

COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
AT YALE UNIVERSITY

Box 2125, Yale Station
New Haven, Connecticut 06520

COWLES FOUNDATION DISCUSSION PAPER NO. 1138

Note: Cowles Foundation Discussion Papers are preliminary materials circulated to stimulate discussion and critical comment. Requests for single copies of a Paper will be filled by the Cowles Foundation within the limits of the supply. References in publications to Discussion Papers (other than mere acknowledgment by a writer that he has access to such unpublished material) should be cleared with the author to protect the tentative character of these papers.

MARKET DIFFUSION WITH TWO-SIDED LEARNING

Dirk Bergemann and Jusso Välimäki

November 1996

Market Diffusion with Two-Sided Learning*

Dirk Bergemann[†] Juuso Välimäki[‡]

First Version: May 1996

This Version: September 1996

Abstract

The diffusion of a new product of uncertain value is analyzed in a duopolistic market in continuous time. The two sides of the market, buyers and sellers, learn the true value of the new product over time as a result of experimentation. Buyers have heterogeneous preferences over the products and sellers compete in prices.

The pricing policies and market shares of the sellers in the unique Markov perfect equilibrium are obtained explicitly. The dynamics of the equilibrium market shares display excessive sales of the new product relative to the social optimum in early stages and too low sales later on. The dynamic resolution of uncertainty implies ex post differentiation and hence *both* sellers value information positively. As information is generated only by experiments with the new product, this relaxes the price competition in the dynamic setting. Finally, the diffusion path of a successful product is shown to be *S*-shaped.

KEYWORDS: Dynamic Duopoly, Learning, Market Shares, Diffusion.

JEL CLASSIFICATION: C73, D43, D83.

* The authors would like to thank Colin Campbell, Ron Goettler, Phillip Leslie, Ariel Pakes, Jim Peck, Mike Riordan and seminar participants at the Institut d'Anàlisi Econòmica, Barcelona, and Ohio State University for many helpful comments.

[†] Department of Economics, Yale University, New Haven, CT 06520-8268, U.S.A., dirk.bergemann@yale.edu; and Institut d'Anàlisi Econòmica, CSIC, Barcelona, Spain.

[‡] Department of Economics, Northwestern University, Evanston, IL 60208-2009, U.S.A., valimaki@nwu.edu.

1. Introduction

The design of a new product or the improvement of an existing product is only the necessary first step in launching a product in the market. Commercial success depends critically on the speed and cost at which buyers learn the relevant characteristics of the product. An eventually successful product may go through a phase of sluggish sales in the beginning of its life cycle simply because buyers are not aware of its true quality. It may then be in the firm's interest to engage in strategies that sacrifice current revenue to generate more information about the product for example through aggressive penetration pricing to capture a larger clientele.

In this paper, we model dynamic competition in a duopolistic market for experience goods. An established firm and a firm with a new product compete in prices in an infinite horizon, continuous time model. Buyers have to try the new product in order to learn how well it suits their needs. We assume that the product incorporates both a common and a private value component to the buyers. To keep the model analytically tractable, we assume that the private value component of every buyer is common knowledge and may reflect idiosyncratic taste, location or the like. In contrast, the common component is learned gradually over time as more experience is accumulated. The information obtained in any single trial with the new product is a noisy signal of product quality or more generally the utility it provides to the consumer. The total available information on the common value component then depends on the cumulative sales. We assume that a statistic on the aggregate performance of the new product (obtained either via consumer reports or an unmodeled process of word of mouth communication) is

available to all parties in the model. This introduces an informational externality into the model. Each individual purchase has effects on all buyers in this model through changes in the market statistic.

The simplest economic example for our story is a linear city with, for example, two fish markets located at the ends of the city. Consumers are uniformly distributed along the segment between the fish markets and incur a transportation cost linear in distance to the stores. One of the stores has been in operation for a long time while the ownership of the other has been changed recently. Buyers choose the store at the beginning of each period depending on their beliefs on the quality of the fish in the new store, their transportation cost and current prices. The quality of the store keeping has to be learned through experience while the idiosyncratic transportation cost to every buyer is common knowledge at the outset. Every purchase in the new store yields an imperfect signal on the quality (taste, freshness, etc.) of the fish in the new market. Based on the buyers' experiences, the two firms and the consumers update their belief on the value the new fish market provides to the buyers in the city.

We characterize the diffusion path of the new product (or the time path of the clientele of the new seller in the fish market example), and derive the equilibrium price path in Markov Perfect Equilibrium. Our continuous time specification allows us to derive analytical solutions to both the price paths of the firms and the associated path of market share evolution. The pricing policies display an interesting asymmetry. Both firms prefer ex post differentiation, as is typical in models of price competition, and hence the value of information is positive for both firms. Only sales of the new product, however, produce more information. Hence, if the firms want to speed up the information transmission, the new firm must make relatively large sales in early periods. Since both firms benefit from

sales by the new firm, competitive pressures in the stage game are reduced and uncertain (and vertical) product differentiation relaxes competition in a similar sense as deterministic product differentiation in Shaked and Sutton (1982). In equilibrium, this results in the new firm collecting higher revenues due to increased market share early on. As buyers and sellers become more convinced about the true quality of the new product, the market shares converge to those of the full information game. The value of accumulating more information has thus two components for the established seller: Due to product differentiation, he expects to get higher profits in expected terms and as the information becomes more accurate, the need to distort sales in favor of the new firm diminishes as well.

This asymmetry is also apparent in the intertemporal behavior of the market share of the new firm. In the myopic case, i.e. the case where the stage game is played repeatedly as a one shot game and all players ignore the informational effects, the *expected* market share of the new firm is constant. In the Markov Perfect Equilibrium, however, the expected market share of the new firm is decreasing (its market share is a supermartingale). High initial sales reflect the value of information to both sellers. Furthermore, the expected revenues of both sellers are increasing over time. Thus firms are sacrificing current profits early on to enhance the accumulation of information.

In equilibrium, the pricing strategies induce sales paths that differ from the socially optimal path. As long as the beliefs about the type of the new product are sufficiently pessimistic, the level of experimentation and hence purchases of the new product exceed socially optimal levels. The intuition behind this result is straightforward: At pessimistic beliefs, the market share of the new firm is small and the cost of attracting new customers, which is the loss of profit on inframarginal buyers from lowering the prices, is small. The established firm has

a large market share and concedes some market share to protect inframarginal profits. On the other hand, if beliefs about the new product are optimistic, then the old firm prices more aggressively and experimentation falls below the social optimum.

The time path of product adoption, or the diffusion curve of a successful new product displays the S-shape, documented in empirical work as in Mansfield (1968) and Gort and Klepper (1982). At relatively pessimistic beliefs, increases in the number of adopting consumers increase the inflow of new information. As beliefs about the product quality become sufficiently optimistic, the growth in the market share of the new firm eventually slows down.

Our model of dynamic competition involves simultaneous determination of two value functions representing the sellers' intertemporal problems in an infinite horizon model. The dynamic programming equations under positive discounting, which describe the equilibrium of the model, are, unfortunately, nonlinear differential equations and can only be solved numerically. To avoid these complications, we use a technique recently employed by Bolton and Harris (1994) by considering the limiting model as discounting in the model becomes small. The limiting model preserves all the desirable features of the original discounted dynamic program, in particular, the optimal policies in the limiting model are the unique limits of optimal policies with discounting.

There are a number of related papers in which issues of strategic pricing are analyzed in a learning environment. Aghion, Espinosa and Jullien (1993) and Harrington (1995) consider product-differentiated duopolies in which firms learn about the substitutability of their products. In their models learning is one-sided. More precisely, the firms try to learn the degree of substitutability between their products. The goods are not experience goods as consumers have perfect infor-

mation about the products at the outset and it is observed demand conditional on the price differential between the products, rather than cumulative sales of a new product that generate the learning. Caminal and Vives (1996) analyze a two-period duopoly model where a sequence of short-lived buyers are uncertain about the quality differential. As consumers are restricted to observe only quantities of past sales, the long-lived firms attempt to signal-jam the learning process of the buyers.

Our techniques are similar to those used by Bolton and Harris (1994) who were the first to consider a continuous time model of strategic experimentation. Bergemann and Välimäki (1996) analyze strategic experimentation in a duopoly with a continuum of identical consumers. The homogeneity excludes market-sharing and the analysis there concentrates on the impact of the informational externalities on the efficiency of the market. Issues of informational externalities in a market for new products were already present in the competitive entry and capacity expansion model of Rob (1991) and its extension to a two-sided learning model by Vettas (1993). Our focus on strategic interaction between the firms results in very different price dynamics and welfare conclusions.

Judd and Riordan (1994) analyze a two-sided learning model in a monopoly with two periods. Their information structure differs from ours as the consumption of the experience good yields a private signal to the individual buyer and the aggregate signal of sales is a private signal to the monopolist. Finally, Schlee (1996) has analyzed the incentives of a monopolist to engage in introductory pricing of a new experience good in a two-period model where the signals are publicly observed.

The model and the Bayesian learning process are introduced in Section 2. The socially efficient allocation policy is described in Section 3. The definition of the

Markov perfect equilibrium is given in Section 4, where we completely describe the pricing policies and market shares in the unique equilibrium. The diffusion path of a successful product is presented in some detail in Section 5. We conclude in Section 6.

2. The model

In a dynamic duopoly, firms with differentiated products compete in prices in an infinite horizon, continuous time setting. The first firm is well established in the market and its product characteristics are common knowledge at the beginning of the game. The second firm has a new product whose value has to be learned over time.

The preferences of the buyers are described by a Hotelling location model. The buyers are uniformly distributed on the interval $[0, 1]$ and they have unit demand at each instant of time. The value of the certain product for individual n is given by s_n with

$$s_n = s + nd, \quad n \in [0, 1]. \quad (2.1)$$

The parameter $d > 0$ represents the substitutability between the products. Symmetrically, the value of the uncertain product for individual n is given by μ_n with

$$\mu_n = \mu + (1 - n)d, \quad n \in [0, 1]. \quad (2.2)$$

The value, μ , of the new product is initially unknown to *all* parties. It can be either low or high:

$$\mu \in \{\mu_L, \mu_H\}, \quad (2.3)$$

with

$$0 < s - d < \mu_L < s < \mu_H < s + d. \quad (2.4)$$

The inner inequalities in (2.4) imply that the new product can be of either lower or higher value than the established one. The outer inequalities assert that in either case, the efficient allocation would assign a positive measure of buyers to both products.¹ The marginal cost of production for both products are normalized to zero.

The size of d determines how much the value of the product to the buyer and ultimately the choice behavior of the buyer is influenced by her location. As such d is a measure of the heterogeneity among the buyers. The model in (2.1)-(2.3) is one of horizontal and vertical differentiation, where the horizontal differentiation is common knowledge at the outset, but the extent of vertical differentiation is uncertain.

The uncertainty about the value of the second product can be resolved only by experimentation, i.e. through purchases of the new product. The performance of the new product is, however, subject to random disturbances and any single experiment with the new product provides only a noisy signal about the true underlying value. The information conveyed by an experiment depends on the size of the experiment. As each buyer is of measure zero, the size of her purchase is negligible and hence the information generated by an individual experiment is also negligible. In consequence all relevant information is contained in the *aggregate outcome*. The aggregate or *market outcome* is the performance of the product over all buyers, which is assumed to be publicly observable.

The market outcome is defined by the following stochastic process in continuous time:

$$dX(n(t)) = n(t) \mu dt + \sigma \sqrt{n(t)} dB(t), \quad t \in [0, \infty), \quad (2.5)$$

where $n(t)$ is the fraction of buyers who use the *new* product at time t and

¹In other words, the innovation is not drastic.

$dB(t)$ is the increment of the one-dimensional standard Brownian motion. The interpretation of (2.5) is most intuitive when μ is regarded as the average success rate of the new product. In this case, $dX(n(t))$ simply counts the number of successes which depends on the number $n(t)$ of experiments and on the random disturbance, $\sigma\sqrt{n(t)}dB(t)$. The market outcome, $dX(n(t))$, thus provides a noisy signal of the true value of the uncertain alternative. The informativeness of the signal is increasing in the sales of the new seller.²

It is immediately verified that the instantaneous mean and the variance of the market outcome are linear in the market share of the new seller. As the value of μ can only be μ_L or μ_H , posterior beliefs about the quality are completely characterized by $\alpha(t)$, with

$$\alpha(t) = \Pr(\mu = \mu_H | \mathcal{F}(t)), \quad (2.6)$$

where $\mathcal{F}(t)$ is the history generated by $X(n(t))$. The conditional expected quality $\mu(\alpha(t))$ of the uncertain product is

$$\mu(\alpha(t)) = (1 - \alpha(t))\mu_L + \alpha(t)\mu_H. \quad (2.7)$$

The market players extract the information provided by the noisy market outcome (2.5) to improve their common prior beliefs $\alpha(0) = \alpha_0$ over time. The game is

²It can be shown that the stochastic model represented in (2.5) is the limit to the following finite model: In an economy with N buyers, each individual experiment X_i^N is normally distributed with unknown mean $\mu_N = \mu/N$ and known variance $\sigma_N^2 = \sigma^2/N$ for fixed μ and σ^2 . As the number of buyers and therefore experiments increases, it is easily verified that the aggregate mean and the aggregate variance of the *market experimentation*:

$$\sum_{i=1}^N X_i^N$$

remain constant at (μ, σ^2) . Hence, the information contained in the market outcome is constant in the number of buyers. The corresponding limit as $N \rightarrow \infty$ in a continuous-time formulation is (2.5).

thus one of incomplete but symmetric information, and no issues of asymmetric information arise.³

The learning process of the market represents a signal extraction problem: Given the information generated by $X(n(t))$, what is the posterior belief about the value of the second alternative. As the beliefs are completely characterized by $\alpha(t)$, the signal extraction problem reduces to the description of the law of motion of the posterior belief $\alpha(t)$.

Proposition 1 (Posterior Belief).

The process $\alpha(t)$ is a Brownian motion with zero drift and variance $n(t)\Sigma^2(\alpha(t))$:

$$n(t)\Sigma^2(\alpha(t)) = n(t) \left[\frac{\alpha(t)(1-\alpha(t))(\mu_H - \mu_L)}{\sigma} \right]^2. \quad (2.8)$$

Proof. See Liptser and Shirayayev (1977), Theorem 9.1. ■

The posterior belief $\alpha(t)$ is a martingale as the posterior belief incorporates all predictable information. The variance of $\alpha(t)$ can be interpreted as the amount of information generated by the market outcome as it indicates how rapidly $\alpha(t)$ can change. The variance increases linearly in the market share of the new firm and depends on the signal to noise ratio $(\mu_H - \mu_L)/\sigma$ and the diffuseness of the prior information $\alpha(t)(1 - \alpha(t))$.

The only payoff relevant source of uncertainty in this model is the value of the new product. All available information in period t is incorporated in the common posterior belief $\alpha(t)$ and the dynamics in the model are driven by the process of belief change. The focus in subsequent sections is the analysis of Markovian policies where the posterior belief α is a natural state variable. We start by

³Because of this simplifying assumption, we do not have to address issues of prices signaling the qualities as in Judd and Riordan (1996). We shall come back to this issue in the conclusion.

solving for the socially efficient experimentation policy to get a benchmark for the equilibrium model.

3. Efficient Experimentation

This section has two objectives. First, we want to introduce the basic technique of dynamic programming with zero discounting to describe intertemporal policies in our model. Second, we apply the technique to characterize the socially efficient experimentation policy, which maximizes the social surplus. The technique is most clearly illustrated here as the efficient allocation is the solution to a single optimization problem without any strategic interaction.

An experimentation policy prescribes for every posterior belief α the shares of buyers allocated to the sellers. Denote by $n(\alpha)$ the market share of the new product and by $1 - n(\alpha)$ the share of the established product. The average flow value from the established product when the $1 - n(\alpha)$ buyers receive the product is with preferences as in (2.1) and (2.2):

$$s(n(\alpha)) \triangleq s + \frac{(1 + n(\alpha))d}{2}, \quad (3.1)$$

and for the new product it is similarly

$$\mu(n(\alpha)) \triangleq \mu + \frac{(2 - n(\alpha))d}{2}. \quad (3.2)$$

The expected flow value of consumption from the two alternatives under posterior belief α and allocation policy $n(\alpha)$ is then given by

$$n(\alpha)\mu(n(\alpha)) + (1 - n(\alpha))s(n(\alpha)).$$

The optimal allocation problem *with discounting* can be written as a dynamic programming problem in continuous time:

solving for the socially efficient experimentation policy to get a benchmark for the equilibrium model.

3. Efficient Experimentation

This section has two objectives. First, we want to introduce the basic technique of dynamic programming with zero discounting to describe intertemporal policies in our model. Second, we apply the technique to characterize the socially efficient experimentation policy, which maximizes the social surplus. The technique is most clearly illustrated here as the efficient allocation is the solution to a single optimization problem without any strategic interaction.

An experimentation policy prescribes for every posterior belief α the shares of buyers allocated to the sellers. Denote by $n(\alpha)$ the market share of the new product and by $1 - n(\alpha)$ the share of the established product. The average flow value from the established product when the $1 - n(\alpha)$ buyers receive the product is with preferences as in (2.1) and (2.2):

$$s(n(\alpha)) \triangleq s + \frac{(1 + n(\alpha))d}{2}, \quad (3.1)$$

and for the new product it is similarly

$$\mu(n(\alpha)) \triangleq \mu + \frac{(2 - n(\alpha))d}{2}. \quad (3.2)$$

The expected flow value of consumption from the two alternatives under posterior belief α and allocation policy $n(\alpha)$ is then given by

$$n(\alpha)\mu(n(\alpha)) + (1 - n(\alpha))s(n(\alpha)).$$

The optimal allocation problem *with discounting* can be written as a dynamic programming problem in continuous time:

$$rV(\alpha) = \max_{n(\alpha)} \left\{ n(\alpha) \mu(n(\alpha)) + (1 - n(\alpha)) s(n(\alpha)) + \frac{1}{2} n(\alpha) \Sigma^2(\alpha) V''(\alpha) \right\}, \quad (3.3)$$

where $V(\alpha)$ is the value function of the allocation problem and $r > 0$ is the discount rate.⁴

The expression (3.3) has a very simple interpretation. The flow benefit $rV(\alpha)$ from an experimentation policy $n(\alpha)$ consists of two parts. The first part is the aggregated flow payoff from the two alternatives and the second part is the flow value of experimentation $\frac{1}{2} n(\alpha) \Sigma^2(\alpha) V''(\alpha)$. The latter term captures the informational gains from the experimentation policy as it improves the quality of future choices. The instantaneous variance $\Sigma^2(\alpha)$ indicates the quantity of information released through a unit of experimentation and $n(\alpha)$ is the current seize of the experiment. The curvature $V''(\alpha)$ of the value function is the shadow price of information. The flow value of the optimal experimentation policy then maximizes the sum of the flow payoff and the flow value of information.

The optimal allocation $n(\alpha)$ can be obtained in principle from the first order conditions of the right hand side term of (3.3). But the differential equation which results when $n(\alpha)$ is replaced by its solution is quadratic in the second order term $V''(\alpha)$. As a consequence, solutions to the differential equation implied by (3.3) can be obtained only through numerical methods and no analytical solutions are available. Note that the nonlinear differential equation would arise with any model of heterogeneous consumers and as such is not an artefact of the preference specification here.

We sidestep these problems by analyzing the optimal allocation problem un-

⁴See Dixit and Pindyck (1994) or Harrison (1985) for a detailed derivation of the dynamic programming equation in continuous time when uncertainty is represented by a Brownian motion.

der zero discounting. With the appropriate optimality criterion, called the *strong long-run average criterion* in Dutta (1991), we preserve the recursive representation of the dynamic programming equation (3.3) as $r \rightarrow 0$. Most importantly for our purposes, the optimal policies under this criterion are the unique limits to the associated policies under *discounting* and as such maintain the intertemporal aspects of the experimentation policies. In other words, all the qualitative properties of the equilibrium we derive in the following will hold also for small discount rates $r > 0$.

The strong long-run average criterion refines the *long-run average criterion* which discriminates insufficiently between alternative intertemporal policies.⁵ The long-run average under the initial belief α_0 is given by

$$v(\alpha_0) = \sup_{n(\alpha)} \lim_{T \rightarrow \infty} \frac{1}{T} E \left[\int_0^T (n(\alpha) \mu(n(\alpha)) + (1 - n(\alpha)) s(n(\alpha))) dt \mid \alpha_0 \right] \quad (3.4)$$

where we suppress dependence on t . The long-run average $v(\alpha)$ is equal to the *expected full-information payoff*

$$v(\alpha) = \alpha v(1) + (1 - \alpha) v(0), \quad (3.5)$$

if $\alpha(t)$ converges almost surely to 0 or 1, as it will be the case here. The full information payoffs $v(0)$ and $v(1)$ are the solutions to the static allocation problems, when μ is known to be either μ_L or μ_H . These values can be computed immediately and so can hence *all* long-run average values $v(\alpha)$.⁶ The strong

⁵See Dutta (1991) for a very careful and detailed analysis on the connection of optimality criteria under discounting and under no discounting. Briefly, the strong-long run average criterion discriminates between policies that differ for a finite time interval, as does of course the usual criterion under discounting, whereas the long run average criterion does not discriminate.

⁶The long-run average values for the efficient and the equilibrium program are recorded in Lemma 1 in the Appendix.

long-run average is defined by the following optimization problem:

$$V(\alpha_0) = \sup_{n(\alpha)} \lim_{T \rightarrow \infty} E \left[\int_0^T (n(\alpha) \mu(n(\alpha)) + (1 - n(\alpha)) s(n(\alpha)) - v(\alpha)) dt \mid \alpha_0 \right], \quad (3.6)$$

where $v(\alpha)$ is as defined in (3.4). The strong long-run average maximizes the expected returns net the long-run average.⁷ As $T \rightarrow \infty$, the limit is well defined and finite. The infinite horizon problem (3.6) can be presented via the dynamic programming equation as

$$\max_{n(\alpha)} \left\{ n(\alpha) \mu(n(\alpha)) + (1 - n(\alpha)) s(n(\alpha)) - v(\alpha) + \frac{1}{2} n(\alpha) \Sigma^2(\alpha) V''(\alpha) \right\} = 0, \quad (3.7)$$

under the condition that $n(\alpha)$ is bounded away from zero for all α . This condition means that even for low values of α , it is optimal to allocate some buyers to the new product. The condition is satisfied with the inequalities (2.4) introduced in the previous section.

The simplicity of the Bellman equation (3.7) in the case of no discounting is transparent as it contains only the second-order term but no lower order terms of the value function. The flow value $rV(\alpha)$ which appeared in the discounted case is replaced by the long-run average $v(\alpha)$ in the undiscounted case.

We proceed to solve for the optimal experimentation policy with the assistance of the Bellman equation (3.7). As the optimal allocation $N(\alpha)$ achieves a maximum value of 0 for the left-hand term of (3.7), and as $N(\alpha)$ is bounded away from zero, we may divide the term inside the maximum operator through

⁷As the flow value $n(\alpha) \mu(n(\alpha)) + (1 - n(\alpha)) s(n(\alpha))$ of the optimal allocation policy under imperfect information is necessarily less than the expected full information payoff $v(\alpha)$, a different interpretation of the strong-long run average criterion is that it minimizes the losses due to imperfect information compared to the maximum achievable under full information.

$n(\alpha)$ and claim that the maximizer $\hat{N}(\alpha)$ of the modification:

$$\max_{n(\alpha)} \left\{ \frac{s + \frac{d}{2} - v(\alpha)}{n(\alpha)} - dn(\alpha) \right\} + \mu(\alpha) - s + d + \frac{1}{2} \Sigma^2(\alpha) V''(\alpha) = 0, \quad (3.8)$$

is identical to the maximizer of the original expression (3.7). The optimal allocation is found by the first order condition of (3.8) which is independent of the second order term of the value function.

Proposition 2 (Efficient Experimentation Policy).

The efficient allocation $N(\alpha)$ is given by:

$$N(\alpha) = \sqrt{\frac{v(\alpha) - s - \frac{d}{2}}{d}}, \quad (3.9)$$

with $N(\alpha) > 0$ for all $\alpha \in [0, 1]$.

Proof. See Appendix. ■

The optimal allocation $N(\alpha)$ is naturally an increasing function of α . The comparative statics are as expected and $N(\alpha)$ is increasing in μ_H, μ_L , and decreasing in s and d . As we compare $N(\alpha)$ with the myopically optimal allocation $M(\alpha)$ which solves the static problem

$$M(\alpha) \in \arg \max_{n(\alpha)} n(\alpha) \mu(n(\alpha)) + (1 - n(\alpha)) s(n(\alpha)),$$

it is verified after some computations that $N(\alpha) > M(\alpha)$ for all $\alpha \in (0, 1)$. This provides additional evidence that the optimal policy under the strong long-run average preserves the intertemporal aspect of the experimentation policy as the optimal seize of the experiment $N(\alpha)$ is larger than the myopically efficient allocation. The difference is of course attributable to the additional incentive to use the new product to generate information for future decisions. This intertemporal benefit is absent in the static decision.

4. Dynamic Equilibrium

We define our solution concept, the Markov Perfect Equilibrium, in Section 4.1. Next, we characterize the unique equilibrium and the associated equilibrium policies. Section 4.2 considers the efficiency of the equilibrium and examines the dynamic properties of the equilibrium.

4.1. Equilibrium and Policies

In this model of dynamic competition, buyers and sellers learn over time more about the true value of the new product. We focus on Markovian strategies with the posterior belief $\alpha(t)$ as the state variable to emphasize the interaction between the changing stage game, pricing and market acquisition.

The sellers choose at each instant of time their prices non-cooperatively. We consider only pricing policies, $p_i(\alpha)$, of the sellers that are measurable with respect to the state variable α . The strategic considerations of the sellers involve both current revenues and the influence that experimentation has on future revenues. Again, we are interested in the limiting case of no discounting and the value functions of the sellers can be constructed in the same way as the value function of the efficient program in (3.7). The dynamic programming equation for the established seller is:

$$0 = \max_{p_1(\alpha)} \left\{ (1 - n(\alpha)) p_1(\alpha) - v_1(\alpha) + \frac{1}{2} n(\alpha) \Sigma^2(\alpha) V_1''(\alpha) \right\}, \quad (4.1)$$

and for the new seller it is:

$$0 = \max_{p_2(\alpha)} \left\{ n(\alpha) p_2(\alpha) - v_2(\alpha) + \frac{1}{2} n(\alpha) \Sigma^2(\alpha) V_2''(\alpha) \right\}, \quad (4.2)$$

where $v_1(\alpha)$ and $v_2(\alpha)$ are the long-run average payoffs of the sellers.⁸ The

⁸The long-run average payoffs are again the expected full information payoffs, which are the

equilibrium market share

$$n(\alpha) \equiv n(\alpha, p_1(\alpha), p_2(\alpha)),$$

naturally depends on the prices offered by the sellers.

The buyers, in turn, make their purchase decisions as a function of the current estimate $\mu(\alpha)$ of the new product, the current prices and their individual preferences. As a single buyer is negligible in the market, the informativeness of the market outcome is independent of any individual purchase. In consequence the buyer's decision is based exclusively on the current (expected) values and prices, as she cannot influence the informativeness of the market outcome. The equilibrium market shares are then determined by the critical consumer $n(\alpha)$ who is indifferent between the two alternatives:

$$s + n(\alpha)d - p_1(\alpha) = \mu(\alpha) + (1 - n(\alpha))d - p_2(\alpha).$$

After rearranging, the equilibrium share of the new seller $n(\alpha)$ emerges as a function of the current values and prices and the degree of horizontal differentiation:

$$n(\alpha) = \frac{(\mu(\alpha) - p_2(\alpha)) - (s - p_1(\alpha)) + d}{2d}. \quad (4.3)$$

As the buyers decide as they would under myopia, the equilibrium share condition (given the prices) is as in the static pricing game. We now formally define the Markov Perfect Equilibrium (Maskin and Tirole (1995)).

static profits when the value of the new product is known. The values are recorded in Lemma 1 in the Appendix.

Definition 1. (Markov Perfect Equilibrium)

A Markov Perfect Equilibrium is a triple $\{p_1(\alpha), p_2(\alpha), n(\alpha)\}$ such that equations (4.1) - (4.3) are satisfied for all $\alpha \in [0, 1]$.

The process of experimentation gradually establishes whether the new product is of low or high value. Experimentation then leads ultimately to an increase in the (vertical) differentiation between the products. The *expected joint profits* for the sellers are higher after differentiation as the superior seller will be able to extract a larger part of the social surplus from the buyers via higher prices. *Ex-ante*, both sellers are equally likely to gain the competitive edge - recall that α is a martingale - and thus sellers attach positive value to additional information. Consequently, the shadow price of information, represented by $V_i''(\alpha)$ is positive to both of them.⁹ This has important, but different implications for the strategies of the sellers.

By inserting the equilibrium share condition (4.3) in the optimization problems of the competing sellers, the strategic considerations of the sellers become more transparent:

$$0 = \max_{p_1(\alpha)} \left\{ \left(1 - \frac{\mu(\alpha) - p_2(\alpha) - s + p_1(\alpha) + d}{2d} \right) p_1(\alpha) - v_1(\alpha) + \frac{\mu(\alpha) - p_2(\alpha) - s + p_1(\alpha) + d}{4d} \Sigma^2(\alpha) V_1''(\alpha) \right\}, \quad (4.4)$$

and

$$0 = \max_{p_2(\alpha)} \left\{ \frac{\mu(\alpha) - p_2(\alpha) - s + p_1(\alpha) + d}{2d} p_2(\alpha) - v_2(\alpha) + \frac{\mu(\alpha) - p_2(\alpha) - s + p_1(\alpha) + d}{4d} \Sigma^2(\alpha) V_2''(\alpha) \right\}. \quad (4.5)$$

The equilibrium conditions (4.4) and (4.5) illustrate that for both firms an increase in their respective market shares requires a decrease in the price at which

⁹The formal result will be stated in Corollary 2.

their product is offered. But the equilibrium conditions also show that the incentives to acquire a larger market share are quite different for the two sellers. The value of information, which is generated by sales of the new product only, operates like an additional revenue source. And while it is associated with sales of the new product, new information actually *benefits both firms*. The additional value from sales will prompt the new seller to price more aggressively and seek a larger market share than he would in a static world. In contrast, the incentives for the established seller to increase his market share are relatively weaker because a larger market share would imply less experimentation and hence reduce the informational gains. This will lead the established firm to adopt a less aggressive pricing policy in equilibrium.

By the same argument we presented for the efficient program we can divide (4.4) and (4.5) by $n(\alpha)$, or its equivalent in (4.3). The resulting equilibrium conditions for the optimal pricing strategies do not involve the second derivatives of the value functions:

$$0 = \frac{1}{2} \Sigma^2(\alpha) V_1''(\alpha) + \max_{p_1(\alpha)} \left\{ \frac{2d(p_1(\alpha) - v_1(\alpha))}{p_1(\alpha) - p_2(\alpha) + \mu(\alpha) - s + d} - p_1(\alpha) \right\}, \quad (4.6)$$

and similarly for the second seller

$$0 = \frac{1}{2} \Sigma^2(\alpha) V_2''(\alpha) + \max_{p_2(\alpha)} \left\{ \frac{-2dv_2(\alpha)}{p_1(\alpha) - p_2(\alpha) + \mu(\alpha) - s + d} + p_2(\alpha) \right\}. \quad (4.7)$$

The advantage of this approach is rather obvious. The optimal pricing policies of the sellers can now be obtained without explicit reference to the second derivatives of the value functions. The optimal pricing policies are the solutions to the first-order conditions of (4.6) and (4.7).

Proposition 3 (Equilibrium Experimentation).

There is a unique Markov Perfect Equilibrium. The equilibrium prices are

$$p_1(\alpha) = \frac{2}{3}(s - \mu(\alpha)) + \sqrt{2dv_2(\alpha)}, \quad (4.8)$$

and

$$p_2(\alpha) = \frac{1}{3}(\mu(\alpha) - s) + d. \quad (4.9)$$

The market share of the new seller is given by

$$n(\alpha) = \sqrt{\frac{v_2(\alpha)}{2d}}. \quad (4.10)$$

Proof. See Appendix. ■

The uniqueness property of the equilibrium is naturally due to the Markovian restriction of the equilibrium strategies, and the set of non-Markovian (perfect) equilibria is obviously much larger.

As the market is always completely shared by the sellers, the market share of the established firm is $1 - n(\alpha)$. The properties of the dynamic pricing equilibrium are most easily illustrated by contrasting the dynamic with the static policies. In the corresponding static equilibrium, buyers and sellers take the current posterior belief as given and make decisions as if they were in a one-shot game without any intertemporal considerations. We denote the static (or myopic) unique equilibrium prices by $p_1^m(\alpha)$ and $p_2^m(\alpha)$, and the market share of the new firm as $n^m(\alpha)$. The following corollary is immediate.

Corollary 1 (Static and Dynamic Prices).

- (i) *The level of prices is $p_1(\alpha) > p_1^m(\alpha)$ and $p_2(\alpha) = p_2^m(\alpha)$.*
- (ii) *The level of sales is $n(\alpha) > n^m(\alpha)$.*

The static prices and market shares are recorded in Lemma 1 in the Appendix. As the established firm adopts a less competitive pricing policy, its price is higher and its market share is smaller in the dynamic equilibrium than in the static one. The new firm is willing to price more aggressively due to the informational gains which come with a larger market share. But the extent to which the new seller lowers his price is limited by the associated losses on the current revenue from the inframarginal buyers. These conflicting incentives for the new seller are resolved at intertemporal prices that are exactly identical to the corresponding static prices. In balance, the new firm then receives a higher market share, but can charge the same price as in the static equilibrium. The established firm gains in terms of current prices, but loses in terms of market shares relative to the static case.

The influence of the information flow on the pricing strategies has a surprising consequence for the optimal choice of innovations. Suppose for the moment that the new seller was currently offering a product of low quality μ_L . Suppose further that he can choose between two forms of innovation: either he gets a product with value $\mu(\alpha)$ with certainty, or he gets an innovation which is of superior quality, μ_H , with probability α , and with probability $1 - \alpha$ it presents no improvement beyond μ_L . Corollary 1 then implies that the second seller will strictly prefer the uncertain innovation. Moreover, the stochastic innovation increases the prices the consumers have to pay on average and thus relaxes the price competition relative to the certain innovation.

The experience with the new product generates information which is valuable for both firms. The equilibrium value of information is given by the dynamic programming equations (4.1) and (4.2) as:

$$\frac{1}{2}n(\alpha)\Sigma^2(\alpha)V_1''(\alpha) = v_1(\alpha) - (1 - n(\alpha))p_1(\alpha),$$

and

$$\frac{1}{2}n(\alpha)\Sigma^2(\alpha)V_2''(\alpha) = v_2(\alpha) - n(\alpha)p_2(\alpha),$$

respectively. For each firm, the equilibrium value of information $\frac{1}{2}n(\alpha)\Sigma^2(\alpha)V_i''(\alpha)$ is thus precisely the difference between its expected full information revenue and its current revenue. As the equilibrium prices and allocations are established in Proposition 3, it can be verified that the value of information is *twice as large for* the established firm than for the new firm. Thus, the established firm values information about the true quality of the new product even higher than the new seller himself. This initially puzzling result can be traced back to the impact the information flow has on the players strategies. To see this better, consider for the moment the value of information if the sellers would follow myopic policies. In this case the value of information for the two sellers is in fact equal:

$$v_1(\alpha) - (1 - n^m(\alpha))p_1^m(\alpha) = v_2(\alpha) - n^m(\alpha)p_2^m(\alpha),$$

as the symmetry of the model and the updating process would suggest. In the dynamic equilibrium, the less aggressive position of the established firm lowers its own revenue flow and raises the revenue flow for the new firm. Thus the current revenue shortfall from the expected full information revenues increases for the established seller relative to the new seller. In consequence, the resolution of uncertainty is valued more highly by the established firm. Next we analyze the efficiency and dynamic features of the equilibrium.

4.2. Inefficiency of the Equilibrium

The pricing policies of the sellers change systematically in the size of their market shares. As the value functions of the sellers indicate, the marginal benefit from persuading an additional buyer *to experiment* at any instant of time are constant at $\frac{1}{2}\Sigma^2(\alpha)V_i''(\alpha)$ for seller i and independent of the market share. But as the marginal buyer has always the lowest valuation for the product among the current buyers, to convince her, each seller would have to decrease the price on all inframarginal buyers. The static revenue loss associated with the acquisition of an additional buyer is therefore increasing in the market share, since the price decrease is granted to a larger measure of inframarginal buyers. For a seller with a small market share, the revenue loss is small and his price policy is mostly determined by intertemporal considerations. Conversely, a seller with a large market share will not be very responsive to a lower price from the competitor as he attempts to maintain the current price level even at the expense of losing some market share.

This suggests that for low values of α , the new firm acquires market share aggressively to generate information and that the established firm will respond only slowly with lower prices. In conjunction, these two strategies will generate excessive experimentation compared to the socially optimal level. As α increases, the market share of the new firm increases and this weakens its aggressive stance. Simultaneously the established firm becomes more responsive in its pricing policy as its market share shrinks. For high values of α , the new seller captures the market to a large extent and his incentives to support additional experimentation become very weak.

Proposition 4 (Inefficiency).

Equilibrium experimentation is excessive for low values of α and insufficient for high values of α . The difference

$$N(\alpha) - n(\alpha)$$

is increasing in α and crosses zero once.

Proof. See Appendix. ■

The dynamic response of the sellers to changes in the market condition as represented in the posterior belief α has several implications for the expected changes in the market shares and revenues of the sellers over time. The following properties of the equilibrium prices and market shares are established by Proposition 3 and the fact that the posterior belief $\alpha(t)$ is a martingale (see Proposition 1).

Corollary 2 (Martingale and Convexity).

- (i) *The price $p_1(\alpha)$ is a supermartingale, the price $p_2(\alpha)$ is a martingale.*
- (ii) *The market share $n(\alpha)$ of the new seller is a supermartingale.*
- (iii) *The revenues $(1 - n(\alpha))p_1(\alpha)$ and $n(\alpha)p_2(\alpha)$ are submartingales.*
- (iv) *The value functions $V_i(\alpha)$ of the sellers are convex.*

For the new seller, the incentives to acquire a larger market share weaken with the size of his current market share. In consequence, marginal changes in the market share of the new seller in response to changes in the posterior belief α are more accentuated for low market shares $n(\alpha)$. The concavity of $n(\alpha)$ together with the martingale property of the posterior belief $\alpha(t)$ turns $n(\alpha)$ into

a strict supermartingale. A symmetric argument leads to a strict submartingale property for the market share of the established seller and thus to the resilience of the established seller, whose market shares are expected to increase over time.

The response of the established seller to a decline in his market share mirrors the behavior of the new firm. His pricing behavior becomes more aggressive and more responsive to changes in the market conditions as his market share decreases. The price $p_1(\alpha)$ is a decreasing and concave function of α and thus the prices of the established seller are expected to decrease over time. This leads to the conclusion that the market share of the established seller is expected to increase, while its price is expected to decrease over time. This highlights the differential in the equilibrium response rates of market shares and prices. The market share of the established firm responds initially, i.e. for low values of α , very fast to changes in the market condition, whereas his price is initially very slow in response to changes in the posterior belief $\alpha(t)$.

We argued earlier that the *joint profits* of the sellers are increasing in the extent of (vertical) differentiation. As experimentation induces the process of differentiation, the revenues of the sellers are expected to increase over time. In probabilistic terms, the revenues of both sellers are strict submartingales with respect to the posterior belief $\alpha(t)$. Finally, the positive value of the experiments for the firms, as it generates ex post differentiation is documented by the convexity of the value functions.

5. Market Diffusion

So far we have described the equilibrium policies as functions of the state variable $\alpha(t)$. But as the changes in the posterior belief $\alpha(t)$ are endogenously determined by the equilibrium policies, one would like to know more about the evolution of

the posterior belief $\alpha(t)$ and the associated policies over *real* time t . We therefore take the analysis one final step further and examine how market shares and prices typically develop over time for a successful or unsuccessful new product. This will give us a more detailed description of the diffusion paths of new products in real time rather than in the state variable $\alpha(t)$ which, after all, only represents the information available at time t .

In the following, we focus on the case when the true value of the product is high, or $\mu = \mu_H$, as only successful products will ultimately capture the market. While the sample path of the market shares and prices is of course stochastic, we start by analyzing the expected changes in the posterior belief $\alpha(t)$ under the equilibrium policies.

The differential equation which describes the market outcome for the high value product is given by

$$dX(n(\alpha(t))) = n(\alpha(t))\mu_H dt + \sqrt{n(\alpha(t))}\sigma dB(t), \quad (5.1)$$

and is similar to the outcome process as defined earlier in (2.5). Here $n(\alpha(t))$ is the equilibrium market share when the market participants hold the belief $\alpha(t)$. The differential equation (5.1) combines two different elements: the true data generating process and the equilibrium policies of the imperfectly informed traders, which were obtained in the previous section. More precisely, in the case of $dX(n(\alpha(t)))$ the signal is generated by the true value μ_H , but as the market participants possess only imperfect information in the form of $\alpha(t)$, the size $n(\alpha(t))$ of the experiment is determined by the equilibrium policies based on $\alpha(t)$. The evolution of the posterior belief $\alpha(t)$ conditional on $\mu = \mu_H$ is given by

$$d\alpha(t) = (\mu_H - \mu_L)\alpha(t)(1 - \alpha(t)) [dX(n(\alpha(t))) - n(\alpha(t))\mu(\alpha(t))dt], \quad (5.2)$$

as an application of the general filtering equation, see again Liptser and Shiriyayev (1977). The change in the posterior belief $\alpha(t)$ is then determined by the difference inside the square bracket, which is simply the difference between the true signal and the expected signal, where the expectation is based on the current imperfect information of the market traders. By inserting the differential equation (5.1) into (5.2) we obtain the evolution of the posterior belief when the process is generated by the true value μ_H :

$$d\alpha(t) = (\mu_H - \mu_L)^2 \alpha(t) (1 - \alpha(t))^2 n(\alpha(t)) dt + (\mu_H - \mu_L) \alpha(t) (1 - \alpha(t)) \sqrt{n(\alpha(t))} dB(t). \quad (5.3)$$

The difference between the evolution of the *conditional posterior belief* $\alpha(t)$ here and the posterior belief in Proposition 1 is that the former receives a continuous push upwards from the conditioning on μ_H . On average, the signal, which is the outcome realization exceeds the expectations of the market, which still puts some positive probability, namely $1 - \alpha(t)$ on the event that the product is of low quality, while in fact the product is of high quality. Hence, $\alpha(t)$ is not a martingale anymore, but a submartingale. The drift of the process $\alpha(t)$ is then given by

$$E[d\alpha(t)] = n(\alpha(t)) (\mu_H - \mu_L)^2 \alpha(t) (1 - \alpha(t))^2 dt. \quad (5.4)$$

The equation (5.4) presents the drift rate of the conditional posterior belief as a function of time. Our interest rests with the description of the posterior belief when its evolution is governed by the expected or mean changes. In other words, we analyze the behavior of the market when the evolution of $\alpha(t)$ is determined only by the drift of the process (5.3). We denote the new and deterministic process by $\hat{\alpha}(t)$, and the evolution of the mean posterior $\hat{\alpha}(t)$ is given by:

$$d\hat{\alpha}(t) = (\mu_H - \mu_L)^2 \hat{\alpha}(t) (1 - \hat{\alpha}(t))^2 n(\hat{\alpha}(t)) dt. \quad (5.5)$$

As the differential equation (5.5) is deterministic, $\hat{\alpha}(t)$ is a function of the initial condition $\hat{\alpha}_0 = \alpha_0$ and time t only.¹⁰ Next we analyze the behavior of the market shares and the prices under the deterministic process $\hat{\alpha}(t)$ rather than under the stochastic process $\alpha(t)$. The market share $n(\hat{\alpha}(t))$ is a composite function of the state variable $\hat{\alpha}(t)$ and time t . The differential equation which describes the changes in the market share of the new seller $dn(\hat{\alpha}(t))$ is given by

$$dn(\hat{\alpha}(t)) = n'(\hat{\alpha}(t)) n(\hat{\alpha}(t)) (\mu_H - \mu_L)^2 \hat{\alpha}(t) (1 - \hat{\alpha}(t))^2 dt. \quad (5.6)$$

As $n(\cdot)$ is given in Proposition 3, we can infer immediately from (5.6) that the market share of the new firm increases over time. The market share $n(\hat{\alpha}(t))$, as $\hat{\alpha}(t)$ before, is a deterministic process, and as such is a function of the initial condition $\hat{\alpha}_0 = \alpha_0$ and time t only. For transparency we relabel

$$\hat{n}(t) \equiv n(\hat{\alpha}(t))$$

for a given $\hat{\alpha}_0 = \alpha_0$. The question we would like to address then is the speed by which the new firm gains market shares. In particular we would like to know whether the acquisition of the new buyers is performed more aggressively over time or whether the acquisition speed slows down with the passage of time.

Proposition 5 (S-shaped diffusion path).

- (i) *The mean posterior belief $\hat{\alpha}(t)$ is increasing over time, first at an increasing, then at a decreasing rate.*
- (ii) *The mean market share $\hat{n}(t)$ is increasing over time, first at an increasing, then at a decreasing rate.*

¹⁰An alternative approach would be to derive the distribution of the state variable $\alpha(t)$ over time by the Kolmogorov forward equation. The nonlinearities present in this model don't allow us to solve the partial differential equation explicitly, but numerical simulations indicate the same qualitative as the one we derive in the following.

Proof. See Appendix. ■

The expected market share of the new product, conditional on it being of high value is therefore an *S*-shaped function, increasing, and initially convex and then concave. A symmetric result holds for the expected market share conditional on the product being of low value, in which case the market share of the new firm is decreasing, initially concave, and then convex.

The evolution of the market shares over time is a *composition* of the behavior of the equilibrium market share $n(\hat{\alpha}(t))$ as a function of $\hat{\alpha}$ and the evolution of the posterior belief $\hat{\alpha}(t)$. The drift in $\hat{\alpha}(t)$ is increasing when starting from low values of $\hat{\alpha}(t)$. This in turn accelerates the growth in market shares and increasingly increases the drift of $n(\hat{\alpha}(t))$ as a function of t . But as $\hat{\alpha}$ approaches one, the learning process slows down and the drift in $\hat{\alpha}(t)$, while remaining positive approaches zero. In conjunction with the concavity of $n(\cdot)$, this decreases the drift $dn(\hat{\alpha}(t))$ and hence the speed of market acquisition by the new seller.

The evolution of the mean market share $\hat{n}(t)$ and an actual sample path of the stochastic share process $n(\alpha(t))$ are displayed in Fig.1. The evolution of the market shares is accompanied by a similar, but somewhat shifted movement in the prices. The differences in the drift rates of market share and price processes reflect again the strategic trade-offs for the sellers. We continue the discussion here for the case when the true value of the new product is high and state the results first. Again, symmetric results hold for the case when the true value of the new product is low. By analogy we denote $\hat{p}_i(t) \equiv p_i(\hat{\alpha}(t))$.

Proposition 6 (Prices over time).

- (i) The price process $\hat{p}_1(t)$ is strictly decreasing in t , first at a decreasing, then at an increasing rate.
- (ii) The price process $\hat{p}_2(t)$ is strictly increasing in t , first at an increasing, then at a decreasing rate.
- (iii) The concavity of $\hat{p}_1(t)$ prevails strictly longer than the convexity of $\hat{p}_2(t)$.
- (iv) The convexity of $\hat{p}_2(t)$ prevails strictly longer than the convexity of $\hat{n}(t)$.

Proof. See Appendix. ■

The movements of the price processes can again be decomposed along the state and time space dimension, and together reveal a rather rich dynamic process. The changes in the prices $\hat{p}_i(t)$ as described in (i) and (ii) follow a pattern parallel to the market shares of the respective sellers. But if we compare the price processes across the sellers and similarly if we compare the price and share processes for a given seller, then the different rates of adjustment elucidate the strategic incentives of the sellers across time.

The drift in the price process $\hat{p}_1(\cdot)$ of the established seller continues to *decrease longer* than the drift in the price $\hat{p}_2(\cdot)$ of the new seller continues to *increase*. In other words, as $\hat{\alpha}(t)$ increases the downward movement in the price $\hat{p}_1(\cdot)$ accelerates even when the upward movement of $\hat{p}_2(\cdot)$ decelerates. This result translates the increasing responsiveness of the established seller to the success of his competitor, documented in Proposition 3 into the time profile of the pricing policy of the established seller.

The dynamic movement of the prices are displayed in Fig.2 in the behavior of the mean prices as well as for an actual sample path of the prices. The underlying sample path of $\alpha(t)$ is the same as the one in Fig.1.

Finally, the trade-off for the new seller between higher market share and higher price is changing over time. As $\alpha(t)$ increases over time, the expected gains in market share $n(\cdot)$ start to decrease earlier than the gains in the prices $p_2(\cdot)$. Hence for a larger market share the new seller is using his competitive edge to increase prices rather than to increase market shares. Again, we display the dynamics for the mean processes as well as for an actual sample path in Fig.3.

6. Conclusion

In this paper we analyzed the dynamic aspects of pricing policies and market shares following the introduction of a new product with uncertain value. The strategic aspect of the learning process lead to an increase in prices in the dynamic competition relative to the static competition with given beliefs. The (uncertain) value of the new product was not a choice variable of its seller and, in particular, it was assumed to be constant over time. While this restriction should clearly be removed in future research on dynamic innovation, the comparison between entry with a certain or uncertain value is nonetheless suggestive. The premium to uncertain innovation for the new firm as well as the positive externality the innovation has on the competing seller, implies that stochastic innovation reduces competition and increases profits in oligopolistic markets. This suggests excessive innovations in differentiated markets as the search for innovations is the *dynamic equivalent* to the static vertical differentiation in Shaked and Sutton (1982).

We recall in this context that the *established* seller valued information about the true quality of the new product higher than the *new* seller and this in spite of the purely strategic role of information. If the established firm could use the information beyond its strategic aspect, say for product improvements of its own, then the value of information would a fortiori even be higher, and may relax

competition even further.

A final remark concerns the information structure in our model. Buyers and sellers make their choices over time under imperfect but symmetric information. As the new product is an experience good, it may be natural to ask what would happen if the individual experiments would provide a private signal and market shares were the only information aggregation device. While sufficient care is required in specifying the signal structure with a continuum of agents, the surprising answer is that a model with private signals would have the same structure and dynamics as the current model. As a first approximation we observe that the private signals create heterogeneity among the buyers similar to the heterogeneity due to the horizontal differentiation. As the strategic weight of each buyer is small relative to the market, the bias in their purchase decision due to the effort to manipulate the sellers beliefs is small, and zero with a continuum of buyers. The sales in the subsequent period then reveal to the sellers the average experience of the buyers, and enable the sellers to update their beliefs about the true quality of the new product.¹¹ Thus many of the basic properties of the model, in particular the structure of the inefficiency and the qualitative behavior of the diffusion rates would be present in the private signal model. The basic difference is the endogenous change in the heterogeneity over time. The influence of a single private signal on the posterior decreases over time, since the posteriors become more concentrated around the true value of the alternative, and so the heterogeneity among buyers decreases over time as well.

¹¹Judd and Riordan (1994) essentially analyze this private information model with a monopolist under the additional complication that the aggregate sales are private information to the monopolist.

7. Appendix

All proofs to the propositions in the text are collected in this appendix.

Lemma 1 (Long-run averages). *The long-run average in the efficient program is given by*

$$v(\alpha) = \frac{(s + \mu(\alpha) + \frac{3}{2}d)}{2} + (1 - \alpha) \frac{(\mu_L - s)^2}{4d} + \alpha \frac{(\mu_H - s)^2}{4d}. \quad (7.1)$$

The long-run averages of the two sellers in the strategic program are given by

$$v_1(\alpha) = (1 - \alpha) \frac{(\frac{1}{3}(s - \mu_L) + d)^2}{2d} + \alpha \frac{(\frac{1}{3}(s - \mu_H) + d)^2}{2d}, \quad (7.2)$$

and

$$v_2(\alpha) = (1 - \alpha) \frac{(\frac{1}{3}(\mu_L - s) + d)^2}{2d} + \alpha \frac{(\frac{1}{3}(\mu_H - s) + d)^2}{2d}. \quad (7.3)$$

Proof. The long-run average values $v(\alpha)$ and $v_i(\alpha)$ are equal to the expected full-information payoffs:

$$v_i(\alpha) = (1 - \alpha)v_i(0) + \alpha v_i(1),$$

as $n(\alpha)$ and $N(\alpha)$ are bounded away from zero under the condition (2.4). But $v_i(0)$ and $v_i(1)$ are the values to the static allocation and equilibrium problems when the value of the new product is known to be either μ_L or μ_H . The composite values are then computed immediately. In particular the equilibrium prices and allocations are

$$p_1 = \frac{1}{3}(s - \mu_i) + d \quad \text{and} \quad p_2 = \frac{1}{3}(\mu_i - s) + d,$$

and the market share of the new firm is

$$n = \frac{\frac{1}{3}(\mu_i - s) + d}{2d}. \blacksquare$$

Proof of Proposition 2: The first order condition for the efficient program:

$$\max_{n(\alpha)} \frac{s + \frac{d}{2} - v(\alpha)}{n(\alpha)} - dn(\alpha),$$

is given by

$$dn(\alpha)^2 + s + \frac{1}{2}d - v(\alpha) = 0,$$

and the positive root of the quadratic form is the solution with

$$N(\alpha) = \sqrt{\frac{v(\alpha) - s - \frac{d}{2}}{d}}. \blacksquare$$

Proof of Proposition 3: The first-order conditions for the sellers are given by:

$$\frac{2d}{p_1(\alpha) - p_2(\alpha) + \mu(\alpha) - s + d} - \frac{2d(p_1(\alpha) - v_1(\alpha))}{(p_1(\alpha) - p_2(\alpha) + \mu(\alpha) - s + d)^2} - 1 = 0, \quad (7.4)$$

and

$$-\frac{2dv_2(\alpha)}{(p_1(\alpha) - p_2(\alpha) + \mu(\alpha) - s + d)^2} + 1 = 0. \quad (7.5)$$

Rearranging the first-order condition (7.5),

$$v_2(\alpha) 2d = (p_1(\alpha) - p_2(\alpha) + \mu(\alpha) - s + d)^2,$$

and by the equilibrium share condition (4.3) we have

$$n(\alpha) = \sqrt{\frac{v_2(\alpha)}{2d}}. \quad (7.6)$$

By rewriting (7.4) and using again the equilibrium share condition (4.3) we obtain

$$2dn(\alpha) + v_1(\alpha) - p_1(\alpha) = 2dn^2(\alpha),$$

which after using (7.6) gives us:

$$p_1(\alpha) = \sqrt{2dv_2(\alpha)} + v_1(\alpha) - v_2(\alpha). \quad (7.7)$$

Finally, the equilibrium share condition (4.3) gives us an explicit expression for the price of the new seller by using (7.6) and (7.7):

$$p_2(\alpha) = \mu(\alpha) - s + v_1(\alpha) - v_2(\alpha) + d. \quad (7.8)$$

As $v_1(\alpha) - v_2(\alpha) = \frac{2}{3}(s - \mu(\alpha))$ by Lemma 1, we obtain the unique equilibrium prices as in (4.8) and (4.9). The martingale and supermartingale characterizations follow directly from the linearity of $v_i(\alpha)$ in α . ■

Proof of Proposition 4: We first prove the strict monotonicity and then the single crossing property. By Proposition 2 and 3 we have

$$N(\alpha) - n(\alpha) = \sqrt{\frac{v(\alpha) - s - \frac{d}{2}}{d}} - \sqrt{\frac{v_2(\alpha)}{2d}}. \quad (7.9)$$

A sufficient condition for $N(\alpha) - n(\alpha)$ to be strictly increasing is that (i) $N(0) < n(0)$, and (ii) $2v'(\alpha) > v_2'(\alpha)$. By Lemma 1, (ii) is equivalent to

$$\mu_H - \mu_L + \frac{(\mu_H - s)^2}{2d} - \frac{(\mu_L - s)^2}{2d} > \frac{(\frac{1}{3}(\mu_H - s) + d)^2}{2d} - \frac{(\frac{1}{3}(\mu_L - s) + d)^2}{2d},$$

which in turn is equivalent to

$$\frac{2}{9}(\mu_H - \mu_L) \frac{2\mu_L + 3d - 4s + 2\mu_H}{d} > 0. \quad (7.10)$$

Since $\mu_H - \mu_L$ is strictly positive,

$$(3d + 2(\mu_H + \mu_L) - 4s) > 0, \quad (7.11)$$

the inequality (7.10) is implied by condition (2.4). The single crossing property is satisfied after verifying that

$$2v(0) - 2s - d < v_2(0), \quad (7.12)$$

and

$$2v(1) - 2s - d > v_2(1), \quad (7.13)$$

which can be computed directly from the long-run averages given in Lemma 1. ■

Proof of Proposition 5: (i) The evolution of the mean posterior $\hat{\alpha}(t)$ as a function of time is given by

$$d\hat{\alpha}(t) = (\mu_H - \mu_L)^2 \hat{\alpha}(t) (1 - \hat{\alpha}(t))^2 n(\hat{\alpha}(t)) dt.$$

By Proposition 3 we have

$$n(\hat{\alpha}(t)) = \sqrt{\frac{v_2(\hat{\alpha}(t))}{2d}}.$$

By Lemma 1, $v_2(\hat{\alpha}(t))$ is linear in $\hat{\alpha}(t)$ and we can express $n(\hat{\alpha}(t))$ for the moment simply as the root of a linear function

$$n(\hat{\alpha}(t)) = \sqrt{b_1 + b_2 \hat{\alpha}(t)}, \quad (7.14)$$

with $b_1, b_2 > 0$ as positive constants. The second time derivative is then given by

$$\frac{d^2 \hat{\alpha}(t)}{dt^2} = (\mu_H - \mu_L)^2 \left((1 - \hat{\alpha}(t))^2 \sqrt{b_1 + b_2 \hat{\alpha}(t)} - 2\hat{\alpha}(t) (1 - \hat{\alpha}(t)) \sqrt{b_1 + b_2 \hat{\alpha}(t)} + \frac{1}{2} \hat{\alpha}(t) \frac{(1 - \hat{\alpha}(t))^2}{\sqrt{b_1 + b_2 \hat{\alpha}(t)}} b_2 \right).$$

It can then be verified that there exists $\bar{\alpha} \in (\frac{1}{3}, \frac{2}{3})$ such that

$$\frac{d^2 \hat{\alpha}(t)}{dt^2} > 0 \Leftrightarrow \hat{\alpha}(t) > \bar{\alpha} \quad \text{and} \quad \frac{d^2 \hat{\alpha}(t)}{dt^2} < 0 \Leftrightarrow \hat{\alpha}(t) < \bar{\alpha}.$$

(ii) The evolution of the mean market share of the new seller as a function of time is given by

$$dn(\hat{\alpha}(t)) = n'(\hat{\alpha}(t)) n(\hat{\alpha}(t)) (\mu_H - \mu_L)^2 \hat{\alpha}(t) (1 - \hat{\alpha}(t))^2 dt. \quad (7.15)$$

The evolution of the market share as a function of time is given in (7.15) and the time derivative of $dn(\hat{\alpha}(t))/dt$ presents the curvature of the market share of the new firm as a function of time. We may replace $n(\hat{\alpha}(t))$ again by (7.14) to obtain:

$$\frac{dn(\hat{\alpha}(t))}{dt} = \frac{1}{2} b_2 (\mu_H - \mu_L)^2 \hat{\alpha}(t) (1 - \hat{\alpha}(t))^2.$$

The second time derivative is given by

$$\frac{d^2n(\hat{\alpha}(t))}{dt^2} = \frac{1}{2} b_2 (\mu_H - \mu_L)^2 \left((1 - \hat{\alpha}(t))^2 - 2\hat{\alpha}(t)(1 - \hat{\alpha}(t)) \right) \frac{d\hat{\alpha}(t)}{dt}.$$

As we have $\frac{d\hat{\alpha}(t)}{dt} > 0$ for all $\hat{\alpha}(t) \in (0, 1)$ the sign of the curvature is determined by the first term from which infer directly that

$$\frac{d^2n(\hat{\alpha}(t))}{dt^2} > 0 \Leftrightarrow \hat{\alpha}(t) \in \left(0, \frac{1}{3} \right). \blacksquare$$

Proof of Proposition 6: The proofs of the four claims rely on the same technique as the ones in Proposition 5 and are therefore omitted. ■

BIBLIOGRAPHY

- AGHION, P., M.P. ESPINOSA, AND B. JULLIEN (1993): "Dynamic Duopoly with Learning through Market Experimentation." *Economic Theory* **3**, 517-539.
- BERGEMANN, D. AND J. VALIMAKI (1996): "Market Experimentation and Pricing." Cowles Foundation Discussion Paper, Yale University.
- BOLTON, P. AND C. HARRIS (1993): "Strategic Experimentation." STICERD TE 93/261, London School of Economics.
- CAMINAL, R. AND X. VIVES (1996): "Why Market Shares Matter: An Information-Based Theory." *Rand Journal of Economics*, **27**, 221-239.
- DIXIT, A.K. AND R.S. PINDYCK (1994): *Investment under Uncertainty*. Princeton: Princeton University Press.
- DUTTA, P. (1991): "What Do Discounted Optima Converge To?" *Journal of Economic Theory* **55**, 64-94.
- GORT, M. AND S. KLEPPER. (1982): "Time Paths in the Diffusion of Product Innovation." *Economic Journal* **17**, 630-653.
- HARRINGTON, J.E. (1995): "Experimentation and Learning in a Differentiated-Products Duopoly." *Journal of Economic Theory* **66**, 275-288.
- HARRISON, J.M. (1985): *Brownian Motion and Stochastic Flow Systems*. New York: John Wiley & Sons.
- KELLER, G. AND S. RADY (1996): "Monopoly Experimentation." mimeo, LSE.
- JUDD, K. AND M. RIORDAN (1994): "Price and Quality in a New Product Monopoly." *Review of Economic Studies* **61**, 773-789.
- LIPTSER, R.S. AND A.N. SHIRYAYEV (1977): *Statistics of Random Processes I*. New York: Springer Verlag.

- MANSFIELD, E. (1968): *Industrial Research and Technological Innovations*. New York: W.W. Norton.
- MASKIN, E. AND J. TIROLE (1995): "Markov Perfect Equilibrium." mimeo, Harvard University.
- ROB, R. (1991): "Learning and Capacity Expansion under Demand Uncertainty." *Review of Economic Studies* 58, 655-675.
- SCHLEE, E. (1996): "Learning about Product Quality and Introductory Pricing." mimeo, Arizona State University.
- SHAKED, A. AND J. SUTTON (1982): "Relaxing Price Competition through Product Differentiation." *Review of Economic Studies* 49, 3-13.
- VETTAS, N. (1993): "Demand and Supply in New Markets: Diffusion with Bilateral Learning." mimeo, Duke University.




Fig. 1: The market share of the new firm

$$(\alpha_0 = 0.1, \mu_L = 2, \mu_H = 6, s = 4, d = 1)$$




Fig. 2: The evolution of the prices

$(\alpha_0 = 0.1, \mu_L = 2, \mu_H = 6, s = 4, d = 1)$




Fig. 3: The evolution of price and market share

$$(\alpha_0 = 0.1, \mu_L = 2, \mu_H = 6, s = 4, d = 1)$$