

MONOPOLY PRICING OF EXPERIENCE GOODS

By

Dirk Bergemann and Juuso Välimäki

July 2004

Revised May 2005

COWLES FOUNDATION DISCUSSION PAPER NO. 1463R



COWLES FOUNDATION FOR RESEARCH IN ECONOMICS

YALE UNIVERSITY

Box 208281

New Haven, Connecticut 06520-8281

<http://cowles.econ.yale.edu/>

Monopoly Pricing of Experience Goods*

Dirk Bergemann[†]

Juuso Välimäki[‡]

First Version: May 2003

This Version: May 2005

Abstract

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.

KEYWORDS: Monopoly, dynamic pricing, learning, experience goods, continuous time, Markov perfect equilibrium.

JEL CLASSIFICATION: D81, D83

* We would like to thank Kyle Bagwell, Ernst Berndt, Steve Berry, Patrick Bolton, Sven Rady, Mike Riordan, Klaus Schmidt and Ferdinand von Siemens for very helpful comments and Colin Stewart for excellent research assistance. We benefited from discussions at the IO Day at New York University and seminars at the University of Darmstadt, EUI, LSE, University of München, University of Nürnberg, IUI, University of Oulu, and UCL. The first author gratefully acknowledges financial support from NSF Grant SES 0095321 and DFG Mercator Research Professorship at the Center of Economic Studies at the University of Munich. The second author gratefully acknowledges support from ESRC through grant R000223448.

[†]Department of Economics, Yale University, New Haven, CT 06520-8268, U.S.A., dirk.bergemann@yale.edu.

[‡]Department of Economics, Helsinki School of Economics and University of Southampton, 00100 Helsinki, Finland, valimaki@hkkk.fi.

1 Introduction

1.1 Motivation

Since the original contribution by Nelson (1970), experience goods have been defined as those goods whose true quality is only learned upon consumption. As a result, any consumer choice model of experience goods must combine the perishable goods nature of repeat purchases with the durable information resulting from previous purchases. In this paper, we develop a simple tractable model of optimal pricing for a monopolist that sells an experience good over time to a population of potential buyers. We find that different types of experience goods induce qualitatively different long run outcomes depending on whether the short run perishable attributes or the long run informational features of the product are the most important in the purchasing decisions.

In an experience goods market, the seller is facing simultaneously two different submarkets. The demand curve in the part of the population that has already learned its preferences is similar to the standard textbook treatment. Those buyers that are uncertain about the true quality of the product must 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 forego purchases in the current period as 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.

Our main result is that the qualitative features of the equilibrium depend on a rather simple intertemporal comparison. When a single new consumer enters the market in its long run full information equilibrium, he calculates his expected consumer surplus from potential future purchases and compares that to the (possible) short run losses from current purchases. We define a market to be a *mass market* if such a buyer is willing to buy at the full information monopoly price. We show that in a mass market, the long run equilibrium price converges to the full information monopoly price and the long run sales also converge to full information sales levels.

In a *niche market*, uninformed buyers do not buy at the monopoly price. In this case,

long run prices also converge to full information levels, but sales fall short of the full information monopoly sales. We also show that the equilibrium price paths take quite different shapes in these two cases. In a mass market, the monopolist skims the more attractive part of the market (i.e. the uninformed buyers). In a niche market, the monopolist offers low initial prices to capture a larger share of the uninformed at the expense of targeting the more attractive informed segment 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 that 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. The key feature of the model is that different buyers become informed at different times. As a result, the market is segmented at each point in time and the degree of segmentation depends on the prices that the monopolist sets. Hence the model incorporates elements of demand management and market building.

To give a concrete example that fits our model, we consider the market for pharmaceuticals. Prior to launch, each new drug undergoes an extensive period of testing to determine its performance in the overall population. The aggregate uncertainty relating to the product has hence been reduced to a large extent at the moment of introduction. At the same time, many drugs have different and unknown effects for different patients. This individual level uncertainty provides a motive for experimentation. A recent empirical study by Crawford & Shum (2005) regarding the dynamic demand behavior in pharmaceutical markets documents the important role of idiosyncratic uncertainty and learning in explaining demand.¹

¹The empirical literature on learning based models in pharmaceutical demand is growing rapidly. Ching (2002) provides structural, dynamic demand estimates when there is learning among patients about a new (generic) pharmaceutical with common values. Coscelli & Shum (2004) estimate the impact of uncertainty and learning for the introduction of a new drug. The role of information is also central in Bhattacharya & Vogt (2003), where a model and preliminary estimates regarding informative marketing for new pharmaceutical products are presented.

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 heterogenous buyers. We provide a tractable analytical framework and demonstrate the flexibility of the framework by outlining extensions of the basic model in Sections 5 and 6. Besides the earlier cited work on pharmaceuticals, there is growing literature on Bayesian learning in consumer markets with experience goods, see e.g. Erdem & Keane (1996), Erdem, Imai & Keane (2003) and Akerberg (2003). All of these contributions focus on the optimal behavior of the buyers facing an exogenously given price. A contribution of our paper is then to present a parsimonious model of equilibrium behavior with fully optimizing and forward agents on both sides of the market. The model has a number of features which facilitate an empirical analysis of experience good markets. In our model, the equilibrium price converges to the static monopoly 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. Our paper therefore provides a modeling framework where it may be possible to extend the empirical analysis of Bayesian learning to an analysis of Bayesian equilibrium behavior. The simple characterization in terms of niche and mass market makes an empirical verification of the fit of the model feasible.

Finally, we should mention that the model allows an alternative interpretation which we spell out in more detail in Section 6. We can rephrase the random arrival of information as a random arrival of a consumption opportunity. With this specification, we may assume that the buyers learn their true preferences upon the first purchase. In addition to being commonly used in other papers on experience goods, this assumption expands the scope of our results to markets such as new brands of cereals, soft drinks, etc. Similarly, the demand for many professional services such as law, and internet services, often arises due to random

events.

1.2 Related Literature

Monopoly models dealing with issues of dynamic pricing include Milgrom & Roberts (1986), Farrell (1986) and Tirole (1988). All of these models make the assumption that the perceived quality is either high or otherwise of no value and the main emphasis is on vertical differentiation between the buyers. Furthermore, these models have only a two period horizon. We view both of these restrictions as unnecessary and unrealistic in many situations. Our model allows for the possibility that the monopolist discriminates intertemporally in the market in a more flexible manner and as a result, our conclusions are quite different from those in the earlier literature. In our model, it is possible that the marginal buyer in the later periods might have a lower willingness to pay for quality than the marginal buyer in the earlier periods and as a result, buyers have an incentive to engage in experimental consumption. An early paper on optimal dynamic pricing of experience goods is Shapiro (1983). In his model, the consumers are differentiated ex-ante in terms of their willingness to pay. Every consumer learns the true quality of the product only through his own experience even though it is assumed to be the same for everybody. Furthermore, it is assumed that (i) each consumer acts myopically and (ii) expected quality is biased with respect to the true quality. Shapiro (1983) derives results for buyers with either an optimistic or a pessimistic bias. The natural benchmark of unbiased buyers would yield 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, Cremer (1984) considers the optimal commitment price path and does not analyze the time consistent pricing policy. In a recent contribution, Villas-Boas (2004) considers the equilibrium 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 subsequent analysis is mostly concerned with brand loyalty, i.e. whether buyers return to the seller they bought from in the past. Villas-Boas (2004) presents condition on the skewness of the distribution under which brand loyalty exists in equilibrium. As the model combines differentiation along two dimensions, it is impossible to obtain clear and general results about the intertemporal price path. Villas-Boas (2005) provides an extension to an infinite horizon model.

Finally, conditions for initially high prices have been obtained in asymmetric information models of entry. In those papers, the monopolist is assumed to know the true value of the product, and the prices chosen serve as signals of the true quality. A prominent example of such models is Bagwell & Riordan (1991) where high and declining prices serve as signals of high product quality. Judd & Riordan (1994) consider a model with initially symmetric information where private signals are received by the monopolist and the buyers after first period choices. The firm then faces a signaling problem in the second period. The results in these models depend on the details of the information revelation mechanism and the cost structure. In our model, the results depend only on the quality difference between the products which can in principle be inferred directly from the realized prices.

The paper is organized as follows. Section 2 sets up the basic model and discusses the appropriate solution concepts. Section 3 presents the problem of optimal demand management for the seller. Section 4 analyzes the properties of the optimal price path. Section 5 discusses the social efficient allocation and the role of idiosyncratic versus social learning. Section 6 provides extensions of the model including inflow of new buyers, and random purchasing opportunities for the buyer. Section 7 concludes. The proofs of all the results are collected in the Appendix.

2 Model

We consider a continuous time model with $t \in [0, \infty)$ and a positive discount rate $r > 0$. A monopolist with a cost of production offers a 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 buyer types 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 quasiconcave 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 , with:

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

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 λdt for all buyers that purchase the product in a time instant of length dt .² In this case, the posterior distribution on θ remains constant at the prior until the signal is observed.³ The most important analytical consequence of this assumption is that conditional on not having observed a signal, the buyers 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.⁴

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 public observable variables in each period 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 each instant t is the fraction of agents that have become informed. We denote this fraction by $\alpha(t) \in [0, 1]$. Even though $\alpha(t)$ is not directly observable to the players, it can be calculated from the equilibrium purchasing strategies. Conditional on the uninformed buying in period t , the state variable

²It might be natural to allow for cases where 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.

³This assumption is made for the ease of exposition only. We have also computed the model for posteriors with positive and negative drift. Again, the qualitative features of the solution are the same as in the constant case.

⁴In section 6, we give an alternative interpretation to the informational structure of our model in terms of random purchasing opportunities for the buyers.

evolves according to:

$$\frac{d\alpha(t)}{dt} = \lambda(1 - \alpha(t)).$$

as in period t there are $1 - \alpha(t)$ currently informed 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 $p(\alpha)$. An uninformed buyer has a Markovian purchasing strategy $d^u(\alpha, p)$ that depends on the state of the system as well as the current price. Finally, the Markovian purchasing strategy $d^\theta(\alpha, p)$ of the informed buyer depends on the state, the current price and his valuation.

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. Note that 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.

3 Demand Management

The basic issue in the introduction of a new product is the dynamic demand management. In the early stages, most 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 focuses the more important segment. More precisely, the type of the buyer whose willingness to pay determines the equilibrium price changes over time. With a new product, the marginal buyer is clearly uninformed in the early stages. As the informed segment grows, the marginal buyers is more likely to come from that segment. Optimal demand management then determines when to switch between these two market segments. The dynamic pricing policy of the seller therefore contains at its core an optimal stopping problem.

After the switch, the marginal buyer is informed. It is not clear whether the uninformed buyers keep on purchasing the product. Either they are priced out of the market or they stay in the market as inframarginal buyers. Whether the uninformed buyers stay in the market or drop out 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 alternatives 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.

3.1 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:

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

Price \hat{p} is also the optimal price in the static monopoly problem where each buyer knows his valuation for the object. The corresponding equilibrium quantity of sales 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} . The expected value v of the product to an uninformed buyer coincides with 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 :

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

The uninformed buyer receives an informative signal at the rate λ . If the information arrives, then buyer becomes informed and purchases the product if and only if $\theta - \hat{p} \geq 0$. Finally, the future benefit of the information is discounted at a 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, $\frac{\lambda}{r} \mathbb{E}_\theta \max \{\theta - \hat{p}, 0\}$.

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} .⁵

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 and/or the discount rate r is small.

We can gain further insight into the notions of niche and mass market by a comparative static analysis. For a given r and λ , let us consider a family of distributions $F(\theta; \sigma^2)$, parametrized by variance σ^2 under a given and constant 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 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. We verified the above intuition exactly for 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

3.2 Optimal Switching

We first describe the intertemporal decision problems in terms of the familiar dynamic programming equations in Markov strategies. We start with the simple decision problem of

⁵In a recent contribution, Johnson & Myatt (2004) use the distinction between niche and mass market to investigate whether a (mean-preserving) spread of the valuation increases the revenue in the static optimal pricing problem. They interpret “Advertising” and “Marketing” as providing information which increases the variance of the valuations by the customers.

⁶We would like to thank the editor, Robert Shimer, who suggested the comparative static result regarding the variance.

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 V^θ of the informed buyer from the Bellman's equation:

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

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 individual buyer and hence the myopic decision 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:

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

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:

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

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:

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

subject to

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

Here $(1 - \alpha)$ is the share of uninformed buyers in the population and $\alpha(1 - F(p(\alpha)))$ is the fraction of informed buyers that 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:

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

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 define:

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 niche market and 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 $\hat{\alpha}$.

3.3 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 which characterize the stopping point $\hat{\alpha}$. After reaching $\hat{\alpha}$, 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} or \hat{w} :

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

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 buyer 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

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

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. $\hat{\alpha} < 1$.
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 5 and 6, we formulate models of social learning and random purchasing opportunities where the separability of the problems with respect to different discount rates is useful.

3.4 Mass Market

Initially, the informed segment does not exist. The monopolist offers prices which leave the uninformed agents just indifferent between buying and not buying. The monopolist thus extracts initially all the surplus from the current purchases of the uninformed agents. As the informed segment grows, any price which 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 point at which she 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 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 or not buying. We can express this in terms of the equilibrium value function of the buyer, using (3):

$$p(\alpha) = v + \lambda \left(\mathbb{E}_\theta V^\theta(\alpha) - V^u(\alpha) \right). \quad (8)$$

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

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

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

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

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. This is because 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:

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

Even though the stopping condition can therefore be expressed in terms of the flow revenues, the problem contains an intertemporal element as the equilibrium price $p(\alpha)$

before and at the stopping point is based on the equilibrium continuation values, see (8)). After $\hat{\alpha}$ is reached, the equilibrium price is computed as the price that maximizes the flow revenue, or

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

This equation is a static optimization condition, but 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. $\hat{\alpha} \leq 1$.
2. $\hat{\alpha}$ is decreasing in λ and increasing in r .

If we consider the comparative static results in Proposition 1 and 2, then it is worth observing that the respective stopping points for niche and mass market move in opposite directions as a function of λ and r . Consider a distribution $F(\theta)$ with the property that the nature of the market depends on the values of λ and r . 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 offers introductory prices for a longer period. Eventually λ will reach a point where 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 is, in fact, constant and $\hat{\alpha} = 1$. For higher λ , 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.

FIGURE 1: STOPPING POINT FOR VARYING $\frac{\lambda}{r}$.

4 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:

$$p(\alpha) = v + \lambda \mathbb{E}(V^\theta(\alpha) - V^u(\alpha)). \quad (10)$$

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. $p(\alpha)$ satisfies for all $\alpha \leq \hat{\alpha}$:

$$\frac{dp}{d\alpha} \frac{d\alpha}{dt} = r(p(\alpha) - v) - \lambda \mathbb{E}_\theta \max\{\theta - p(\alpha), 0\}. \quad (11)$$

2. $p(\alpha)$ is decreasing in α for all $\alpha < \hat{\alpha}$.

3. $q(\alpha)$ is initially decreasing and convex in α .

The differential equation (11) describes the evolution of the price in the early market. By rearranging the equality (11), it becomes apparent that the differential equation represents the trade-off of the uninformed buyer in his current 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. This benefit has two components. First, buying tomorrow allows the buyer to postpone the net cost of a purchase, $p(\alpha) - v$, and second, the price to be paid changes.

Proposition 4 (Mature Market)

1. In the niche market, $p(\alpha)$ 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 $\lim_{\alpha \rightarrow 1} p(\alpha) = \hat{p}$.

The value of information before the stopping point, is decreasing over time. While the price dynamics is governed by the same differential equation for the niche and the mass market, the source of the decrease in the value of information is different in the niche and the mass market. In the niche market, a phase of introductory prices ends at the stopping point. After this, the seller increases her price to \hat{p} and a large fraction of surplus 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 path 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 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 niche and mass market, 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.

INSERT HERE FIGURE 2: EQUILIBRIUM PRICE FOR NICHE MARKET

INSERT HERE FIGURE 3: EQUILIBRIUM PRICE FOR MASS MARKET

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 estimates a static demand function for a product using data that includes observations of prices for $\alpha < \hat{\alpha}$, it is clear that the resulting estimators are biased.

5 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.

5.1 Social Efficiency

The socially efficient policy maximizes the sum of the expected discounted value across all agents. As learning and the resolution of uncertainty is purely idiosyncratic, the socially optimal policy can be determined as the solution for a representative consumer. The buyers and the seller have quasilinear preferences and hence socially efficient policy simply maximizes expected social surplus. The optimal consumption policy can be determined for informed and uninformed consumer 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:

$$rW^u = \max \left\{ v + \lambda \mathbb{E}_\theta \left[W^\theta - W^u \right] - c, 0 \right\}. \quad (12)$$

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, $\lambda \mathbb{E} [W^\theta - W^u]$, exceeds zero. The expected value of the informed consumer on the other hand is

$$\mathbb{E} [W^\theta] = \frac{1}{r} \mathbb{E}_\theta \max \{\theta - c, 0\}. \quad (13)$$

By inserting the value function for the informed consumer, or (13) into (12), we obtain the critical value v so that the social value of consuming an additional unit of the new good is larger 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 is 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 , are negative.

The monopoly position of the seller introduces static as well as intertemporal distortions away from the socially optimal level of consumption. The static element of the distortion come from the standard revenue consideration 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 market 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 both for the informed and the uninformed consumer. The important ingredient of the model which leads to an agreement of the competitive and the efficient outcome 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 market will now be different.

5.2 Social Learning

We introduce social learning into our model by replacing each individual buyer with a group of k buyers who agree on value of the good. The group can be interpreted as a family, a department, or a neighborhood where individuals have similar preferences and the ability to communicate. At each point in time, one buyer is selected from the group of k members to consider a purchase. The buyer bears the cost and enjoys the benefits from the purchase privately. The information resulting from his purchase is assumed to be shared within the group.

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

$$rV^\theta(\alpha) = \frac{1}{k} \max\{\theta - p(\alpha), 0\} + \frac{dV^\theta}{d\alpha} \frac{d\alpha}{dt}. \quad (14)$$

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:

$$rV^u(\alpha) = \frac{1}{k} \max\left\{v - p(\alpha) + \lambda \left(\mathbb{E}_\theta V^\theta(\alpha) - V^u(\alpha)\right), 0\right\} + \frac{k-1}{k} \lambda \left(\mathbb{E}_\theta V^\theta(\alpha) - V^u(\alpha)\right) + \frac{dV^u}{d\alpha} \frac{d\alpha}{dt}. \quad (15)$$

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

Social learning has two effects on optimal pricing. First, the uninformed buyers are less willing to pay a premium for the information, as there is a chance to learn from the experiences of others. Second, the seller is willing to sponsor the experiment of each agent as she knows that the information will spread among all members of the group. If they value the object sufficiently highly, the seller can extract some of their surplus in the future. For the individual buyer, the probability $1/k$ of making the purchase, acts as an 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{a} are determined by using a modified discount rate $r \cdot k$ 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 competition and monopoly becomes quite different. The monopolist partially internalizes the informational externality as she understands that new 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 only takes into account 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.⁷

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 $v - c + \frac{\lambda}{r} \hat{p} (1 - F(\hat{p})) > 0$, then it is optimal for the monopolist to launch the product for all k .*

6 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 that 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.

⁷The proof of this result is standard and therefore omitted.

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 to consumer goods such as cereals, cosmetics, etc. 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 where the buyers' discount rate is set to $\frac{r+\lambda}{\lambda}$ and the arrival rate of information is set to 1. As was explained in section 3, 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.

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, either because of the ageing of the customers or other systematic changes to the agents' preferences. In Bergemann & Valimaki (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 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 as 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 for example Sobel (1991)). In our setting, the product is effectively a bundle, consisting of perishable element in form of the immediate consumption benefit, but also a durable element in form of the information obtained with the purchase. From this perspective, the relationship to the durable good pricing problem appears inherently.

Uniqueness of Equilibrium The equilibrium analysis in this paper focused on the notion of a Markov Perfect Equilibrium and derived the unique equilibrium in this class. The uniqueness result extends to a much larger class of equilibria. In Bergemann & Valimaki (2004) we show that the MPE 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.

7 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 as 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 where 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 consider the role of idiosyncratic learning in a strategic environment against either a known or similarly unknown product. An interesting variation of the current model and of special importance for the pharmaceutical market would be to consider optimal pricing when a competitor will only appear in T periods hence, induced by the expiration of the patent, see Berndt, Ling & Kyle (2003) for an empirical analysis of this situation.

8 Appendix

The 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. (1.) The optimal stopping point is given as a solution to (6) by:

$$\hat{\alpha} = \frac{\hat{w} + \frac{\lambda}{r}\hat{p}(1 - F(\hat{p}))}{\hat{w}F(\hat{w}) + \hat{p}(1 - F(\hat{p}))\left(1 + \frac{\lambda}{r}\right)}. \quad (16)$$

For simplicity we define $\sigma \triangleq \frac{\lambda}{r}$, and with this we rewrite equation (16) as:

$$\hat{\alpha} = \frac{\hat{w} + \sigma\hat{p}(1 - F(\hat{p}))}{\hat{w} + (1 + \sigma)\hat{p}(1 - F(\hat{p})) - \hat{w}(1 - F(\hat{w}))}, \quad (17)$$

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

(2.) After replacing \hat{w} by its explicit expression given in (1), we rearrange equation (17) to get:

$$\begin{aligned} & \hat{\alpha} \left((v + \sigma\mathbb{E}(\theta - \hat{p})^+) F(v + \sigma\mathbb{E}(\theta - \hat{p})^+) + \hat{p}(1 - F(\hat{p})) (1 + \sigma) \right) \\ &= v + \sigma\mathbb{E}(\theta - \hat{p})^+ + \sigma\hat{p}(1 - F(\hat{p})). \end{aligned} \quad (18)$$

Differentiating the equality (18) implicitly with respect to σ yields:

$$\begin{aligned} & \frac{d\hat{\alpha}}{d\sigma} (\hat{w}F(\hat{w}) + \hat{p}(1 - F(\hat{p})) (1 + \sigma)) + \hat{\alpha} ((\hat{w}f(\hat{w}) + F(\hat{w}))\mathbb{E}(\theta - \hat{p})^+ + \hat{p}(1 - F(\hat{p}))) \\ &= \mathbb{E}(\theta - \hat{p})^+ + \hat{p}(1 - F(\hat{p})), \end{aligned}$$

and hence

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

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 quasiconcavity of $p(1 - F(p))$ together with the fact that $\hat{w} < \hat{p}$. Therefore, the numerator is also positive, as needed. ■

Proof of Proposition 2. (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:

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

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

$$\int_t^T e^{-r(\tau-t)} (v - p(\alpha(\tau))) d\tau + \int_{\theta_l}^{\theta_h} \left[\int_T^\infty e^{-r(\tau-t)} (\theta - p(\alpha(\tau)))^+ d\tau \right] dF(\theta). \quad (19)$$

The value function of the uninformed buyer is obtained from (19) 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:

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

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

$$\mathbb{E}_\theta V^\theta(\alpha) - V^u(\alpha) = \int_t^\infty \int_{\theta_l}^{\theta_h} \int_t^T e^{-r(\tau-t)} (p(\alpha(\tau)) - \theta)^+ d\tau dF(\theta) \lambda e^{-\lambda T} dT. \quad (20)$$

The gain of the informed vis-a-vis the uninformed buyer, arises in all those instances where 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) as $p(\alpha)$ is decreasing in α (and t). We can then run $w(\alpha)$ backwards as long as $w(\alpha) \geq p(\alpha)$.

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

$$w(\alpha) = p(\alpha). \quad (21)$$

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

$$w(1) > p(1). \quad (22)$$

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 (20). It therefore follows that, provided $p(\alpha)$ is monotone,

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

which together with (22) is sufficient to guarantee existence and uniqueness of the stopping point.

(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 (21) 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. ■

Proof of Proposition 3. (1.) We obtain the differential equation by differentiating $p(\alpha)$ as given by (8) with respect to t to obtain

$$\frac{dp}{d\alpha} \frac{d\alpha}{dt} = \lambda \mathbb{E} \left(\frac{dV^\theta(\alpha)}{d\alpha} \frac{d\alpha}{dt} - \frac{dV^u(\alpha)}{d\alpha} \frac{d\alpha}{dt} \right).$$

We can then replace the first derivatives of the value function of informed and uninformed agent by using the dynamic programming equations, (2) and (3). After cancelling redundant terms we arrive at the differential equation (11) for the price $p(\alpha)$. The equation (11) for the full extraction prices has a unique rest point, $\dot{p}(t) = 0$, at $p(t) = w$ as w uniquely solves:

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

(2.) We show the monotonicity of $p(\alpha)$ separately for $\hat{p} < \hat{w}$ and $\hat{p} \geq \hat{w}$. We start with the later 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

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

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 foregone 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 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 (11).

(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

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

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

References

- Akerberg, D. 2003. "Advertising, Learning and Consumer Choice in Experience Good Markets: A Structural Empirical Examination." *International Economic Review* 44:1007–1040.
- Bagwell, K. & M. Riordan. 1991. "High and Declining Prices Signal Product Quality." *American Economic Review* 81:224–239.
- Bergemann, D. & J. Valimaki. 2004. Monopoly Pricing of Experience Goods. Technical Report 1463 Cowles Foundation for Research in Economics, Yale University.
- Berndt, E.R., D. Ling & M.K. Kyle. 2003. The Long Shadow of Patent Expiration: Do Rx to OTC Switches Provide an Afterlife. In *Scanner Data and Price Indexes*, ed. R. Feenstra & M. Shapiro. Chicago: University of Chicago Press.
- Berry, Steve. 1994. "Estimating Discrete Choice Models of Product Differentiation." *RAND Journal of Economics* 25:242–262.
- Bhattacharya, J. & W.B. Vogt. 2003. "A Simple Model of Pharmaceutical Price Dynamics." *Journal of Law and Economics* 46:599–626.
- Ching, A.T. 2002. Consumer Learning and Heterogeneity: Dynamics of Demand for Prescription Drugs After Patent Expiration. Technical report University of Toronto.
- Coscelli, A. & M. Shum. 2004. "An Empirical Model of Learning and Spillovers in New Drug Entry." *Journal of Econometrics* 122:213–246.
- Crawford, G.S. & M. Shum. 2005. "Uncertainty and Learning in Pharmaceutical Demand." *Econometrica* forthcoming.
- Cremer, J. 1984. "On the Economics of Repeat Buying." *Rand Journal of Economics* 15:396–403.
- Erdem, T. & M. Keane. 1996. "Decision Making under Uncertainty: Capturing Dynamic Brand Choice Processes in Turbulent Consumer Goods Markets." *Marketing Science* 15:1–20.

- Erdem, T., S. Imai & M. Keane. 2003. "Brand and Quantity Choice Dynamics under Price Uncertainty." *Quantitative Marketing and Economics* 1:5–64.
- Farrell, J. 1986. "Moral Hazard as an Entry Barrier." *Rand Journal of Economics* 17:440–449.
- Johnson, J.P. & D.P. Myatt. 2004. On the Simple Economics of Advertising, Marketing, and Product Design. Technical report Cornell University and Oxford University.
- Judd, K. & M. Riordan. 1994. "Price and Quality in a New Product Monopoly." *Review of Economic Studies* 61:773–789.
- Milgrom, P. & J. Roberts. 1986. "Prices and Advertising Signals of Product Quality." *Journal of Political Economy* 94:796–821.
- Nelson, P. 1970. "Information and Consumer Behaviour." *Journal of Political Economy* 78:311–329.
- Shapiro, C. 1983. "Optimal Pricing of Experience Goods." *Bell Journal of Economics* 14:497–507.
- Sobel, J. 1991. "Durable Goods Monopoly with Entry of New Consumers." *Econometrica* 59:1455–1485.
- Tirole, J. 1988. *The Theory of Industrial Organization*. Cambridge: MIT Press.
- Villas-Boas, J.M. 2004. "Consumer Learning, Brand Loyalty, and Competition." *Marketing Science* 23:134–145.
- Villas-Boas, J.M. 2005. "Dynamic Competition with Experience Goods." *Journal of Economics and Management Strategy* forthcoming.

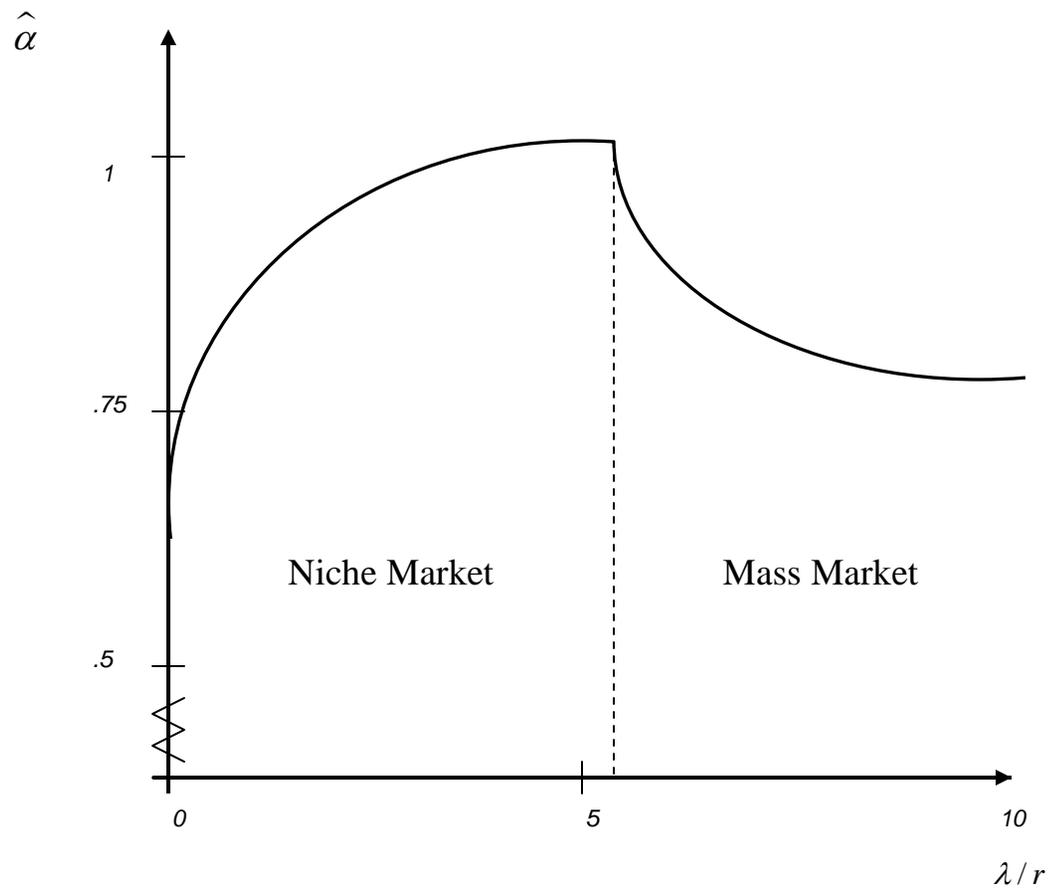


Figure 1: Stopping Point $\hat{\alpha}$

Niche Market: $\lambda/r < 5.25$; Mass Market: $\lambda/r \geq 5.25$; $F(\theta) = U[-0.75, 1]$, $c=0$.

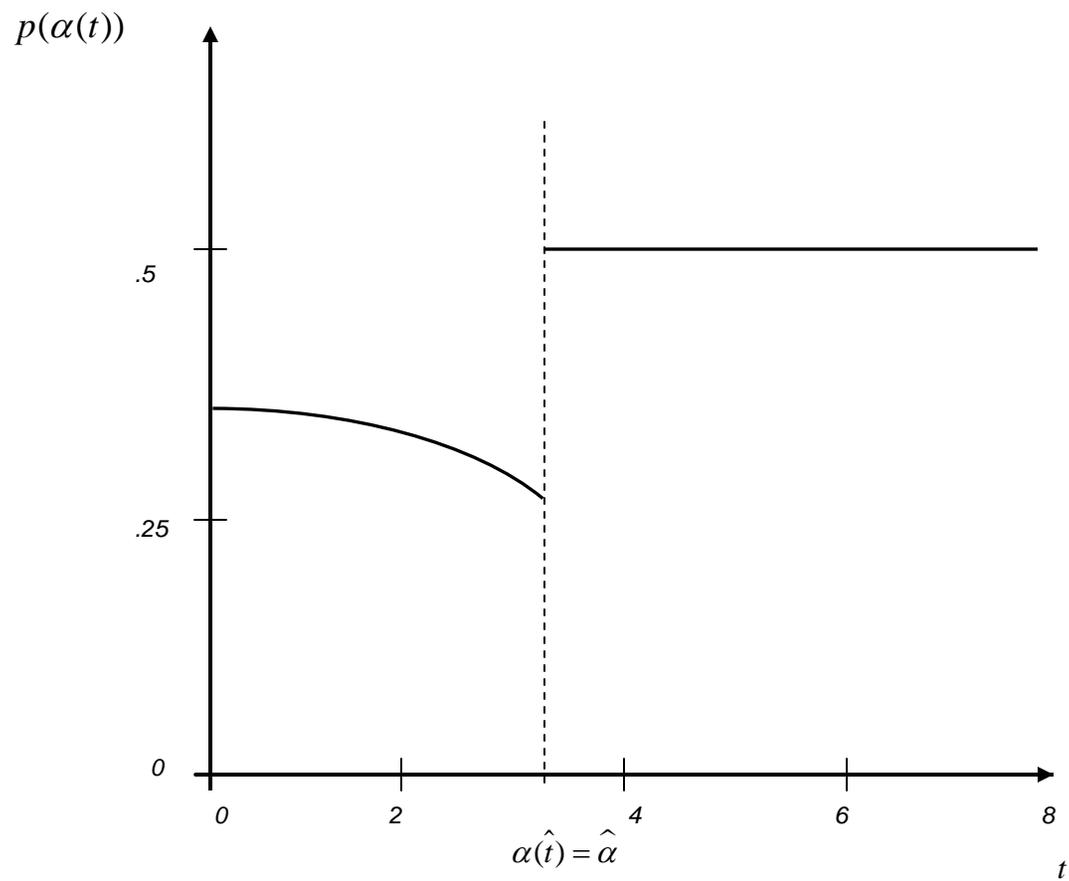


Figure 2: Intertemporal Price in Niche Market

$$F(\theta) = U[-0.75, 1], \quad c=0, \quad \lambda=0.4, \quad r=0.05$$

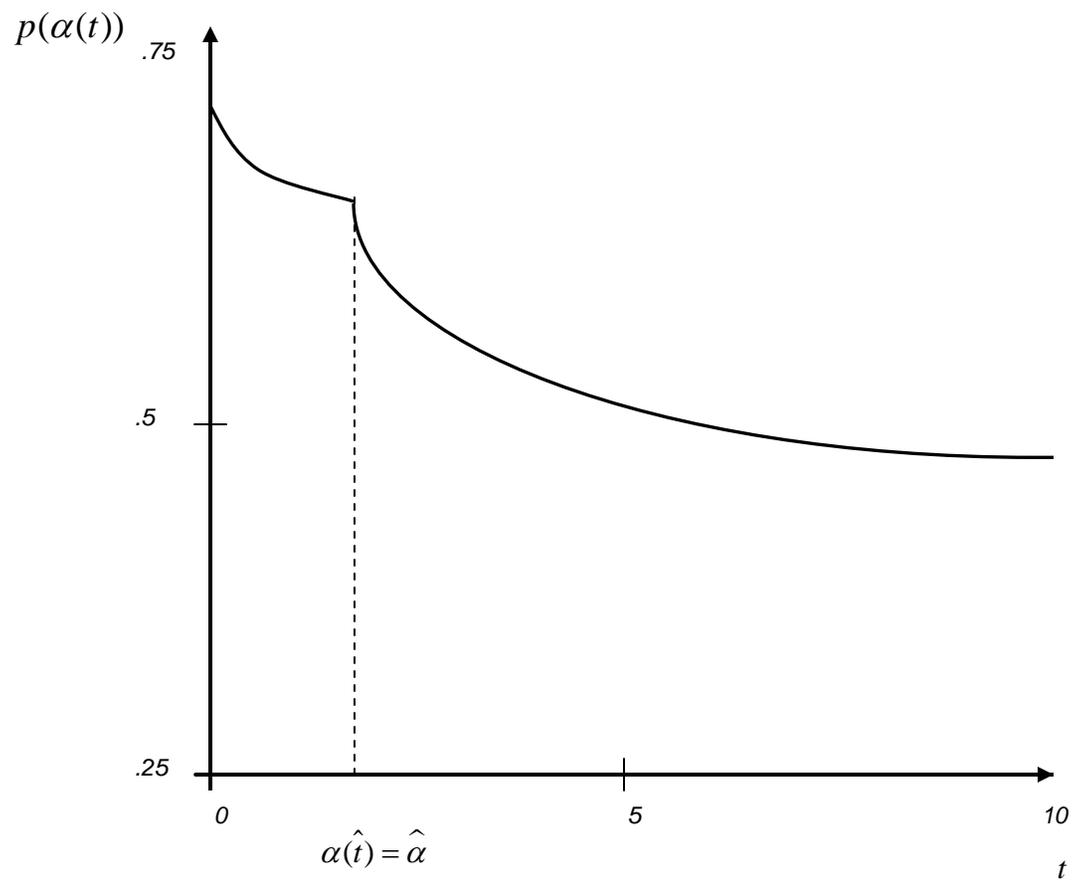


Figure 3: Intertemporal Price in Mass Market

$$F(\theta) = U[-0.75, 1], \quad c=0, \quad \lambda=0.1, \quad r=0.05$$