Continuous-Time Screening Contracts

Alessandro Bonatti*
Yale University

February 14, 2008

Abstract

This paper studies the evolution of nonlinear pricing schemes for new experience goods using a continuous-time screening model. It combines a traditional price discrimination model with a Brownian motion information structure by assuming that the monopolist can control the rate at which consumers learn about the product's quality. The paper characterizes the optimal nonlinear tariff as a function of consumer beliefs about quality, and derives predictions about price and quantity dynamics.

The optimal policy reflects not only the need for experimentation, but also the dynamic trade-offs between efficiency, the value of information, and informational rents. In a linear-quadratic setting, these three effects imply that equilibrium quantity provision need not be a monotonic function of beliefs; in particular, the monopolist initially adopts an aggressive sales policy to accelerate the rate of learning, then reduces the quantity supplied to low-valuation buyers in order to limit informational rents. Similarly, with constant price elasticity of demand, unit prices are nonmonotonic in market beliefs (first decreasing, then increasing), reflecting the use of introductory pricing by the monopolist.

Keywords: Nonlinear pricing, menus of contracts, experience goods, Bayesian learning, experimentation.

J.E.L. codes: D42, D82, D83, L12

*Department of Economics, Yale University, New Haven CT 06520-8268, U.S.A. alessandro.bonatti@yale.edu.
I wish to thank Rossella Argenziano, Dirk Bergemann, Ian Larkin, Marco Ottaviani, and Maher Said for helpful discussions and suggestions, as well as participants in the Micro Theory Lunch at Yale University. All errors are my own. This work was generously supported by a Cowles Foundation Anderson Fellowship.

1 Introduction

When new products are introduced on the market, consumers often have the option of selecting among several versions. These may differ in their qualitative features (think of the Home and Business versions of a new operating system), or simply in the number of units (such as the number of channels in satellite television, the number of movies in a Netflix DVD rentals subscription, or the number of minutes in a cellular phone plan). A common feature to these markets is that consumers have heterogeneous tastes for the product (watching a movie, obtaining a license for a software), thereby creating an incentive for firms not to price all versions uniformly, for example by offering quantity discounts. Another common feature is that the inherent quality of the product may be not be perfectly known at the time of its introduction. However, the product's diffusion enables consumers and firms to gradually obtain information about its quality. In this setting, the arrival of information may induce both buyers and firms to adjust their sales policies and purchasing decisions.

The goal of this paper is to study the evolution of the optimal nonlinear tariffs as the market receives information about product quality. It addresses questions such as: will firms offer increasing discounts for large orders? Will the product line display more or less variety? In other words, we are interested in the dynamics of the prices and quantities combinations available within a firm's product line. To be more concrete, consider the case of the online DVD rental company Netflix. Their initial menu offer consisted of five different plans and their product soon proved to be successful.[1] In the following two years, Netflix first raised its prices (2004) and then reduced them slightly, while at the same time adding more options (2005). Our model aims to explain the evolution of these menus, relating changes in the offered plans to the diffusion of information about product quality. Figure 1 reports the price-quantity bundles over the years 2003-2005, that is, immediately before Blockbuster established itself as a serious competitor.

In order to address these issues, we develop a continuous-time screening model in which, at each instant, a monopolist firm offers a menu of contracts to a population of buyers. As purchases are made, both the firm and the consumers observes a sequence of public signals about the product's quality, and revise their beliefs. We consider buyers who purchase repeatedly and a seller who may modify the tariffs at any point in time. We assume further that buyers with higher tastes for the product benefit more from an increase in quality, and that the precision of the signals observed by the market is directly related to the number of units sold in each period.

As an application of our model, consider the case of DVD rental plans. The key feature of these contracts is that consumers have the option of choosing the number of movies they may rent out at the same time. While buyers differ in their personal tastes for watching movies on DVD, the quality of the recommender system (or also of the delivery system) is a key component in determining the overall quality of the service. With this interpretation, each rented movie constitutes an informative experiment on the product's quality. It is reasonable to imagine that customers with higher tastes

[1] The total number of users grew from 900,000 in 2002 to 3.3 million at the end of 2004 according to the company's Investor Relations website (http://ir.netflix.com).

Figure 1: Netflix Rental Plans 2003-2005

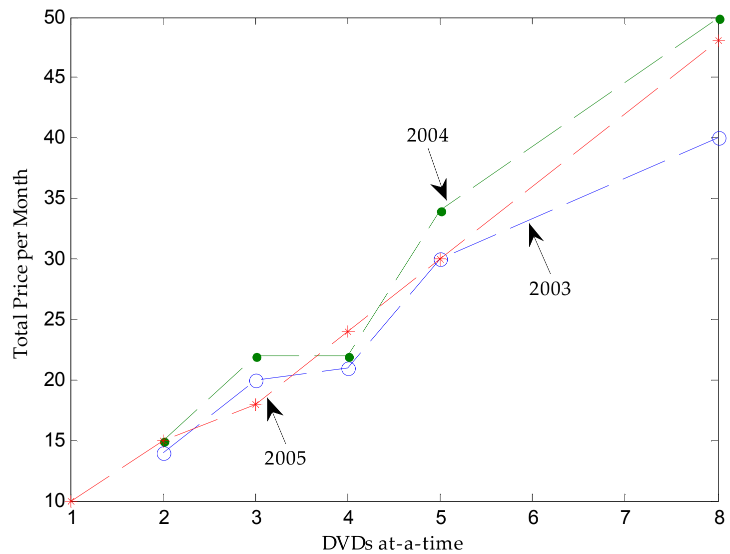

for movies also care particularly about the fit of the recommendation to their own tastes (or about the speed of delivery). Furthermore, both the prices for each plan and the plan choice itself are easily adjustable, making the repeated purchases framework realistic.

The market for enterprise software provides an alternate application of our model. An emerging contractual arrangement in this industry is given by software-as-a-service (SaaS). Under this contract, firms have the option of renting a given number of use licenses ("seats") for a given software (e.g. customer management or database programs). Larger firms need to rent more seats and arguably benefit more from a higher quality product. In this market, each employee using the software constitutes an experiment for product quality, so that the number of seats may be directly tied to the arrival speed of information. Moreover, the license renting contracts, and the corresponding unit prices are also easily adjustable. Finally, network externalities are not as big an issue in enterprise software (which is designed for internal use), making the private values framework quite realistic.

Our model characterizes the optimal dynamic nonlinear pricing policy under two different sets of functional form assumptions. In particular, we consider a linear-quadratic specification analogous to that of Mussa and Rosen (1978) and a constant price elasticity demand specification, as in Maskin and Riley (1984). In both cases, we focus on the undiscounted version of the firm's problem. In particular, we utilize the strong long-run average criterion to derive the limiting optimal policy function as the discount factor goes to one.[2] This method provides increased tractability and allows us to solve the model analytically. At the same time, the solution preserves all the qualitative features of the dynamic (discounted) optimal policies.

[2] For an extensive description of this criterion, and its relation to other undiscounted optimality criteria, the reader is referred to Dutta (1991).

The results show that the seller's optimal policy is influenced by several factors operating simultaneously: since information is generated via product diffusion, selling large numbers of units provides an additional value from experimentation. More importantly, a trade-off emerges between efficiency, information generation, and informational rents. In brief, a long series of positive signals about the product's quality raises each consumer type's valuation for the good. At the same time, it reduces the value to experimenting and it raises the informational rents demanded by the highest types. In a linear-quadratic setting, this implies the quantity levels offered to each type are not necessarily monotonic in the beliefs about product quality. In a constant price elasticity model, quantity is a strictly increasing function of the market's beliefs, but unit prices are not. In particular, unit prices are set at their lowest levels when information is valuable (uncertainty is high) and rents are not excessively high (meaning the beliefs are quite pessimistic), in order to induce larger purchases. Thus, the characteristics of technology and demand functions influence the instruments through which the firm induces the market to experiment with its product. Furthermore, under both sets of assumptions, we obtain predictions on the dynamics of product line variety and on the expected time path of quantity and price levels. We find results consistent with introductory pricing: positive signals about product quality increase variety, while quantities are expected to decrease, and unit prices to increase over time.

This paper builds upon both the screening literature (e.g. Mussa and Rosen (1978), Maskin and Riley (1984)) and the strategic experimentation literature (e.g. Bolton and Harris (1999)). It is also tightly connected to the literature on dynamic pricing. In particular, Bergemann and Välimäki (2002) consider a duopoly game of price competition in which consumers have heterogeneous tastes for quality and the two firms' products are vertically differentiated. We adopt a similar demand structure and we characterize the solution to the firm's undiscounted problem. In contrast to their model, our paper allows consumers to have multi-unit demands and firms to price discriminate. At this stage, however, our analysis is limited to the case of monopoly. Finally, the present paper is also related to papers on dynamic regulation, such as Lewis and Yildirim (2002). In fact, an easy reformulation of our model addresses the case of procurement under uncertainty over production costs.

The rest of this paper is organized as follows: Section 2 describes the model in detail; Section 3 solves the firm's (undiscounted) optimization problem; Sections 4 and 5 characterize the solution in the linear-quadratic and constant price elasticity frameworks respectively; and Section 6 summarizes the results and suggests the future developments of our work.

2 The Model

We consider a model with a monopolist firm and a continuum of consumers. Each consumer's valuation for q units of a product is assumed to be a separable function of the product's inherent quality and of the consumer's idiosyncratic taste parameter. In other words, the consumer's type is given by the product of two components, θ and μ, where μ is a common component representing the quality of the entire product line, while θ is an idiosyncratic component. With reference to the Netflix or enterprise software examples, one may think of μ as the quality of the software or of the recommender system and of θ as the consumer's particular taste for movies, or need for "seats". The idiosyncratic component θ is the consumer's private information. It is distributed in the population according to a distribution

$$F(\theta) : [\theta_L, \theta_H] \to [0,1]$$

with a continuously differentiable density $f(\theta)$. For simplicity, we also assume that $F(\theta)$ satisfies the nondecreasing likelihood ratio property. The parameter μ may only take two values, $\mu \in \{\mu_L, \mu_H\}$ and it is initially unknown to both the firm and the consumers. We allow the firm to offer consumers different quantity bundles at different prices. The total utility of a type θ consumer purchasing q units of a good of quality μ at price $p(q)$ is given by:

$$U = \mu\theta u(q) - p(q).$$

Thus, our model can be interpreted as a screening problem with an unknown support of consumers' valuations. In fact, depending on product quality, the type space is effectively either $[\mu_L\theta_L, \mu_L\theta_H]$ or $[\mu_H\theta_L, \mu_H\theta_H]$.

We focus on the continuous time version of the model. Furthermore, we assume that, at each instant, the market observes a public signal π about the inherent quality of the product (μ). The evolution of this signal is given by a Brownian motion with drift μ and variance $\sigma^2/\bar{q}$,

$$d\pi = \mu dt + \frac{\sigma}{\sqrt{\bar{q}}} dz,$$

where $\bar{q} = \int_{\theta_L}^{\theta_H} q(\theta) f(\theta) d\theta$ is the total number of units purchased in a given period by the consumers and dz is the standard Wiener process. With this structure for the public signals, one can show (see theorem 9.1 in Liptser and Shiryaev (1977)) that the evolution of the common belief $\alpha = \Pr(\mu = \mu_H)$ is also given by a Brownian motion:

$$d\alpha = \alpha(1-\alpha)\frac{\mu_H - \mu_L}{\sigma}\sqrt{\bar{q}} dz.$$

Thus, the belief process has a zero drift and a variance term equal to $\alpha^2(1-\alpha)^2\left(\frac{\mu_H-\mu_L}{\sigma}\right)^2\bar{q}$. This term depends positively on the current degree of uncertainty, on signal-to-noise ratio, and on the total purchased quantity. More specifically, the rate at which the firm and the consumers learn about the product's quality proportional to the total quantity purchased.

2.1 The Buyer's Problem

We model each buyer as making a purchase decision in every period.[3] In a dynamic setting without commitment on the firm's side, buyers have an incentive not to reveal their private information. Since the focus of our analysis is on the dynamics of prices as public information is released, and not on the revelation of private information, we introduce an anonymity constraint on the firm's pricing decisions. In other words, we literally require the firm to post a menu in every period, and do not allow for buyer-specific offers, which keep track of each individual's purchases history. Under this assumption, the buyer's problem in every period simply consists of choosing an item from a menu. More specifically, consider a nonlinear tariff scheme derived from a direct mechanism $p(\theta) = p(q(\theta))$. Denote by $\hat{U}(\alpha, \theta, \theta')$ buyer type θ's utility when holding belief α and purchasing the item designed for buyer type θ'.

$$\begin{aligned} \hat{U}(\alpha, \theta, \theta') &= \mathbb{E}_\mu \left[\mu\theta u\left(q(\theta')\right)\right] - p(\theta') \\ &= m(\alpha)\,\theta u\left(q(\theta')\right) - p(\theta'), \\ \text{where } m(\alpha) &= \alpha\mu_H + (1-\alpha)\,\mu_L. \end{aligned}$$

It follows that each buyer chooses the item that maximizes his total utility $\hat{U}(\alpha, \theta, \theta')$. Therefore, let

$$U(\alpha, \theta) = \max_{\theta'} \hat{U}(\alpha, \theta, \theta').$$

Note that the public signals assumption implies each consumer's participation decision does not affect the amount of information she receives. Similarly, the existence of a continuum of buyers implies each individual quantity choice does not affect the total number of units being consumed at each instant, which determines the precision of the information. Therefore, each consumer's willingness to pay for the good does not reflect any learning consideration. It follows that, from the firm's perspective, the buyers' types are effectively distributed on $[m(\alpha)\,\theta_L, m(\alpha)\,\theta_H]$.

2.2 The Firm's Problem

The Static Benchmark As a static benchmark, we modify the monopolistic screening model (Mussa and Rosen (1978), Maskin and Riley (1984)) to take our model's particular type space into consideration. Denote by $q(\alpha, \theta)$ and $U(\alpha, \theta)$ respectively the quantity and indirect utility schedules offered by the monopolist when the market holds belief α. The firm's static profit function may then be written as

$$\Pi(\alpha, q, U) = \int_{\theta_L}^{\theta_H} \left(m(\alpha) \cdot \theta \cdot u(q(\alpha, \theta)) - c(q(\alpha, \theta)) - U(\alpha, \theta)\right) f(\theta)\, d\theta. \tag{1}$$

[3] Equivalently, one may reinterpret the buyer's as a rental decision. The repeat purchases setting is then equivalent to assuming the customer can easily switch between different rental plans (as in the Netflix example).

For each value of α, the firm's problem is to maximize $\Pi(\alpha, q, U)$ with respect to the functions $q(\alpha, \cdot)$ and $U(\alpha, \cdot)$, subject to the Incentive Compatibility and Individual Rationality constraints.

The Dynamic Problem We now turn to the dynamic version of the problem. Given our assumption $\mu \in \{\mu_L, \mu_H\}$, at each point in time the value of α is a sufficient statistic for the firm's problem. It is therefore natural to use it as a state variable. Using Ito's lemma, the firm's value function is given by

$$rV(\alpha) = \max_{q(\alpha,\cdot),U(\alpha,\cdot)} \left[\Pi(\alpha, q, U) + \frac{1}{2}\bar{q}\left(\alpha(1-\alpha)\frac{\mu_H - \mu_L}{\sigma}\right)^2 V''(\alpha)\right].$$

We then define a learning component as $\Lambda(\alpha)$. More specifically,

$$\Lambda(\alpha) = \frac{1}{2}\left(\alpha(1-\alpha)\frac{\mu_H - \mu_L}{\sigma}\right)^2 V''(\alpha).$$

With this notation, one may write the firm's Bellman equation as

$$\begin{aligned} rV(\alpha) &= \max_{q(\alpha,\cdot),U(\alpha,\cdot)} \left[\Pi(\alpha, q, U) + \bar{q}\Lambda(\alpha)\right] \qquad (2) \\ &= \max_{q(\alpha,\cdot),U(\alpha,\cdot)} \int_{\theta_L}^{\theta_H} \left(m(\alpha)\theta u(q(\alpha,\theta)) - c(q(\alpha,\theta)) - U(\alpha,\theta) + \Lambda(\alpha) q(\alpha,\theta)\right) f(\theta)\, d\theta. \end{aligned}$$

This maximization problem is again subject to the Incentive Compatibility and Individual Rationality constraints.

3 Equilibrium

We now review the solution to the static problem and derive some properties of this solution that will be useful later on. Since buyers solve a static decision problem in each period, when offered a menu $(q(\alpha,\theta), U(\alpha,\theta))$, the necessary and sufficient conditions for incentive compatibility are given by:

$$\frac{\partial U(\alpha,\theta)}{\partial \theta} = m(\alpha) u(q(\alpha,\theta)), \qquad \text{(IC1)}$$

$$\frac{\partial q(\alpha,\theta)}{\partial \theta} \geq 0, \text{ for all } \alpha \text{ and } \theta. \qquad \text{(IC2)}$$

Furthermore, we normalize outside options to zero for all buyer types and all beliefs α. Therefore, the individual rationality constraint is given by

$$U(\alpha,\theta) \geq 0 \text{ for all } \alpha, \theta. \qquad \text{(IR)}$$

Substituting (IC1) and (IR) into the profit function (1), one may rewrite the firm's expected profits as follows:

$$\Pi(\alpha,q,U)=\int_{\theta_L}^{\theta_H}\left(m(\alpha)\left(\theta-\frac{1-F(\theta)}{f(\theta)}\right)u(q(\alpha,\theta))-c(q(\alpha,\theta))-U(\alpha,\theta)\right)f(\theta)\,d\theta.$$

Pointwise maximization then yields the following first order condition for quantity provision:

$$c'(q(\alpha,\theta))=m(\alpha)\left(\theta-\frac{1-F(\theta)}{f(\theta)}\right)u'(q(\alpha,\theta)).$$

Remark 1 (The Effect of (Good) News) .
The static quantity provision $q(\alpha,\theta)$ is increasing in α for all θ.

This result states that, in the static case, positive signals about the product's quality increase efficiency and lead to larger quantity provision for each type. It will be useful to contrast this result with the solution to the firm's dynamic problem.

In the dynamic problem, an analogous substitution of the incentive compatibility and individual rationality constraints (IC1) and (IR) into the objective function (2) allows us to write the firm's value function as

$$rV(\alpha)=\max_{q(\alpha,\cdot),U(\alpha,\cdot)}\int_{\theta_L}^{\theta_H}(m(\alpha)\phi(\theta)u(q(\alpha,\theta))-c(q(\alpha,\theta))-U(\alpha,\theta)+\Lambda(\alpha)q(\alpha,\theta))f(\theta)\,d\theta,$$

where

$$\phi(\theta)=\theta-\frac{1-F(\theta)}{f(\theta)}$$

represents type θ's virtual valuation. The first order condition for the optimal quantity schedule $q(\alpha,\theta)$ may then be written as

$$c'(q(\alpha,\theta))=m(\alpha)\phi(\theta)u'(q(\alpha,\theta))+\Lambda(\alpha). \tag{3}$$

This first order condition differs from the statical problem's only for the presence of the learning component $\Lambda(\alpha)$. However (as we show in the following sections), $\Lambda(\alpha)$ is typically not monotonic in α. This implies that, in the dynamic solution, the RHS of (3) - and therefore the quantity provision - need not be monotonic in the beliefs α.

Whenever condition (3) may be solved explicitly for $q(\alpha,\theta)$ (such as in the following functional form examples), it is possible to then substitute the optimal policy back into the value function and obtain a full characterization of the solution.

4 Linear-Quadratic Model

In this section, we specify the model to the case of linear utility and quadratic costs. More precisely, we adopt the Mussa and Rosen (1978) functional forms $u(q)=q$ and $c(q)=\frac{1}{2}q^2$. The convex costs assumption is generally used in models which interpret q as product quality. However, we stress that this formulation is also equivalent to a model with constant marginal costs and a non separable gross utility function of the form $u(\alpha,\theta,q)=m(\alpha)\theta q-\frac{1}{2}q^2$. This formulation is more realistic for the quantity interpretation given in this paper.

Under these functional form assumptions, for all types θ served in equilibrium, the first order condition (3) may be written as

$$q(\alpha,\theta)=m(\alpha)\phi(\theta)+\Lambda(\alpha). \tag{4}$$

When the entire market is covered, substituting back into the objective function, one obtains:

$$\begin{aligned} rV(\alpha) &= \int_{\theta_L}^{\theta_H}\frac{1}{2}\left(m(\alpha)\phi(\theta)+\Lambda(\alpha)\right)^2 f(\theta)\,d\theta \\ &= \frac{1}{2}m^2(\alpha)\,\mathbb{E}_\theta\left[\phi^2\right]+m(\alpha)\,\mathbb{E}_\theta[\phi]\,\Lambda(\alpha)+\frac{1}{2}\Lambda^2(\alpha). \end{aligned} \tag{5}$$

It can be verified that the first term in (5) is the expression for the firm's equilibrium expected profits in the static case. We denote this static benchmark by $\Pi^*(\alpha)$:

$$\Pi^*(\alpha)=\frac{1}{2}m^2(\alpha)\,\mathbb{E}_\theta\left[\phi^2\right]. \tag{6}$$

We can then show that the firm's dynamic problem with discounting reduces to an ordinary differential equation. In fact, solving equation (5) for $\Lambda(\alpha)$, one obtains

$$\Lambda(\alpha)=-m(\alpha)\,\mathbb{E}_\theta[\phi]+\sqrt{\left(m(\alpha)\,\mathbb{E}_\theta[\phi]\right)^2+2\left(rV(\alpha)-\Pi^*(\alpha)\right)}.$$

Then, using the definition $\Lambda(\alpha)=\frac{1}{2}\left(\alpha(1-\alpha)\frac{\mu_H-\mu_L}{\sigma}\right)^2V''(\alpha)$, one obtains the following second-order (nonlinear) differential equation for the firm's value function

$$V''(\alpha)=2\left(-m(\alpha)\,\mathbb{E}_\theta[\phi]+\sqrt{m^2(\alpha)\,\mathbb{E}_\theta[\phi]^2+2\left(rV(\alpha)-\Pi^*(\alpha)\right)}\right)\left(\alpha(1-\alpha)\frac{\mu_H-\mu_L}{\sigma}\right)^{-2}$$

with boundary conditions given by

$$V(0) = \frac{1}{r}\Pi^*(0)\text{, and} \tag{7}$$

$$V(1) = \frac{1}{r}\Pi^*(1). \tag{8}$$

Unfortunately, this differential equation does not have an analytical solution. As anticipated, we therefore adopt a different procedure. In particular, we analyze the undiscounted version of the firm's problem by adopting the *strong long run average* criterion.[4] This approach provides more

[4] This criterion was pioneered by Ramsey (1928), then analyzed in detail by Dutta (1991) and used in applications, among others, by Bergemann and Valimaki (1997), Bolton and Harris (2000), and Bergemann and Välimäki (2002).

tractability, by allowing us to solve a static optimization problem, while preserving the qualitative properties of the dynamic solution. It may be synthesized as follows: in the lung run, we know beliefs will convergence (either to μ_L or to μ_H). In the absence of discounting, given current belief α, the expected long-run per-period average payoff is given by

$$rV(\alpha) \to v(\alpha) = \alpha\Pi^*(1) + (1-\alpha)\Pi^*(0).$$

Furthermore, many policy functions may attain the average value $v(\alpha)$ in the long run, independently of their finite time behavior. The main contribution of Dutta (1991) is to prove that the limiting (for $r \to 0$) optimal policy function maximizes the undiscounted stream of payoffs, net of their long run averages:

$$V(\alpha) = \sup_{q(\alpha(t),\cdot),U(\alpha(t),\cdot)} \lim_{T\to\infty} \mathbb{E}\left[\int_0^T \left[\Pi(\alpha(t), q, U) - v(\alpha(t))\right] dt\right].$$

Furthermore, the strong long run average criterion allows the following recursive representation:

$$\begin{aligned} v(\alpha) \quad &= \max_{q(\alpha,\cdot),U(\alpha,\cdot)} \left[\Pi(\alpha, q, U) + \bar{q}\Lambda(\alpha)\right] \\ &\quad \text{s.t. (IC1), (IC2), and (IR).} \end{aligned}$$

This equation is identical to the one in the discounted version of the problem, except for the absence of $V(\alpha)$ from the left hand side. Therefore, the first order condition may still be written as in (4). Substituting into the objective function one obtains:

$$v(\alpha) = \frac{1}{2}m^2(\alpha)\mathbb{E}_\theta\left[\phi^2\right] + m(\alpha)\mathbb{E}_\theta[\phi]\Lambda(\alpha) + \frac{1}{2}\Lambda^2(\alpha).$$

The learning component $\Lambda(\alpha)$ is then given by

$$\Lambda(\alpha) = -m(\alpha)\mathbb{E}_\theta[\phi] + \sqrt{m^2(\alpha)\mathbb{E}_\theta[\phi]^2 + 2(v(\alpha) - \Pi^*(\alpha))}. \tag{9}$$

Note that $\Lambda(\alpha) \neq 0$, unless $\Pi^*(\alpha)$ is linear. Moreover, as long as $\Lambda(\alpha) \neq 0$, the solution to the undiscounted problem is different from the static optimum. We can now summarize these derivations in the following proposition.

Proposition 1 (Equilibrium Quantity Provision) .

1. *The undiscounted optimal (second-best) quantity provision in the linear-quadratic model is given by*

$$q(\alpha,\theta) = m(\alpha)(\phi(\theta) - \mathbb{E}_\theta[\phi]) + \sqrt{2v(\alpha) - m^2(\alpha)\operatorname{Var}[\phi]}. \tag{10}$$

2. *The undiscounted optimal (second-best) quantity provision is always greater than the static (second-best) quantity provision.*

3. *The undiscounted first-best optimal quantity provision is given by:*

$$q(\alpha,\theta) = m(\alpha)(\theta - \mathbb{E}[\theta]) + \sqrt{2v(\alpha) - m^2(\alpha)\,\mathrm{Var}[\theta]}.$$

Proof. See the Appendix. ■

The main implication of this result is that the effect of market beliefs on quantity depends on the consumer's type θ. Types with a virtual valuation above the average will benefit more from an improvement of the market belief than those with below-average virtual valuations. This observation is largely driving the nonmonotonicity of quantity provision, a result which we show in the next proposition.

Proposition 2 (Nonmonotonic Quantity Provision) .

1. *A sufficient condition for quantity provision $q(\alpha,\theta)$ to be increasing in α is $\phi(\theta) > \mathbb{E}_\theta[\phi]$ and $\alpha < \frac{1}{2}$.*

2. *Quantity provision $q(\alpha,\theta)$ is nonmonotonic in α for all types $\theta \leq \tilde{\theta}$, where $\tilde{\theta}$ satisfies*

$$\phi\left(\tilde{\theta}\right) = \frac{\mathbb{E}_\theta\left[\phi^2\right]\left(1 - \frac{\mu_L}{\mu_H}\right)}{2\mathbb{E}_\theta[\phi]}.$$

3. *A sufficient condition for the first best quantity provision to be increasing in α for all θ is*

$$\mathrm{Var}[\theta] \leq 2\mathbb{E}[\theta]\,\theta_L \frac{\mu_H}{\mu_H - \mu_L} - \mathbb{E}[\theta]^2.$$

Proof. See the Appendix. ■

This result shows that the set of types receiving a nonmonotonic quantity supply is increasing in the ratio μ_H/μ_L. This ratio constitutes a measure of how relevant the inherent product quality is in determining buyers' utility. More generally, as good news about product quality are announced, gains from trade increase, since consumers are willing to pay more for each unit. This effect is stronger for high consumer types, as they benefit the most from a quality increase. The opposite is clearly true for bad news. We will refer to this as the "mean" effect. At the same time, as the market belief moves away from $\frac{1}{2}$, the value of information attached to each unit is reduced. In other words, once the consumers are almost certain they know the true μ, it will be hard to influence their belief by supplying additional units. This variance-type effect affects all types in the same way, since each type's contribution to the informational content of purchases proportionally to her consumption. These two effects characterize both the first-best and the second-best policies. However, under adverse selection, the mean effect (increased valuations) bears an additional cost to the seller, given by the need to provide increasing informational rents. Remember that in the two-type static adverse selection model, the high type's rent is given by $U(\theta_H) = U(\theta_L) + (\theta_H - \theta_L)\,q_L$. The equivalent formulation for this model would be

$U(\theta_H) = U(\theta_L) + m(\alpha)(\theta_H - \theta_L) q_L$. In other words, positive signals provide the seller an incentive to increase the distortions in the quantity consumed by lower types. As seen from the equilibrium quantity provision, this effect is stronger the higher the variance of the distribution of types. The balance of the mean, variance and informational rent effects determines a set of types for which quantity provision is nonmonotonic in α. These types consume the largest quantities for intermediate values of α, where the value of information is highest and rents are not too costly. Finally, the last result in Proposition 2 and Remark 1 on the static quantity provision suggest that, if the distribution of types θ has a low enough variance (in other words, if product quality has an important enough effect on the type space), then both experimentation and asymmetric information are necessary elements for the nonmonotonicity result.

Figure 2 displays the quantity levels supplied to each buyer type, as a function of the market belief α. Types θ are assumed to be uniformly distributed over $[1, 2]$, while $\mu_L = 1$ and $\mu_H = 5$. The critical type $\tilde{\theta}$ is therefore given by $\tilde{\theta} = \frac{7}{6}$.

Figure 2: Equilibrium Quantity Provision

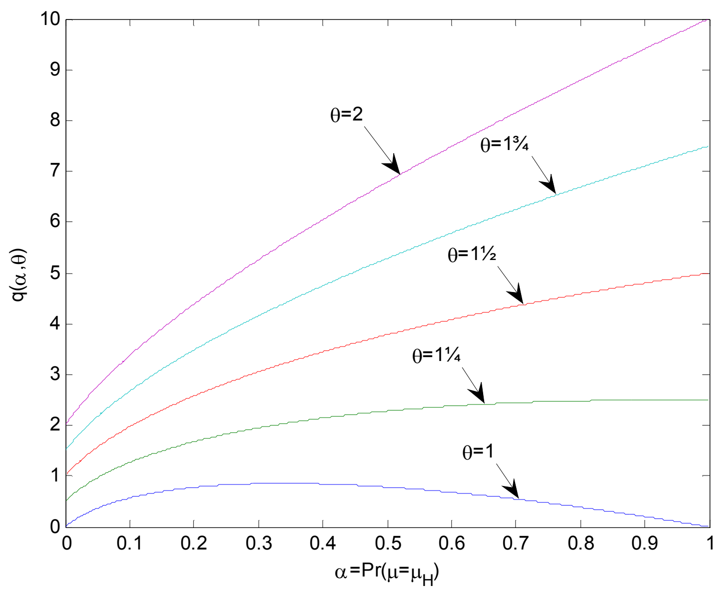

While Proposition 2 analyzes the evolution of quantity schedules as a function of the market beliefs, the next two results are more informative of the expected time pattern of equilibrium consumption.

Proposition 3 (Concave Quantity Provision) .
In the linear-quadratic model, the limiting optimal quantity provision schedule is concave in α.

Proof. The first term in expression (10) is linear in α. The term inside the square root is concave, since $v(\alpha)$ is linear, $m(\alpha)$ is convex, and $\text{Var}[\phi] \geq 0$. Therefore, $q(\cdot, \theta)$ is an increasing, concave function of a concave function of α, and it is itself concave. ■

Note that the concavity of quantity provision depends positively on the variance of the distribution of virtual valuations $\phi(\theta)$. More importantly, the concavity of $q(\alpha,\theta)$ immediately allows to conclude that the expected variation in quantity supply ($\mathbb{E}[dq(\alpha,\theta)]$) is negative, since we know $\mathbb{E}[d\alpha]=0$. This means we can expect the supplied quantity levels to decrease over time. In contrast to the absolute quantity provision, we can obtain the following result about the relative evolution of menu items by inspection of the first order condition (4).

Remark 2 *Differences between the quantity levels of any two items in the optimal menu are given by* $q(\alpha,\theta)-q(\alpha,\theta')=m(\alpha)(\phi(\theta)-\phi(\theta'))$. *Hence they are linear in* α *and constant in expectations.*

We now turn to the analysis of the equilibrium tariffs. The total tariff charged by the monopolist for $q(\alpha,\theta)$ units of the product is obtained from the expression for the informational rents (IC1), and may be written as follows:

$$p(\alpha,\theta)=m(\alpha)\,\theta q(\alpha,\theta)-U(\alpha,\theta).$$

Using expression (10) for the equilibrium quantity provision, we can express $p(\alpha,\theta)$ as

$$\begin{aligned} p(\alpha,\theta) &= m(\alpha)\,\theta\,(m(\alpha)\,\phi(\theta)+\Lambda(\alpha))-m(\alpha)\int_{\theta_L}^{\theta}(m(\alpha)\,\phi(s)+\Lambda(\alpha))\,ds \\ &= m^2(\alpha)\left(\theta_L\phi(\theta_L)+\int_{\theta_L}^{\theta}s\phi'(s)\,ds\right)+m(\alpha)\,\theta_L\Lambda(\alpha), \end{aligned}$$

where $\Lambda(\alpha)$ is given by (9). We can then prove the following result.

Proposition 4 (Equilibrium Prices) .
In the linear-quadratic model, the second derivative of the price function, for a given α, *is increasing in* θ.

Proof. Since $\phi(\theta)$ is an increasing function, the coefficient multiplying $m^2(\alpha)$ is positive. Furthermore, the second term does not depend on θ. ■

In other words, for a given α_0, if $p(\theta_0;\alpha_0)$ is convex, so is $p(\alpha,\theta)$ for all $\theta>\theta_0$. When the distribution of types is given, we can identify a threshold function $\hat{\theta}(\alpha)$ such that, given α, all types $\theta>\hat{\theta}$ are offered convex price schedules. When this is the case, we can expect total prices to increase over time ($\mathbb{E}[dp]>0$) for high types, even though quantity is expected to decrease.

We again consider an example with types θ uniformly distributed over $[1,2]$, $\mu_L=1$ and $\mu_H=5$. Figure 3(a) displays the equilibrium prices for low types θ, as a function of α. These types receive concave prices, while higher types (in figure 3(b)) receive convex prices.

The analysis of unit prices allows us to better evaluate the experimentation and rent extraction

Figure 3: Equilibrium Prices

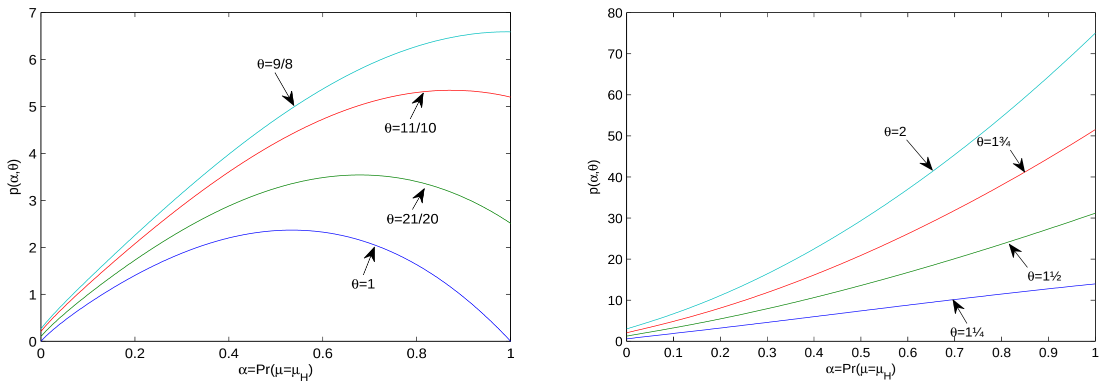

(a) (b)

motives. Unit prices are given by:

$$\frac{p(\alpha,\theta)}{q(\alpha,\theta)} = \frac{m^2(\alpha)\left(\theta_L\phi(\theta_L) + \int_{\theta_L}^{\theta} s\phi'(s)\,ds\right) + m(\alpha)\,\theta_L\Lambda(\alpha)}{m(\alpha)\,\phi(\theta) + \Lambda(\alpha)},$$

with $\Lambda(\alpha)$ again given by (9). The next result follows directly from the previous propositions:

Corollary 1 *For all types θ for which $p(\alpha,\theta)$ is convex in α, unit prices $\frac{p(\alpha,\theta)}{q(\alpha,\theta)}$ are also convex in α.*

Proof. Since quantity is always a positive concave function of α, the convexity of the price function implies that unit prices $\frac{p(\alpha,\theta)}{q(\alpha,\theta)}$ are the product of two convex functions. ■

We now analyze the joint dynamics of quantity provision and prices, to determine the evolution of the menus of contracts offered by the firm. Figure 4 displays the price-quantity bundles for several values of α.

The arrival of good news increases unit prices uniformly and extends the range of offered products at the top. Conversely, the variety of products offered at the bottom only increases as uncertainty is revealed, that is, when α approaches either zero or one. In fact, as the value of experimentation decreases, there is less of a need to increase quantity provision for high types and to consequently distort consumption downwards for lower types. This effect is best seen in Figure 5, which details the dynamics of the "lower tail" of the equilibrium menus.

In terms of predicted time patterns, increasing marginal costs and linear utility[5] suggest that successful product lines should be characterized by increasing differences in prices and increasing variety: low types receive smaller bundles and high types receive increasing numbers of units. Moreover, as consumers learn about the quality of the product, the differences between the unit prices charged on different bundles become larger. Therefore, the model predicts the firm will charge increasing markups on the most profitable bundles.

[5] Or, equivalently, constant marginal costs and quadratic utilities.

Figure 4: Equilibrium Nonlinear Tariffs

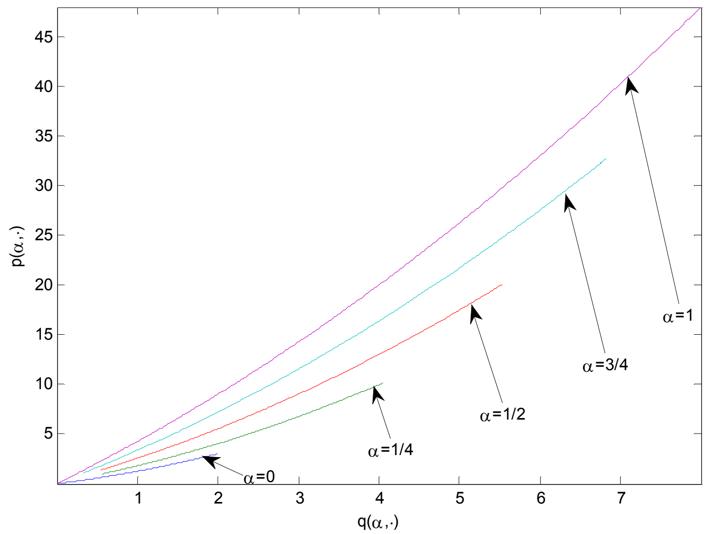

Extension: Variations in Market Coverage We now relax the full market coverage assumption and extend the analysis to the case in which, for any value of μ, it is not optimal for the monopolist to cover the entire market. At this stage, we must introduce a functional form assumption on the distribution of types θ, for tractability reasons. For ease of exposition, we assume that types θ are uniformly distributed on $[\theta_L, \theta_H]$. Under this functional form assumption, the optimal policy rule (from the first order conditions) is unchanged.

$$\begin{aligned} q(\alpha, \theta) &= m(\alpha)\,\phi(\theta) + \Lambda(\alpha) \\ &= m(\alpha)\,(2\theta - \theta_H) + \Lambda(\alpha)\,. \end{aligned}$$

The only difference is the presence of a lowest type $\theta^*(\alpha)$ being served in equilibrium. The firm's value function in equilibrium is given by

$$\begin{aligned} rV(\alpha) &= \int_{\theta^*(\alpha)}^{\theta_H} \frac{1}{2}\left(m(\alpha)\,\phi(\theta) + \Lambda(\alpha)\right)^2 f(\theta)\,d\theta \\ &= \int_{\theta^*(\alpha)}^{\theta_H} \frac{1}{2}\left(m(\alpha)\,(2\theta - \theta_H) + \Lambda(\alpha)\right)^2 \frac{1}{\theta_H - \theta_L} d\theta, \end{aligned}$$

where the critical type is determined through the equation $q(\theta^*(\alpha), \alpha) = 0$. Therefore,

$$\theta^*(\alpha) = \frac{1}{2}\left(\theta_H - \frac{\Lambda(\alpha)}{m(\alpha)}\right).$$

Figure 5: Equilibrium Nonlinear Tariffs (detail)

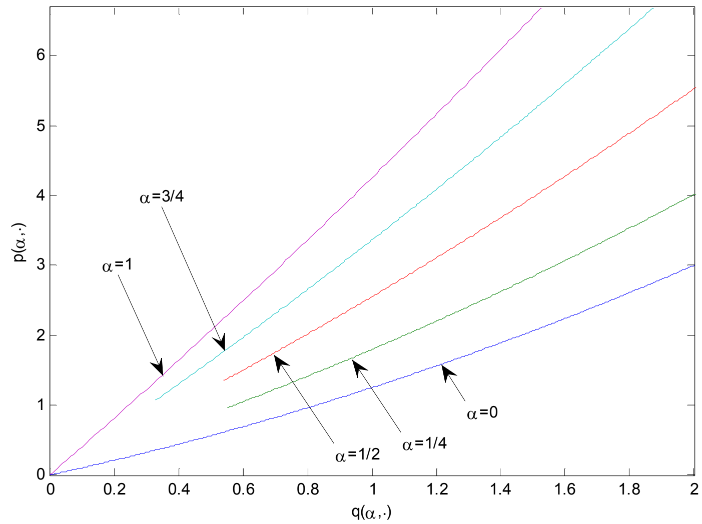

Simple algebra yields the following expression for equilibrium profits and for the learning component

$$
\begin{aligned}
rV(\alpha) &= \frac{(m(\alpha)\theta_H + \Lambda(\alpha))^3}{12m(\alpha)(\theta_H - \theta_L)}, \\
\Lambda(\alpha) &= -m(\alpha)\theta_H + (12m(\alpha)(\theta_H - \theta_L) rV(\alpha))^{\frac{1}{3}}.
\end{aligned}
$$

We again consider the undiscounted case and apply the long-run average criterion. We find that the qualitative features of the solution are unchanged from the full market coverage case. It is worthwhile noticing how the value of information generally leads to more types being served, as compared to the static benchmark. This effect is clearly strongest when uncertainty is most severe. Figure 6 displays the quantity provision schedules $q(\alpha;\theta)$ for $\theta \in \left\{\frac{1}{3}, \frac{1}{2}, \frac{3}{5}, \frac{3}{4}\right\}$.

Figure 6: Equilibrium Quantity Provision

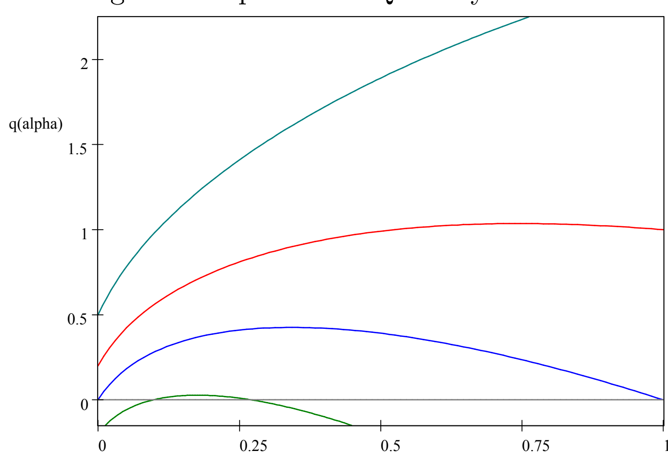

The parameter values are $(\theta_L = 0, \theta_H = 1, \mu_H = 5, \mu_L = 1)$. Note that the quantity supplied to $\theta = \frac{1}{3}$ would be equal to zero for all values of α in the static case. In the dynamic model, $q\left(\alpha, \frac{1}{3}\right)$ intersects the horizontal axis twice, meaning this type is served for intermediate values of α - when the value of information is highest.

Figure 7 shows the evolution of market coverage as a function of α. Market coverage is higher for intermediate values of α, where information is more valuable. It does not achieve its maximum at $\alpha = \frac{1}{2}$ due to the increasing cost of providing quantity to low types, as α increases. The increased distortion in the cutoff type's consumption level determines a decline in market coverage even for values of α lower than $\frac{1}{2}$.

Figure 7: Fraction of the Market Covered

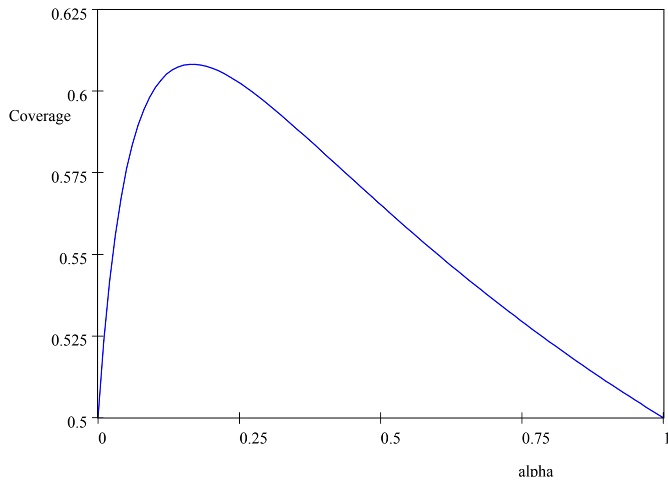

5 Constant Price Elasticity Model

We now analyze a model with a different demand and cost function specification. In particular, we adopt the Maskin and Riley (1984) model with constant price elasticity of demand and constant marginal cost. More specifically, we consider the case of $\eta = 2$, which corresponds to square-root utility. These assumptions provide a better fit our leading examples (DVDs and software licenses), in which duplication costs are small and constant (delivery) or even decreasing. Even in its static version, this model differs sharply from the Mussa and Rosen (1978) quality pricing framework. In fact, decreasing willingness to pay for additional units, combined with constant marginal costs, induces the firm to offer substantial quantity discounts for large purchases.[6] This contrasts with the linear quadratic case. In that model, increasing production costs and constant per unit willingness to pay determine an increasing pattern of unit prices. The comparison is made clear in Figure 8, which displays the Mussa and Rosen (1978) and the Maskin and Riley (1984) allocation (the latter with cost parameter $c = \frac{1}{\sqrt{2}}$) for the case of $\mu = 1$ and uniformly distributed types θ over the

[6] Quantity discounts for high valuation buyers (*i.e.* those who make large purchases) is the main finding of Maskin and Riley (1984).

interval $[1,2]$.

Figure 8: Nonlinear Tariffs - Static Case

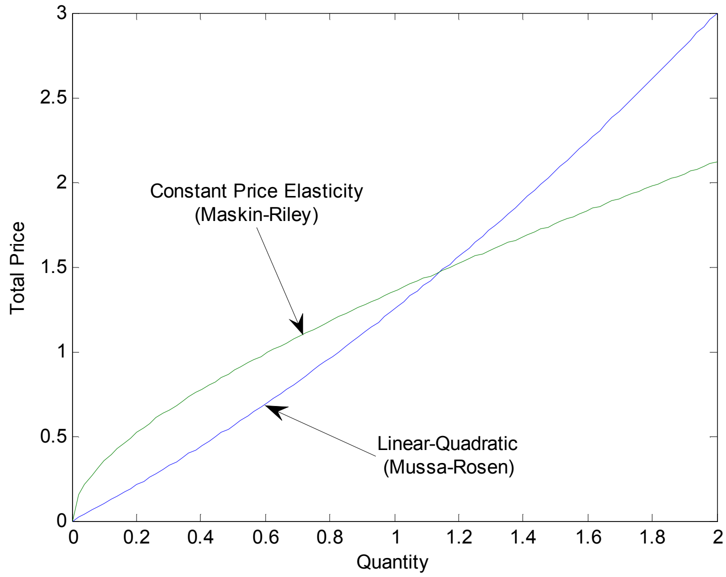

Throughout the rest of this section, we assume that the entire market is covered,[7] and we skip those derivations that just replicate the ones of the previous section.

The functional form and full coverage assumptions allow us to solve first order condition (3) and to obtain an explicit solution for the optimal quantity provision:

$$q(\alpha,\theta)=\left(\frac{m(\alpha)\,\phi(\theta)}{2\,(c-\Lambda(\alpha))}\right)^2. \tag{11}$$

Substituting the expression for $q(\alpha,\theta)$ back into the objective functions, one may write

$$\begin{aligned} \Pi^*(\alpha) &= \frac{m^2(\alpha)\,\mathbb{E}_\theta\left[\phi^2\right]}{4c}, \\ rV(\alpha) &= \frac{m^2(\alpha)\,\mathbb{E}_\theta\left[\phi^2\right]}{4\,(c-\Lambda(\alpha))} \\ &= \Pi^*(\alpha)\,\frac{c}{c-\Lambda(\alpha)}. \end{aligned} \tag{12}$$

In the discounted case, solving (12) for $\Lambda(\alpha)$, one obtains

$$\Lambda(\alpha)=c\left(1-\frac{\Pi^*(\alpha)}{rV(\alpha)}\right).$$

Substituting the definition of $\Lambda(\alpha)$, one obtains a second-order nonlinear ordinary differential

[7] An extension to allow imperfect market coverage is possible along the lines of the previous section.

equation for the firm's value function

$$V''(\alpha) = 2c\left(1 - \frac{\Pi^*(\alpha)}{rV(\alpha)}\right)\left(\alpha(1-\alpha)\frac{\mu_H - \mu_L}{\sigma}\right)^{-2}$$

with boundary conditions (7) and (8). As for the linear-quadratic model, we consider the undiscounted case ($r \to 0$) and we adopt the strong long run average criterion to derive an analytical solution. Under this criterion, we obtain the following characterization for the second-best quantity provision

Proposition 5 (Equilibrium Quantity Provision) .
The equilibrium quantity provision in the (undiscounted) constant price elasticity model is given by

$$q(\alpha,\theta) = \frac{\phi^2(\theta)\left(\alpha\mu_H^2 + (1-\alpha)\mu_L^2\right)^2}{4c^2 m^2(\alpha)}.$$

Proof. From the undiscounted equivalent of condition (12), one obtains the following expression for the learning component:

$$\Lambda(\alpha) = c\left(1 - \frac{\Pi^*(\alpha)}{v(\alpha)}\right) = c\left(1 - \frac{m^2(\alpha)}{\alpha\mu_H^2 + (1-\alpha)\mu_L^2}\right).$$

Substitution into condition (11) delivers the result. ∎

Figure 9 displays the equilibrium quantity supply schedules, as a function of α. The chosen parameter values are $\{\theta_L = 2, \theta_H = 3, \mu_H = 2 + \sqrt{3}, \mu_L = 1, c = 1\}$, with types θ distributed uniformly.

Figure 9: Equilibrium Quantity Provision

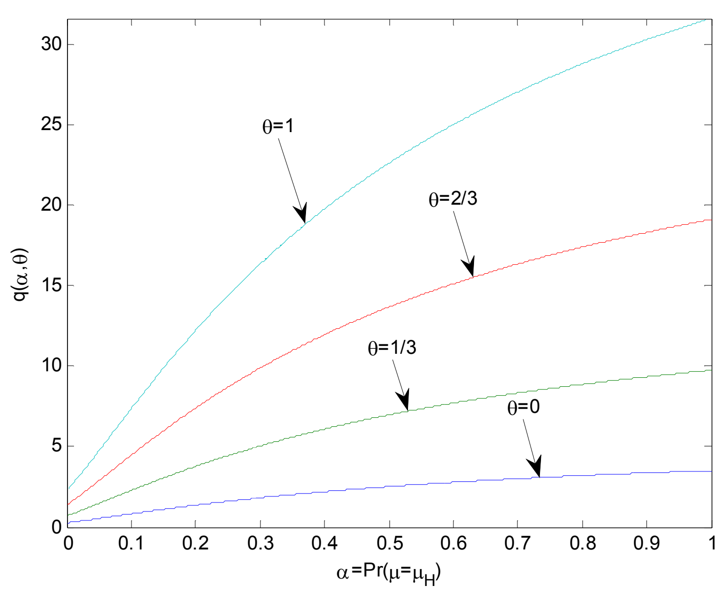

The solution to the constant price elasticity model shares some of the qualitative features of the linear quadratic model. However, the two models differ substantially in terms of their predictions on the monotonicity of quantity supply (as a function of α). In fact, in the constant price elasticity model, quantity is directly proportional to each type's virtual valuation. It follows that variations in α either increase or decrease quantity provision for all types θ. The following results highlight some of the qualitative properties of this model.

Proposition 6 (Monotonicity, Concavity) .

1. *In the constant price elasticity model, quantity provision is strictly increasing in α, for all θ and all α.*
2. *There exists a critical value α^* above which the limiting optimal quantity provision $q(\alpha, \theta)$ is concave in α for all θ.*

Proof. Consider the first derivative of the quantity provision schedule

$$\frac{\partial q(\alpha, \theta)}{\partial \alpha} \propto \frac{(\mu_H - \mu_L)\left(\alpha \mu_H^2 + (1-\alpha)\mu_L^2\right)}{(\alpha \mu_H + (1-\alpha)\mu_L)^3} > 0.$$

Then consider the second derivative:

$$\frac{\partial^2 q(\alpha, \theta)}{(\partial \alpha)^2} \propto \frac{\mu_H \mu_L - 2\left(\alpha \mu_H^2 + (1-\alpha)\mu_L^2\right)}{m^4(\alpha)}.$$

Therefore,

$$\frac{\partial^2 q(\alpha, \theta)}{(\partial \alpha)^2} \leq 0 \iff \alpha \geq \alpha^* = \frac{\mu_L(\mu_H - 2\mu_L)}{2\left(\mu_H^2 - \mu_L^2\right)}.$$

■

The next two results are immediate, yet useful, consequences of the previous proposition.

Corollary 2 *A sufficient condition for global concavity of $q(\alpha, \theta)$ is $\mu_H \leq 2\mu_L$. Moreover, the critical value α^* is bounded from above by $\bar{\alpha}^* = \frac{1}{2} - \frac{1}{4}\sqrt{3} \approx 0.067$ for all $\mu_H > \mu_L$.*

This result, roughly speaking, means quantity provision is convex only in "extreme bad news" cases. The next result follows from the separable (multiplicative) form of the equilibrium quantity provision in α and θ.

Corollary 3 *Absolute differences in the quantity levels of two items in a menu are given by $\left|q(\alpha, \theta) - q(\theta', \alpha)\right| = \left(\left|\phi^2(\theta) - \phi^2(\theta')\right|\right)(2cm(\alpha))^{-2}\left(\alpha \mu_H^2 + (1-\alpha)\mu_L^2\right)^2$. Hence they are concave (convex) in α (depending on $(\alpha - \alpha^*)$).*

As in the linear quadratic model, the optimal quantity provision is affected by considerations of efficiency, informational rents and extraction of information. In the square-root linear model, each type θ's virtual valuation enters multiplicatively, not additively, in the expression for $q(\alpha, \theta)$.

As a result, quantity is increasing in α for all θ, though at different rates. This constitutes an important difference between the two models. However, as in the linear-quadratic model, we find that quantity provision is expected to decrease over time, reflecting an experimentation pattern. Furthermore, the differences between the quantities supplied to different types are predicted to shrink: this implies we expect to see more similar items within a product line as time progresses. We now turn to the analysis of equilibrium prices.

Corollary 4 *In the constant price elasticity model, total equilibrium prices are linear in α and they are given by*

$$\begin{aligned} p(\alpha,\theta) &= m(\alpha)\,\phi(\theta)\sqrt{q(\alpha,\theta)} - m(\alpha)\int_{\theta_L}^{\theta}\sqrt{q(\alpha,\theta)}d\theta \\ &= \frac{1}{2c}\left(\theta_L\phi(\theta_L) + \int_{\theta_L}^{\theta} s\phi'(s)\,ds\right)\left(\alpha\mu_H^2 + (1-\alpha)\,\mu_L^2\right). \end{aligned}$$

Intuitively, since prices are linear and quantities concave, we can expect consumers to pay more per unit as time progresses. Figure 10 confirms this intuition. It displays the behavior of the unit prices charged to several consumer types, as a function of α. The parameter values are again given by $\left\{\theta_L = 2, \theta_H = 3, \mu_H = 2+\sqrt{3}, \mu_L = 1, c = 1\right\}$, with types θ distributed uniformly.

Figure 10: Equilibrium Unit Prices

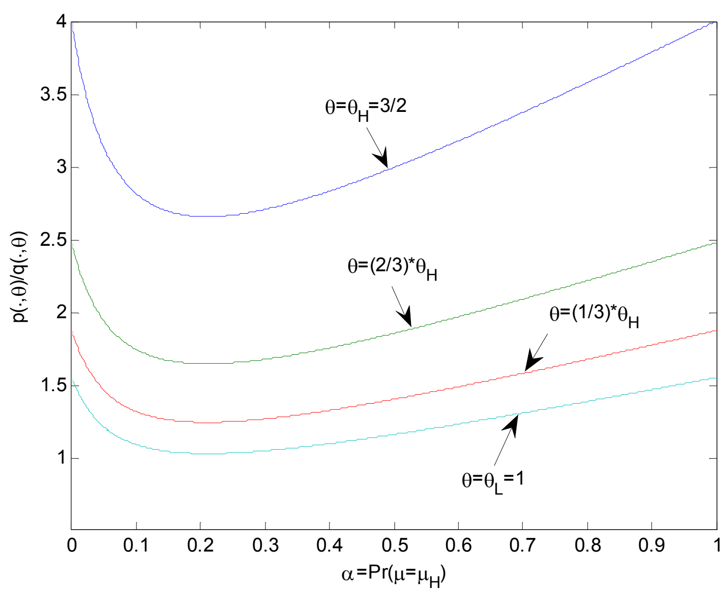

We can now state the final result of this section - a characterization of the behavior of unit prices - in the following proposition.

Proposition 7 (Unit Prices) .

1. *In the constant price elasticity model, the static equilibrium unit prices do not depend on market beliefs* α.

2. *In the undiscounted optimum, unit prices are lower than in the static case for all* $\alpha \in (0, 1)$.

3. *In the undiscounted optimum, unit prices are decreasing in* α *for all* θ *if and only if* $\alpha \leq \hat{\alpha} = \frac{\mu_L}{\mu_H + \mu_L}$

4. *In the undiscounted optimum, unit prices are convex in* α, *for all* α *and* θ.

Proof. See the Appendix. ■

These results inform us of the way the optimal screening contracts are translated into pricing policies. Consumers are offered quantity discounts whenever information is most valuable, while more rents are extracted as the product is increasingly believed to be of good quality. This finding is in sharp contrast with the static case, in which unit prices do not depend on α. In terms of predicted time pattern, unit prices are expected to increase over time. Thus, our model predicts that firms will adopt introductory pricing policies to diffuse information early on, and will subsequently increase prices, once the market beliefs over the product's quality have improved. This practice differs sharply from the information diffusion and rent extraction techniques found in the linear-quadratic model. In that setting, both experimentation and rent extraction were achieved through the quantity provision schedules.

Figure 11 summarizes the equilibrium of the constant price elasticity model, by displaying the menus of contracts offered by the firm.

Figure 11: Equilibrium Nonlinear Tariffs

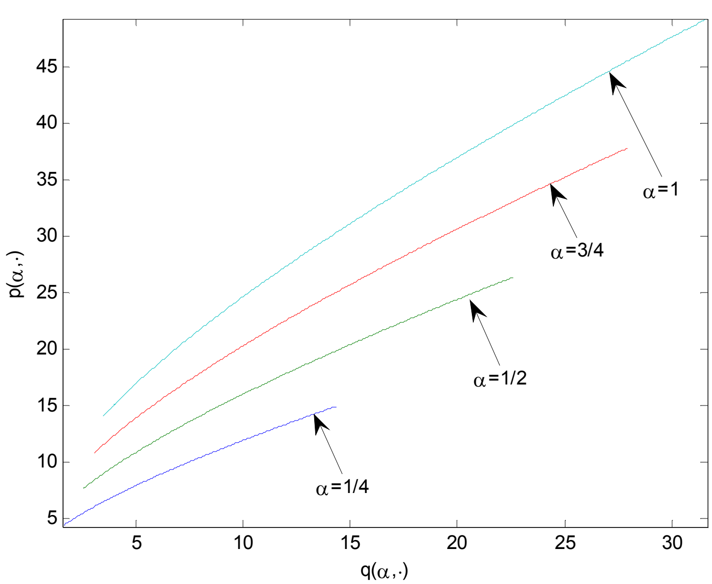

6 Conclusions and Further Research

This paper has analyzed a monopoly dynamic screening problem for experience goods, where the amount of information obtained by the market is proportional to the total quantity consumed in each period. It has focused on the undiscounted version of the firm's problem, in order to provide closed-form solutions that are at the same time tractable and informative of the qualitative properties of the optimal contracts, as well as of their predicted evolution through time.

In a linear utility - quadratic cost model, we have shown that quantity is nonmonotonic and concave in the market beliefs. The nonmonotonicity property is a consequence of the interaction of the learning component with the informational rents. Under conditions on the distribution of the buyer's private information, removing either one of these components (*i.e.* considering the dynamic first-best or the static second-best benchmarks) restores increasing quantities for each type. Firms supply additional units to the market in order to induce more experimentation, in particular when information is valuable and rents are not too high. Upon arrival of sufficiently positive signals, the firm begins extracting more rents by reducing quantities for the lowest types more than the corresponding prices. In terms of predicted time-patterns, quantities are expected to decline, prices to increase, and differences between menu items to remain constant over time.

In a model with constant price elasticity of demand and linear production cost, quantity is again concave but always increasing in the market belief. Conversely, unit prices are nonmonotonic in α, demonstrating how the firm uses introductory volume discounts to obtain more information. Similarly, as the expected product quality increases, the firm extracts more rents by increasing both quantities and unit prices for all types. In terms of predictions, quantity is expected to decrease over time, as well as differences between menu items. In contrast, prices remain constant in expectation, which implies that unit prices must increase over time.

The future developments of this paper include a discussion of the possibility of adopting this model to explain the evolution of nonlinear tariffs from some case studies. Examples are given by the movie rentals, satellite television, and possibly enterprise software industries. From the theoretical point of view, this model naturally lends itself to incorporate the strategic interaction between a well-known incumbent and an entrant whose product is of unknown quality. The tools developed for the monopoly case may then be applied to a model of imperfect competition with vertically differentiated products. A different, but closely related question is, when will firms introduce a free "basic version" of the product? Addressing this question requires some modifications to this model, which we pursue in parallel work.

References

Bergemann, D. and J. Valimaki (1997). Market Diffusion with Two-Sided Learning. *The RAND Journal of Economics 28*(4), 773–795.

Bergemann, D. and J. Välimäki (2002). Entry and vertical differentiation. *Journal of Economic Theory 106*(1), 91–125.

Bolton, P. and C. Harris (1999). Strategic Experimentation. *Econometrica 67*(2), 349–374.

Bolton, P. and C. Harris (2000). Strategic Experimentation: The Undiscounted Case. *Incentives and Organization: Essays in Honour of James Mirrlees*.

Dutta, P. K. (1991). What do discounted optima converge to?: A theory of discount rate asymptotics in economic models. *Journal of Economic Theory 55*(1), 64–94.

Lewis, T. and H. Yildirim (2002). Learning by Doing and Dynamic Regulation. *The RAND Journal of Economics 33*(1), 22–36.

Liptser, R. and A. Shiryaev (1977). *Statistics of Random Processes, I.* Springer-Verlag, New York.

Maskin, E. and J. Riley (1984). Monopoly with Incomplete Information. *The RAND Journal of Economics 15*(2), 171–196.

Mussa, M. and S. Rosen (1978). Monopoly and Product Quality. *Journal of Economic Theory 18*(2), 301–317.

Ramsey, F. (1928). A Mathematical Theory of Saving. *The Economic Journal 38*(152), 543–559.

Appendix

Proof of Proposition 1

1. Using the definition of $\Pi^*(\alpha)$ in (6), and substituting (9) into (4), we obtain:

$$q(\alpha,\theta) = m(\alpha)(\phi(\theta) - \mathbb{E}_\theta[\phi]) + \sqrt{2v(\alpha) - m^2(\alpha)\left(\mathbb{E}_\theta\left[\phi^2\right] - \mathbb{E}_\theta[\phi]^2\right)},$$

which can be more effectively written as (10).

2. The static optimal quantity provision is given by

$$q(\alpha,\theta) = m(\alpha)\phi(\theta),$$

The difference between the two expressions is given by the term

$$\Lambda(\alpha) = -m(\alpha)\mathbb{E}_\theta[\phi] + \sqrt{2v(\alpha) - m^2(\alpha)\left(\mathbb{E}_\theta\left[\phi^2\right] - \mathbb{E}_\theta[\phi]^2\right)}.$$

Notice that $\Lambda(0) = \Lambda(1) = 0$. Moreover, we show in the proof of Proposition 2 that the function $\Lambda(\alpha)$ is strictly concave. Therefore, it must be the case that $\Lambda(\alpha) > 0$ for all $\alpha \in (0,1)$.

3. The first best policy is simply obtained by replacing virtual valuations $\phi(\theta)$ in equation (10) with types θ.

Proof of Proposition 2

Using the fact that $\mathbb{E}_\theta\left[\phi^2\right] = \mathbb{E}_\theta[\phi]^2 + \mathrm{Var}[\phi]$, the optimal quantity provision for the linear-quadratic model may be rewritten as

$$q(\alpha,\theta) = m(\alpha)(\phi(\theta) - \mathbb{E}_\theta[\phi]) + \sqrt{\left(\alpha\mu_H^2 + (1-\alpha)\mu_L^2\right)\mathbb{E}_\theta\left[\phi^2\right] - m^2(\alpha)\,\mathrm{Var}[\phi]}.$$

1. The first partial derivative $\frac{\partial q(\alpha,\theta)}{\partial \alpha}$ may be written as

$$\frac{\partial q(\alpha,\theta)}{\partial \alpha} \propto \phi - \mathbb{E}_\theta[\phi] + \frac{\mathbb{E}_\theta\left[\phi^2\right]((\mu_H - \mu_L)(1-2\alpha)) + 2(\mu_L + \alpha(\mu_H - \mu_L))\mathbb{E}_\theta[\phi]^2}{2\sqrt{\mathbb{E}_\theta[\phi]^2(\mu_L + \alpha(\mu_H - \mu_L))^2 + \mathbb{E}_\theta\left[\phi^2\right]\alpha(\mu_L - \mu_H)^2(1-\alpha)}}.$$

The first part of the statement (sufficiency of $\phi > \mathbb{E}_\theta[\phi]$ and $\alpha < \frac{1}{2}$) follows immediately form inspection of the term in parentheses.

2. Evaluating this expression at $\alpha = 0$ and $\alpha = 1$, one obtains

$$\left(\frac{\partial q(\alpha,\theta)}{\partial \alpha}\right)\Bigg|_{\alpha=0} = \phi(\theta) + \frac{\mathbb{E}_\theta\left[\phi^2\right]}{2\mathbb{E}_\theta[\phi]}\left(\frac{\mu_H}{\mu_L} - 1\right) > 0,\ \forall\theta$$

and

$$\left(\frac{\partial q(\alpha,\theta)}{\partial \alpha}\right)\Bigg|_{\alpha=1} = \phi(\theta) - \frac{\mathbb{E}_\theta\left[\phi^2\right]}{2\mathbb{E}_\theta[\phi]}\left(1 - \frac{\mu_L}{\mu_H}\right).$$

Since $\frac{1}{2\mathbb{E}_\theta[\phi]}\mathbb{E}_\theta\left[\phi^2\right]\left(1-\frac{\mu_L}{\mu_H}\right) > 0$, if $\phi(\theta_L)$ is low enough, for some θ quantity q is decreasing in α around $\alpha = 1$. Furthermore, since $\frac{1}{2\mathbb{E}_\theta[\phi]}\mathbb{E}_\theta\left[\phi^2\right]\left(1-\frac{\mu_L}{\mu_H}\right) > \frac{1}{2}\mathbb{E}_\theta[\phi]\left(1-\frac{\mu_L}{\mu_H}\right)$, it follows that a sufficient condition for the critical type θ^* to be larger than θ_L is given by

$$\frac{1}{2}\mathbb{E}_\theta[\phi]\left(1 - \frac{\mu_L}{\mu_H}\right) > \phi(\theta_L).$$

Considering that $\mathbb{E}_\theta[\phi(\theta)] = \theta_L$, the previous condition may be written as

$$\theta_L f(\theta_L)\left(1 + \frac{\mu_L}{\mu_H}\right) < 2.$$

3. Finally, the first best quantity provision is given by

$$q^*(\alpha,\theta) = m(\alpha)(\theta - \mathbb{E}[\theta]) + \sqrt{2v(\alpha) - m^2(\alpha)\operatorname{Var}[\theta]},$$

which is also concave. It follows that its first derivative (at $\alpha = 1$) may be written as

$$\begin{aligned}\frac{\partial q^*(\alpha,\theta)}{\partial \alpha} &= (\mu_H - \mu_L)(\theta - \mathbb{E}[\theta]) + \frac{\left(\mathbb{E}[\theta]^2 + \operatorname{Var}[\theta]\right)\left(\mu_H^2 - \mu_L^2\right) - 2m(\alpha)(\mu_H - \mu_L)\operatorname{Var}[\theta]}{2\sqrt{2v(\alpha) - m^2(\alpha)\operatorname{Var}[\theta]}}\\ \left(\frac{\partial q^*(\alpha,\theta)}{\partial \alpha}\right)\Bigg|_{\alpha=1} &\propto \theta - \frac{\mathbb{E}\left[\theta^2\right]}{2\mathbb{E}[\theta]}\left(1 - \frac{\mu_L}{\mu_H}\right).\end{aligned}$$

Since $\mathbb{E}\left[\theta^2\right] = \mathbb{E}[\theta]^2 + \operatorname{Var}[\theta]$, letting $\theta = \theta_L$, the last part of the proposition follows directly.

Proof of Proposition 7

1. The static benchmark quantity and total prices are given by

$$q(\alpha,\theta) = \left(\frac{m(\alpha)\phi(\theta)}{2c}\right)^2$$

and

$$p(\alpha,\theta) = m(\alpha)\,\theta\frac{m(\alpha)\,\phi(\theta)}{2c} - \frac{m^2(\alpha)}{2c}\int_{\theta_L}^{\theta}\phi(x)\,dx$$
$$= \frac{m^2(\alpha)}{2c}\left(\phi(\theta)\,\theta - \int_{\theta_L}^{\theta}\phi(x)\,dx\right).$$

Therefore unit prices are given by

$$\frac{p(\alpha,\theta)}{q(\alpha,\theta)} = 2c\frac{\phi(\theta)\,\theta - \int_{\theta_L}^{\theta}\phi(x)\,dx}{\phi^2(\theta)}$$

and they are independent of α.

2. In the undiscounted optimum, are given by

$$\frac{p(\alpha,\theta)}{q(\alpha,\theta)} = \frac{\frac{1}{2c}\left(\phi(\theta)\,\theta - \int_{\theta_L}^{\theta}\phi(x)\,dx\right)\left(\alpha\mu_H^2 + (1-\alpha)\,\mu_L^2\right)}{\frac{\phi^2(\theta)\left(\alpha\mu_H^2+(1-\alpha)\mu_L^2\right)^2}{4c^2m^2(\alpha)}},$$

which may be simplified into

$$\frac{p(\alpha,\theta)}{q(\alpha,\theta)} = 2c\frac{\phi(\theta)\,\theta - \int_{\theta_L}^{\theta}\phi(x)\,dx}{\phi^2(\theta)}\frac{m^2(\alpha)}{\alpha\mu_H^2 + (1-\alpha)\,\mu_L^2}. \tag{13}$$

The behavior of unit prices as a function of α is determined, for all θ, by the component $\frac{m^2(\alpha)}{\alpha\mu_H^2+(1-\alpha)\mu_L^2}$. Note that $\frac{m^2(\alpha)}{\alpha\mu_H^2+(1-\alpha)\mu_L^2} < 1$ for all $\alpha \in (0,1)$, so that the dynamic solution is always below the static one.

3. Moreover, note that this function is increasing in α over the range $[\hat{\alpha}, 1]$, with $\hat{\alpha} = \frac{\mu_L}{\mu_H+\mu_L}$. Its first derivative may be written as

$$\frac{d\left(\frac{(\alpha\mu_H+(1-\alpha)\mu_L)^2}{\alpha\mu_H^2+(1-\alpha)\mu_L^2}\right)}{d\alpha} = \frac{(\alpha\mu_H + (1-\alpha)\,\mu_L)\,(\mu_H - \mu_L)^2}{\left(\alpha\mu_H^2 + (1-\alpha)\,\mu_L^2\right)^2}\,(\alpha\mu_H - (1-\alpha)\,\mu_L)\,.$$

4. The function $\frac{m^2(\alpha)}{\alpha\mu_H^2+(1-\alpha)\mu_L^2}$ is always convex in α. Its second derivative may be written as

$$\frac{d^2\left(\frac{(\alpha\mu_H+(1-\alpha)\mu_L)^2}{\alpha\mu_H^2+(1-\alpha)\mu_L^2}\right)}{(d\alpha)^2} = \frac{2\,(\mu_H - \mu_L)^2\,\mu_H^2\mu_L^2}{\left(\alpha\mu_H^2 + (1-\alpha)\,\mu_L^2\right)^3} > 0.$$