DYNAMIC PRICING OF NEW EXPERIENCE GOODS

BY

DIRK BERGEMANN AND JUUSO VÄLIMÄKI

Dynamic Pricing of New Experience Goods

Dirk Bergemann

Yale University

Juuso Välimäki

Helsinki School of Economics and University of Southampton

We develop a dynamic model of experience goods pricing with independent private valuations. We show that the optimal paths of sales and prices can be described in terms of a simple dichotomy. In a mass market, prices are declining over time. In a niche market, the optimal prices are initially low followed by higher prices that extract surplus from the buyers with a high willingness to pay. We consider extensions of the model to integrate elements of social rather than private learning and turnover among buyers.

I. Introduction

In this paper, we develop a simple and tractable model of optimal pricing for a monopolist that sells a new experience good over time to a population of heterogeneous buyers. Our main result shows that it is possible

to classify all markets to be either *mass markets* or *niche markets* with qualitatively different equilibrium patterns for prices and quantities. These two different markets lead to different pricing strategies that have been called skimming and penetration pricing, respectively, in the earlier literature.

To classify the market as a mass or a niche market, it is sufficient to consider the intertemporal incentives of a new buyer who is uncertain about her tastes for the product. For this calculation, we *assume* that the market price is at its static monopoly level. The buyer compares (potential) current losses to future informational benefits from the purchase. In a mass market, she is willing to purchase at the static monopoly price, whereas in a niche market she refuses such trades. We show that in a mass market, the dynamic equilibrium prices decrease over time and uninformed buyers purchase in all periods. In a niche market, uninformed buyers do not buy at the static monopoly price. As a result, the monopolist must offer low initial prices to capture a larger share of the uninformed at the expense of targeting the more attractive informed segment of the market. This is in contrast to the mass market, where the monopolist skims in the early stages the more attractive part of the market (i.e., the uninformed buyers). In both cases, prices converge in the long run to static monopoly price levels. Yet in a mass market, the sales also converge to the static monopoly level, but in a niche market, quantities remain below that level as some uninformed buyers drop out of the market.

The model in this paper is an infinite-horizon, continuous-time model of monopoly pricing. There is a continuum of ex ante identical consumers who have a unit demand per period for the purely perishable good. At each instant of time, the monopolist offers a spot price and the buyers decide whether to purchase or not. In the beginning, as the new product is introduced, all the buyers are uncertain about their valuation. By consuming the product, they learn their true valuation for the product in a stochastic manner. For analytical convenience, we assume that a perfectly revealing signal arrives according to a Poisson process to the active buyers in the market. We also assume that the aggregate distribution of preferences in the population is common knowledge.

In an experience goods market, the seller is facing two different submarkets simultaneously. The demand curve in the part of the population that has already learned its preferences is similar to the standard textbook case. Those buyers who are uncertain about the true quality of the product behave in a more sophisticated manner. Each purchase incorporates an element of information acquisition that is relevant for future decisions. The value of this information is endogenously determined in the market. If future prices are high, purchases are unlikely

to yield information that results in future consumer surpluses. If future prices are low, it may be in the buyers' best interest to forgo purchases in the current period since future prices are attractive regardless of the true value of the product. As a result, current and future prices determine simultaneously the sales in the informed segment of the market and the value of information in the uninformed segment.

The market for new pharmaceuticals illustrates the key features of our model. In the pharmaceutical industry, each new drug undergoes an extensive period of prelaunch testing to determine its performance in the overall population. The aggregate uncertainty relating to the product is therefore negligible at the moment of introduction. Nevertheless, many drugs differ in their effectiveness and the incidence of side effects across patients. This idiosyncratic uncertainty provides a motive for the individual agent to experiment. A recent empirical study by Crawford and Shum (2005) regarding the dynamic demand behavior in pharmaceutical markets documents the important role of idiosyncratic uncertainty and learning in explaining demand.[1] For a data set of anti-ulcer descriptions, they observe substantial uncertainty about the idiosyncratic effectiveness of the individual drugs and high precision in the signals received through consumption experience. We model the effectiveness of the new treatment to an individual new patient as a random event. The time of response to the drug is random, and the response may be either positive or negative (successful recovery from the illness or severe side effects).

To our knowledge, the current paper is the first to address the issue of experimental consumption in a fully dynamic model with a population of heterogeneous buyers. We provide a tractable analytical framework and demonstrate the flexibility of our framework by outlining extensions of the basic model in Sections V and VI. We should mention that our model allows an alternative interpretation, which we spell out in Section VI. We can rephrase the random arrival of information as a random arrival of a consumption opportunity. With this specification, we assume that the buyers learn their true preferences upon their first purchase. In addition to being commonly used in other papers on experience goods, this assumption allows us to apply the model to a wider range of markets with experience goods such as the market for many professional services (e.g., hospitality, law, and Internet services).

Besides the earlier cited work on learning in markets for new pharmaceutical products, there is a growing literature on Bayesian learning

[1] In closely related work on learning in pharmaceutical markets, Coscelli and Shum (2004) estimate the impact of uncertainty and learning for the introduction of a new drug, and Ching (2002) provides structural, dynamic demand estimates when there is learning among patients about a new (generic) pharmaceutical with common values.

in consumer markets with experience goods.[2] Of particular interest for the current model is the recent work by Israel (2005), who uses the random arrival of information as an identification strategy in the context of the automobile insurance market. In his empirical analysis, automobile insurance is an experience good that is provided continuously to the buyer, punctuated by distinct "learning events," namely, the occurrence of an accident. If the insuree files a claim following the accident, he effectively has a random opportunity to learn about the quality of the service, the automobile insurance. The pattern of consumer departures from the insurer before and after such events therefore reveals the impact of learning.

The estimation of Bayesian models of experience goods is inherently difficult since it requires a dynamic model *with* substantial heterogeneity to estimate the Bayesian learning process by idiosyncratic agents. A common limitation in the above empirical contributions is their focus on the optimal behavior of the buyers paired with an exogenous pricing policy. A contribution of our paper is to present a parsimonious model of joint equilibrium behavior with forward-looking buyers and seller. The model has a number of features that will facilitate the empirical analysis of experience good markets. In our model, the equilibrium converges to the static optimal price. The true underlying demand can therefore be estimated as in standard static discrete-choice models (see Berry 1994). This in turn plausibly allows the identification and estimation of the learning rate λ, given the discount rate r, or the identification of the ratio λ/r from the dynamics of the price path. The result here may provide an important step to extend the estimation from models of Bayesian learning to models of Bayesian equilibrium behavior. The simple characterization in terms of niche and mass markets makes an empirical verification of the fit of the model feasible.

The empirical work on dynamic pricing is surprisingly sparse, even in the absence of the experience good aspect.[3] An important exception is the paper by Lu and Comanor (1998), which investigates the deter-

---

[2] A partial list of recent papers includes Erdem and Keane (1996), Ackerberg (2003), Erdem, Imai, and Keane (2003), Israel (2005), Osborne (2005), and Goettler and Clay (2006). The empirical evidence that emerges from these papers is that the initial information and the rate of learning differ substantially across markets. Ackerberg (2003) and Crawford and Shum (2005) find a fast learning rate for yogurt and pharmaceuticals, respectively. Erdem and Keane (1996) find little evidence of experimental consumption for laundry detergents. In the same market, Osborne (2005) finds evidence of initially similar expectations but different individual experiences in a model that allows for heterogeneous buyers. In Goettler and Clay (2006) and Israel (2005), learning is slow for an online grocery and automobile insurance, respectively.

[3] Recently Gowrisankaran and Rysman (2005) and Nair (2005) made substantial progress toward estimating dynamic price policies. They suggest structural models to estimate the dynamic pricing policies of firms in the context of durable good markets. They provide estimates for console video games and digital video disc players and digital cameras.

minants and the intertemporal evolution of prices for a large panel of
new pharmaceutical products. The authors find that pharmaceuticals
that represent a large therapeutic improvement (as measured by Food
and Drug Administration ranking and surveys) over previous products
tend to follow a skimming strategy. They start with a high price and
then lower it over time. On the other hand, pharmaceuticals with a
small improvement over previous products tend to follow a penetration
strategy by first offering low prices and eventually switching to higher
prices. These two different price patterns correspond to our distinction
between mass and niche markets. A product with a substantial improve-
ment provides value to the average consumer, and hence we expect the
mass market condition to hold. On the other hand, a product with a
small improvement naturally constitutes a niche market and hence fol-
lows a penetration strategy. The panel in Lu and Comanor's study is
split between products designed for acute and chronic ailments. They
show that products for chronic ailments are more likely to follow a
skimming strategy. This empirical finding is consistent with our theo-
retical prediction since more frequent purchases, all else equal, lead to
a faster arrival of information and hence a higher λ. Alternatively, more
frequent consumption acts imply that the time elapsed between uses is
smaller and hence the effective discount rate r is smaller, again a factor
that contributes positively to a mass market.

Related literature.—We now briefly discuss the relationship between
prior theoretical contributions and the current model. An important
early paper on the dynamic pricing of experience goods is Shapiro
(1983). The author investigates the optimal price policy of a monopolist
in a two-period model when each consumer learns the true value of the
product through experience. The analysis and the results differ sub-
stantially from the current contribution since Shapiro imposes two spe-
cific assumptions: (i) each consumer acts myopically, and (ii) the ex-
pectation of the quality is biased with respect to the true expected
quality. The myopic consumer fails to evaluate the option value of the
experiment, and with the natural benchmark of expectationally unbi-
ased buyers, his model predicts constant prices.

Cremer (1984) considers a model with initially identical buyers and
idiosyncratic experience to explain the use of coupons for repeat buyers
in a two-period setting. In contrast to our analysis, he considers the
optimal commitment price path and does not analyze the time-consis-
tent pricing policy.

In subsequent work by Farrell (1986), Milgrom and Roberts (1986),
and Tirole (1988), the dynamic selling policy for experience goods has
been investigated without biased expectations. A common feature of
these three models is that the buyers differ ex ante in their willingness
to pay for quality and the true quality has either a zero or a positive

value. When one of the possible quality values is zero, the identity of
the marginal buyer is constant across the two periods. In consequence,
the marginal buyer does not receive a positive surplus in either period
and never associates an option value to early purchases. As a result, the
marginal buyer acts myopically, and his decision does not reflect any
intertemporal trade-off. Further, the equilibrium price path is always
increasing, independent of the shape of the demand or the discount
factor. In contrast, our model allows for the possibility that the monop-
olist discriminates intertemporally in the market in a more flexible man-
ner, and as a result, our conclusions are quite different from those in
the earlier literature.

Villas-Boas (2004) considers the equilibrium pricing of experience
goods in a *duopoly model* with differentiation along a location and a taste
dimension. The location is known at the outset, whereas tastes are
learned through experience. The analysis is concerned mostly with
brand loyalty, that is, whether buyers return to the seller they bought
from in the past. It presents a sufficient condition on the skewness of
the distribution under which brand loyalty exists in equilibrium.

In a recent and complementary contribution, Johnson and Myatt
(2006) use the distinction between niche and mass markets to investigate
under which conditions a (mean-preserving) spread of the valuation
increases the revenue in the *static* optimal pricing problem. They in-
terpret "advertising" and "marketing" as providing information that in-
creases the variance of the valuations by the customers. In contrast, in
our model the shape of the distribution is influenced by the pricing
policy itself rather than being determined through separate instruments.

The paper is organized as follows. Section II sets up the basic model
and discusses the appropriate solution concepts. Section III presents
the problem of optimal demand management for the seller. Section IV
analyzes the properties of the optimal price path. Section V discusses
the social efficient allocation and the role of idiosyncratic versus social
learning. Section VI provides extensions of the model, including inflow
of new buyers, commitment to price paths, and random purchasing
opportunities for the buyer. Section VII presents conclusions. The proofs
of all the results are collected in the Appendix.

II.   Model

We consider a continuous-time model with $t \in [0, \infty)$ and a positive
discount rate $r > 0$. A monopolist with a zero marginal cost of production
offers a single product for sale in a market consisting of a continuum
of consumers. For analytical simplicity, we assume that the buyers have
unit demand for the product within periods and that the product is not
storable. We also abstract from the possibility of price differentiation

within periods. At each instant, the monopolist offers a spot price. Upon seeing the price, each consumer decides whether to purchase or not.

Every consumer is characterized by his idiosyncratic willingness to pay for the product, denoted by θ. The good is an experience good, and the true value of θ is initially unknown to the buyer and the seller. The ex ante distribution of valuations is given by a continuously differentiable distribution function $F(\theta)$ with support $[\theta_l, \theta_h] \subset \mathbb{R}$. This distribution is assumed to be common knowledge and reflects our assumption that there is no aggregate uncertainty in the model. As the focus in this paper is on private individual experiences, we abstract from possible common sources of uncertainty. To simplify the analysis, we also require that $\theta[1 - F(\theta)]$ be strictly quasi-concave in θ. This assumption guarantees that the full-information profit maximization problem is well behaved.

All buyers are ex ante identical, and their expected utility from consuming the product prior to learning their type is given by v:

$$v \triangleq \int_{\theta_l}^{\theta_h} \theta dF(\theta).$$

Throughout the paper, we assume that a perfectly informative signal (e.g., the emergence of side effects in a drug therapy) arrives at a constant Poisson rate λ for all buyers who purchase the product in a time interval of length dt.[4] In this case, the posterior distribution on θ remains constant at the prior until the signal is observed.[5] The most important analytical consequence of this assumption is that if the buyers have not observed a signal, they remain identical. After observing the signal, the buyers are heterogeneous, and the monopolist's key objective is to manage the endogenous composition of these two market segments.

As we analyze the dynamic behavior of the market, it is natural to use dynamic programming tools to derive the equilibrium conditions for the model. We assume that the only publicly observable variables are the prices and aggregate quantities. This is in line with the assumption that each individual buyer is small and has no strategic impact on the aggregate outcomes. The state variable of the model at time t is the fraction of informed buyers at t, denoted by $\alpha(t) \in [0, 1]$. Even though $\alpha(t)$ is not directly observable to the players, it can be calculated from

---

[4] It might be natural to allow for cases in which λ depends on t. In the pharmaceutical example, such a time-varying arrival rate might reflect, e.g., the decline in the probability of a treatment being eventually successful given a number of unsuccessful trials. We have analyzed this possibility, but given that the qualitative features of the model remain the same, we report only the constant case.

[5] This assumption is made for ease of exposition only. We have computed the model for posteriors with positive and negative drift. The qualitative features of the equilibrium remain as in the case of a constant prior.

the equilibrium purchasing strategies. If the uninformed buy in period t, the state variable $\alpha(t)$ evolves according to

$$\frac{d\alpha(t)}{dt} = \lambda[1 - \alpha(t)],$$

since in period t there are $1 - \alpha(t)$ currently uninformed buyers and a fraction λdt of them become informed in a time interval of length dt.

A Markovian pricing strategy for the seller is denoted by $p(\alpha)$. The uninformed buyer has a Markovian purchasing strategy $d^u(\alpha, p)$ that depends on the state variable α as well as the current price p. Similarly, the informed buyer with valuation θ adopts a Markovian purchasing strategy $d^\theta(\alpha, p)$.

The monopolist maximizes her expected discounted profit over the horizon of the game. The buyers maximize the expected discounted value of their utilities from consumption net of price. As there is no aggregate uncertainty, the price and aggregate sales processes are deterministic. The individual buyer, however, faces uncertainty regarding his true valuation and the random time at which he will receive the information.

III.   Demand Management

The basic issue in the introduction of a new product is the dynamic demand management. In the early stages, the majority of the buyers are inexperienced and uninformed. Over time, the segment of informed buyers grows. As the relative sizes of these two market segments change, the seller adapts her policy and shifts her attention to the more important segment. More precisely, the type of the marginal buyer whose willingness to pay determines the equilibrium price changes over time. With a new product, the marginal buyer is inevitably uninformed in the early stages. As the informed segment grows, the marginal buyer is more likely to come from that segment. Optimal demand management then determines the switch between these two market segments. The dynamic pricing policy of the seller therefore contains at its core an optimal stopping problem. The stopping point identifies the time at which the marginal buyer ceases to be uninformed.

After stopping, the marginal buyer is informed. What needs to be determined, however, is whether the uninformed buyers keep on purchasing the product. Either the uninformed buyers are priced out of the market or they stay in the market as inframarginal buyers. Whether the uninformed buyer will eventually stay in or drop out of the market is essentially a question of the size of the market in equilibrium.

The demand management problem of the seller is more subtle than

a pure optimal stopping problem. In the canonical optimal stopping problem, the alternative payoffs do not depend on future policy choices. In the optimal pricing problem here, buyers are forward looking, and their willingness to pay today depends on future prices. The current revenue of the seller therefore depends on his future prices.

A.   Market Size

As the number of informed buyers increases in the market, the optimal price is determined by the distribution of the valuations, $F(\theta)$. When the seller ignores the uninformed buyers, the optimal monopoly price $\hat{p}$ maximizes the flow revenues from the informed buyers:

$$\hat{p} = \arg\max_{p \in \mathbb{R}_+} \{p[1 - F(p)]\}.$$

Of course, the price $\hat{p}$ is also the optimal price in the static monopoly problem in which each buyer knows his valuation for the object. The corresponding equilibrium quantity is denoted by $\hat{q} = 1 - F(\hat{p})$.

The key comparison for the analysis is between the willingness to pay of the uninformed buyers and the static equilibrium price $\hat{p}$. In a static setting the expected value v of the product is the *average* willingness to pay in the market. It follows that if the optimal price $\hat{p}$ is below the average willingness to pay, then the uninformed buyers stay in the market and eventually become informed.

In the intertemporal setting, the willingness to pay of the uninformed actually exceeds v in most cases. In addition to the expected flow value from consumption, the uninformed buyer also has a chance to learn more about his true valuation for the product. If the future price stays constant and equal to $\hat{p}$, then his willingness to pay is the value of a purchase today, or

$$\hat{w} = v + \frac{\lambda}{r} \mathbb{E}_\theta \max\{\theta - \hat{p}, 0\}. \tag{1}$$

The uninformed buyer becomes informed at rate λ. When informed, she purchases the product if and only if $\theta - \hat{p} \geq 0$. Finally, the future benefits are discounted at rate r. The value of a purchase today is then simply the sum of the expected value of the flow consumption, v, and the expected value of information, $(\lambda/r)\mathbb{E}_\theta \max\{\theta - \hat{p}, 0\}$.

The monopoly price in the informed segment $\hat{p}$ and the expected value of information both depend on the distribution $F(\theta)$. We now distinguish between a *niche market* and a *mass market* by comparing the willingness to pay of the uninformed buyers, $\hat{w}$, with the optimal static price $\hat{p}$.

DEFINITION 1 (Niche market and mass market).

1.   The market is said to be a niche market if $\hat{w} < \hat{p}$.
2.   The market is said to be a mass market if $\hat{w} \geq \hat{p}$.

In a mass market, the price $\hat{p}$ is so low that new buyers are willing to enter the market. The monopoly price $\hat{p}$ is independent of λ and r, and hence the mass market condition is more likely to occur if the rate of information arrival λ is large or the discount rate r is small.

We can gain further insight into the notions of niche and mass markets by a comparative static analysis. For a given r and λ, let us consider a family of distributions $F(\theta; \sigma^2)$, parameterized by variance σ^2 with a fixed mean. In this environment, the market is more likely to be a mass market if the variance is small and more likely to be a niche market if the variance is large. For a small variance, the seller can increase the sales by lowering his price just below the mean. In contrast, if the variance is large, the seller prefers to sell at a price above the mean to the upper tail of the market. The size of this particular segment is sufficiently large whenever the variance is large. This intuition is exact within the classes of binary, uniform, and normal distribution (with constant mean). In other words, for any such family of distributions, there exists a critical value $\bar{\sigma}^2$ such that for all $\sigma^2 < \bar{\sigma}^2$ the market is a mass market and for all $\sigma^2 > \bar{\sigma}^2$ the market is a niche market.[6]

B.   Optimal Switching

We first describe the intertemporal decision problems in terms of the familiar dynamic programming equations in Markov strategies. The solution of the optimal stopping problem will then follow from the optimality conditions. The size of the segment of informed buyers, $\alpha(t) \in [0, 1]$, in period t is the state variable of the model. We omit the indexation with respect to time and simply write all value functions as a function of α rather than $\alpha(t)$. We start with the simple decision problem of the informed buyers. These buyers have complete information about their true valuation θ of the object. For a given price policy $p(\alpha)$ by the monopolist, we can determine the value function

---

[6] We would like to thank the editor, Robert Shimer, who suggested the variance analysis. For arbitrary distributions, the comparison between the willingness to pay $\hat{w}$ and the optimal price $\hat{p}$ is more difficult since $\hat{w}$ and $\hat{p}$ are determined by different properties of the distribution function: $\hat{p}$ is determined by the first-order condition of the static revenue function and hence by the hazard rate at $\hat{p}$, whereas $\hat{w}$ is (partially) determined by the expectation over all valuations exceeding $\hat{p}$. This conditional expectation of the upper tail of the distribution can change in either direction with an increase in variance or related measures, such as dispersion or second-order stochastic dominance. Within the above parameterized classes of distributions, the behavior of the tail expectation is in balance with the change in the local hazard rate condition.

V^θ of the informed buyer from the Bellman equation

$$rV^\theta(\alpha) = \max\{\theta - p(\alpha),\ 0\} + \frac{dV^\theta}{d\alpha}\frac{d\alpha}{dt}. \tag{2}$$

The decision whether to buy or not to buy is solved by the myopic decision rule: buy whenever θ exceeds the current price $p(\alpha)$. The only intertemporal component in this equation (the second term) reflects the effect of a change in the composition of the market segments, represented by α, on the future utilities. Future utilities are affected by changes in α as future prices respond to changes in aggregate demand. These changes are beyond the control of any single (informed) buyer, and hence the myopic rule characterizes optimal behavior.

For the uninformed buyers, a purchase of the new product represents a bundle, consisting of the flow of consumption *and* information. Their value function $V^u(\alpha)$ is given by

$$rV^u(\alpha) = \max\{v - p(\alpha) + \lambda[\mathbb{E}_\theta V^\theta(\alpha) - V^u(\alpha)],\ 0\} + \frac{dV^u}{d\alpha}\frac{d\alpha}{dt}. \tag{3}$$

The main difference between these two value functions reflects the value of information to the uninformed buyers. A purchase in the current period generates an inflow of information at rate λ. Conditional on receiving the signal, the uninformed becomes informed. In consequence the new value function becomes $V^\theta(\alpha)$ for some θ. From the point of a currently uninformed buyer, there is uncertainty about his true valuation θ. He estimates the expected gain from the information by taking the expectation with respect to θ. The informational gain attached to a current purchase is given by

$$\lambda[\mathbb{E}_\theta V^\theta(\alpha) - V^u(\alpha)].$$

The value function of the seller is denoted by $V(\alpha)$. We describe the seller's dynamic programming equation in two parts to separate the intertemporal considerations as cleanly as possible. The basic trade-off facing the firm is that sales are made at a single price in two separate market segments. If the firm decides to sell to the uninformed buyers as well as some informed ones, the relevant equation is given by

$$rV(\alpha) = \max_{p(\alpha) \in \mathbb{R}_+} \{p(\alpha)\{1 - \alpha + \alpha[1 - F(p(\alpha))]\}\} + \frac{dV}{d\alpha}\frac{d\alpha}{dt} \tag{4}$$

subject to

$$p(\alpha) \leq v + \lambda\mathbb{E}_\theta[V^\theta(\alpha) - V^u(\alpha)].$$

Here $1 - \alpha$ is the share of uninformed buyers in the population and $\alpha[1 - F(p(\alpha))]$ is the fraction of informed buyers who are willing to buy

at prices $p(\alpha)$. The constraint on the price $p(\alpha)$ guarantees that the uninformed buyers are indeed willing to purchase at prices $p(\alpha)$.

If the monopolist sells to the informed segment only, then her value function satisfies

$$rV(\alpha) = \max_{p(\alpha) \in \mathbb{R}_+} \{p(\alpha)\alpha[1 - F(p(\alpha))]\}. \qquad (5)$$

In this latter case, the size of the informed segment, α, remains constant and $d\alpha/dt = 0$, since the flow of information to the uninformed buyers has stopped. The Markovian prices in this regime must hence remain constant in all future periods. With these preliminaries, we can state the following definition.

DEFINITION 2 (Markov-perfect equilibrium).   A Markov-perfect equilibrium of the dynamic game is a triple (d^u, d^θ, p) such that the problems (2)–(5) are simultaneously solved for all α and θ.

We now employ the dichotomy between a niche market and a mass market to find the optimal launch strategy as the solution to a specific stopping problem. We denote the size of the informed market segment at the stopping point by $\hat{\alpha}$.

C.   Niche Market

In the niche market the willingness to pay of the uninformed buyers is below the static optimal price: $\hat{w} < \hat{p}$. It follows that if the seller sets prices optimally in the informed segment, then the uninformed stop buying. In consequence, the seller has to decide how long she wishes to serve the uninformed market segment.

We now describe the marginal conditions that characterize the stopping point $\hat{\alpha}$. After $\hat{\alpha}$ is reached, the optimal dynamic price equals $\hat{p}$. At the stopping point, the uninformed buyers purchase the new product for the last time. Their willingness to pay at the stopping point is therefore exactly equal to $\hat{w}$. At the stopping point, the seller must be indifferent between charging $\hat{p}$ and $\hat{w}$:

$$\hat{\alpha}[1 - F(\hat{p})]\hat{p} = \{(1 - \hat{\alpha}) + \hat{\alpha}[1 - F(\hat{w})]\}\hat{w} + \frac{\lambda(1 - \hat{\alpha})}{r}[1 - F(\hat{p})]\hat{p}. \qquad (6)$$

The indifference condition compares the revenue from $\hat{p}$ relative to revenue from $\hat{w}$. If the seller were to offer $\hat{p}$, then only those informed buyers who have a true valuation $\theta \geq \hat{p}$ purchase the product, leading to a sales volume of $\hat{\alpha}[1 - F(\hat{p})]$. On the other hand, if the seller were to offer $\hat{w}$, then all uninformed buyers would stay in the market and all informed buyers with $\theta \geq \hat{w}$ would also buy the object, leading to a larger sales volume of $(1 - \hat{\alpha}) + \hat{\alpha}[1 - F(\hat{w})]$. At price $\hat{w}$, $\lambda(1 - \hat{\alpha})$ currently uninformed customers become informed, and hence they will add to

the revenue from the informed customers for all future periods. If we denote by $\pi(p, \alpha)$ the flow profit to the monopolist from price p when α is the fraction of informed buyers, then the above equation can be written as

$$\pi(\hat{p}, \hat{\alpha}) - \pi(\hat{w}, \hat{\alpha}) = \frac{\lambda(1 - \hat{\alpha})[1 - F(\hat{p})]\hat{p}}{r}. \tag{7}$$

The left-hand side represents the differential gains from extracting surplus from the informed agents, and the right-hand side represents the benefits from building up future demand. The latter is the long-term gain from an additional inflow of $\lambda(1 - \alpha)$ informed buyers of whom $1 - F(\hat{p})$ are willing to purchase at price $\hat{p}$. As the right-hand side is positive, we conclude that with niche markets, the monopolist sacrifices current profits to build up future demands.

PROPOSITION 1 (Equilibrium stopping in the niche market).   If $\hat{w} < \hat{p}$, then

1. $\hat{\alpha} < 1$ and

2. $\hat{\alpha}$ is increasing in λ and decreasing in r.

In the calculation of $\hat{\alpha}$, the buyers' optimality conditions are reflected only through $\hat{w}$. As a result, it is quite straightforward to extend the model to allow for different discount factors for the buyers and the seller. The discount rate of the buyer determines $\hat{w}$ through (1) and the seller's discount rate determines the long-run gains from additional goodwill customers. In Sections V and VI, we formulate models of social learning and random purchasing opportunities in which the separability of the problems with respect to different discount rates is useful.

D.   Mass Market

Initially, the informed segment does not yet exist. The monopolist then offers high prices, which leave the uninformed agents just indifferent between buying and not buying. The monopolist can thus extract initially all the surplus from the current purchases of the uninformed agents. As the informed segment grows, any price that leaves the uninformed indifferent results in revenue losses in the informed segment relative to pricing at $\hat{p}$. The monopolist's problem is therefore to determine the stopping point at which he starts to leave surplus to the uninformed buyers.

In contrast to the niche market, the uninformed buyers continue to purchase in the mass market. After the stopping point, they become

inframarginal rather than marginal buyers. As a result, the optimal stopping condition can (almost) exclusively be described in terms of the flow revenue for the seller. Until the stopping point $\hat{\alpha}$, uninformed buyers are marginal. In consequence, the equilibrium price makes the uninformed buyer just indifferent between buying and not buying. We can express this in terms of the equilibrium value function of the buyer, using (3):

$$p(\alpha) = v + \lambda[\mathbb{E}_\theta V^\theta(\alpha) - V^u(\alpha)]. \tag{8}$$

The flow revenue of the seller at a given price p and a fraction α of informed buyers is

$$\pi(p, \alpha) = (1 - \alpha)p + \alpha[1 - F(p)]p \quad \text{for } p \leq p(\alpha). \tag{9}$$

The seller sets prices to make the uninformed customers marginal as long as the marginal revenue from increasing the price at $p(\alpha)$ is nonnegative, or

$$\frac{\partial \pi(p(\alpha), \alpha)}{\partial p} \geq 0. \tag{10}$$

As long as the uninformed buyer is the marginal buyer, the marginal flow revenue $\partial \pi / \partial p$ at $p = p(\alpha)$ can well be strictly positive. The reason is that the true payoff function has a discontinuity at $p = p(\alpha)$ reflecting the positive mass of uninformed buyers that drop out of the market at prices above $p(\alpha)$. The optimal stopping point $\hat{\alpha}$ is hence derived from

$$\frac{\partial \pi(p(\hat{\alpha}), \hat{\alpha})}{\partial p} = 0. \tag{11}$$

In a mass market, the stopping condition can therefore be expressed in terms of the flow revenues. Still, the problem contains an intertemporal element since the equilibrium price $p(\alpha)$ before and at the stopping point is based on the equilibrium continuation values; see equation (8). After the stopping point $\hat{\alpha}$, the equilibrium price is computed as the price that maximizes the flow revenue, or

$$\frac{\partial \pi(p, \alpha)}{\partial p} = 0. \tag{12}$$

Even though this equation has the flavor of a static optimization condition, the dynamics of the model still enter into the determination of prices through the evolution of α.

PROPOSITION 2 (Equilibrium stopping in the mass market).   If $\hat{w} \geq \hat{p}$, then

1. $\hat{\alpha} \leq 1$ and

2. $\hat{\alpha}$ is decreasing in λ and increasing in r.

If we consider the comparative statics results in propositions 1 and 2, then it is worth observing that the respective stopping points for niche and mass markets move in opposite directions as a function of λ and r. Suppose that we consider a distribution $F(\theta)$ with the property that it may be either a niche or a mass market depending on the values of λ and r and consider a comparative statics exercise in λ. For very low values of λ, the willingness to pay by the uninformed is low, the market is a niche market, learning stops early, and few buyers become informed. As λ increases, the uninformed buyers are willing to pay more. In turn the seller will offer introductory prices for a longer period. Eventually λ will reach a point at which the willingness to pay of the uninformed exactly equals the static optimal price, or $\hat{w} = \hat{p}$. At this knife-edge case the optimal dynamic price will in fact be constant for all t and $\hat{\alpha} = 1$. At this point, the market turns from a niche market into a mass market. The willingness to pay of the uninformed is now high enough for them to stay in the market until they become informed. In consequence, the uninformed buyer becomes an inframarginal buyer. The stopping point $\hat{\alpha}$ now decreases in λ as the uninformed customers become inframarginal earlier. The comparative statics are simply reversed for r. In figure 1, the increasing part of the graph corresponds to the niche market case and the decreasing part belongs to the mass market.

IV.   Equilibrium Pricing

We are now in a position to characterize the complete equilibrium pricing policy on the basis of the equilibrium stopping point. Before the stopping point $\hat{\alpha}$, the marginal buyer is the uninformed buyer. We showed earlier that the price before stopping is

$$p(\alpha) = v + \lambda \mathbb{E}[V^\theta(\alpha) - V^u(\alpha)].$$

We first discuss the pricing policy in the early market and then describe the equilibrium conditions for the mature market.

PROPOSITION 3 (Early market).

1.   The price $p(\alpha)$ satisfies for all $\alpha \leq \hat{\alpha}$

$$\frac{dp}{d\alpha}\frac{d\alpha}{dt} = r[p(\alpha) - v] - \lambda \mathbb{E}_\theta \max\{\theta - p(\alpha), 0\}. \tag{13}$$

2.   The price $p(\alpha)$ is decreasing in α for all $\alpha < \hat{\alpha}$.
3.   The equilibrium sales $q(\alpha)$ are decreasing for small α and convex in α.

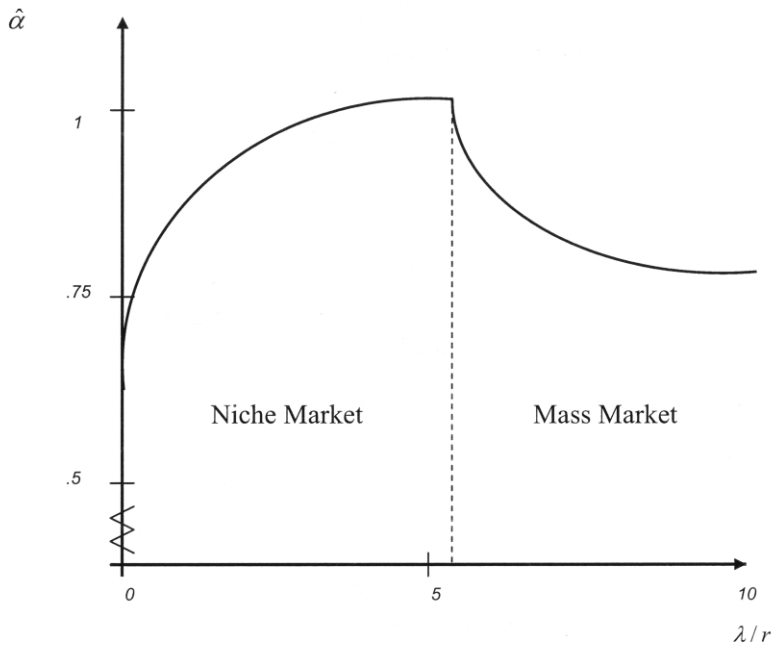The differential equation (13) describes the evolution of the price in

FIG. 1.—Stopping point for varying λ/r

the early market. When we rearrange the equation, it becomes apparent that it represents the trade-off that the uninformed buyer is facing in his purchase decision:

$$\lambda \mathbb{E}_\theta \max \{\theta - p(\alpha), \, 0\} = r[p(\alpha) - v] - \frac{dp}{d\alpha}\frac{d\alpha}{dt}.$$

The left-hand side represents the net benefit of buying the new product today rather than tomorrow. In particular, a purchase today generates an informative signal at rate λ and allows the buyer to make an informed decision. The right-hand side represents the net benefits of buying tomorrow rather than today. The net benefit has two components. First, buying tomorrow allows the buyer to postpone the net cost of a purchase, $p(\alpha) - v$, by an instant; and second, it may change the price the buyer has to pay for the information.

In the early market, the price contains a premium for the option value generated by the information. As time goes by, this option value decreases, and in consequence the price will have to decline to offset the lower option value. The rate at which the price decreases depends directly on the variation of valuations in the market. In particular, the

price decreases at a faster rate if there is more riskiness in the distri-
bution in the sense of second-order stochastic dominance. For a given
p and v, the rate at which the price decreases is determined by the
second term, namely $-\lambda \mathbb{E}_\theta \max\{\theta - p(\alpha), 0\}$, which is a concave function
in θ. Consequently, its expected value decreases if the riskiness in the
sense of second-order stochastic dominance increases. Thus we expect
to see a more rapid price decline in niche markets relative to mass
markets. In niche markets the seller initially pursues a more aggressive
pricing strategy to build up a clientele before he eventually switches to
extract the surplus from his target market, namely the niche of high-
valuation buyers.[7]

PROPOSITION 4 (Mature market).

1. In the niche market, the price $p(\alpha)$ jumps up and stays at
   $p(\hat{\alpha}) = \hat{p}$ at the stopping point $\hat{\alpha}$.
2. In the mass market, $p(\alpha)$ is decreasing for all $\alpha \geq \hat{\alpha}$ and
   $\lim_{\alpha \to 1} p(\alpha) = \hat{p}$.

The value of information before the stopping point, as expressed by
$p(\alpha) - v$, is decreasing over time. While the evolution of the price is
governed by the same differential equation for the niche and the mass
markets, the source of the decrease in the value of information is dif-
ferent in the niche and the mass markets. In the niche market, the
stopping point is the end of a phase of introductory prices. After this,
the seller increases her price to $\hat{p}$ and a large fraction of surplus is
extracted from the informed buyers. Hence the initially high value of
information results from the relatively low initial prices.

In the mass market, the seller stops extracting all the surplus from
the uninformed buyers and lowers her price to attract more informed
buyers with lower valuations for the object. The value of information is
now decreasing because with lower future prices, the option value that
arises from the possibility of rejecting the product when θ is low is
smaller.

Our interpretation of the two qualitatively different price paths goes
as follows. In the niche market, the monopolist makes introductory
offers to increase the number of goodwill customers once the price is
raised. In the mass market, the monopolist skims the high-valuation
buyers in the market (the uninformed buyers) with a high and declining
price. It should be noted that in the niche as well as in the mass market,
the prices do not change by large amounts before $\hat{\alpha}$, and as a result,
adjustment costs to changing prices might well force the monopolist to
adopt a two-price regime with low initial prices followed by higher prices

---

[7] We would like to thank the editor and a referee for inquiring about the relationship
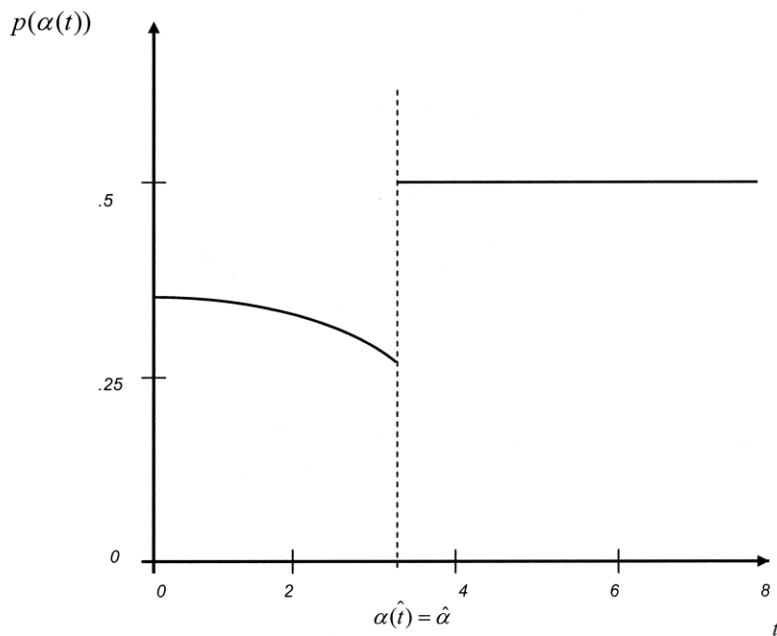between the rate of price change and the variance in the valuations.

FIG. 2.—Equilibrium price for the niche market

in the niche market and high prices followed by low prices in the mass market.

The intertemporal pricing policies are graphically depicted in figures 2 and 3 for the niche and mass markets, respectively. With the niche market, the introductory price slowly decreases until it reaches a value equal to the willingness to pay, and at that point, the seller ceases to pursue new customers and sells only to informed customers with sufficiently high valuations. In the mass market, the discount factor r is small, and hence the option value for the uninformed buyer is almost constant. In consequence, the price declines very slowly until the seller begins to seek sales more aggressively from the informed customers. At this point, the price begins to decrease more rapidly and eventually converges to the static monopoly price.

Notice also that our model provides a theoretical prediction for the joint movements of prices and equilibrium quantities. These effects should be taken into account when estimating the demand for new products. If one were to estimate a static demand function for a product using data that include observations of prices for $\alpha < \hat{\alpha}$, then it is immediate that the resulting estimators would be biased.

The equilibrium analysis in this paper focused on the notion of a

$p(\alpha(t))$

.75

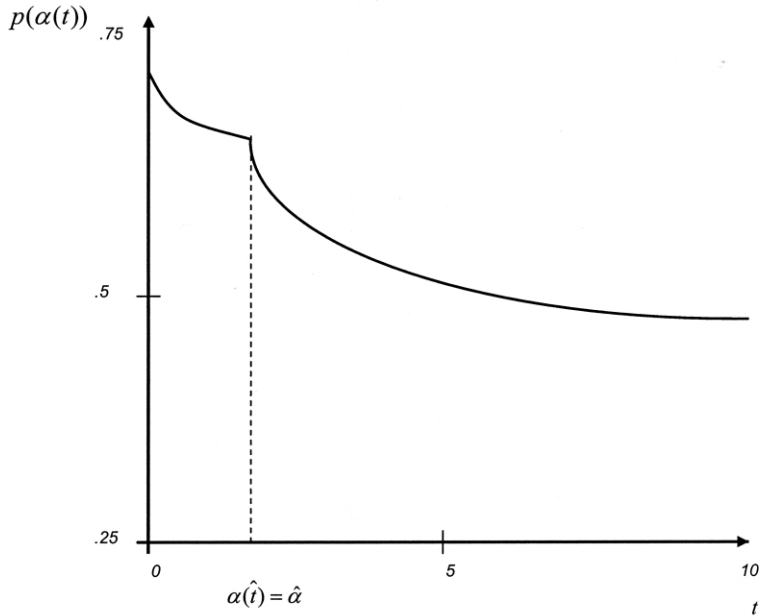.5

.25

0

5

10

$\alpha(\hat{t}) = \hat{\alpha}$

t

Fig. 3.—Equilibrium price for the mass market

Markov-perfect equilibrium and derived the unique equilibrium in this class. The uniqueness result extends to a much larger class of equilibria. In Bergemann and Välimäki (2004), we show that the Markov-perfect equilibrium remains the only sequential equilibrium outcome as long as the information sets of the players include only their own past actions, observations, and the aggregate market data. With this much weaker restriction, the continuation paths of play are still independent of the choices of any individual buyer, but they may depend on past prices in an arbitrary manner.

V.   Idiosyncratic versus Social Learning

We now contrast the equilibrium allocation with the socially efficient allocation. In the presence of idiosyncratic learning, the socially efficient outcome can be implemented in a competitive equilibrium with marginal cost pricing. We then augment the analysis of idiosyncratic learning with an element of social learning. The social learning naturally introduces informational externalities among the buyers, and we show that the earlier welfare ranking between competitive and monopolistic market structure may be reversed in the presence of social learning.

A.   Social Efficiency

The socially efficient policy maximizes the sum of the expected discounted value across all agents. As learning and the resolution of uncertainty are purely idiosyncratic, the socially optimal policy can be determined as the solution for a representative consumer. The buyers and the seller have quasi-linear preferences, and hence the socially efficient policy simply maximizes expected social surplus. The optimal consumption policy can be determined for informed and uninformed consumers separately. For a given constant marginal cost c of producing the object, the (social) value function W^θ of the informed customer is

$$rW^\theta = \max\{\theta - c, 0\},$$

and for the uninformed consumer it is given by W^u, or

$$rW^u = \max\{v + \lambda\mathbb{E}_\theta[W^\theta - W^u] - c, 0\}. \qquad (14)$$

For the informed consumer, the socially efficient decision is simply to consume if and only if his net value $\theta - c$ is positive. For the uninformed consumer, he should buy if the *current* social net benefit, $v - c$, and the *future* social net benefit, $\lambda\mathbb{E}[W^\theta - W^u]$, exceed zero. The expected value of the informed consumer, on the other hand, is

$$\mathbb{E}[W^\theta] = \frac{1}{r}\mathbb{E}_\theta\max\{\theta - c, 0\}. \qquad (15)$$

By inserting the value function for the informed consumer, or (15), into (14), we obtain the critical value v so that the social value of consuming an additional unit of the new good is larger than or equal to zero:

$$v + \frac{\lambda}{r}\mathbb{E}_\theta\max\{\theta - c, 0\} \geq c.$$

The socially optimal policy is therefore to consume the new good as long as the expected net value today, $v - c$, and the value of information from current consumption are positive. It is worth emphasizing that it may be efficient to try the new product even if the expected net value of the object, $v - c$, or even the expected gross value v, is negative.

The monopoly position of the seller introduces static as well as intertemporal distortions away from the socially optimal level of consumption. The static distortionary element comes from the standard revenue considerations of the seller. Her objective is to maximize the revenue rather than the social welfare. In consequence, she will typically set the price above marginal cost and hence fail to offer the product to some buyers who would have a positive contribution to the social surplus. The dynamic element of the distortion arises in the niche mar-

ket as well as in the mass market, but with different consequences. In the niche market, the willingness to pay of the uninformed is low, and hence the seller will eventually stop selling to the uninformed in pursuit of higher per unit revenue from the informed buyers. In the mass market, efficient learning takes place. But as the uninformed agents are willing to pay a premium over and above their expected value for the object, the seller maintains a high price even relative to static monopoly price $\hat{p}$. This implies that during the launch phase, many informed buyers will not purchase the object even though their purchase would generate a positive social surplus.

In contrast to this, a competitive market would support the efficient allocation in this model of idiosyncratic learning. In a competitive market, the object is offered at the marginal cost c. In consequence, the objective function of the buyer in a competitive market coincides with the social objective for both the informed and the uninformed consumers. The important ingredient of the model that leads to an agreement of the competitive and the efficient outcomes is the idiosyncratic nature of the learning experience. We next discuss a minimal extension of our model to introduce an element of social learning. We shall see that the qualitative insights of the equilibrium analysis will essentially carry over but that the welfare comparison between competitive and monopoly markets will now be different.

B. Social Learning

We introduce social learning into our model by extending every single buyer to a group of k buyers who consume individually but share the same true valuation for the object and share the information about the valuation among the members of the group. We may think of the group as a family, a department, or a neighborhood with similar preferences and the ability to communicate among each other. At each point in time, we randomly select one buyer out of the group of k members to consider a purchase. If that member decides to purchase the object, then he receives the consumption value of the object privately but shares all resulting information from his purchase with the other members of his group.

The equilibrium conditions of the idiosyncratic learning model are easily adapted to accommodate social learning. The value function $V^{\theta}(\alpha)$ of an informed agent is given by

$$rV^{\theta}(\alpha) = \frac{1}{k}\max\{\theta - p(\alpha),\ 0\} + \frac{dV^{\theta}}{d\alpha}\frac{d\alpha}{dt}. \tag{16}$$

The only change from the earlier formulation is that the frequency of

purchases is reduced by factor $1/k$. Similarly, the value function of an uninformed buyer is given by

$$rV^u(\alpha) = \frac{1}{k}\max\{v - p(\alpha) + \lambda[\mathbb{E}_\theta V^\theta(\alpha) - V^u(\alpha)],\ 0\}$$

$$+ \frac{k-1}{k}\lambda[\mathbb{E}_\theta V^\theta(\alpha) - V^u(\alpha)] + \frac{dV^u}{d\alpha}\frac{d\alpha}{dt}. \qquad (17)$$

The element of social learning enters the dynamic programming equation of the individual buyer since he now has two sources of information: either he can directly purchase the product as an occasion arises (at the rate $1/k$) or he can benefit from the information generated by the other members of the group (at the rate $k - [1/k]$).

The social learning has two effects on the pricing policy. First, the uninformed are less willing to pay a premium for the information, since there is a chance to learn from the experiences of others. On the other hand, the seller is willing to sponsor the experiment of each agent since she knows that the information will spread among all members of the group. If they value the object sufficiently highly, the seller can extract their surplus in the future. For an individual buyer, the frequency $1/k$ by which an opportunity for a purchase arrives acts as a (stochastic) increase in the discount rate. From the point of view of the seller, the actual market size in each period remains unchanged. In consequence, the stopping conditions for $\hat{\alpha}$ are determined by using a modified discount rate, namely $r \cdot k$, for the buyer. In consequence, the basic distinctions in terms of launch strategies remain unchanged, and the equilibrium policies of the seller are exactly as if he would face buyers with a larger discount rate.

With this informational externality the comparison between the competitive market outcome and the monopoly outcome becomes quite different. The monopolist partially internalizes the informational externality since he is aware that the release of information today leads to a higher revenue in future periods. In a competitive setting, the product is offered at marginal cost in each period. As each buyer takes into account only the private benefits from learning, the social optimum may not be reached in the competitive market. In particular, if the expected value of the current consumption v falls below marginal cost c, then the only reason to purchase the product is to acquire more information. If the informational externality between the buyers is increased by increasing k, then the uninformed will not purchase at the marginal cost pricing. The socially optimal policy is independent of k, and the size of the market for the monopolist is independent of k as well. It follows that the monopolist may now sustain the market and

induce socially beneficial learning even though the competitive market cannot do so.

PROPOSITION 5 (Monopoly versus competitive market).   If $v < c$ and it is socially efficient to adopt the new product, then

1. there exists $\bar{k} > 0$ such that, for all $k \geq \bar{k}$, the new product is never sold in the competitive market;
2. if $v - c + (\lambda/r)\hat{p}[1 - F(\hat{p})] > 0$, then it is optimal for the monopolist to launch the product for all k.

VI.   Extensions

Random purchases.—We first describe an alternative interpretation for the informational structure of our basic model. We assume as before that there is a continuum of buyers who are initially uncertain about their tastes for the new product. In contrast to the basic model, we now assume that the consumption opportunities arrive at random time intervals. For analytical convenience, we assume that these arrivals follow a Poisson process with parameter λ.

In this reformulation of the model, it makes sense to assume that the buyers learn their true tastes upon consuming the first unit of the good. Even though this assumption is less realistic in the context of the pharmaceuticals market, it may fit better consumer goods such as cereals, cosmetics, and so forth. In the initial periods, the monopolist is facing demand mostly from uninformed buyers. The fraction of repeat buyers increases as the good stays in the market, and eventually most of the buyers are repeat buyers. The only change in the dynamic programming formulations is the random purchase rate, which acts like an increase in the discount rate. The problem of the informed buyers is now

$$(r + \lambda)V^{\theta}(\alpha) \; = \; \lambda \max\{\theta - p(\alpha),\, 0\} + \frac{dV^{\theta}(\alpha)}{d\alpha}\frac{d\alpha}{dt}.$$

Correspondingly, the value function of the uninformed buyers is given by

$$(r + \lambda)V^{u}(\alpha) \; = \; \lambda \max\{v - p(\alpha) + \mathbb{E}_{\theta}V^{\theta}(\alpha) - V^{u}(\alpha),\, 0\} + \frac{dV^{u}(\alpha)}{d\alpha}\frac{d\alpha}{dt}.$$

These two value functions correspond to the original model in which the buyers' discount rate is set to $(r + \lambda)/\lambda$ and the arrival rate of information is set to one. As was explained in Section III, the analysis of the original model can be carried out with different interest rates for the buyers and the seller, and hence this alternative interpretation is included as a special case of the original model.

Commitment.—New distortionary effects arise from the intertemporal trade-offs faced by a time-consistent monopolist. In order to understand these dynamic effects, it is interesting to ask whether they could be overcome if the seller had the ability to commit himself to an entire future price path at the beginning of the game.

The analytical techniques of the solution with commitment are quite different from the equilibrium solution, since we must abandon the requirement of sequential rationality or time consistency in characterizing the optimal commitment paths. For this reason, we simply report the basic results, which are formally stated in an earlier working paper version (Bergemann and Välimäki 2004).

The optimal commitment solution is qualitatively different from the time-consistent equilibrium solution in the niche market. With the ability to commit, the seller can promise future prices below $\hat{p}$ in exchange for higher prices today. This allows her to maintain future prices that are low enough so that *all* uninformed customers will stay in the market until they are informed. The optimal price path is then a constant price $\bar{w}$, which can be obtained as the solution from the willingness to pay of the uninformed (see eq. [1]):

$$\bar{w} = v + \frac{\lambda}{r}\mathbb{E}_\theta \max\{\theta - \bar{w},\, 0\}.$$

In consequence, the uninformed customers receive zero expected surplus in the commitment solution, but their purchasing policy is now socially efficient. It follows that commitment increases social welfare in the niche market. The consumers do not benefit since the ability to commit enables the seller simply to extract more surplus over time.

In the mass market, the ability to commit has negative consequences on social welfare. In the time-consistent policy, the seller eventually cuts her price to sell more to the informed buyers. With the ability to commit, the seller can commit to higher prices in the future. As such higher prices increase the value of information for the uninformed, prices can be higher in the initial periods as well. Thus in a mass market, commitment allows the seller to set higher prices everywhere and hence further depress socially efficient purchases.

Stationary model.—The current model describes the optimal pricing for the introduction of a new product. The buyers started with ex ante identical information regarding the new product, and over time their personal experiences lead them to have heterogeneous and idiosyncratic valuations. It is then natural to expand the scope of our analysis to a market with a constant inflow and outflow of consumers. In fact, many consumer products face a constant renewal in their customer base, because of either the aging of the customers or other systematic changes

to the agents' preferences. In Bergemann and Välimäki (2004), we analyze the steady-state equilibrium in a market with idiosyncratic learning. We model the change in the population by constant entry rate γ of new customers and equal exit rate γ of old customers. The new customers are all initially uninformed and become informed according to the same information technology as in the current paper.

The steady-state equilibrium with a constant renewal of customers displays the same basic features as the current model of the launch of a new good to an entirely new market. The willingness to pay of the new customers is modified in the obvious way by the birth and death rate γ, or

$$\hat{w} = v + \frac{\lambda}{r + \gamma} \mathbb{E}_\theta \max\{\theta - \hat{p}, \, 0\}.$$

With this modification, the optimal policy can again be described in terms of mass market versus niche market. If, as before, the complete information price $\hat{p}$ is below $\hat{w}$, then the market is a mass market, and the seller will offer a price in between the statically optimal price $\hat{p}$ and the willingness to pay of the new customers, $\hat{w}$, so as to balance his revenue management. Importantly, all new customers will enter the market and learn more about the new product.

If, on the other hand, the willingness to pay is below the statically optimal price, or $\hat{p} > \hat{w}$, then we are again in the situation of a niche market. Now, the seller does not want to sell to the new customers all the time, but rather extract surplus from the informed agents. Yet, with a constant inflow of new customers, the seller cannot abandon the new customers altogether since this would imply a diminishing customer base. The resolution of the trade-off is that the seller offers probabilistically both low and high prices. The high price is given by $\hat{p}$ and optimally extracts surplus from the informed agents, and the low price allows the new customers to learn and become acquainted with the new product. In this way, the seller balances revenue objectives yet maintains a constant informed clientele for his product. The dispersed prices in the niche market are common in models of durable goods with entry of new buyers (see, e.g., Sobel 1991). In our setting, the product is effectively a bundle, consisting of a perishable element in the form of the immediate consumption benefit, but also a durable element in the form of the information obtained with the purchase. From this perspective, the relationship to the durable good pricing problem appears inherently.

VII.   Conclusion

In this paper, we have shown that the optimal sales policy of a monopolist in a model of experience goods is qualitatively different depending on whether the market is a mass market or a niche market. In a niche market, it is in the monopolist's best interest to build up a sufficient base of goodwill clientele. To achieve this, the monopolist sacrifices current profit by pricing low in order to find new future buyers. The durability of information about the product quality thus plays a key role in this situation. In a mass market, managing information is less important since the uninformed buyers are willing to buy at the static optimal prices. In this case the monopolist's optimal price path can be seen as an attempt to skim the uninformed segment of the market until the informed segment becomes large.

We have kept the model as simple as possible in order to highlight the dynamics of price setting. The modeling strategy of the current paper could be used to investigate models with more general specifications for either the buyers' valuations of the product or the strategic environment. An interesting instance of a more general demand structure would be one in which the buyers differ in their willingness to pay for quality, but the perceived quality is idiosyncratic and must be learned over time. Regarding the competitive structure of the model, the natural next step would be to consider the role of idiosyncratic learning in a strategic environment against either a known or a similarly unknown product. An interesting variation of the current model and of special importance for the pharmaceutical market would be to consider optimal pricing in which a competitor will appear only in T periods hence, induced by the expiration of the patent (see Berndt, Kyle, and Ling [2003] for an empirical analysis of this situation).

Appendix

This appendix collects the proofs of all the results in the main body of the paper. For notational convenience, we shall adopt a standard notation from probability theory by writing

$$(\theta - p)^{+} \triangleq \max\{\theta - p, \, 0\}.$$

Proof of Proposition 1

Part 1. The optimal stopping point is the solution to (6) and is given by

$$\hat{\alpha} = \frac{\hat{w} + (\lambda/r)\hat{p}[1 - F(\hat{p})]}{\hat{w}F(\hat{w}) + \hat{p}[1 - F(\hat{p})][1 + (\lambda/r)]}. \tag{A1}$$

For simplicity we define $\eta \triangleq \lambda/r$, and with this we rewrite equation (A1) as

$$\hat{\alpha} = \frac{\hat{w} + \eta\hat{p}[1 - F(\hat{p})]}{\hat{w} + (1 + \eta)\hat{p}[1 - F(\hat{p})] - \hat{w}[1 - F(\hat{w})]}, \tag{A2}$$

which shows that $\hat{\alpha} < 1$ since $\hat{p}[1 - F(\hat{p})] - \hat{w}[1 - F(\hat{w})] > 0$.

Part 2. After replacing $\hat{w}$ by its explicit expression given in (1), we rearrange (A2) to get

$$\hat{\alpha}\{[v + \eta\mathbb{E}(\theta - \hat{p})^+]F(v + \eta\mathbb{E}(\theta - \hat{p})^+) + \hat{p}[1 - F(\hat{p})](1 + \eta)\} =$$
$$v + \eta\mathbb{E}(\theta - \hat{p})^+ + \eta\hat{p}[1 - F(\hat{p})]. \tag{A3}$$

Differentiating the equality (A3) implicitly with respect to η yields

$$\frac{d\hat{\alpha}}{d\eta}\{\hat{w}F(\hat{w}) + \hat{p}[1 - F(\hat{p})](1 + \eta)\} + \hat{\alpha}\{[\hat{w}f(\hat{w}) + F(\hat{w})]\mathbb{E}(\theta - \hat{p})^+ + \hat{p}[1 - F(\hat{p})]\} =$$
$$\mathbb{E}(\theta - \hat{p})^+ + \hat{p}[1 - F(\hat{p})],$$

and hence

$$\frac{d\hat{\alpha}}{d\eta} = \frac{\mathbb{E}(\theta - \hat{p})^+[1 - \hat{\alpha}F(\hat{w}) - \hat{\alpha}\hat{w}f(\hat{w})] + (1 - \hat{\alpha})\hat{p}[1 - F(\hat{p})]}{\hat{w}F(\hat{w}) + \hat{p}[1 - F(\hat{p})](1 + \eta)}.$$

The denominator is clearly positive. For the numerator, we observe that $1 - \hat{\alpha}F(\hat{w}) - \hat{\alpha}\hat{w}f(\hat{w})$ is the derivative of the profit function $p\{1 - \hat{\alpha} + \hat{\alpha}[1 - F(p)]\}$ evaluated at $\hat{w}$, which is positive by the assumed quasi concavity of $p[1 - F(p)]$ together with the fact that $\hat{w} < \hat{p}$. Therefore, the numerator is also positive, as needed. QED

Proof of Proposition 2

Part 1. Suppose that, for all $\alpha \geq \hat{\alpha}$, the marginal buyer is an informed buyer. The equilibrium price $p(\alpha)$ is then given as the solution of the static revenue maximization problem (9). The value function of an informed buyer at $\alpha(t) = \alpha$ is

$$V^\theta(\alpha) = \int_t^\infty e^{-r(\tau - t)}[\theta - p(\alpha(\tau))]^+ d\tau.$$

The expected value of the informed buyer at $\alpha(t) = \alpha$ is

$$\mathbb{E}_\theta[V^\theta(\alpha)] = \int_{\theta_l}^{\theta_h}\left\{\int_t^\infty e^{-r(\tau - t)}[\theta - p(\alpha(\tau))]^+ d\tau\right\} dF(\theta).$$

In contrast, the value function of the uninformed buyer for a particular realization $T \geq t$ of the signal arrival time is given by

$$\int_t^T e^{-r(\tau - t)}[v - p(\alpha(\tau))]d\tau + \int_{\theta_l}^{\theta_h}\left\{\int_T^\infty e^{-r(\tau - t)}[\theta - p(\alpha(\tau))]^+ d\tau\right\} dF(\theta). \tag{A4}$$

The value function of the uninformed buyer is obtained from (A4) by taking

the expectation over all signal arrival times $T \geq t$, or $V^u(\alpha)$ is given by

$$\int_t^\infty \int_{\theta_l}^{\theta_h} \left\{ \int_t^T e^{-r(\tau-t)}[v - p(\alpha(\tau))]d\tau + \int_T^\infty e^{-r(\tau-t)}[\theta - p(\alpha(\tau))]^+ d\tau \right\} dF(\theta) \lambda e^{-\lambda T} dT.$$

The willingness to pay for all $\alpha \geq \hat{\alpha}$ is given by

$$w(\alpha) = v + \lambda[\mathbb{E}_\theta V^\theta(\alpha) - V^u(\alpha)].$$

From the above expressions, the difference in the value functions, $\mathbb{E}_\theta V^\theta(\alpha) - V^u(\alpha)$, can therefore be written as

$$\mathbb{E}_\theta V^\theta(\alpha) - V^u(\alpha) = \int_t^\infty \int_{\theta_l}^{\theta_h} \int_t^T e^{-r(\tau-t)}[p(\alpha(\tau)) - \theta]^+ d\tau dF(\theta) \lambda e^{-\lambda T} dT. \qquad (A5)$$

The gain of the informed vis-à-vis the uninformed buyer arises in all those instances in which the uninformed buyer accepts the offer by the seller even though his true valuation is below the equilibrium price. It follows that $w(\alpha)$ is decreasing in α (and t) since $p(\alpha)$ is decreasing in α (and t). We can then run $w(\alpha)$ backward as long as $w(\alpha) \geq p(\alpha)$.

The stopping point $\hat{\alpha}$ is the smallest α at which

$$w(\alpha) = p(\alpha). \qquad (A6)$$

We next argue that there is a unique stopping point $\hat{\alpha}$ by showing that $w(\alpha)$ and $p(\alpha)$ are single crossing. By the hypothesis of $\hat{w} > \hat{p}$, we have

$$w(1) > p(1). \qquad (A7)$$

Observe further that

$$\lim_{\alpha \downarrow 0} p(\alpha) = +\infty.$$

The maximal willingness to pay, $w(\alpha)$, is a constant v and a discounted average over future prices $p(\alpha)$, represented by (A5). It therefore follows that, provided that $p(\alpha)$ is monotone,

$$|p'(\alpha)| > w'(\alpha),$$

which together with (A7) is sufficient to guarantee existence and uniqueness of the stopping point.

Part 2. We observe first that $p(\alpha)$ is independent of r. It follows that r affects the expression $\mathbb{E}_\theta V^\theta(\alpha) - V^u(\alpha)$ only through discounting. As an increase in r decreases $w(\alpha)$, it follows that the intersection (A6) is reached later and thus at a higher value of $\hat{\alpha}$. The argument for λ is similar except for the obvious reverse in the sign. QED

Proof of Proposition 3

Part 1. The differential equation (13) for the full extraction prices has a unique rest point, $\dot{p}(t) = 0$, at $p(t) = w$ since w uniquely solves

$$0 = r(w - v) - \lambda \mathbb{E}_\theta \max\{\theta - w, 0\}.$$

Part 2. We show the monotonicity of $p(\alpha)$ separately for $\hat{p} < \hat{w}$ and $\hat{p} \geq \hat{w}$. We start with the latter case and argue by contradiction. Thus suppose that

$p(0) > w$. Then $p(\alpha) > w > \hat{w}$ for all α. It follows that at $\alpha = \hat{\alpha}$, we have $p(\hat{\alpha}) > w > \hat{w}$; but at $a(t) = \hat{\alpha}$, we have to have $p(\hat{\alpha}) = \hat{w}$ for the uninformed buyer to be willing to buy, and this leads to the desired contradiction.

Consider then $\hat{p} < \hat{w}$ and consequently $\hat{p} < w < \hat{w}$. Suppose that $p(0) > w$ and hence by the differential equation $p(\alpha) > w$ for all t with $\alpha < \hat{\alpha}$. We also recall that as the equilibrium price path is continuous, it follows that at α such that $p(\alpha) = w > \hat{p}$, we have $w(\alpha) > p(\alpha)$. From proposition 2 we recall that the equilibrium during the full extraction phase satisfies

$$p(\alpha) = v + \lambda[\mathbb{E}_\theta V^\theta(\alpha) - V^u(\alpha)],$$

or, more explicitly,

$$p(\alpha) = v + \lambda \int_t^\infty \int_{\theta_l}^{\theta_h} \int_t^T e^{-r(\tau - t)} [p(\alpha(\tau)) - \theta]^+ d\tau dF(\theta) \lambda e^{-\lambda T} dT.$$

For notational ease we shall denote by $R(t)$ the forgone utility benefit from being uninformed in period t (i.e., the regret) or

$$R(t) \triangleq \int_{\theta_l}^{\theta_h} [p(\alpha(t)) - \theta]^+ dF(\theta),$$

and the equilibrium price is then given by

$$p(t) = v + \lambda \int_t^\infty \int_t^T e^{-r(\tau - t)} R(\tau) d\tau \lambda e^{-\lambda T} dT.$$

By hypothesis, $p(t)$ is strictly increasing until $t = \hat{t}$, where $\alpha(\hat{t}) = \hat{\alpha}$ and is decreasing thereafter. It is immediate that $R(t)$ shares the monotonicity properties with $p(t)$. We next show that $p(t)$ cannot be monotone increasing for all $t < \hat{t}$. After integrating with respect to T, we get

$$p(t) = v + \lambda \int_t^\infty e^{-(r+\lambda)(\tau - t)} R(\tau) d\tau.$$

Differentiating with respect to t, we get

$$p'(t) = \lambda \left[-R(t) + (r + \lambda) \int_t^\infty e^{-(r+\lambda)(\tau - t)} R(\tau) d\tau \right],$$

which has to turn negative as $t \uparrow \hat{t}$ by the hypothesis of an increasing price for all $t < \hat{t}$ and the continuity of the price path. This delivers the desired contradiction. The concavity in t follows immediately from $p(0) < w$ and the differential equation (13).

Part 3. The equilibrium sales are given by

$$q(t) = (1 - \alpha) + \alpha[1 - F(p(t))],$$

as long as the uninformed buyers are participating. The equilibrium sales are governed by the following differential equation:

$$q'(t) = -(1 - \alpha) F(p(t)) \lambda - \alpha f(p(t)) p'(t).$$

It follows that, even though $p'(t) < 0$, for all α sufficiently small, $q'(t) < 0$. The

second derivative is given by

$$q''(t) = (1 - \alpha)F(p(t))\lambda^2 - 2(1 - \alpha)\lambda f(p(t))p'(t)$$
$$- \alpha\{f'(p(t))[p'(t)]^2 + f(p(t))p''(t)\},$$

and again for all α sufficiently close to zero, the convexity of the sales follows directly from the decreasing price. QED

References

Ackerberg, Daniel A. 2003. "Advertising, Learning, and Consumer Choice in Experience Good Markets: An Empirical Examination." *Internat. Econ. Rev.* 44 (August): 1007–40.

Bergemann, Dirk, and Juuso Välimäki. 2004. "Monopoly Pricing of Experience Goods." Tech. Report no. 1463, Yale Univ., Cowles Found. Res. Econ.

Berndt, Ernst R., Margaret K. Kyle, and Davina C. Ling. 2003. "The Long Shadow of Patent Expiration: Generic Entry and Rx-to-OTC Switches." In *Scanner Data and Price Indexes*, edited by Robert C. Feenstra and Matthew D. Shapiro. Chicago: Univ. Chicago Press.

Berry, Steven T. 1994. "Estimating Discrete-Choice Models of Product Differentiation." *Rand J. Econ.* 25 (Summer): 242–62.

Ching, Andrew T. 2002. "Consumer Learning and Heterogeneity: Dynamics of Demand for Prescription Drugs after Patent Expiration." Manuscript, Univ. Toronto.

Coscelli, Andrea, and Matthew Shum. 2004. "An Empirical Model of Learning and Patient Spillovers in New Drug Entry." *J. Econometrics* 122 (October): 213–46.

Crawford, Gregory S., and Matthew Shum. 2005. "Uncertainty and Learning in Pharmaceutical Demand." *Econometrica* 73 (July): 1137–73.

Cremer, Jacques. 1984. "On the Economics of Repeat Buying." *Rand J. Econ.* 15 (Autumn): 396–403.

Erdem, Tulin, Susumu Imai, and Michael P. Keane. 2003. "Brand and Quantity Choice Dynamics under Price Uncertainty." *Quantitative Marketing and Econ.* 1 (1): 5–64.

Erdem, Tulin, and Michael P. Keane. 1996. "Decision-Making under Uncertainty: Capturing Dynamic Brand Choice Processes in Turbulent Consumer Goods Markets." *Marketing Sci.* 15 (1): 1–20.

Farrell, Joseph. 1986. "Moral Hazard as an Entry Barrier." *Rand J. Econ.* 17 (Autumn): 440–49.

Goettler, Ronald L., and Karen Clay. 2006. "Price Discrimination with Experience Goods: A Structural Econometric Analysis." Manuscript, Carnegie Mellon Univ.

Gowrisankaran, Gautam, and Marc Rysman. 2005. "Determinants of Price Declines for New Durable Consumer Goods." Manuscript, Washington Univ. and Boston Univ.

Israel, Mark. 2005. "Service as Experience Goods: An Empirical Examination of Consumer Learning in Automobile Insurance." *A.E.R.* 95 (December): 1444–63.

Johnson, Justin P., and David P. Myatt. 2006. "On the Simple Economics of Advertising, Marketing, and Product Design." *A.E.R.*, forthcoming.

Lu, Z. John, and William S. Comanor. 1998. "Strategic Pricing of New Pharmaceuticals." *Rev. Econ. and Statis.* 80 (February): 108–18.

Milgrom, Paul, and John Roberts. 1986. "Price and Advertising Signals of Product Quality." *J.P.E.* 94 (August): 796–821.

Nair, Harikesh. 2005. "Intertemporal Price Discrimination with Forward Looking Consumers: Application to the US Market for Console Video Games." Manuscript, Stanford Univ.

Osborne, Martin. 2005. "Consumer Learning, Habit Formation and Heterogeneity: A Structural Examination." Manuscript, Univ. Chicago.

Shapiro, Carl. 1983. "Optimal Pricing of Experience Goods." *Bell J. Econ.* 14 (Autumn): 497–507.

Sobel, Joel. 1991. "Durable Goods Monopoly with Entry of New Consumers." *Econometrica* 59 (September): 1455–85.

Tirole, Jean. 1988. *The Theory of Industrial Organization.* Cambridge, MA: MIT Press.

Villas-Boas, J. Miguel. 2004. "Consumer Learning, Brand Loyalty, and Competition." *Marketing Sci.* 23 (Winter): 134–45.