convergence of the payoff function to the limit u_i^0 as $\epsilon \rightarrow 0$, and this limit has precisely the properties stated in the main text. In a last step, one then shows that the law of motion of beliefs converges in distribution to a diffusion process with instantaneous mean and variance as described under (ii) above.

Note that 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 *finite* arrival rate κ , so results based on the properties of the limiting belief process can easily be interpreted in terms of price processes for large κ .

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

We shall need the following notation:

$$\begin{aligned}\tau(\pi) &= \pi^2(1-\pi)^2 \frac{(c_1 - c_0)^2}{\sigma^2}, \\ \check{m}(\pi) &= m(\pi) + [b + c(\pi)]\xi, \\ g(\pi) &= \frac{\tilde{\pi}(1-\pi)}{\pi} + \frac{(1-\tilde{\pi})\pi}{1-\pi}.\end{aligned}$$

PROOF OF PROPOSITION 4.1: If condition (15) holds, existence is trivial: we can simply take $u^* = u^m$. Suppose therefore that (15) 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 (10), 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

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

where

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

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 (10) satisfies Nagumo's condition²⁰ 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 (10) 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 (10) 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 subsequence which converges uniformly on all compact subintervals of $]0, 1[$ to a solution u^* with the desired properties.

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

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

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

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

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

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

and therefore

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

²⁰See Bernfeld and Lakshmikantham (1974, page 25).

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

Next, define

$$v^*(\pi) = u^*(\pi) - \lambda(\pi) (u^*)'(\pi)/r.$$

Suppose that the set $\{\pi : v^*(\pi) > \check{m}(\pi)\}$ is not empty. A straightforward calculation using (10) shows that on this set,

$$\frac{\tau(\pi)}{2} (v^*)''(\pi) = r [b + c(\pi)] + \Lambda \left\{ 2g(\pi) [b + c(\pi)] + (\pi - \tilde{\pi})(c_1 - c_0) \right\} =: F(\pi). \quad (\text{B.1})$$

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

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

by (B.1). The same inequality holds at π_3 . On the other hand, we have $v^*(\pi_2) \leq \check{m}(\pi_2)$, $(v^*)'(\pi_2) = \check{m}'(\pi_2) = m'(\pi_2) + (c_1 - c_0)\xi$ and $(v^*)''(\pi_2) \geq \check{m}''(\pi_2) = m''(\pi_2)$, hence

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

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 $\{\pi : m''(\pi) \geq 0\}$.

The remaining parts of the proposition follow as indicated in the main text. ■

For the following two proofs, recall from Section 5 that u^m , the firms' payoff function under price matching at the myopic belief, is the unique solution to the ODE (14) subject to the boundary conditions (11)–(12). Define

$$v^m(\pi) = u^m(\pi) - \lambda(\pi) (u^m)'(\pi)/r.$$

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 $v(0) = m(0)$ and $v(1) = m(1)$.

PROOF OF PROPOSITION 5.1: 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 (15).

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

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

Whenever $\sigma^{-2} \leq r/\rho_m + \Lambda/\ell_m$, (B.3) holds, 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 5.2: Define $\overline{m}(\pi) = (1 - \pi)m(0) + \pi m(1)$ on $[0, 1]$, and set $\xi_d = \min\{\xi > 0 : \check{m}(\pi) \geq \overline{m}(\pi) \text{ for all } \pi\}$. Fix $\xi < \xi_d$, 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

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

Whenever $\sigma^{-2} \geq r/\rho_d + \Lambda/\ell_d$, (B.4) holds, hence $v^m \geq \underline{m}$; in particular, $v^m(\pi) > \check{m}(\pi)$ for some π , so (15) does not hold. \blacksquare

PROOF OF PROPOSITION 6.1: Consider v^* as defined in the proof of Proposition 4.1. This is the unique function that has a continuous first derivative and solves

$$\frac{\tau(\pi)}{2} v''(\pi) \xi = \begin{cases} [r + 2\Lambda g(\pi)] [v(\pi) - m(\pi)] + \Lambda (\pi - \tilde{\pi}) [v'(\pi) - m'(\pi)] & \text{if } v(\pi) < \check{m}(\pi) \\ [r + 2\Lambda g(\pi)] [b + c(\pi)] \xi + \Lambda (\pi - \tilde{\pi}) (c_1 - c_0) \xi & \text{if } v(\pi) > \check{m}(\pi) \end{cases} \quad (\text{B.5})$$

subject to the boundary conditions $v(0) = m(0)$ and $v(1) = m(1)$. Given that m is strictly convex by assumption, it is a strict subsolution of (B.5), so we have $v^* > m$ on the open unit interval. Next, a simpler version of the argument given in Keller and Rady (1999, Appendix B) shows that v^* is itself strictly convex.

Given these properties of v^* , straightforward sub- and supersolution arguments show that on the open unit interval, $v^*(\pi)$ is strictly decreasing in r and σ , and in Λ if r is sufficiently small. By equation (13), this translates into the stated comparative statics of price dispersion with respect to these three parameters.

As to the quality of background information, consider $\xi' > \xi$ and $\bar{v}(\pi) = v^*(\pi) + [b + c(\pi)](\xi' - \xi)$. It is easy to see that the function \bar{v} is a strict supersolution to (B.5) with ξ replaced by ξ' if and only if

$$r [b + c(\pi)] + \Lambda \left\{ 2g(\pi) [b + c(\pi)] + (\pi - \tilde{\pi})(c_1 - c_0) \right\} > 0 \quad \text{for } 0 < \pi < 1,$$

which we know to hold from the proof of Proposition 5.1. By equation (13), this implies that price dispersion decreases as the quality of background information increases from ξ to ξ' . \blacksquare

PROOF OF PROPOSITION 7.2: Comparing the ODEs (10) and (19), and arguing along the same lines as in the uniqueness part of the proof of Proposition 4.1, one sees that $u^\dagger - u^*$ cannot assume a negative minimum, which implies the weak inequality $u^\dagger \geq u^*$.

Next, define

$$v^\dagger(\pi) = u^\dagger(\pi) - \lambda(\pi) (u^\dagger)'(\pi)/r$$

and

$$m^\dagger(\pi) = m(\pi) + [2b + 3c(\pi)]\xi/4.$$

Note that v^\dagger is the unique function that has a continuous first derivative and solves

$$\frac{\tau(\pi)}{2} v''(\pi) \xi = \begin{cases} [r + 2\Lambda g(\pi)] [v(\pi) - m(\pi)] + \Lambda (\pi - \tilde{\pi}) [v'(\pi) - m'(\pi)] & \text{if } v(\pi) < m^\dagger(\pi) \\ [r + 2\Lambda g(\pi)] [2b + 3c(\pi)] \xi / 4 + \Lambda (\pi - \tilde{\pi}) 3(c_1 - c_0) \xi / 4 & \text{if } v(\pi) > m^\dagger(\pi) \end{cases} \quad (\text{B.6})$$

subject to the boundary conditions $v(0) = m(0)$ and $v(1) = m(1)$. It is now easy to see that the function $\underline{v}(\pi) = v^*(\pi) - [2b + c(\pi)] \xi / 4$ (with v^* as defined in the proof of Proposition 4.1) is a strict subsolution to (B.6) if and only if

$$r [2b + c(\pi)] + \Lambda \left\{ 2g(\pi) [2b + c(\pi)] + (\pi - \tilde{\pi})(c_1 - c_0) \right\} > 0 \quad \text{for } 0 < \pi < 1.$$

The coefficient of r is clearly positive, and the same is true for the coefficient of Λ since it is larger than its counterpart in (B.3), which we already know to be positive. So \underline{v} is indeed a strict subsolution, hence $v^\dagger > \underline{v}$ on the open unit interval by the usual argument.

Now assume that condition (21) does not hold. Then, $I = \{\pi : v^\dagger(\pi) > m^\dagger(\pi)\}$ is a nonempty open interval, and since $v^\dagger > \underline{v}$, this interval is strictly larger than $\{\pi : v^* > \check{m}(\pi)\}$. Moreover, we have $V^\dagger(\pi) > V^*(\pi)$ at all beliefs in I .

We are now ready to show the strict inequality $u^\dagger > u^*$. Suppose that u^\dagger and u^* assume the same value at some belief π_0 in the interior of the unit interval. Comparing the relevant ODEs again, one sees that π_0 must lie outside of I . Suppose without loss of generality that $\pi_0 \leq \pi_1$, where π_1 is the left boundary of I . Since u^\dagger and u^* have identical values and first derivatives at π_0 and solve the same second-order ODE outside of I , they must coincide on the entire closed interval $[0, \pi_1]$.

Immediately to the right of π_1 , we must have $v^*(\pi) \leq m^\dagger(\pi)$; otherwise, we would have

$$\frac{\tau(\pi)}{2r} (u^*)''(\pi) \xi = v^*(\pi) - m(\pi) > \frac{2b + 3c(\pi)}{4} \xi = \frac{\tau(\pi)}{2r} (u^\dagger)''(\pi) \xi,$$

hence $u^\dagger(\pi) < u^*(\pi)$ immediately to the right of π_1 , which is impossible. As $v^*(\pi_1) = m^\dagger(\pi_1)$, this implies $(v^*)'(\pi_1) \leq (m^\dagger)'(\pi_1)$. Now, a strict inequality is impossible here since we also have $(v^*)'(\pi_1) = (v^\dagger)'(\pi_1)$ and $v^\dagger(\pi) > m^\dagger(\pi)$ immediately to the right of π_1 . So we can conclude that $(v^*)'(\pi_1) = (v^\dagger)'(\pi_1) = (m^\dagger)'(\pi_1)$.

Next, a comparison of the ODEs for v^* and v^\dagger shows that the second derivatives of these functions at π_1 are identical. As $v^*(\pi) \leq m^\dagger(\pi)$ to the right of π_1 , this common second derivative cannot exceed $(m^\dagger)''(\pi_1)$; as $v^\dagger(\pi) > m^\dagger(\pi)$ to the right of π_1 , on the other hand, the common second derivative cannot be strictly lower than that of m^\dagger . So we conclude that $(v^*)''(\pi_1) = (v^\dagger)''(\pi_1) = (m^\dagger)''(\pi_1)$.

Differentiating the ODEs for v^* and v^\dagger with respect to π and using the fact that the values, first derivatives and second derivatives of these functions at π_1 coincide with those of m^\dagger , one now sees that the third derivatives of v^* and v^\dagger are also identical at π_1 . By the same arguments as before, they are equal to the third derivative of m^\dagger at that belief. Repeating this step, one proves by induction that *all* derivatives of v^* and v^\dagger are identical at π_1 , coinciding with the respective derivatives of m^\dagger . As solutions to ODEs where the second derivative is an analytic function of the independent variable, the dependent variable and its first derivative, v^* and v^\dagger are themselves analytic, hence equal to their Taylor series, on some interval $[\pi_1, \pi_1 + \epsilon]$ with $\epsilon > 0$. Therefore, v^* and v^\dagger coincide in a neighborhood of π_1 , which implies in particular that $v^\dagger(\pi) \leq m^\dagger(\pi)$ immediately to the right of π_1 . This is the desired contradiction.

We still have to rule out equality of the payoff functions u^\dagger and u^* at the boundary of the unit interval. Suppose for example that $u^\dagger(0) = u^*(0)$. By the boundary condition (11), this implies $(u^\dagger)'(0) = (u^*)'(0)$. Next, note that v^\dagger and v^* both solve the ODE (B.2) on an open interval J with left boundary 0. We will use this below to show that $(v^\dagger)'(0) = (v^*)'(0) = m'(0)$, which in turn implies $(u^\dagger)''(0) = (u^*)''(0)$ by the definitions of v^\dagger and v^* and the fact that $(u^\dagger)'(0) = (u^*)'(0)$. On J , the relevant ODE for u^* and u^\dagger is (14). Thus, the function $z = u^\dagger - u^*$ solves the ODE

$$\pi^2 (1 - \pi)^2 \frac{z''(\pi)}{2r} \frac{(c_1 - c_0)^2}{\sigma^2} \xi = z(\pi) - \lambda(\pi) \frac{z'(\pi)}{r} \quad (\text{B.7})$$

on J with $z(0) = z'(0) = z''(0) = 0$; moreover, z is strictly positive on the open unit interval. Now, choose ϵ in J so that

$$\left[(1 - \pi)^2 \frac{(c_1 - c_0)^2}{r\sigma^2} \xi - 1 \right] \pi > -2\lambda(\pi)/r$$

for all $\pi < \epsilon$. The strictly convex function $\underline{z}(\pi) = z(\epsilon)\pi^2/\epsilon^2$ coincides with z at the boundaries of the interval $[0, \epsilon]$; in the interior of this interval, it is a strict subsolution of (B.7). By the usual argument, therefore, $z > \underline{z}$ at beliefs strictly between 0 and ϵ . This is a contradiction, however, since $z''(0) = 0 < \underline{z}''(0)$.

The proof is complete once we show that any solution v to (B.2) with $v(0) = m(0)$ also has $v'(0) = m'(0)$. Suppose for instance that $v'(0) > m'(0)$. Using (B.2) and the fact that $g(\pi)[v(\pi) - m(\pi)] \rightarrow \tilde{\pi}[v'(0) - m'(0)]$ as $\pi \rightarrow 0$, we can find $K > 0$ and $\epsilon > 0$ such that $v''(\pi) \geq K\pi^{-2}$ on $[0, \epsilon]$. But then

$$v'(\pi) = v'(\epsilon) - \int_{\pi}^{\epsilon} v''(\eta) d\eta \leq v'(\epsilon) - K \int_{\pi}^{\epsilon} \frac{d\eta}{\eta^2} = v'(\epsilon) - K \left[\frac{1}{\pi} - \frac{1}{\epsilon} \right] \rightarrow -\infty$$

as $\pi \rightarrow 0$, which contradicts the assumption from which we started. The same contradiction arises if we assume $v'(0) < m'(0)$. ■

C Stationary Distribution

In all the MPE of Proposition 4.1, the process of posterior beliefs is a diffusion whose increment $d\pi$ has mean $\lambda(\pi) dt$ and variance $\tau(\pi)(2V^*(\pi) + \xi) dt$ with $\tau(\pi)$ as defined in Appendix B. According to Karlin and Taylor (1981, p. 221), the stationary density of this process, if it exists, has the form

$$\psi(\pi) = \mu(\pi)[C_1 S(\pi) + C_2]$$

where

$$\mu(\pi) = \frac{1}{s(\pi)\tau(\pi)(2V^*(\pi) + \xi)}$$

and

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

with

$$s(\pi) = \exp \left\{ - \int_{\frac{1}{2}}^{\pi} \frac{2\lambda(x)}{\tau(x)(2V^*(x) + \xi)} dx \right\}.$$

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, $V^*(\pi)$ 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 6.2: Given sufficiently small r and ξ , Propositions 5.2 and 6.1 imply the existence of a cutoff level $\bar{\Lambda}$ such that the interval of beliefs where the price variance $V^*(\pi) > 0$ is non-empty if and only if $\Lambda < \bar{\Lambda}$. 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. ■

References

- AGHION, P., ESPINOZA, M.P. and JULLIEN, B. (1993): “Dynamic Duopoly with Learning through Market Experimentation”, *Economic Theory*, **3**, 517–539.
- BAGWELL, K. and STAIGER, R.W. (1997): “Collusion over the Business Cycle”, *RAND Journal of Economics*, **28**, 82–106.
- BERGEMANN, D. and VÄLIMÄKI, J. (1997): “Market Diffusion with Two-Sided Learning”, *RAND Journal of Economics*, **28**, 773–795.
- BERGEMANN, D. and VÄLIMÄKI, J. (2000): “Experimentation in Markets”, *Review of Economic Studies*, **67**, 213–234.
- BERNFELD, S.R. and LAKSHMIKANTHAM, V. (1974): *An Introduction to Nonlinear Boundary Value Problems* (New York and London: Academic Press).
- BLANKENSHIP, G. and PAPANICOLAOU, G.C. (1978): “Stability and Control of Stochastic Systems with Wide-Band Noise Disturbances”, *SIAM Journal of Applied Mathematics*, **34**, 437–476.
- 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 ROB, R. (1998): “Experimentation and Competition”, *Journal of Economic Theory*, **78**, 299–320.
- FOUQUE, J.-P., PAPANICOLAOU, G. and SIRCAR, K.R. (2000): *Derivatives in Financial Markets with Stochastic Volatility* (Cambridge, UK: Cambridge University Press).
- GREEN, E.J. and PORTER, R.H. (1984): “Noncooperative Collusion under Imperfect Price Information”, *Econometrica*, **52**, 87–100.
- 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**, 475–507.
- 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.
- PAPANICOLAOU, G.C. (1978): “Asymptotic Analysis of Stochastic Equations”, in M. Rosenblatt (ed.), *Studies in Probability Theory* (Washington, DC: Mathematical Association of America).
- PAPANICOLAOU, G.C., STROOCK, D. and VARADHAN, S.R.S. (1977): “Martingale Approach to Some Limit Theorems”, in M. Reed (ed.), *Statistical Mechanics, Dynamical Systems and the Duke Turbulence Conference* (Duke University Mathematics Series, vol. 3) (Durham, NC: Duke University Press).
- REINGANUM, J. (1979): “A Simple Model of Equilibrium Price Dispersion”, *Journal of Political Economy*, **87**, 851–858.
- RIORDAN, M.H. (1985): “Imperfect Information and Dynamic Conjectural Variations”, *RAND Journal of Economics*, **16**, 41–50.
- ROB, R. (1991): “Learning and Capacity Expansion under Demand Uncertainty”, *Review of Economic Studies*, **58**, 655–675.
- ROTEMBERG, J.J. and SALONER, G. (1986): “A Supergame-Theoretic Model of Price Wars During Booms”, *American Economic Review*, **76**, 390–407.
- ROTHSCHILD, M. (1974): “A Two-Armed Bandit Theory of Market Pricing”, *Journal of Economic Theory*, **9**, 185–202.
- SALOP, S. and STIGLITZ, J. (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.
- SOBEL, J. (1984): “The Timing of Sales”, *Review of Economic Studies*, **51**, 353–368.




Figure 1: Standard deviation of price for $\alpha = 0.1$, $\Lambda = 0.1$; independent randomization




Figure 2: Standard deviation of price for $\alpha = 0.1$, $\Lambda = 0.2$; independent randomization





Figure 3: Sample paths for $\alpha = 0.1$, $\Lambda = 0.1$; 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 4: Sample paths for $\alpha = 0.1$, $\Lambda = 0.2$; 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 5: Price dispersion for $x = 0.1$, $\Lambda = 0, 0.1, 0.2$ & 0.4 ; independent randomization
 The straight line is the threshold x .
 The curves are the graphs of $[u^*(p) - I(p)(u^*)'(p)/r - m(p)] / [b+c(p)]$ for the different values of Λ .




Figure 6: Price dispersion for $x = 0.1, 0.2$, $\Lambda = 0.1$; independent randomization
 The straight lines are the thresholds, high and low x .
 The curves are the graphs of $[u^*(p) - I(p)(u^*)'(p)/r - m(p)] / [b+c(p)]$ for the two values of x .




Figure 7: Density functions of the invariant distribution of posterior beliefs for $\Lambda = 0.1, 0.2$ & 0.4 ; $\mathbf{x} = 0.1$; independent randomization

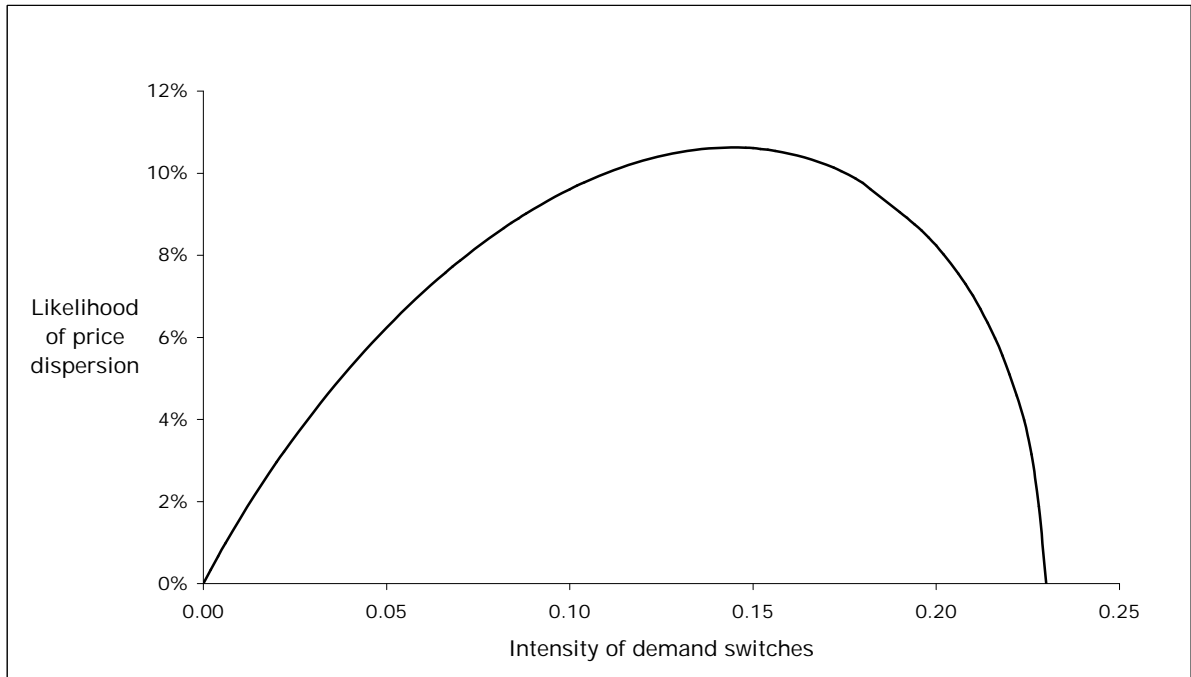


Figure 8: Likelihood of price dispersion (under the invariant distribution of posterior beliefs) as a function of Λ ; $\mathbf{x} = 0.1$; independent randomization




Figure 9: Price dispersion for $x = 0.1$, $\Lambda = 0.2$; independent *v.* joint randomization

The straight lines are the thresholds x and $[2b+3c(\mathbf{p})] x / 4[b+c(\mathbf{p})]$.

The lower curve is the graph of $[u^*(\mathbf{p}) - \mathbf{I}(\mathbf{p})(u^*)'(\mathbf{p})/r - m(\mathbf{p})] / [b+c(\mathbf{p})]$ (independent randomization); the upper curve is that of $[u^+(\mathbf{p}) - \mathbf{I}(\mathbf{p})(u^+) '(\mathbf{p})/r - m(\mathbf{p})] / [b+c(\mathbf{p})]$ (joint randomization).