Review of Economic Studies (2000) 67, 213–234

Experimentation in Markets

DIRK BERGEMANN
Yale University

and

JUUSO VÄLIMÄKI
University of Southampton

First version received April 1996; final version accepted March 1999 (Eds.)

We present a model of entry and exit with Bayesian learning and price competition. A new product of initially unknown quality is introduced in the market, and purchases of the product yield information on its true quality. We assume that the performance of the new product is publicly observable. As agents learn from the experiments of others, informational externalities arise.

We determine the Markov Perfect Equilibrium prices and allocations. In a single market, the combination of the informational externalities among the buyers and the strategic pricing by the sellers results in excessive experimentation. If the new product is launched in many distinct markets, the path of sales converges to the efficient path in the limit as the number of markets grows.

1. INTRODUCTION

In multi-agent learning situations, informational externalities may reduce the number of experiments undertaken below the socially efficient level. As buyers choose among new experience goods or firms decide whether to adopt a new technology, the availability of information from others' decisions gives rise to a free rider problem. Rather than perform a costly experiment herself, a buyer may opt to wait and see how the market evaluates the new product. In this paper, we develop a simple market model of experimentation and analyse the informational effects in a model with buyers and sellers. In contrast to the one-sided experimentation problems, we find that equilibrium experimentation often exceeds the Pareto optimal level in two-sided models.

The first model we consider is a dynamic duopoly in continuous time with homogeneous buyers. Two firms with differentiated products engage in price competition over time. One product has a known value while the other is new and its true value is initially uncertain to all the parties in the model. The value of the product is determined by the quality of the match between consumer preferences and product characteristics. Additional information is acquired only through repeat purchases. In each period, buyers observe a noisy signal of the true value of the product. We assume that all signals are publicly observable and that all uncertainty is about a common value component. With this assumption, all buyers and sellers condition their behaviour on the same information and we can abstract from individual differences in past observations. This is justified in economic situations where public data on other buyers' choices is available and informative. For example, when choosing among providers of communication or transportation services, such as mail delivery firms or airlines, it is reasonable to rely on consumer reports and other published data in addition to one's private experience. We also show that our model has an alternative interpretation as one where a flow of buyers make a once and for all purchase of a durable good. Under that scenario, signals from earlier purchases are the only source of information for a given buyer. In both interpretations of the model, the public observability of signals gives rise to an informational externality in the market.

We determine the paths of sales and prices in Markov Perfect Equilibrium and compare them to the Pareto optimal paths.[1] In a model without sellers (or equivalently with fixed prices), informational externalities among the buyers result in too little investment in information acquisition as in Bolton and Harris (1999). When we allow the firms to set their prices optimally, a new source of potential inefficiency is discovered: experimentation determines the future competitive positions of the two firms. The new firm extracts benefits from successful experiments through higher future prices while buyers bear the risk of unsuccessful experiments. When buyers become more pessimistic about the new firm's product, the established firm is able to charge higher prices on its product in equilibrium. The new seller has to compensate each buyer for the costs resulting from higher future prices charged by the established seller. The individual buyer fails, however, to take into account the effect of her own purchases on all other buyers. This makes sales relatively inexpensive for the new firm and the equilibrium displays excessive experimentation. A well known result on optimal learning states that the *ex post* efficient alternative is not always selected in the long run along the *ex ante* efficient path (*e.g.* Rothschild (1974)). Our results therefore imply that introducing price competition on the supply side of the market may achieve *ex post* optimality with a larger probability.

The second model analyses price competition in multiple markets. Each market has a separate established producer, but the new product is a competitor in all markets. One may think of each single market as a different geographic location or, alternatively, as the customer base of a small firm producing under a capacity constraint. Multiple markets introduce a new externality. Sales of the new product in any given market provide information to all of the small producers. As a result, they are more aggressive in their own pricing decisions as undercutting the new firm no longer results in a complete stop in the flow of information. As the number of markets increases, the equilibrium experimentation path converges to the efficient path.

The continuous time techniques we use allow us to derive the equilibria in closed form. In particular, we can analyse the prices in a more detailed manner than would be possible in a discrete time model. By examining the price paths, we gain new insights into the strength of the rivalry between the firms at various points in the game.

There are a number of papers examining learning and experimentation in multi-agent settings. The most relevant papers for our purposes are Bolton and Harris (1999), Bergemann and Välimäki (1996) and Felli and Harris (1996). Bolton and Harris analyse a continuous-time game of strategic experimentation and our methodology follows theirs closely. They focus on the pure informational externality between a set of identical agents and show that all the equilibria involve too little experimentation. Bergemann and Välimäki present a model with a single buyer, and consequently the issue of informational externalities between the buyers does not arise. Finally, Felli and Harris study the wage dynamics in a continuous time learning model about an employee's firm-specific productivity. Again, the model considers only a single employee and hence informational externalities do not arise.

The first paper to address the potential free rider problems in the presence of informational externalities is the Hendricks and Kovenock (1989) analysis of oil exploration

1. There is a multiplicity of equilibrium prices. The situation is similar to the static Bertrand model with differentiated products. We concentrate on equilibria in trembling hand perfect prices.

games. They showed that experimentation can be insufficient or excessive due to the externalities. Rob (1991) studies the informational impact of entry decisions into an industry with uncertain profitability. Successful entry attracts more entry which reduces the long-run profits of early successful entrants. As a result the amount of entry in each period is inefficiently low. Chamley and Gale (1994) present the free-riding aspect in a similar timing game in which each agent has to make a single investment decision. Vettas (1998) allows for two-sided learning in a market of new products. The firms learn about the market size while the buyers make inferences about product quality. Bergemann and Välimäki (1997) consider the diffusion of a new product in a model of horizontal differentiation. Keller and Rady (1999) consider experimentation in a changing environment by a monopolist. The public observability of the utility signals, or a subset of those signals, is central to some recent models of world-of-mouth communication and social learning such as McFadden and Train (1996) and Banerjee and Fudenberg (1997).

Section 2 introduces the elements of the dynamic game in a two-period example. Its structure is kept as simple as possible to convey the basic aspects of the two-sided market environment. The continuous-time model is introduced in Section 3. The single market model is analysed in Section 4. In Section 5 we present the model with many markets. Section 6 concludes.

2. A TWO-PERIOD EXAMPLE

Two sellers provide quality differentiated products to a unit mass of identical buyers with unit demand in each period.[2] The incumbent supplies a product with known quality, while the quality supplied by the entrant is initially unknown. The value of the established product is s per period and the new product has a value of either μ_L or μ_H with $\mu_L < s < \mu_H$. Let α be the common prior probability that the product has value μ_H and denote the expected quality by $\mu(\alpha) \triangleq \alpha\mu_H + (1-\alpha)\mu_L$. The marginal costs of production are identical and normalized to zero.

Firm j chooses price p_j^t in period $t \in \{1, 2\}$, where $j = 0$ indexes the entrant, and $j = 1$ the incumbent. The net utility of a purchase to the buyer is the (expected) quality of the product minus its current price. Buyers and sellers maximize the sum of their per-period payoffs.

The revelation of uncertainty takes an extremely simple form. If a fraction x of the buyers experiment with the new product, then its true quality is revealed to all agents in the second period with probability x. With the complementary probability, no new information arrives.

Consider first the continuation games after zero or complete experimentation: $x \in \{0, 1\}$. Under complete experimentation in the first period ($x = 1$), the new product is worth μ_H with probability α in the second period. The second period prices are given by Bertrand competition: $p_0^2 = \mu_H - s$ and $p_1^2 = 0$, and all buyers purchase from the entrant. With probability $1 - \alpha$, the quality is low and the second-period prices are $p_0^2 = 0$ and $p_1^2 = s - \mu_L$, and all buyers purchase from the incumbent. Conditional on complete experimentation in the first period, the expected second-period profits for the two firms are $\pi_0^2 = \alpha(\mu_H - s)$ and $\pi_1^2 = (1-\alpha)(s - \mu_L)$. If there is zero experimentation in the first period, then second-period prices are given by $p_0^2 = \max\{\mu(\alpha) - s, 0\}$ and $p_1^2 = \max\{s - \mu(\alpha), 0\}$, and the firm with a positive price sells to the entire market. We assume for the rest of this

2. We are grateful to Eric Maskin for suggesting this example.

section that $\mu(\alpha)<s$. This implies that from a myopic point of view, experiments are costly.

The first-period equilibrium prices, p_0^1 and p_1^1, are found by backward induction. Since each consumer is of measure zero, the future payoff of an individual buyer is independent of her current product choice. The equilibrium condition under Bertrand pricing then requires the buyer to be indifferent between the two offers

$$\mu(\alpha)-p_0^1=s-p_1^1, \tag{1}$$

and hence the price differential has to be equal to the (expected) quality difference. Moreover, we require that the non-selling firm be indifferent between selling and not selling at equilibrium prices. Prices satisfying this requirement are called cautions. With the linearity of the payoffs in x, either all buyers or none buy from the new firm in equilibrium. The values of α at which experimentation occurs in equilibrium are characterized by two conditions. First, the incumbent must prefer to concede the market in the first period and to make sales in the second period if the new good fails in the first period

$$p_1^1+s-\mu(\alpha)\leqq(1-\alpha)(s-\mu_L).$$

With cautious pricing, this holds as an equality and

$$p_1^1=\alpha(\mu_H-s). \tag{2}$$

Second, the entrant has to make nonnegative expected profits by selling today and betting on a favourable resolution of uncertainty tomorrow

$$p_0^1+\alpha(\mu_H-s)\geqq 0. \tag{3}$$

The values of α that satisfy (1)–(3) induce experimentation in the first period. The conditions (1)–(3) imply that

$$p_0^1=\alpha(\mu_H-s)+\mu(\alpha)-s\geqq-\alpha(\mu_H-s).$$

Hence experimentation occurs in equilibrium whenever

$$\alpha\geqq\alpha^*=\frac{s-\mu_L}{(\mu_H-\mu_L)+2(\mu_H-s)}.$$

On the other hand, the socially efficient policy requires experimentation whenever current costs of experimentation are outweighed by future gains

$$\alpha(\mu_H-s)\geqq s-\mu(\alpha),$$

or

$$\alpha\geqq\hat{\alpha}=\frac{s-\mu_L}{(\mu_H-\mu_L)+(\mu_H-s)}.$$

As $\alpha^*<\hat{\alpha}$, we conclude that the cautious equilibrium exhibits excessive experimentation.

This inefficiency can be traced to the divergence of the private cost from the social cost of experiments in equilibrium. The social benefit, $\alpha(\mu_H-s)$, coincides with the entrant's private benefit. The social cost is given by the myopic losses, $s-\mu(\alpha)$. The private cost of supporting the experiment, *i.e.* the negative price that the entrant has to quote, is $p_0^1=\mu(\alpha)-s+\alpha(\mu_H-s)$. The additional term $\alpha(\mu_H-s)$ is the price of the incumbent, and thus reflects his informational gain through cautious pricing. The failure of the buyers to take the future surplus extraction into account reduces the private cost to finance experimentation. In contrast to the duopoly, where the identity of the benefitting seller depends on the outcome of the experiment, a monopoly would extract the social surplus at every stage and the equilibrium would be efficient. More insight into the discrepancy between the efficient and the equilibrium allocation may be obtained by considering the case where all buyers act collectively and make purchases as a cooperative.[3]

In equilibrium, the cooperative is indifferent between the two products at current prices. Hence the price differential is equal to the sum of the quality differential *and* the change in the continuation payoff resulting from experimentation,

$$p_0^1 - p_1^1 = \mu(\alpha) - s + \alpha s + (1-\alpha)\mu_L - \mu(\alpha) = \mu(\alpha) - s - \alpha(s - \mu_H). \tag{4}$$

By cautious pricing, p_1^1 equals the expected gain from experimentation for the incumbent when the entrant is selling in the first period. Notice that with the cooperative, the expected losses from experimentation for the buyer equal exactly the incumbent's expected gain. The equilibrium condition (4) then shows that the private cost of experimentation for the entrant coincides with the social cost, or $p_0^1 = \mu(\alpha) - s$, and efficiency follows. The dynamic implications of the informational externality are considered next in a more gradual process of information revelation and with a finite number of buyers and sellers.

3. THE MODEL

The market consists of buyers and sellers. The buyers are indexed by $i \in \{1, \ldots, N\}$. Two sellers, $j = 0, 1$, offer differentiated products and compete in prices in a continuous time model with an infinite horizon. Firm 0 is called the new firm or the entrant. The value, μ, of its product is initially unknown to *all* parties in the market and we refer to it as the new or uncertain product. It can be either low or high

$$\mu \in \{\mu_L, \mu_H\},$$

with $0 < \mu_L < s < \mu_H$. The value of firm 1's product is s, and the firm is called the established firm or the incumbent. All players share a common prior $\Pr(\mu = \mu_H) = \alpha_0$. Let $n_j(t)$ denote the number of buyers that purchase firm j's product in period t.

3.1. *Bayesian learning*

The uncertainty about the value of the new product can be resolved only by experimenting. The performance of the new product is, however, subject to random disturbances. The information resulting from any single purchase provides a noisy signal of the true underlying quality. The flow utility from the uncertain alternative is

$$du_i(t) = \mu dt + \sigma dW_i(t),$$

where $dW_i(t)$ is the increment of the one-dimensional standard Wiener process. We assume that dW_i and $dW_{i'}$ are independent for $i \neq i'$. The flow utility of the established product is given by

$$du_i(t) = s dt.$$

The aggregate performance of the product over all buyers is the sum of flow utility realizations, which we assume to be publicly observable. Let $n_0(t)$ be the number of the new firm's customers and w.l.o.g. assume that for a given $n_0(t)$, the buyers with the smallest

3. This is the case analysed in Bergemann and Välimäki (1996).

indices buy from the new firm. The aggregate performance of the new product is given by

$$\sum_{i=1}^{n_0(t)} du_i(t) = n_0(t)\mu dt + \sigma \sum_{i=1}^{n_0(t)} dW_i(t). \tag{5}$$

All relevant information is contained in the aggregate outcome.[4] As the value of μ is either μ_L or μ_H, the posterior beliefs are given by $\alpha(t)$, with

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

where $\mathcal{F}(t)$ is the history (or more accurately, the filtration) generated by the past observations. The conditional expected quality $\mu(\alpha(t))$ of the uncertain product is

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

The players extract the information contained in the noisy market outcome (5) to update their beliefs. The game is thus one of incomplete but symmetric information, and no issues of asymmetric information arise. As the beliefs are characterized by $\alpha(t)$, the inference problem reduces to the description of the law of motion of $\alpha(t)$.[5] It can be shown that $\alpha(t)$ is a diffusion process with zero drift and instantaneous variance $n_0(t)\Sigma(\alpha(t))$, where

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

The process of the posterior $\alpha(t)$ has zero drift since posterior beliefs form a martingale and any change in $\alpha(t)$ has zero expectation. The variance $\Sigma(\alpha(t))$ measures the additional information obtained through an experiment. An increase in the variance causes a more rapid change in the posterior. The variance of $\alpha(t)$ is linear in the size of the aggregate experiment, $n_0(t)$, and in the "signal to noise" ratio $((\mu_H - \mu_L)^2)/\sigma^2$. A large market share for the new firm results in a faster change of the posterior $\alpha(t)$ as more information is generated.

The key assumption in our model is the public observability of the utility signals. This allows us to abstract from the issue of idiosyncratic differences in posterior beliefs that might arise from different experiences. While this is clearly a strong assumption, we feel that it approximates well two important economic situations. In the first, a large number of buyers are faced with the same dynamic decision problem. Even an imperfect aggregate measure of performance by the new product is valuable to the buyers' choices. This situation is common in the provision of services, and examples include the data collected by consumer agencies on the percentage of flights arriving on time on a new airline or average waiting times for internet services on a new provider. The second class of problems involves once and for all purchases of durable goods. In Section 4, we argue that our model can be reinterpreted as one of durable goods sales to an inflow of new buyers. In that case, no buyer has individual information on the products at the moment when the choice is made and public information is the only basis for assessing the products.

4. Recall that in Bayesian learning models with normal distributions, the sample mean and the number of observations constitute a sufficient statistic for the observation of n i.i.d. draws from a distribution with known variance and unknown mean.

5. See Liptser and Shirayayev (1977), Chapter 9, for the filtering equations of the Brownian motion in continuous time.

3.2. *Strategies and equilibrium*

The dynamic game consists of two components: the pricing strategies of the sellers and the acceptance decisions of the buyers. The strategies depend on information available to the agents at the instant of decision. Since we want to analyse how the resolution of uncertainty affects the pricing game between the sellers, we concentrate on equilibria in Markovian strategies. This allows us to rule out collusive equilibria with continuation strategies that depend on information that is not payoff relevant.

We view the model as a continuous-time analogue of the repeated extensive form game where the sellers set prices at the beginning of each stage, and buyers then choose where to buy. With this in mind, the natural state variable in period t is $\alpha(t)$ for the sellers, and for the buyers it is $\alpha(t)$ together with the prevailing prices, $p_0(t)$ and $p_1(t)$. A *pricing policy* p_j is a measurable function p_j: $[0, 1]\rightarrow\mathbb{R}$. In the equilibrium analysis we show that the restriction to pure strategies is without loss of generality. Each buyer has unit demand at every instant and her *acceptance policy*, a_i, determines where to purchase the product: a_i: $[0, 1]\times\mathbb{R}\times\mathbb{R}\rightarrow\{0, 1, R\}$. The acceptance policy specifies whether the buyer accepts seller 0 or 1 or rejects both (R).

The firms offer prices $p_j(t)$ at each instant of time and their instantaneous revenues evolve as a function of market share $n_j(t)$ and price $p_j(t)$

$$d\pi_j(t) = n_j(t)p_j(t)dt.$$

The marginal cost of production is constant and normalized to zero. The flow utility of a buyer is determined by her choice among the competing products. It is her flow utility net of the current price

$$du_i(t) - p_j(t)dt.$$

The intertemporal profits of each firm j are a function of the pricing and acceptance policies of the market players

$$V_j(a, p_j, p_{-j}; \alpha) = \mathbb{E}_\alpha \int_0^\infty e^{-rt} d\pi_j(t),$$

where $r>0$ is the discount rate and $a = (a_1, \ldots, a_N)$. The intertemporal utility functional of each buyer is the discounted flow of net utilities

$$V_i(a_i, a_{-i}, p; \alpha) = \mathbb{E}_\alpha \int_0^\infty e^{-rt}(du_i(t) - p_j(t)dt),$$

where $a_i(\alpha, p_0(t), p_1(t)) = j$ and $p = (p_0, p_1)$.

Definition 1. (*Markov Perfect Equilibrium, MPE*). A collection of strategies, $\{a^*, p^*\}$, is a Markov Perfect Equilibrium if

(i) $V_j(a^*, p_j^*, p_{-j}^*; \alpha) \geqq V_j(a^*, p_j, p_{-j}^*; \alpha), \forall j, \forall p_j, \forall \alpha,$

(ii) $V_i(a_i^*, a_{-i}^*, p^*; \alpha) \geqq V_i(a_i, a_{-i}^*, p^*; \alpha), \forall i, \forall a_i, \forall \alpha.$

Notice that the deviation strategies a_i and p_j are not required to be Markovian. The equilibrium analysis presented in Sections 4 and 5 involves the solution of stopping problems derived from the equilibrium conditions. The basic technique is most clearly illustrated in the solution of the socially efficient allocation.

3.3. *Efficiency*

The efficient allocation policy in the market model can be obtained by solving a specific multi-armed bandit problem. Here we provide only the essentials of the solution technique and refer the reader to Karatzas (1984) for details. Since we have assumed quasilinear utilities for all of the players in the game, finding the set of Pareto-efficient allocations is equivalent to solving the total surplus maximization problem.

An allocation policy n is a measurable and adapted process with values $n_j(t) \in \{1, \dots, N\}$. Clearly, an optimal policy allocates all buyers to one of the firms. Let $n(t)$ be the number of the new firm's customers and w.l.o.g. assume that for a given $n(t)$, the buyers with the smallest indices buy from the new firm. The expected value of an allocation policy n is given by

$$V(n; \alpha) = \mathbb{E}_\alpha \int_0^\infty e^{-rt}\left(\sum_{i=1}^{n(t)} (\mu dt + \sigma dW_i(t)) + \sum_{i=n(t)+1}^{N} s dt\right). \tag{6}$$

The problem is then to find a policy n^* so as to maximize expected value

$$V(\alpha) \triangleq \max_n V(n; \alpha), \qquad \text{with } V(\alpha) = V(n^*; \alpha), \qquad \forall \alpha \in [0, 1].$$

The controls in the allocation problem are continuous, but due to the convexity of the value function, it can be shown that (6) is equivalent to the optimal stopping problem[6]

$$V(\alpha) = \max_\tau V(\tau; \alpha) = \mathbb{E}_\alpha\left[\int_0^\tau e^{-rt}\left(\sum_{i=1}^{N} \mu dt + \sigma dW_i(t)\right) + e^{-r\tau}\frac{Ns}{r}\right],$$

where τ is a stopping time.[7] As the instantaneous payoff $\mu(\alpha)$ is increasing in α, one might expect a solution that chooses the uncertain alternative in a half-open interval $(\hat{\alpha}, 1]$ and stops the process at $\hat{\alpha}$. Using Itô's lemma and the stochastic differential equation governing the posterior process, the Hamilton–Jacobi–Bellman (HJB) equation of this problem for $\alpha > \hat{\alpha}$ is the following differential equation[8]

$$r V(\alpha) = N(\mu(\alpha) + \tfrac{1}{2}\Sigma(\alpha)V''(\alpha)), \tag{7}$$

with the initial boundary conditions given by the value matching and the smooth pasting conditions

$$r V(\hat{\alpha}) = Ns,$$
$$V'(\hat{\alpha}) = 0.$$

The differential equation (7) represents the flow benefit during the experimentation phase. It consists of the expected payoff $N\mu(\alpha)$ and the informational gains $(N/2)\Sigma(\alpha)V''(\alpha)$ which improve the intertemporal policy. The instantaneous variance $\Sigma(\alpha)$ indicates the

6. The assumed convexity is shown to hold in the solution of the stopping problem and the uniqueness of the value function then justifies the initial hypothesis.

7. Recall that a stopping time is a real valued, $\mathcal{F}(t)$-measurable random variable. In other words, stopping at t has to be decided based upon history at time t.

8. The HJB equation is the dynamic programming equation in continuous time. It can be shown that the HJB equation admits a unique solution that is piecewise twice continuously differentiable. See Dixit and Pindyck (1994) or Harrison (1985) for the details.

quantity of information released by a unit of experimentation and the curvature of the value function $V''(\alpha)$ is the shadow price of information for the planner. It is then optimal to experiment as long as the flow payoff and the flow value of information exceed the value of the safe alternative.

Theorem 1. (Efficient Stopping). *The Pareto efficient stopping point* $\hat{\alpha}$ *is*

$$\hat{\alpha} = \frac{(s-\mu_L)(\gamma-1)}{(\mu_H-\mu_L)(\gamma-1)+2(\mu_H-s)},$$

with

$$\gamma = \sqrt{1+\frac{8r\sigma^2}{N(\mu_H-\mu_L)^2}}.$$

Proof. See Appendix. ||

For notational brevity, we define ρ by

$$\rho \triangleq \frac{8r\sigma^2}{(\mu_H-\mu_L)^2}. \tag{8}$$

The stopping point $\hat{\alpha}$ is strictly increasing and $V(\cdot)$ is strictly decreasing in ρ/N.

4. SINGLE MARKET

The basic model with a finite number of buyers in a single market is presented in Subsection 4.1. The limiting case with a continuum of small buyers is analysed in Subsection 4.2. The limiting model is also interpreted as a model of durable goods sales with a flow of incoming buyers.

4.1. *Finitely many buyers*

Since we are looking for equilibria in Markovian strategies, we can use dynamic programming techniques directly. The value functions of the sellers are functions of the current posterior $\alpha(t)\in[0, 1]$. Using Itôs lemma and the equation governing the posterior process, we get the Hamilton–Jacobi–Bellman equations

$$rV_0(\alpha) = \max_{p_0}\left\{np_0+\frac{n}{2}\Sigma(\alpha)V_0''(\alpha)\right\}, \tag{9}$$

and

$$rV_1(\alpha) = \max_{p_1}\left\{(N-n)p_1+\frac{n}{2}\Sigma(\alpha)V_1''(\alpha)\right\}, \tag{10}$$

where n is the number of buyers who purchase from the entrant. Observe that the objective functions are linear in sales for both firms.

Each $V_j(\cdot)$ can be decomposed into the flow revenue resulting from sales, n_jp_j, and the expected change in the competitive position generated by the sales of the new product, captured by $(n/2)\Sigma(\alpha)V_j''(\alpha)$. As the posterior belief $\alpha(t)$ forms a martingale, $\mathbb{E}[d\alpha(t)] = 0$, only the second order term in the expected change remains. With Markovian strategies, each firm's future competitive position is influenced only by the arrival of new information. Hence $\frac{1}{2}\Sigma(\alpha)V_j''(\alpha)$ can be interpreted as the value of information to firm j. Experiments lead to more differentiation (in *ex post* terms) between the two competitors as information is accumulated. In a more differentiated environment, competitive pressure between the firms is reduced and they are able to extract more surplus from the buyers. As a result, one would expect that $V''(\alpha) > 0$.

The value function of buyer i is

$$rU_i(\alpha) = \max\left\{s - p_1 + \frac{n}{2}\Sigma(\alpha)U_i''(\alpha), \mu(\alpha) - p_0 + \frac{n+1}{2}\Sigma(\alpha)U_i''(\alpha)\right\}, \tag{11}$$

when n other buyers choose the uncertain object. The choice of the buyer is determined by the expected flow payoff and the flow value of information released through the choice. By selecting the uncertain firm, her future payoffs are changed by the amount of $\frac{1}{2}\Sigma(\alpha)U_i''(\alpha)$. Equation (11) of buyer i is representative of all buyers as they have identical preferences and they have access to the same information. We can therefore omit the index i entirely and consider the representative buyer. Differentiation leads to less competitive prices and hence enables the sellers to extract more surplus from the buyers. As a consequence, it is to be expected that the buyers have a negative value of information, or $U''(\alpha) < 0$.

Due to price competition, each buyer has to be indifferent between the alternatives, or formally

$$s - p_1 = \mu(\alpha) - p_0 + \tfrac{1}{2}\Sigma(\alpha)U''(\alpha). \tag{12}$$

It is never optimal to leave any buyer with more net surplus than she would obtain from the alternative seller, nor could any seller ever expect to make sales if he would offer strictly less utility than the competitor.

As the objective functions of the sellers are linear in market shares and (12) holds, we conclude that nontrivial market sharing can happen only if both sellers are indifferent between selling any amount $n_j > 0$ or not selling at all. But at any such state α, stopping the experiments must also be an equilibrium outcome. We then conjecture that the equilibrium allocation has the following simple structure: Since the payoff of the second seller is increasing in α, there is a half-open interval $(\alpha^*, 1]$, called the continuation region, where experiments occur at the maximal rate $n_0 = N$, and a closed interval $[0, \alpha^*]$, called the stopping region, in which no experiments take place and $n_1 = N$.

Before we verify this conjecture, a qualification for the pricing policies is made. We require that any price quoted by a firm which is not selling in a given period would make the firm at least weakly better off if accepted. In an earlier paper, Bergemann and Välimäki (1996), we called prices which satisfy this property "cautious".[9] the value functions in (9) and (10) imply that cautious prices satisfy

$$p_0(\alpha) \geqq -\tfrac{1}{2}\Sigma(\alpha)V_0''(\alpha), \tag{13}$$

9. This requirement captures the logic behind trembling hand prefection in this infinite time horizon framework. Notice that prices p_j at which sales, $n_j > 0$, occur always satisfy the cautious property. Without cautiousness, any switching point between $\hat{\alpha}$ and the cautious equilibrium switching point α^* can be supported.

and

$$p_1(\alpha) \geq \tfrac{1}{2}\Sigma(\alpha)V_1''(\alpha). \tag{14}$$

Condition (13) states that the entrant is willing to sell only if the price is at least offset by the value of the information flow. In contrast, condition (14) states that the incumbent is willing to sell only if he receives at least enough revenue to compensate for the foregone informational gains. A simple undercutting argument establishes that the prices p_j need to satisfy the appropriate inequality as an equality if $n_j = 0$. With positive sales, the prices are obtained by using (12)

$$p_0(\alpha) = \mu(\alpha) - s + \tfrac{1}{2}\Sigma(\alpha)(U''(\alpha) + V_1''(\alpha)), \quad \text{if } n_0 > 0, \tag{15}$$

and

$$p_1(\alpha) = s - \mu(\alpha) - \tfrac{1}{2}\Sigma(\alpha)(U''(\alpha) + V_0''(\alpha)), \quad \text{if } n_1 > 0. \tag{16}$$

The price of each seller has two components. First, the price extracts or subsidizes the difference in the current expected value of the alternatives. The second component reflects the intertemporal incentives of the competitor and the individual buyer. The pricing policies (15) and (16) display an important asymmetry in the influence the value of information has on the pricing policies. If the new firm sells its product, then the experiments generate information in the market. In contrast, if the established firm intends to make a sale, it has to recognize that it reduces the information flow in the market, which is reflected in the sales prices in (16).

Since any single experiment provides relevant information not only to the buyer who purchases the new good, but to all buyers, the experiment of an individual buyer generates an externality among all buyers. This effect is not properly reflected in the equilibrium prices. If, as was previously argued, the value of information, $U''(\alpha)$, is negative to the buyers, then the equilibrium price overstates the value of an experiment. In fact, if the new seller were to absorb the cost of the negative externality imposed on all buyers by a single experiment, then the price would have to be

$$p_0(\alpha) = \mu(\alpha) - s + \tfrac{1}{2}\Sigma(\alpha)(NU''(\alpha) + V_1''(\alpha)). \tag{17}$$

The difference between (15) and (17) indicates the divergence between the market price and the social price of the experiment, which increases in the number of buyers.

Using (15) and (16), the sellers' optimality conditions can be written as

$$rV_0(\alpha) = N\max\{\mu(\alpha) - s + \tfrac{1}{2}\Sigma(\alpha)(V_0''(\alpha) + V_1''(\alpha) + U''(\alpha)), 0\}, \tag{18}$$

and

$$rV_1(\alpha) = N\max\{\tfrac{1}{2}\Sigma(\alpha)V_1''(\alpha), s - \mu(\alpha) - \tfrac{1}{2}\Sigma(\alpha)(V_0''(\alpha) + U''(\alpha))\}. \tag{19}$$

It is easily verified that the two value functions represent the same stopping problem. The indifference condition (12) and cautious pricing reduce the two-dimensional control problem in $p_0(\alpha)$ and $p_1(\alpha)$ into a one-dimensional stopping problem in α which can be stated as: How long can the entrant afford a pricing policy that captures the entire market? The value function $V_0(\alpha)$ indicates that the extent of experimentation depends on the benefits to the sellers and the costs to the buyers. The equilibrium stopping point is obtained by continuity conditions on the value functions of all players and a smoothness condition associated with the stopping problem of the new seller.[10]

10. The explicit derivation of the value functions is presented in Lemma 1 in the Appendix.

Theorem 2. (Equilibrium Stopping). *There is a unique MPE in cautious strategies. The equilibrium path displays excessive experimentation. The stopping point is given by*

$$\alpha^* = \frac{(s-\mu_L)(\gamma N - \lambda(N-1)-1)}{(\mu_H-\mu_L)(\gamma N-\lambda(N-1)-1)+2(\mu_H - s)} < \hat{\alpha},$$

with $\lambda = \sqrt{1+\rho/(N-1)}$, *and* γ *and* ρ *as in Theorem* 1.

Proof. See Appendix. ||

The threshold α^* at which stopping occurs is decreasing in the number of buyers. An increase in the market size N increases the signal to noise ratio of the outcome when all buyers experiment and thus increases the return from experimentation. For $N=1$, the equilibrium coincides with efficient stopping: $\alpha^* = \hat{\alpha}$, as shown in Bergemann and Välimäki (1996). The equilibrium price policies $p_j(\alpha)$ and the curvatures of the value functions $V_j(\alpha)$ and $U(\alpha)$ follow directly from the solution of the equilibrium stopping problem.

Corollary 1. (Submartingale and Convexity).

1. *The pricing policies* $p_j(\alpha)$ *are submartingales.*
2. *The value functions of the sellers are convex.*
3. *The value functions of the buyers are concave.*

The submartingale characterization of the prices illustrates the dilemma facing the buyers. As the expected quality of each product follows a martingale, the submartingale prices imply that buyers expect decreasing net utilities over time. In fact, the instantaneous utilities, $du_i(\alpha) = \mu(\alpha) - p_0(\alpha)$, of the buyers form a strict supermartingale in the continuation region. As the established seller is indifferent between selling and not selling in the experimental phase, we have

$$p_1(\alpha) = \tfrac{1}{2}\Sigma(\alpha)V_1''(\alpha),$$

and in consequence

$$p_1(\alpha) = \frac{r}{N}V_1(\alpha).$$

The foregone revenue at any instant of time must be equal to the expected increase in discounted future revenue, or

$$p_1(\alpha)dt = \frac{\mathbb{E}[dp_1(\alpha)]}{r} \Leftrightarrow \mathbb{E}[dp_1(\alpha)] = rp_1(\alpha)dt,$$

from which it follows that the price of the incumbent has a positive drift.

The selling price $p_0(\alpha)$ has two interesting features. It is negative between α^* and some α with $\mu(\alpha) - s < 0$. The negative price compensates the buyers for their purchases of the (myopically) lower quality product. But $p_0(\alpha)$ is not monotone increasing in α as one might have expected. As α approaches α^* from the right, the value function of the incumbent firm increases as the likelihood of stopping increases and the expected time to stopping (conditional on eventual stopping) decreases. Through cautious pricing, this leads to higher prices posted by the incumbent. The competitive pressure on the entrant is thus relieved and he can charge higher prices. Thus a segment of decreasing prices (as a function of α) is observed. Less aggressive prices by the incumbent allow the entrant to shift the cost of the experiments to the buyers as beliefs become more pessimistic.

The curvatures of the value functions reflect the attitudes towards information. As the posterior belief α approaches the stopping point α^*, the value of information for the entrant declines and at α^*, we observe $V_0''(\alpha^*)=0$ and by implication $p_0(\alpha)$ converges to 0 from the right. For posterior beliefs close to α^*, stopping becomes almost certain and the buyers become less averse to experimentation, and at α^*, we find $U''(\alpha^*)=0$. A typical pair of equilibrium price policies $p_j(\alpha)$ and flow utilities $du(\alpha)-p_j(\alpha)dt$ as a function of α are presented in Figure 1.

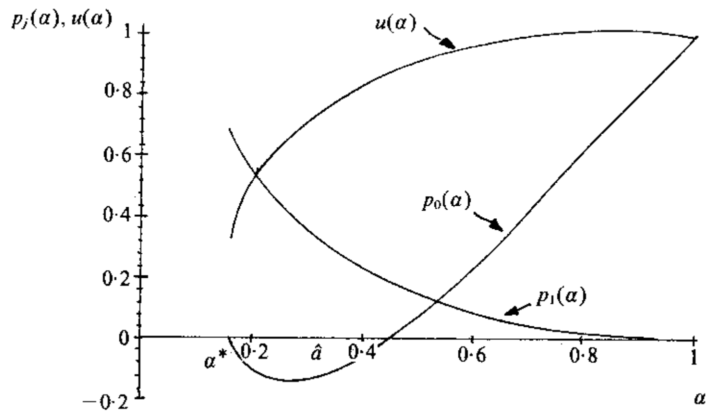

FIGURE 1

Prices $p_j(\alpha)$ and flow utilities $u(\alpha)$ in a single market: $\mu_H=2$, $\mu_L=0$, $s=1$, $N=2$, $\sigma=4$, $r=\frac{1}{2}$

4.2. *Infinitely many buyers*

The previous analysis suggests that the inefficiency increases as the size of a single buyer becomes small relative to the size of the market. The experiments should then extend to threshold levels $\alpha^*(N)$ which decline in the number N of buyers. To make this intuition precise, we have to be careful when changing N. If we were to simply increase the number of buyers, we would also change the informativeness of the aggregate outcome in the continuation region. In fact, as the number of buyers increases, the instantaneous aggregate outcome would become completely informative by the law of large numbers, and even the efficient stopping point $\hat{\alpha}$ would converge to 0.

It is therefore necessary to keep the informativeness of the aggregate outcome independent of the number of buyers if we want to analyse the effect of making the buyers small without eliminating the aggregate uncertainty. This can be accomplished by decreasing the size of the individual experiment as we increase the total number of experiments. By normalizing the flow utility from the experiment to

$$du_i(t)=\frac{\mu}{N}dt+\frac{\sigma}{\sqrt{N}}dW_i(t),$$

the size of the aggregate experiment remains unchanged as we increase N. Intuitively, in a market with fixed aggregate variance σ^2, the effect of increasing N then reflects solely the increase in the externality between the buyers. Note in particular that the socially-efficient allocation policy is independent of N.

The inefficiency due to the free-riding aspect increases as each buyer has an ever smaller influence on the release of information. In consequence, the entrant has to compensate each individual buyer less and less for the participation in the experiment. It then becomes less costly for the entrant to finance the experiments and the stopping point $\alpha^*(N)$ decreases as a function of N. The equilibrium stopping point $\alpha^*(N)$ is as in Theorem 2 given by

$$\alpha^*(N)=\frac{(s-\mu_L)(\tilde{\gamma}N-\tilde{\lambda}(N-1)-1)}{(\mu_H-\mu_L)(\tilde{\gamma}N-\tilde{\lambda}(N-1)-1)+2(\mu_H-s)}, \tag{20}$$

where the only changes are introduced through the parameters $\tilde{\gamma}$ and $\tilde{\lambda}$ which reflect the normalization of the market size, with

$$\tilde{\gamma}=\sqrt{1+\rho}, \tag{21}$$

and

$$\tilde{\lambda}=\sqrt{1+\rho\frac{N}{N-1}}. \tag{22}$$

The parameter $\tilde{\gamma}$ is now independent of N as the aggregate market size is constant and equal to 1. The parameter $\tilde{\lambda}$ represents the share of the variance in the posterior belief for which the individual consumer is not compensated, namely $(N-1)/N$. The social value of the game is decreasing in N due to the excessive information acquisition. The monotonicity of the allocation in N is associated with the monotonicity of the value functions of the agents. As N increases, the seller of the unknown product receives a successively larger share of the social surplus to the detriment of the known seller and the aggregate value of the buyers. The comparative static results follow directly from the equilibrium value functions.

Corollary 2. (Convergence and Monotonicity).

1. *The equilibrium stopping point $\alpha^*(N)$ is strictly decreasing in N and converges to the stopping point with infinitely many buyers*

$$\alpha^*_\infty=\frac{(s-\mu_L)(\tilde{\gamma}-1)^2}{(\mu_H-\mu_L)(\tilde{\gamma}-1)^2+4\tilde{\gamma}(\mu_H-s)}.$$

2. *The value function $V_0(\cdot)$ is strictly increasing in N for all $\alpha>\alpha^*(N)$.*
3. *The value functions $NU(\cdot)$ and $V_1(\cdot)$ are strictly decreasing in N for all $\alpha>\alpha^*(N)$.*

Proof. See Appendix. ||

The extent of the inefficiency as N increases is depicted in Figure 2, which shows the price path of $p_0(\alpha)$ as N varies. As N increases the new firm succeeds in shifting the costs of the experiments to the buyers. Moreover the plateau with $p_0(\alpha)=0$ recedes as N increases.

In the limit as $N\to\infty$, any single experiment carries no information about the value of the uncertain alternative. The single buyer is now infinitesimally small relative to the market and her purchase decision has no impact on the informational content of the market outcome. Consequently, neither the individual buyer nor the sellers attach any strategic importance to her current decision. In the limit the representative buyer then

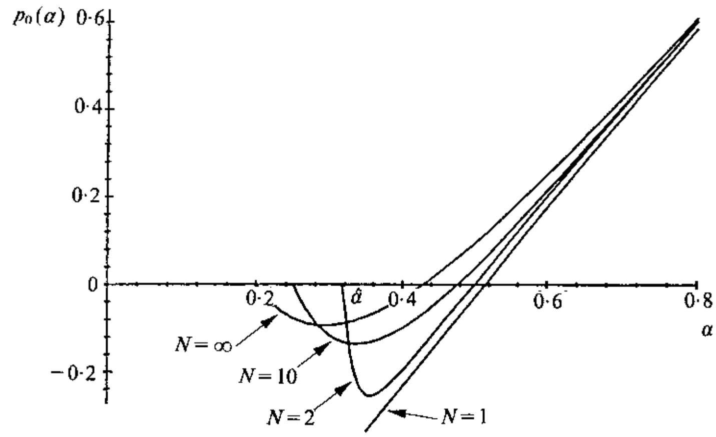

FIGURE 2

Prices $p_0(\alpha)$ in a single market with a variable number of buyers: $\mu_H=2$, $\mu_L=0$, $s=1$, $\sigma=4$, $r=\frac{1}{2}$

behaves as if she were completely myopic. The indifference condition of each buyer simplifies to

$$\mu(\alpha)-p_0=s-p_1, \tag{23}$$

as the size of the individual experiment is infinitesimal. We can therefore strengthen Corollary 1 and obtain symmetry in the expected rates of change of the prices quoted by the sellers.

Corollary 3. *In the continuation region, the pricing policies of both firms are submartingales with identical and positive drift*

$$\mathbb{E}[dp_j(\alpha)]=rp_1(\alpha)dt, \quad \textit{for } j=0,1.$$

Interestingly, the model with a continuum of buyers can be reinterpreted as one with a constant inflow of buyers that make a once and for all purchase of a durable good. With perfect durables, it is natural that the experiences of past buyers have a large impact on current buyers' beliefs on product quality. In the absence of repeat purchases, it is also incentive compatible for all buyers to report their experiences truthfully *even* in the case of private observability. The only new constraint that appears in the durable goods case is that we need to make sure that no buyer wants to delay her purchase. As prices are always below the expected value of the product and the expected change in prices is nonnegative, waiting entails a loss to the buyers in all periods. Notice also that it is never in the firms' interest to delay sales. Hence the equilibrium as described above yields an equilibrium in this durable goods model as well. Instead of having the same buyers switch from one product to the other, different buyers choose different products in the durable goods interpretation. As each buyer purchases only once, the informational effects on future prices are not internalized.

5. MANY MARKETS

In the single market model, the sellers internalize the effect from all information flows in the model while the buyers are not able to do this. In this section, we consider the case where the informational externality affects both sellers and buyers.

We start by considering M distinct markets, each having a separate incumbent offering an established product of value s to the buyers in his market. The M markets may be thought of as local markets and the entrant as a global competitor, who can introduce his new product into all local markets. An alternative interpretation of this model would be that the incumbents are using a technology with a capacity constraint while the entrant has a technology with unlimited capacity available. The true value of the new product, $\mu \in \{\mu_L, \mu_H\}$, is the same across all markets. The main difference compared to the previous sections is that there is now an informational externality across the markets in addition to the one within the markets.

We assume that the buyers in all markets are identical and that each incumbent is serving the same number of buyers, N/M with $N \geq M$. We normalize the size of the sum of all markets to 1 and as a consequence, each incumbent sells to a market of size $1/M$.[11] The main result in this section is a characterization of the equilibrium stopping for different market structures in terms of M and N. We consider only the symmetric equilibria of the model. The representation of the equilibrium conditions is similar to the previous section and we focus on the differences. By symmetry, it is sufficient to consider a representative incumbent and a representative buyer.

Denote the representative incumbent by j and the entrant by 0. The equilibrium in this model can again be characterized as an optimal stopping problem. The indifference condition of the buyer is, as before

$$s - p_j(\alpha) = \mu(\alpha) - p_0(\alpha) + \tfrac{1}{2}\Sigma(\alpha)U''(\alpha). \tag{24}$$

The optimality condition of the incumbent in the continuation region reflects the trade-off between sales and information

$$rV_j(\alpha) = \max\left\{\frac{1}{M}p_j(\alpha) + \frac{1}{2}\frac{M-1}{M}\Sigma(\alpha)V_j''(\alpha), \tfrac{1}{2}\Sigma(\alpha)V_j''(\alpha)\right\}. \tag{25}$$

The cautious price per unit of sale is again given by setting the price equal to the value of information generated by a sale of the new product

$$p_j(\alpha) = \tfrac{1}{2}\Sigma(\alpha)V_j''(\alpha).$$

Despite the similarity to the single-market behaviour, note how M appears in the decision of the individual incumbent in the Bellman equation (25). The trade-off between sales and information is now relaxed by the fact that even if the incumbent makes a sale in his local market, information may still be generated in the remaining markets. This leads the incumbent to adopt a more aggressive pricing policy which in turn makes it more expensive for the entrant to sell in that particular local market.

The failure of the individual incumbent to take the value of information to the other local incumbents into account leads to less experimentation in equilibrium. This is perhaps most clearly reflected in the equilibrium stopping problem for the entrant. After inserting (24) and (25) into the value function of the entrant, we obtain

$$rV_0(\alpha) = \max\{\mu(\alpha) - s + \tfrac{1}{2}\Sigma(\alpha)(V_0''(\alpha) + V_j''(\alpha) + U''(\alpha)), 0\}. \tag{26}$$

11. This normalization allows us to keep the efficient stopping point constant and independent of the size of the individual buyers and sellers. However, the convergence results to be presented are independent of the normalization.

In contrast, the stopping problem which reflects the social value of all the players and hence would lead to the efficient stopping point is given by

$$\max\{\mu(\alpha)-s+\tfrac{1}{2}\Sigma(\alpha)(V_0''(\alpha)+MV_j''(\alpha)+NU''(\alpha)),0\}.$$

The symmetry in the way that the externalities among buyers and incumbents effect the equilibrium stopping is now apparent. The qualitative difference is that these externalities work in opposite directions in terms of determining α as $V_j''(\alpha)>0$ and $U''(\alpha)<0$. This symmetry is also evident in the solution to the equilibrium stopping problem in (26).

Theorem 3. (Many Market Equilibrium).

1. *The equilibrium stopping point in the unique cautious equilibrium is*

$$\alpha^*(M,N)=\frac{(N\tilde{\gamma}-\tilde{\lambda}(N-M)-M)(s-\mu_L)}{(N\tilde{\gamma}-\tilde{\lambda}(N-M)-M)(\mu_H-\mu_L)+2M(\mu_H-s)},$$

with $\tilde{\gamma}$ *and* $\tilde{\lambda}$ *as in* (21) *and* (22).

2. *If* $M=N$, *the equilibrium stopping is efficient*: $\alpha^*(M,N)=\hat{\alpha}$.
3. *The stopping point* $\alpha^*(M,N)$ *is decreasing in* N *and increasing in* M; *and*

$$\lim_{M,N\to\infty}\alpha^*(M,N)=\hat{\alpha}.$$

Proof. See Appendix. ||

The case of $M=N$ has a few interesting properties. If there is a single buyer in each market, then the local incumbent and buyer can trade intertemporal payoffs on a one-to-one basis. They both have control of $1/M$ of the total market. Hence the future gains from experiments for the incumbent equal exactly the future losses of the buyer. Cautious pricing then allows the new firm to make sales at the myopic quality differential $\mu(\alpha)-s$.

As the number of markets M increases, the strategic aspect in the acceptance and pricing policies vanishes. In consequence, the sales prices in the experimental phase converge to the myopic quality differential

$$\lim_{M,N\to\infty}p_0(\alpha)=\mu(\alpha)-s,\qquad \lim_{M,N\to\infty}p_j(\alpha)=0,\qquad \text{for all } \alpha>\alpha^*(M,N),$$

and $\alpha^*(M,N)$ converges to the efficient stopping point. The convergence result holds also if instead of the single large entrant, each market had a small entrant with access to the same new product as all other local entrants. This may represent a situation where local franchisees share a common franchising product.

With many markets, the prices at the switching point display a discontinuity which is linked to the externality among the incumbents. For $\alpha>\alpha^*(M,N)$, the value $V_j(\alpha)$ of each incumbent is given by the value of information generated in all markets, or

$$rV_j(\alpha)=\tfrac{1}{2}\Sigma(\alpha)V_j''(\alpha),\quad \text{for } \alpha>\alpha^*.$$

After the switch the value is determined by sales in the local market, or

$$rV_j(\alpha^*)=\frac{1}{M}p_j(\alpha^*).$$

We can immediately infer the size of the jump in prices at the switching point by the continuity of the value functions. We denote the limiting point from the right by

$$p_j^+(\alpha^*) \triangleq \lim_{\alpha \downarrow \alpha^*} p_j(\alpha).$$

Corollary 4. *The equilibrium prices at the stopping point display an upward jump for* $M > 1$

$$Mp_j^+(\alpha^*) = p_j(\alpha^*) = s - \mu(\alpha^*),$$

and

$$p_0^+(\alpha^*) = \frac{N(M-1)}{(N-1)M}(\mu(\alpha^*) - s) < 0 = p_0(\alpha^*).$$

The size of the jump is positively related to the number of markets M and disappears for $M = 1$ and $N > 1$. The discontinuity of the prices implies that the flow utility of the buyers decreases discontinuously when the new firm stops selling in the markets.

6. CONCLUSION

This paper shows that the conventional wisdom that informational externalities lead to inefficiently low levels of experimentation may be reversed in a two-sided learning model. The introduction of sellers into the multi-agent learning model creates a market where experiments are priced. The new seller *sponsors* the uncertain alternative and rewards buyers for experiments through low prices. In contrast to one-sided learning models, the seller provides direct incentives for the buyers to experiment. The ownership of the product allows the seller to extract the future benefits of current experimentation which would have evaporated without the assignment of property rights.

The main theme of the paper is the importance of the market structure for efficiency conclusions in a model of informational externalities. To abstract from other forms of distortions we analyse a stage game without static distortions. The cost of this modelling choice is that it is very difficult to generate genuine market sharing between firms. In an earlier working paper (Bergemann and Välimäki (1998)), we considered a variation of the current model where buyers in two markets differ in terms of the volume of their purchases. For simplicity, the first market segment consists of a continuum of identical small buyers and the second is formed by a single large buyer. Buyers in both segments have the same flow valuation for the products and the sellers can price discriminate between the segments. In contrast to the small buyers, the large buyer controls a sizeable part of the current information flow and thus internalizes the impact of her current purchases on her future utilities. The willingness to pay for the new good is then different across the market segments. As a result, there is an intermediate range of the values of α where the new firm sells to the small buyers in equilibrium while the established firm caters to the large buyer.

When interpreting the basic model as a durable goods model, an interesting question arises. Since the buyers have a strict preference to purchase the product at the moment of their arrival rather than wait for a period, it might be useful to analyse a model where the entry of new buyers is endogenous. For example, one could imagine a model with a constant population of buyers that own a physically depreciating durable good. The possibility of early purchases at an endogenously determined cost might have important implications on the shape of the price process and the speed of information transmission.

APPENDIX

The appendix contains the proofs of all propositions and theorems presented in the main body of the paper.

Proof of Theorem 1.

We describe the solution procedure for the inhomogeneous second-order differential equation in some detail as it reappears in the construction of the various equilibrium value functions. All solutions of the inhomogeneous equation (7) permit the following representation

$$V(\alpha) = c_1 H_1(\alpha) + c_2 H_2(\alpha) + \psi(\alpha),$$

where $H_1(\alpha)$ and $H_2(\alpha)$ are two linearly independent solutions of the corresponding homogeneous equation and $\psi(\alpha)$ is a particular solution of the inhomogeneous equation. The complete solution to (7) is established by the variation of parameters method, see Chapter 2 in Birkhoff and Rota (1978). A particular solution to the inhomogeneous differential equation is given by

$$\psi(\alpha) = u_1(\alpha)H_1(\alpha) + u_2(\alpha)H_2(\alpha),$$

where the parameters $u_1(\alpha)$ and $u_2(\alpha)$ are determined by

$$u_1'(\alpha) = -\frac{G(\alpha)H_2(\alpha)}{W(H_1(\alpha), H_2(\alpha))},$$

and

$$u_2'(\alpha) = \frac{G(\alpha)H_1(\alpha)}{W(H_1(\alpha), H_2(\alpha))},$$

and $G(\alpha)$ is the forcing term of the inhomogeneous differential equation. The Wronskian determinant, $W(H_1(\alpha), H_2(\alpha))$, is given by

$$W(H_1(\alpha), H_2(\alpha)) = H_1(\alpha)H_2'(\alpha) - H_1'(\alpha)H_2(\alpha).$$

The solution to the homogeneous version of the differential equation (7) is

$$H(\alpha) = b_1 \alpha^{(\gamma+1)/2}(1-\alpha)^{(\gamma-1)/2} + b_2\alpha^{-(\gamma-1)/2}(1-\alpha)^{(\gamma+1)/2}, \tag{27}$$

with $\gamma = \sqrt{1+\rho/N}$. The general solution for the value function $V(\alpha)$ is

$$V(\alpha) = \frac{N\mu(\alpha)}{r} + b_1\alpha^{(\gamma+1)/2}(1-\alpha)^{-(\gamma-1)/2} + b_2\alpha^{-(\gamma-1)/2}(1-\alpha)^{(\gamma+1)/2}.$$

The efficient switching point $\hat{\alpha}$ is obtained by requiring that the value matching and the smooth pasting conditions are satisfied:[12]

$$\begin{aligned} rV(\hat{\alpha}) &= Ns, \\ V'(\hat{\alpha}) &= 0. \end{aligned} \tag{28}$$

Boundedness of the value function as $\alpha \to 1$ implies $b_1 = 0$. The conditions in (28) determine $\hat{\alpha}$ and b_2

$$\hat{\alpha} = \frac{(s-\mu_L)(\gamma-1)}{2(\mu_H - s) + (\mu_H - \mu_L)(\gamma - 1)},$$

$$b_2 = \frac{2N(\mu_H - s)}{(\gamma-1)r}\left(\frac{\hat{\alpha}}{1-\hat{\alpha}}\right)^{(\gamma+1)/2}.$$

12. See Shiryayev (1978), Chapter 3, for a formal statement of this principle.

In particular $b_2 > 0$ implies the convexity of the value function. Further computation shows that the efficient cut-off $\hat{\alpha}$ is strictly increasing in ρ/N and $V(\cdot)$ is strictly decreasing in ρ/N. ||

To characterize the equilibrium stopping point, we need the solutions to the differential equations describing the value functions, which is obtained after inserting the equilibrium prices (15)–(16) in the value functions.

Lemma 1. (Value functions). *The value functions in the experimentation region display the general form*

$$V_0(\alpha) = \frac{N}{r}(\mu(\alpha) - s) - cN\alpha^{-(\lambda-1)/2}(1-\alpha)^{(\lambda+1)/2} + d\alpha^{-(\gamma-1)/2}(1-\alpha)^{(\gamma+1)/2},$$

$$V_1(\alpha) = b\alpha^{-(\gamma-1)/2}(1-\alpha)^{(\gamma+1)/2},$$

and

$$U(\alpha) = \frac{s}{r} + c\alpha^{-(\lambda-1)/2}(1-\alpha)^{(\lambda+1)/2} - b\alpha^{-(\gamma-1)/2}(1-\alpha)^{(\gamma+1)/2},$$

with $\lambda = \sqrt{1 + \rho/(N-1)}$, *and* γ *as in Theorem* 1.

Proof. The construction of the value functions proceeds from the homogenous part of the differential equation (19) of the incumbent to the inhomogeneous differential equation of the individual buyer (11) and the entrant (18). The solution to (19) is as in (27), where the second term vanishes as $V_1(\alpha) \to 0$ as $\alpha \to 1$

$$V_1(\alpha) = b\alpha^{(\gamma-1)/2}(1-\alpha)^{(\gamma-1)/2}.$$

The general solution to the homogeneous version $U_h(\alpha)$

$$rU_h(\alpha) = \tfrac{1}{2}(N-1)\Sigma(\alpha)U_h''(\alpha),$$

of the differential equation $U(\alpha)$ is

$$U_h(\alpha) = c_1\alpha^{(\lambda+1)/2}(1-\alpha)^{-(\lambda-1)/2} + c_2\alpha^{-(\lambda-1)/2}(1-\alpha)^{(\lambda+1)/2}.$$

A particular solution of the equation (11) is

$$\psi(\alpha) = \frac{s}{r} - b\alpha^{(\gamma-1)/2}(1-\alpha)^{(\gamma+1)/2}.$$

The complete solution to the inhomogeneous equation (11) is

$$U(\alpha) = \frac{s}{r} + c\alpha^{-(\lambda-1)/2}(1-\alpha)^{(\lambda+1)/2} - b\alpha^{(\gamma-1)/2}(1-\alpha)^{(\gamma+1)/2}.$$

As before we set $c_1 = 0$, since the value function of the buyer has to be bounded as α goes to one. For simplicity set $c_2 = c$. The final step is to construct $V_0(\alpha)$. A particular solution to the inhomogeneous equation is

$$\psi(\alpha) = \frac{N}{r}(\mu(\alpha) - s) - cN\alpha^{-(\lambda-1)/2}(1-\alpha)^{(\lambda+1)/2}. \tag{29}$$

The fundamental solutions together with the particular solution $\psi(\alpha)$ gives us the format of all possible solutions of the value function $V_0(\alpha)$. Boundedness of the value function forces one term of the homogenous solution to vanish

$$V_0(\alpha) = \frac{N}{r}(\mu(\alpha) - s) + d\alpha^{-(\gamma-1)/2}(1-\alpha)^{(\gamma+1)/2} - cN\alpha^{-(\lambda-1)/2}(1-\alpha)^{(\lambda+1)/2},$$

which completes the construction of the value functions. ||

Proof of Theorem 2.

The existence of an equilibrium is established by construction. Given the guess on the shape of the continuation and stopping regions, we construct the value functions. Following the derivation of the value functions, it is then verified that the initial guess is satisfied.

The equilibrium conditions for the stopping point α^* is the smooth-pasting condition of the optimal stopping problem for the new seller and the continuity of the value functions for the sellers and the buyers at the stopping point α^*

$$\begin{aligned} V_0(\alpha^*) &= 0, \\ V_0'(\alpha^*) &= 0, \\ rV_1(\alpha^*) &= N(s-\mu(\alpha^*)), \\ rU(\alpha^*) &= \mu(\alpha^*). \end{aligned} \tag{30}$$

The conditions in (30) yield the stopping point α^* and the values of the parameters b, c, and d, which determine the curvature of the value functions

$$\alpha^* = \frac{(s-\mu_L)(\gamma N - \lambda(N-1)-1)}{(\mu_H-\mu_L)(\gamma N-\lambda(N-1)-1)+2(\mu_H-s)} < \hat{\alpha},$$

and

$$\begin{aligned} b &= \frac{2N}{(N\gamma-(N-1)\lambda-1)}\frac{(\mu_H-s)}{r}\left(\frac{\alpha^*}{1-\alpha^*}\right)^{(\gamma+1)/2}, \\ c &= \frac{2(N-1)}{(N\gamma-(N-1)\lambda-1)}\frac{(\mu_H-s)}{r}\left(\frac{\alpha^*}{1-\alpha^*}\right)^{(\gamma+1)/2}, \\ d &= \frac{2N^2}{(N\gamma-(N-1)\lambda-1)}\frac{(\mu_H-s)}{r}\left(\frac{\alpha^*}{1-\alpha^*}\right)^{(\gamma+1)/2}. \end{aligned} \tag{31}$$

To show uniqueness, it is sufficient to prove that no other shape of the stopping region is possible. To see this, let $C\subset[0,1]$ denote the continuation region and $S\subset[0,1]$ the stopping region of an arbitrary equilibrium. By cautiousness and Markovian strategies, it is immediate that sales are made at all $\alpha\in[0,1]$, *i.e.* $S\cup C=[0,1]$.

We need to show that there is an $\alpha^*\in[0,1]$ such that $S=[0,\alpha^*]$, and $C=(\alpha^*,1]$. First note that $p_1(\alpha)\geqq 0$ for all α and therefore, $\alpha\in S\Rightarrow\alpha\leqq\hat{\alpha}$. By cautiousness, $p_0(\alpha)=0$ for all $\alpha\in S$. By the continuity of the value functions in α we may take S to be closed. Therefore C is a union of pairwise disjoint (relatively) open intervals. We need to show that if $(\alpha_1,\alpha_2)\subset C$, then $\alpha_2>\hat{\alpha}$. But this follows immediately from cautiousness as $p_0(\alpha_1)=p_0(\alpha_2)=0$ and there must be an $\alpha\in(\alpha_1,\alpha_2)$ with $p_0(\alpha)\geqq 0$. This yields the desired conclusion. ||

Proof of Corollary 2.

The equilibrium value functions are as in Theorem 2, with the exception of the normalization as transparent in the stopping point $\alpha^*(N)$ through the parameters $\tilde{\gamma}$ and $\tilde{\lambda}$. The parameters b, c, d are as in (31) only to be divided through N and inserting $\tilde{\gamma}$ and $\tilde{\lambda}$. The comparative statics follow after straightforward algebra. ||

Proof of Theorem 3.

(1) The value functions are derived as in Lemma 1 with the obvious modifications due to the normalization of the size of the market. The solutions to the differential equations are given by

$$V_0(\alpha)=\frac{1}{r}(\mu(\alpha)-s)+d\alpha^{-(\gamma-1)/2}(1-\alpha)^{(\gamma+1)/2}-cN\alpha^{-(\lambda-1)/2}(1-\alpha)^{(\lambda+1)/2},$$

$$V_j(\alpha)=b\alpha^{-(\gamma-1)/2}(1-\alpha)^{(\gamma+1)/2},$$

and

$$U(\alpha)=\frac{s}{rN}+c\alpha^{-(\lambda-1)/2}(1-\alpha)^{(\lambda+1)/2}-b\alpha^{-(\gamma-1)/2}(1-\alpha)^{(\gamma+1)/2},$$

with $\gamma=\sqrt{1+\rho}$ and $\lambda=\sqrt{1+\rho N/(N-1)}$. The equilibrium stopping conditions are modified only by the normalization in the market size to

$$\begin{aligned} V_0(\alpha^*) &= 0, \\ V_0'(\alpha^*) &= 0, \\ rMV_j(\alpha^*) &= (s-\mu(\alpha^*)), \\ rNU(\alpha^*) &= \mu(\alpha^*). \end{aligned} \tag{32}$$

The stopping point is given as a solution to the equations in (32)

$$\alpha^* = \frac{(N\gamma - \lambda(N-M) - M)(s-\mu_L)}{(N\gamma - \lambda(N-M) - M)(\mu_H - \mu_L) + 2M(\mu_H - s)},$$

and the parameters of the value functions are given by

$$b = \frac{2(\mu_H - s)}{(N\gamma - \lambda(N-M) - M)r}\left(\frac{\alpha^*}{1+\alpha^*}\right)^{(\gamma+1)/2},$$

$$c = \frac{2(N-M)(\mu_H - s)}{(N\gamma - \lambda(N-M) - M)Nr}\left(\frac{\alpha^*}{1-\alpha^*}\right)^{(\lambda+1)/2},$$

$$d = \frac{2N(\mu_H - s)}{(N\gamma - \lambda(N-M) - M)r}\left(\frac{\alpha^*}{1-\alpha^*}\right)^{(\gamma+1)/2}.$$

Uniqueness is proved as in Theorem 2. (2) and (3) follow directly. ||

Acknowledgements. The authors express their gratitude to Drew Fudenberg, Bengt Holmström, Al Klevorick, George Mailath, Eric Maskin, Georg Nöldeke, Andy Postlewaite, Mike Riordan, Larry Samuelson, Karl Schlag and Avner Shaked for many helpful discussions. Detailed comments by two anonymous referees and the editor, Patrick Bolton, greatly improved the paper. We benefited from seminar participants at Bonn, Harvard, Mannheim, MIT, Northwestern and Paris. The first author would like to thank Avner Shaked for his hospitality during a stay at the SFB 303 at Bonn University. The authors acknowledge support from NSF Grant SBR 9709887 and 9709340 respectively.

REFERENCES

BANERJEE, A. and FUDENBERG, D. (1997), "A Simple Model of Word-of-Mouth Learning" (Mimeo).

BERGEMANN, D. and VÄLIMÄKI, J. (1996), "Learning and Strategic Pricing", *Econometrica*, **64**, 1125–1149.

BERGEMANN, D. and VÄLIMÄKI, J. (1997), "Market Diffusion with Two-Sided Learning", *Rand Journal of Economics*, **28**, 773–795.

BERGEMANN, D. and VÄLIMÄKI, J. (1998), "Experimentation in Markets" (Northwestern University CMSEMS DP. 1220).

BIRKHOFF, G. and ROTA, G. C. (1978) *Ordinary Differential Equations*. 3rd ed. (New York: Wiley).

BOLTON, P. and HARRIS, C. (1999), "Strategic Experimentation", *Econometrica*, **67**, 349–374.

CHAMLEY, C. and GALE, D. (1994), "Information Revelation and Strategic Delay in a Model of Investment", *Econometrica*, **62**, 1065–1086.

DIXIT, A. and PINDYCK, R. (1994) *Investment under Uncertainty* (Princeton: Princeton University Press).

FELLI, L. and HARRIS, C. (1996), "Job Matching, Learning and Firm-Specific Human Capital", *Journal of Political Economy*, **104**, 838–868.

HARRISON, J. (1985) *Brownian Motion and Stochastic Flow Systems* (Malabar: Krieger).

HENDRICKS, K. and KOVENOCK, D. (1989), "Asymmetric Information, Information Externalities, and Efficiency: the Case of Oil Exploration", *Rand Journal of Economics*, **20**, 164–182.

KARATZAS, I. (1984), "Gittins Indices in the Dynamic Allocation Problem for Diffusion Processes", *Annals of Probability*, **12**, 173–192.

KELLER, G. and RADY, S. (1999), "Optimal Experimentation in a Changing Environment", *Review of Economic Studies*, **66**, 475–509.

LIPTSER, R. S. and SHIRYAYEV, A. N. (1977) *Statistics of Random Process I* (New York: Springer-Verlag).

McFADDEN, D. and TRAIN, K. (1996), "Consumers' Evaluation of New Products: Learning from Self and Others", *Journal of Political Economy*, **104**, 683–703.

ROB, R. (1991), "Learning and Capacity Expansion under Demand Uncertainty", *Review of Economic Studies*, **58**, 655–675.

ROTHSCHILD, M. (1974), "A Two-Armed Bandit Theory of Market Pricing", *Journal of Economic Theory*, **9**, 185–202.

SHIRYAYEV, A. N. (1978) *Optimal Stopping Rules* (Berlin: Springer-Verlag).

VETTAS, N. (1998), "Demand and Supply in New Markets: Diffusion with Bilateral Learning", *Rand Journal of Economics*, **29**, 215–233.