

MENU PRICING AND LEARNING

ALESSANDRO BONATTI*

YALE UNIVERSITY

FEBRUARY 11, 2009

JOB MARKET PAPER

Abstract

We address the question of designing dynamic menus to sell experience goods. A dynamic menu consists of a set of price-quantity pairs in each period. The quality of the product is initially unknown, and more information is generated through experimentation. The amount of information in the market is increasing in the total quantity sold in each period, and the firm can control the information flow to the market by adjusting the level of sales. We derive the optimal menu as a function of consumers' beliefs about product quality, and characterize the changes in prices and quantities resulting from information diffusion and its effects on beliefs.

The equilibrium menu prices are the result of a dynamic trade-off between immediate gains from trade, information production, and information rents. The firm initially charges lower prices, in order to increase sales above the static optimum, sacrificing short-term gains in order to invest in information. As the market obtains more information, the firm gradually shifts to a policy designed to extract revenue from high-valuation buyers. This policy may eventually exclude low-valuation buyers from the market, even if the product's underlying quality is in fact high.

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

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

*Department of Economics, Yale University, New Haven CT 06520-8268. alessandro.bonatti@yale.edu. I am indebted to my advisor, Dirk Bergemann, for his invaluable help and encouragement throughout this project. I also wish to thank Johannes Hörner and Ben Polak for many helpful discussions. This paper benefited from comments from Rossella Argenziano, Eduardo Faingold, Dino Gerardi, William Hogan, Sven Rady, Maher Said, Larry Samuelson, Jiwoong Shin, K. Sudhir, Richard Zeckhauser, as well as participants at various seminars and conferences. All errors, of course, are my own. This work was generously supported by a Cowles Foundation Carl A. Anderson Fellowship and a Yale University Leylan Fellowship in the Social Sciences.

1 Introduction

1.1 Motivation

Learning plays a crucial role in many markets and other strategic environments. In particular, in markets for new products and services, sellers often face uncertainty over the product's fit for consumers' needs. Consider, for example, new software products and new online services such as DVD rentals, data backup, Internet telephony, and Internet access itself. The quality of these products is only revealed to market participants through consumption, as buyers learn from their own experience and from that of others. Another salient feature of these markets is that consumers' willingness to pay is heterogeneous. This creates the opportunity for firms to profitably adopt price discrimination techniques (such as menu pricing) to elicit the buyers' private information and extract more surplus.

In markets for new products and services, demand conditions evolve over time. As consumers gradually learn about the quality of the product, they modify their purchasing behavior. In consequence, the diffusion of information induces firms to modify their menu prices. In particular, in markets for experience goods, the diffusion of information is endogenous to the behavior of market participants. Buyers and sellers influence the amount of information conveyed to the market through the level of sales, which is determined by consumers' purchasing behavior and firms' pricing strategies. Moreover, information about these products' performance is widely and publicly available through an increasing number of channels.¹ The availability of aggregate information provides firms with an additional instrument in a dynamic environment. Firms are able to condition their product lines and prices on the aggregate opinion of their customers.

In this scenario, the forward looking firm's problem is to screen consumers in order to maximize revenues, while taking into account the informational value of sales. By selling additional units of the product (by offering introductory discounts for example), the firm accelerates the buyers' learning process, which can lead to higher profits in the long run.

In this paper, we address the question of designing dynamic menus to sell experience goods. We focus on the dynamics of prices and quantities to answer questions such as: should firms offer quantity discounts to high-valuation buyers? How does the range of offered prices and quantities vary over time? Which buyers benefit from the diffusion of information?

We develop a dynamic screening model in which a monopolist in each period offers a menu of contracts to a population of buyers. These buyers have private information about their willingness to pay, providing the firm with an incentive to charge discriminatory prices.

¹For example, <http://www.consumerreports.org>, <http://www.cnet.com>, or feedback reports on large retailers' websites such as Amazon.com.

The quality of the product is initially unknown, and more information is generated through experimentation. As purchases are made, both the firm and the consumers observe signals about the product's quality and, as a result, revise their beliefs. The amount of information in the market is increasing in the total quantity sold in each period. The firm can therefore control the information flow to the market by adjusting the level of sales. Learning about the product occurs faster as more units are sold, and so the firm may use low introductory prices.

The uncertainty about the quality of the product introduces a new dynamic element into the standard trade-off between efficiency and rent extraction. More specifically, the determination of the quantity levels supplied to each buyer is the combination of three effects. The first of these components is one of information generation. Since learning occurs through consumption, each unit sold provides additional information value. This leads the firm to sell additional units when uncertainty about quality is high and beliefs are more responsive to news. The second component is related to efficiency. As consumers' beliefs improve, their willingness to pay is higher, hence creating the opportunity to realize larger gains from trade. Therefore the firm offers larger quantities in these cases. The third effect and final component is an adverse selection effect. Positive signals about quality increase the spread in buyers' valuations for the product. This makes the incentive compatibility constraints more difficult to satisfy, and induces the firm to offer lower quantity levels to buyers with low willingness to pay.

While the firm simultaneously pursues the dual objectives of generating information and screening consumers, the balance between the two shifts over time. Initially, the firm increases the level of sales to all buyers above the static optimum: it sacrifices short-term gains in order to invest in information production. As more information is revealed, the firm gradually shifts to a policy that targets the consumers with the highest valuations, in order to extract more surplus. This policy may eventually exclude low-valuation buyers from the market, even if the product's underlying quality is in fact high. As consumers' beliefs improve, the cost of providing incentives to high-valuation buyers increases due to the adverse selection effect. This leads the firm to reduce the supply of its product to low-valuation buyers. Consequently, the quantity levels offered to a low-valuation buyer need not be monotonic in the beliefs about product quality. The model also predicts that successful products should be characterized by a greater price dispersion and a wider variety of available quantities. In particular, the firm expands the range of offered quantities through the addition of new options both at the top and at the bottom of the menu.

In the model, learning takes place based on aggregate information. More precisely, we assume that each consumer's action (quantity choice) and payoff (experienced quality level)

is observable to other buyers and to the firm. In other words, all information is publicly available to the market. While this is an important assumption, it suits the purpose of this paper for two reasons. First, in large markets, consumers realize that others' experience is also indicative of the underlying product quality and take public information into account. More importantly, this paper is interested in modeling the firm's optimal response to variations in demand arising from the arrival of new information. As such, it focuses only on information the firm that can condition its strategies upon. In an alternative model, demand for the product would be determined by consumers' private experiences, while the firm only observes the market's average experience. In the context of this paper, the introduction of private information would add noise to the demand process, but would not alter the qualitative properties of the firm's behavior. We therefore abstract away from further heterogeneity in demand, and only consider the market's observable aggregate experience.

1.2 Examples

The model is well suited to analyze several different markets. Consider first the market for online DVD rentals. Companies like Netflix or Blockbuster allow consumers to subscribe to plans specifying the number of movies they may rent at the same time. While buyers differ in their personal willingness to pay for watching DVD movies, the quality of the recommender system is a common component in determining the overall quality of the service. With this interpretation, each rented movie constitutes an informative experiment about the product's quality. It is reasonable to assume that customers with higher willingness to pay also care more about the fit of the recommendation to their own preferences. Furthermore, both the prices for each plan and the plan choice by the consumer are easily adjustable. Finally, Netflix subscribers exchange information about their experience through a surprisingly large number of channels.² This means that information about the overall performance of the service circulates very rapidly.

Netflix launched the rental service in 2001 and held a near-monopoly position for several years. Figure 1 reports the menus offered by Netflix over the years 2002 through 2005, that is, immediately before Blockbuster established itself as a serious competitor. In 2002, the Netflix menu offer consisted of two plans, allowing for the simultaneous rental of two and four titles, respectively. The variety of the plans offered increased over time, as the service soon proved to be a clear success.³ In 2003, Netflix modified its offer of plans to a four-item

²For example, <http://www.hackingnetflix.com> and <http://blog.netflix.com> are two of the most popular blogs among Netflix customers.

³The 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 DVD Rental Plans 2002-2005

menu, while raising unit prices across the product line. It then raised its prices further in 2004, and added several more options in 2005. At the same time, it reduced all prices slightly, possibly due to competitive pressures. Consistent with the model’s predictions, the range of total charges (in dollars per month) went from a minimum of \$12 and a maximum of \$20 in 2002 to \$5 and \$48, respectively, in 2005. At the same time, the set of available quantities increased from the two options offered in 2002 to the 2005 menu, which allows consumers to choose any number of simultaneous rentals, with a minimum of one and a maximum of eight. Finally, notice that the lowest quantity offered decreased from two rentals at a time in 2002 to one in 2005.⁴

The market for enterprise software provides an alternate application. An emerging contractual arrangement in this industry is given by software-as-a-service (SaaS). Under this contractual form, firms have the option of renting a given number of licenses for the use of a given software product (for example, a customer database system or an online backup program). Larger firms need to rent more licenses seats and 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 rate of information. Moreover, the rental contracts and their corresponding prices are easily adjustable. Finally, network externalities between firms are not a significant issue in enterprise software (as it is designed for internal use), making the private values framework realistic.⁵

⁴We are ignoring plans which impose a limit to the number of monthly rentals. If we were to include them, then the lowest quantity would be given by a one DVD at a time, up to four per month.

⁵Larkin (2008) provides a detailed analysis of some frequently used contractual arrangements in this industry.

1.3 Related Literature

This paper enriches the screening literature by extending nonlinear pricing techniques beyond the canonical, static environment to a model where information is revealed over time. It therefore builds upon the classic studies in price discrimination, such as Mussa and Rosen (1978) and Maskin and Riley (1984). At the same time, it is tightly connected to continuous time models in which the information flow is controlled by one or more of the agents. Examples of these papers include the works of Bolton and Harris (1999), Keller and Rady (1999), and Moscarini and Smith (2001) on strategic experimentation, as well as the continuous-time reputation game analyzed by Faingold and Sannikov (2007). In particular, we adopt the method of Faingold and Sannikov (2007) for showing existence of and uniqueness of a solution to the firm's problem.⁶

Our analysis also complements several models of introductory and dynamic pricing under product quality uncertainty. The main work in this area is due to Bergemann and Välimäki (1997, 2002, 2006) and to Villas-Boas (2004, 2006). In particular, Bergemann and Välimäki (2002) analyze a duopoly model of price competition with vertically differentiated products. In Bergemann and Välimäki (2002), market participants are uncertain about the degree of vertical differentiation of the two firm's products. In this paper, as in Bergemann and Välimäki (2002), consumers are heterogeneous in terms of their willingness to pay. However, we allow consumers to have multi-unit demands and the firm to price discriminate.

The problem of generating information through sales, was first studied, in the context of a screening model, by Braden and Oren (1994). In their model, a monopolist is uncertain about the probability distribution over each buyer's willingness to pay. One buyer arrives in each period, and her choice from the firm's menu provides information about this distribution. In contrast, in our paper both sides of the market are learning, and the inference procedure is based on aggregate information. This framework provides a more realistic description of markets with a large number of customers (such as the DVD rentals or the enterprise software markets).

The techniques used in this paper also relate to the model of dynamic regulation of Lewis and Yildirim (2002). In their model, a planner offers a menu of contracts to a firm whose production costs decrease by a deterministic amount, but where innovation follows a stochastic process. Compared to Lewis and Yildirim (2002), our analysis emphasizes learning and the buyer-specific effects of information. With a slight change in interpretation, our framework could also be used to study repeated procurement contracts, when the regulator and the firm share uncertainty over a production cost parameter.

⁶The main differences between the two papers are discussed at the end of subsection 2.2.

In addition to this theoretical body of work, a recent empirical literature attempts to quantify the importance of learning considerations on consumers' dynamic purchasing behavior. In these studies, consumers learn from their individual experience, revise their beliefs about product quality, and consequently modify their choices. A non-exhaustive list of empirical papers on learning and dynamic consumer choice includes Akerberg (2003), Akçura, Gonul, and Petrova (2004), Crawford and Shum (2005), Erdem and Keane (1996), Goettler and Clay (2006), Gowrisankaran and Rysman (2007), and Israel (2005). From a different perspective, Hitsch (2006) and Song and Chintagunta (2003) analyze learning about the demand on the firm's side, but focus on investment decisions such as product adoption or exit, not on pricing strategies. This paper complements this literature with a theoretical framework for nonlinear pricing, in which firms' learning is just as important as buyers', and in which information is obtained from aggregate experience.

The remainder of this paper is organized as follows: section 2 introduces the model, section 3 derives the equilibrium conditions and discusses the comparative statics properties of the optimal menu; section 4 specifies the model to a linear utility - quadratic cost structure and fully characterizes the equilibrium menu; section 5 derives predictions for the intertemporal patterns of prices and quantities; section 6 extends our analysis to a competitive setting through a duopoly example; and section 7 concludes. All proofs are in the Appendix.

2 The Model

2.1 Payoffs

We consider a dynamic model with a monopolist firm and a continuum of small consumers. Consumers purchase repeatedly and have multi-unit demands in each period. Each consumer's valuation for the firm's product depends both on a private value and a common value component. We denote by θ an idiosyncratic, private value component, representing the buyer's personal willingness to pay for the product. For each buyer, θ belongs to the interval $\Theta = [\theta_L, \theta_H]$. The idiosyncratic component θ is the consumer's private information. It is distributed in the population according to a distribution

$$F : \Theta \rightarrow [0, 1].$$

Assumption 1 (Monotone Hazard Rate)

$F(\theta)$ satisfies the monotone hazard rate condition: $(1 - F(\theta)) / f(\theta)$ is decreasing.

We denote by μ a common value component representing the inherent quality of the product. This parameter may only take one of two values, $\mu \in \{\mu_L, \mu_H\}$ with $0 < \mu_L < \mu_H$. Each consumer's valuation for q units of a product is a separable function of the product's quality μ and of the consumer's willingness to pay θ . The complete information utility of a consumer with willingness to pay θ , who purchases q units of a product of quality μ , for a total charge of p , is given by

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

The function $u(q)$ is assumed to be strictly increasing. As a consequence, the consumer's utility function $U(\mu, \theta, q, p)$ displays the single crossing property in (θ, q) . Furthermore, quality and personal taste are complement goods for the consumer. We assume that each buyer makes a purchase decision in every period, and that she can freely switch between different quantity levels. We normalize each buyer's outside option to zero. Finally, we assume that production costs are given by a strictly increasing function $c(q)$.

Product quality μ is initially unknown to both the firm and the consumers, and all market participants share the common prior belief

$$\alpha_0 = \Pr(\mu = \mu_H).$$

At each time t , the expected product quality, given current beliefs α_t , is denoted by

$$\mu(\alpha_t) \triangleq \mathbb{E}_{\alpha_t} \mu = \alpha_t \mu_H + (1 - \alpha_t) \mu_L.$$

In each period, a monopolist posts a menu of price-quantity pairs. We require the firm to price anonymously, and we allow for prices and quantities to be flexibly adjusted. In a direct mechanism, the firm's strategy is a pair of functions $q_t : \Theta \rightarrow \mathbb{R}_+$ and $p_t : \Theta \rightarrow \mathbb{R}_+$ in each period. These functions determine the quantity level and the total charges assigned to each buyer θ . Suppose each buyer purchases quantity level $q_t(\theta)$ and pays total charges of $p_t(\theta)$. The firm then obtains flow profits of

$$\Pi(q_t, p_t) \triangleq \int_{\theta_L}^{\theta_H} (p_t(\theta) - c(q_t(\theta))) f(\theta) d\theta.$$

In our environment, the social gains from trade realized by selling quantity q to type θ , when product quality is μ , are given by $\mu\theta u(q) - c(q)$. In order to ensure that the firm's problem has an interior solution, we introduce the following assumptions on the social gains from trade.

Assumption 2 (Social Gains from Trade)

1. Gains from trade $\mu\theta u(q) - c(q)$ are strictly concave in q for all μ and θ .
2. For all μ and θ , $\mu\theta u'(0) - c'(0) > 0$.
3. For all μ and θ , $\lim_{q \rightarrow \infty} [\mu\theta u'(q) - c'(q)] = -\infty$.

2.2 Information and Learning

Information about product quality is obtained through consumption. The level of aggregate sales determines the total number of experiments with the product. In turn, the level of experimentation influences the rate at which the firm and the consumers learn about its quality. We assume that beliefs α_t evolve according to the following diffusion process:

$$d\alpha_t = \alpha_t(1 - \alpha_t) \frac{\mu_H - \mu_L}{\sigma} \sqrt{Q_t} dz_t, \quad (1)$$

where dz_t denotes the standard Wiener process. We also define the total quantity sold at time t as:

$$Q_t = \int_{\theta_L}^{\theta_H} q_t(\theta) f(\theta) d\theta.$$

In the Appendix, we present a model which provides a micro foundation for the belief process based on individual signals. At this stage, note that beliefs follow a martingale, and thus the process has a zero drift. We define following function:

$$\Sigma(\alpha_t) \triangleq \frac{1}{2} \left(\alpha_t(1 - \alpha_t) \frac{\mu_H - \mu_L}{\sigma} \right)^2. \quad (2)$$

The function $\Sigma(\alpha_t)$ captures the marginal contribution of each unit sold to the variance of the belief process $(d\alpha_t)^2 = Q_t \Sigma(\alpha_t) dt$. The variance is increasing in the degree of dispersion $\alpha_t(1 - \alpha_t)$ and in the signal-to-noise ratio $(\mu_H - \mu_L)/\sigma$. Beliefs evolve more quickly when current uncertainty is high, and when signals are precise. Finally, we stress that the evolution of the belief process is endogenously determined, since the total quantity sold Q_t depends on the firm's pricing and on the consumers' purchasing decisions.

In this environment, the firm can therefore control the rate of information flow to the market, by adjusting the level of sales. In this sense, the paper is linked to the strategic experimentation literature, as well as to the continuous-time reputation model of Faingold and Sannikov (2007). In that paper, the long-lived player's actions affect the information flow received by the short-lived players. Our screening model differs from Faingold and Sannikov (2007) along two essential dimensions. First, in our model the firm does not know its own type (unlike the long lived player in their reputation game), and is learning about it along

with the short-lived players. Second, the firm cannot manipulate the buyers' beliefs, but only affect the speed at which learning takes place. Nevertheless, we will be able to apply the method of proof of Faingold and Sannikov (2007) to show existence and uniqueness of a solution in our model, as we discuss in the next section (and in the Appendix).

3 Equilibrium Analysis

We now turn to the equilibrium analysis, starting with a characterization of the incentive compatible menus of contracts. In our model, each individual buyer has a negligible impact on the information flow. Therefore, each buyer chooses the price-quantity pair that maximizes her expected utility, given beliefs α_t and the firm's menu offer. The two-point prior on quality μ ensures that, at every point in time, the value of α_t is a sufficient statistic for the firm's problem. Therefore, we denote by $(q(\alpha_t, \theta), p(\alpha_t, \theta))$ the menu offered by the firm when the market belief is given by α_t . We also denote by $U(\alpha_t, \theta, \theta')$ the expected utility of a buyer with beliefs α_t and willingness to pay θ who purchases the item $(q(\alpha_t, \theta'), p(\alpha_t, \theta'))$ intended for a buyer of type θ' :

$$U(\alpha_t, \theta, \theta') = \mu(\alpha_t) \cdot \theta \cdot u(q(\alpha_t, \theta')) - p(\alpha_t, \theta'). \quad (3)$$

Let $U(\alpha_t, \theta) = U(\alpha_t, \theta, \theta)$ denote buyer θ 's indirect utility when reporting truthfully. The incentive compatibility constraints for the firm's problem are then given by the consumer's first and second order conditions for truthful revelation:

$$\frac{\partial U(\alpha_t, \theta)}{\partial \theta} = \mu(\alpha_t) \cdot u(q(\alpha_t, \theta)), \quad (4)$$

$$\frac{\partial q(\alpha_t, \theta)}{\partial \theta} \geq 0, \text{ for all } \alpha_t \text{ and } \theta. \quad (5)$$

Each buyer's valuation for q units of the product depends positively on her beliefs α_t . Positive signals therefore allow the firm to charge higher prices. However, when beliefs become more optimistic, the firm must concede higher information rents to the buyers, as shown by equation (4). This effect is due to the complementarity between product quality and buyers' willingness to pay. As beliefs α_t increase, the difference between any two buyers' willingness to pay also increases, creating stronger incentives to misreport one's type. This means that for high values of α_t the incentive compatibility constraints are harder to satisfy. Finally, the buyers' participation constraints are given by

$$U(\alpha_t, \theta) \geq 0, \text{ for all } \alpha_t \text{ and } \theta. \quad (6)$$

3.1 Myopic Benchmark

We now provide a benchmark for the firm's problem. We consider an impatient (myopic) firm, who only maximizes the current flow profits. By expressing total charges $p(\alpha_t, \theta)$ in terms of the buyers' indirect utilities $U(\alpha_t, \theta)$, we can rewrite the firm's flow profits as

$$\Pi(\alpha_t, q, U) \triangleq \int_{\theta_L}^{\theta_H} (\mu(\alpha_t) \cdot \theta \cdot u(q(\alpha_t, \theta)) - c(q(\alpha_t, \theta)) - U(\alpha_t, \theta)) f(\theta) d\theta.$$

The myopic firm maximizes $\Pi(\alpha_t, q, U)$ subject to the incentive compatibility constraints (4) and (5) and to the participation constraint (6). As in many screening problems, Assumption 1 ensures that (5) holds in equilibrium. We then substitute the incentive compatibility constraint (4) in the objective, and integrate by parts. As a result, we can express the firm's flow profits as a function of only beliefs α_t and quantities $q(\alpha_t, \theta)$. Denote by $\phi(\theta) \triangleq \theta - (1 - F(\theta)) / f(\theta)$ the virtual valuation of buyer θ . The firm's flow profits are then given by

$$\Pi(\alpha_t, q) \triangleq \int_{\theta_L}^{\theta_H} (\mu(\alpha_t) \phi(\theta) u(q(\alpha_t, \theta)) - c(q(\alpha_t, \theta))) f(\theta) d\theta, \quad (7)$$

and the myopic equilibrium profit function is defined as

$$\Pi_m(\alpha_t) \triangleq \max_{q: \Theta \rightarrow \mathbb{R}_+} \Pi(\alpha_t, q).$$

The myopically optimal menu is determined by maximizing (7) pointwise. The first order condition for quantity provision is given by

$$\mu(\alpha_t) \phi(\theta) u'(q) - c'(q) = 0. \quad (8)$$

The myopic equilibrium quantity level $q_m(\alpha_t, \theta)$ is then given by the solution to (8), whenever this solution is positive, and by zero otherwise. The firm equalizes marginal cost and the buyer's marginal utility. Notice that the expected product quality $\mu(\alpha_t)$ acts as scale parameter for marginal utilities, and hence for equilibrium quantity provision. The following proposition describes the key properties of the myopic solution.

Proposition 1 (Myopic Solution)

1. The myopic quantity $q_m(\alpha_t, \theta)$ is strictly increasing in both α_t and θ .
2. The myopic profit function $\Pi_m(\alpha_t)$ is strictly increasing and strictly convex in α_t .

Higher beliefs about product quality improve every buyer's willingness to pay, and the firm finds it profitable to sell a larger number of units. Convexity of the myopic profit function implies, that as beliefs improve, the firm can also charge higher unit prices. The convexity of the myopic profit function has implications for the firm's incentives to learn about the quality of its product. In the myopic world, the firm would be willing to pay in order to enter a fair bet between the two states $\mu = \mu_L$ and $\mu = \mu_H$. Define the complete information average profit as

$$v(\alpha_t) \triangleq \alpha_t \Pi_m(1) + (1 - \alpha_t) \Pi_m(0). \quad (9)$$

For all interior α_t , we then have $v(\alpha_t) > \Pi_m(\alpha_t)$, with $\Pi_m(1) = v(1)$ and $\Pi_m(0) = v(0)$.

3.2 Dynamic Solution

We now turn to the dynamic version of the problem. A strategy for the firm consists of a sequence of quantity supply functions $q_t : \Theta \rightarrow \mathbb{R}_+$. The corresponding total charges are determined by the incentive compatibility constraints. Given the prior belief α_0 , the firm's objective function may then be written as

$$\Pi^*(\alpha_0) \triangleq \sup_{q_t : \Theta \rightarrow \mathbb{R}_+} \lim_{T \rightarrow \infty} \mathbb{E}_\alpha \left[\int_0^T e^{-rt} \Pi(\alpha_t, q_t) d\alpha_t \mid \alpha_0 \right], \quad (10)$$

where the evolution of beliefs α_t is given by the law of motion (1). Using the law of motion for beliefs and Itô's Lemma, the Hamilton-Jacobi-Bellman (HJB) equation for this maximization problem is given by

$$rV(\alpha_t) = \max_{q : \Theta \rightarrow \mathbb{R}_+} [\Pi(\alpha_t, q) + Q\Sigma(\alpha_t) V''(\alpha_t)]. \quad (11)$$

The expression for $rV(\alpha_t)$ differs from the flow profits $\Pi(\alpha_t, q)$ only through the term $Q\Sigma(\alpha_t) V''(\alpha_t)$. This term is proportional to the total quantity sold Q , and its sign depends on whether $V''(\alpha_t) \gtrless 0$. Each unit sold provides an informative signal whose effect on beliefs depends on the variance $\Sigma(\alpha_t)$. Therefore, the term $Q\Sigma(\alpha_t) V''(\alpha_t)$ can be interpreted as the informational value generated through sales. In particular, the marginal value of information is given by the second derivative of the value function $V''(\alpha_t)$. The convexity of the value function is related to the firm's incentives to amplify the variance of the belief process, and hence to learn faster about the product's quality. Finally, note that information has no value when $\alpha_t = 0$ and $\alpha_t = 1$, since beliefs no longer change in those cases. Writing the HJB

equation (11) more explicitly, we obtain an expression that may be maximized pointwise:

$$rV(\alpha_t) = \max_{q: \Theta \rightarrow \mathbb{R}_+} \left[\int_{\theta_L}^{\theta_H} (\mu(\alpha_t) \phi(\theta) u(q(\alpha_t, \theta)) - c(q(\alpha_t, \theta))) f(\theta) d\theta \right. \\ \left. + \int_{\theta_L}^{\theta_H} q(\alpha_t, \theta) \Sigma(\alpha_t) V''(\alpha_t) f(\theta) d\theta \right]. \quad (12)$$

We now prove existence of a solution to the firm's dynamic problem. We then return to the optimal menu of contracts and illustrate the role of the value of information in determining the equilibrium prices and quantities. Our approach to proving existence of a solution consists of turning the HJB equation into a second order differential equation with the two boundary conditions $rV(0) = \Pi_m(0)$ and $rV(1) = \Pi_m(1)$. Since α_t is the independent variable in our boundary value problem, we drop time subscripts. Our existence and uniqueness result is stated in the following theorem.

Theorem 1 (Existence and Uniqueness)

1. *There exists a unique solution $V(\alpha)$ to the HJB equation (12). $V(\alpha)$ is \mathcal{C}^2 and satisfies $\Pi_m(\alpha) \leq rV(\alpha) \leq v(\alpha)$ for all α .*
2. *The policy function $q(\alpha, \theta)$ maximizing the right hand side of (12) pointwise is the unique optimal control. It is continuous and differentiable in α and θ .*
3. *The solution $V(\alpha)$ coincides with the supremum value $\Pi^*(\alpha)$ of (10).*

The logic of the proof of (1.) and (2.) follows the one in Faingold and Sannikov (2007), who first developed a method to prove existence and uniqueness of a solution given the singularity points of the differential equation. The proof of (3.) uses a verification theorem from Fleming and Soner (2006).

We now derive some elementary properties of the policy function. For this purpose, a prior question to be addressed is whether the firm assigns positive value to information.

Theorem 2 (Convexity of the Value Function)

The firm's value function $V(\alpha)$ is convex in α .

Theorem 2 shows that the firm's value function inherits the convexity property of the myopic profit function. A convex value function implies that the forward-looking firm is willing to give up some revenue in the short run (i.e. to depart from $\Pi_m(\alpha)$) in order to generate more information through sales. We relate the degree of patience and the incentives to invest in information in the following comparative statics result. We normalize the firm's payoffs by focusing on the return (or annuity) function $rV(\alpha)$.

Proposition 2 (Value of Information)

1. *The return function $rV(\alpha)$ and the value of information $\Sigma(\alpha)V''(\alpha)$ are decreasing in σ and in r , for all α .*
2. *For all μ_H and μ_L such that $\alpha\mu_H + (1-\alpha)\mu_L = \mu$ for some $\mu > 0$, $rV(\alpha)$ and $\Sigma(\alpha)V''(\alpha)$ are increasing in the difference $(\mu_H - \mu_L)$.*

These results show that the firm's patience level, the precision of the available signals, and the relevance of the learning process provide the firm with greater incentives to experiment. A measure of the relevance of the learning process is given by the difference between the two possible quality levels.

3.3 Properties of the Equilibrium Menus

We now turn to the properties of the equilibrium menus of contracts. Pointwise maximization of the right-hand side of the HJB equation (12) yields an expression for the optimal quantity provision, as a function of the value of information $V''(\alpha_t)$. The equilibrium $q(\alpha_t, \theta)$ is given by the solution to the first order condition

$$\mu(\alpha_t)\phi(\theta)u'(q(\alpha_t, \theta)) - c'(q(\alpha_t, \theta)) + \Sigma(\alpha_t)V''(\alpha_t) = 0, \quad (13)$$

whenever this solution is positive, and by zero otherwise. This condition differs from that of the myopic firm because of the marginal value of information. In particular, the forward-looking firm equalizes marginal cost to the buyer's marginal utility, augmented by the marginal value of information $\Sigma(\alpha_t)V''(\alpha_t)$. Notice that the firm's incentives to experiment, captured by $\Sigma(\alpha_t)V''(\alpha_t)$, are uniform across buyers, as this term does not depend on the buyer's type θ . We summarize our comparative statics results for the forward-looking firm's problem in the following proposition.

Proposition 3 (Quantity Supply)

1. *The quantity level $q(\alpha, \theta)$ is everywhere higher than the myopic quantity $q_m(\alpha, \theta)$.*
2. *The quantity level $q(\alpha, \theta)$ is increasing in the value of information $\Sigma(\alpha)V''(\alpha)$.*

Proposition 3 shows that the comparative statics results of Proposition 2 have important implications for the properties of the equilibrium menus. In particular, the firm experiments by selling quantity levels larger than the myopic optima for all α and θ . This leads, for example, to a (weakly) larger set of types receiving positive quantities in the dynamic solution than in the myopic one. Combining the results of propositions 2 and 3, we obtain that the

number of additional units sold is increasing in the firm's patience level, and in the precision of the signals. However, the value of information $\Sigma(\alpha) V''(\alpha)$, as well as quantities and market coverage levels, are typically not monotonic in beliefs. In particular, the firm has no incentive to experiment when beliefs are degenerate and $\alpha \in \{0, 1\}$.

The quantity levels $q(\alpha, \theta)$ in the direct mechanism can be linked to the actual price-quantity menus offered by the firm in an indirect mechanism. We can represent the firm's strategy through a nonlinear price function $\hat{p}(\alpha, q)$. This function defines the total amount charged by the firm for q units of the product, when the market beliefs are given by α . Similarly, we denote by $\theta(\alpha, q)$ the buyer who purchases quantity level q in equilibrium, as a function of the beliefs. Consumers maximize their utility given the firm's current menu offer. This allows us to characterize the marginal prices charged by the firm on each unit. The equilibrium marginal prices must in fact solve the buyer's first order condition

$$\mu(\alpha) \theta(\alpha, q) u'(q) - \hat{p}_q(\alpha, q) = 0.$$

Since any quantity sold is defined by equation (13) for some type θ , the equilibrium marginal prices are given by

$$\hat{p}_q(\alpha, q) \triangleq \mu(\alpha) u'(q) \phi^{-1} \left(\frac{c'(q) - \Sigma(\alpha) V''(\alpha)}{\mu(\alpha) u'(q)} \right).$$

Proposition 4 (Marginal Prices)

1. *Marginal prices $\hat{p}_q(\alpha, q)$ are everywhere lower than in the myopic benchmark.*
2. *Marginal prices $\hat{p}_q(\alpha, q)$ are decreasing in the value of information $\Sigma(\alpha) V''(\alpha)$.*

As in all screening problems, a precise characterization of prices requires knowledge of the distribution of types $F(\theta)$. However, regardless of the distribution of types, Proposition 4 shows that experimentation reduces the marginal prices paid by each consumer. The firm is willing to give up revenue (by lowering prices) to induce experimentation, while the consumer has no incentives to pay for information.

To summarize our results so far, the solution to the firm's dynamic optimization problem implies higher sales and lower marginal prices, compared to the myopic benchmark. The level of experimentation depends positively on the firm's patience level and on the precision of the available signals. It also depends positively on the difference between the two possible quality levels of the product. In the next section, we examine how the optimal menus adjust to the evolution of beliefs for the case of linear utility and quadratic costs.

4 Linear-Quadratic Model

In this section, we adopt the Mussa and Rosen (1978) functional form assumptions: $u(q) = q$ and $c(q) = q^2/2$. We characterize the solution for a setting in which all buyers participate, and discuss the properties of the equilibrium menu with positive discounting. We turn to the undiscounted limit to describe the effects of information more in detail. We then discuss the properties of the equilibrium menu that extend to the case of small positive discounting. Finally, we extend the analysis to the case of imperfect market coverage.

The first order condition (13) can now be written as

$$q(\alpha, \theta) = \max \{ \mu(\alpha) \phi(\theta) + \Sigma(\alpha) V''(\alpha), 0 \}. \quad (14)$$

This explicit expression for quantity provision allows us to separately identify the role of the value of information in determining the evolution of the equilibrium menus as a function of beliefs.

4.1 Full Market Coverage and Positive Discounting

Full market coverage is obtained in equilibrium when $\phi(\theta_L) > 0$. In this case, the myopic solution, in terms of quantities and total charges, is given by

$$\begin{aligned} q_m(\alpha, \theta) &= \mu(\alpha) \phi(\theta), \\ p_m(\alpha, \theta) &= \mu^2(\alpha) \left(\theta \phi(\theta) - \int_{\theta_L}^{\theta} \phi(s) ds \right). \end{aligned}$$

The following proposition relates the equilibrium menus to the myopic benchmark.

Proposition 5 (Full Market Coverage)

1. *The equilibrium quantities and prices are*

$$\begin{aligned} q(\alpha, \theta) &= q_m(\alpha, \theta) + \Sigma(\alpha) V''(\alpha), \\ p(\alpha, \theta) &= p_m(\alpha, \theta) + \mu(\alpha) \theta_L \Sigma(\alpha) V''(\alpha). \end{aligned}$$

2. *The marginal value of information is*

$$V''(\alpha) = \Sigma(\alpha)^{-1} \left(-\mu(\alpha) \mathbb{E}_{\theta}[\phi] + \sqrt{(\mu(\alpha) \mathbb{E}_{\theta}[\phi])^2 + 2(rV(\alpha) - \Pi_m(\alpha))} \right). \quad (15)$$

A few remarks are in order. First, each type receives $\Sigma(\alpha) V''(\alpha)$ units over and above the myopic quantity level. These additional units constitute the marginal level of experi-

mentation. In this setting, the marginal level of experimentation is constant across types. The number of additional units $\Sigma(\alpha)V''(\alpha)$ need not, however, be monotonic in beliefs α . Second, quantities increase above q_m by $\Sigma(\alpha)V''(\alpha)$, but prices only exceed p_m by $\mu(\alpha)\theta_L\Sigma(\alpha)V''(\alpha)$. This means that each additional unit sold is priced uniformly at $\mu(\alpha)\theta_L$. In other words, the firm charges the lowest type's willingness to pay. Hence it cannot extract any more surplus on the additional units sold. This is a consequence of the fact that buyers are not willing to pay for experimentation, and need to be offered a low enough price in order to be convinced to purchase more. Third, unlike many experimentation models, the value of information does not only depend on the difference $rV(\alpha) - \Pi_m(\alpha)$, but also directly on the current level of demand, captured by $\mu(\alpha)\mathbb{E}_\theta[\phi]$. This term, in fact, is equal to the total quantity sold by the myopic seller. As such, it is a measure of the “default” level of experimentation. It influences the speed of the learning process in the absence of any additional investment in information production by the firm.

Another implication of Proposition 5 is that the effects of new information on the supplied quantity levels depend on the consumer's willingness to pay θ . In particular, types with a virtual valuation above the average $\mathbb{E}_\theta[\phi]$ benefit more from an improvement in beliefs than those with below-average virtual valuations. At the same time, differences between the quantity levels offered to different buyers do not depend on the level of experimentation. The following proposition focuses on the variations in the price-quantity pairs offered to each consumer. These variations constitute the basis for our results on the dynamics of the equilibrium menu's variety.

Proposition 6 (Contract Variety)

1. For all $\theta > \theta'$, quantity differences $q(\alpha, \theta) - q(\alpha, \theta')$ are increasing and linear in α .
2. For all $\theta > \theta'$, differences in total prices charged $p(\alpha, \theta) - p(\alpha, \theta')$ are increasing and convex in α .

If we let $\theta = \theta_H$ and $\theta' = \theta_L$, this result implies that higher values of beliefs α bring about a wider range of options, in terms of offered quantities, and a higher dispersion of total charges.

We would now like to explicitly characterize the behavior of the equilibrium menus as a function of beliefs α . This requires solving the differential equation (15) for the firm's value function. Unfortunately, this differential equation does not have an analytical solution, as it is a second order, nonlinear problem. However, we are able to obtain closed form solutions by analyzing the undiscounted version of the firm's problem.

4.2 No Discounting

For the analysis of the undiscounted version of the problem, we adopt the *strong long run average* criterion.⁷ This approach identifies the limit of the discounted policy functions as the discount rate approaches zero. The solution provided through the strong long run average criterion therefore preserves the qualitative properties of the optimal solution for small discount rates. This criterion also allows us to preserve the recursive formulation of the problem, and to obtain analytical solutions for the policy function. More importantly, the analysis in the previous sections has established a clear link between the firm's patience level and the effects of learning. As the latter are the focus of this paper, considering high degrees of patience removes the risk of assuming learning away because of a high discount rate.

With reference to our model, the strong long run average criterion may be informally summarized as follows: we know beliefs will converge either to μ_L or to μ_H . In the limit for $r \rightarrow 0$, given beliefs α , the return function $rV(\alpha)$ converges to the long run average payoff $v(\alpha)$ defined in (9). However, many policy functions attain the long run average value $v(\alpha)$, independently of their finite time properties.⁸ The main contribution of Dutta (1991) is to prove that the policy function maximizing the undiscounted stream of payoffs, net of their long run averages

$$V(\alpha_0) \triangleq \sup_{q: \Theta \rightarrow \mathbb{R}_+} \mathbb{E}_\alpha \left[\int_0^\infty (\Pi(\alpha_t, q_t) - v(\alpha_t)) d\alpha_t \mid \alpha_0 \right],$$

represents the limit (for $r \rightarrow 0$) of the policy functions maximizing the discounted stream of payoffs (10). The strong long run average solution combines the finite time properties of catching-up optimality and the recursive representation of criteria like the limit of the means. We can therefore write the undiscounted analog of the HJB equation (11) as

$$v(\alpha) = \max_{q: \Theta \rightarrow \mathbb{R}_+} [\Pi(\alpha, q) + Q\Sigma(\alpha) V''(\alpha)], \quad (16)$$

where now $\Sigma(\alpha) V''(\alpha)$ represents the limit marginal value of information. This value does not vanish as $r \rightarrow 0$. On the contrary, Proposition 2 showed that the value of information is increasing in the firm's patience level. In the undiscounted case, we directly solve for this value in closed form. In particular, in the linear quadratic case, the undiscounted equilibrium

⁷This criterion was pioneered by Ramsey (1928). In more recent work, Dutta (1991) discusses the relationship between the strong long run average and other undiscounted optimization criteria.

⁸In particular, any policy that prescribes myopically optimal behavior after beliefs have converged achieves an average payoff of $v(\alpha)$ in the absence of discounting.

quantity supply function is given by

$$q(\alpha, \theta) = \mu(\alpha) (\phi(\theta) - \mathbb{E}_\theta[\phi]) + \sqrt{(\mu(\alpha) \mathbb{E}_\theta[\phi])^2 + (\text{Var}[\phi] + \mathbb{E}_\theta[\phi]^2)\alpha(1-\alpha)(\mu_H - \mu_L)^2}. \quad (17)$$

Theorem 3 (Undiscounted Equilibrium Quantities)

1. The equilibrium quantity $q(\alpha, \theta)$ is strictly concave in α for all θ .
2. There exists a threshold type $\tilde{\theta}$ such that $q(\alpha, \theta)$ is first increasing then decreasing in α for all types $\theta \leq \tilde{\theta}$, and strictly increasing in α for all types $\theta > \tilde{\theta}$.

The main result of Theorem 3 is that experimentation has buyer-dependent implications for the evolution of equilibrium quantities. Contrary to the myopic case, a set of types $[\theta_L, \tilde{\theta}]$ does not always receive a larger number of units as beliefs improve. In particular, the threshold type $\tilde{\theta}$ satisfies the following equation:

$$\phi(\tilde{\theta}) = \frac{\text{Var}[\phi] + \mathbb{E}_\theta[\phi]^2}{2\mathbb{E}_\theta[\phi]} \frac{\mu_H - \mu_L}{\mu_H}. \quad (18)$$

The set of types whose consumption levels are nonmonotonic in beliefs is increasing in the relative quality difference $(\mu_H - \mu_L) / \mu_H$ and in the dispersion of buyers' valuations. This is a consequence of the fact that the firm's profit $p(\alpha, \theta) - c(q(\alpha, \theta))$ is convex in θ . Therefore, an increase in the spread of the distribution $F(\theta)$ improves the firm's profits makes the learning process more significant. The concavity of equilibrium quantities suggests that experimentation is higher for intermediate values of the beliefs. Figure 2(a) confirms this intuition. In this figure, we show the quantity levels supplied to three different buyers $\theta < \theta' < \theta''$ as a function of the market belief α . The lowest type has a zero virtual valuation, and would never be served in the static solution. Figure 2(b) illustrates the equilibrium total charges. Consistent with the result from Proposition 6, the differences between total charges paid by different buyers are increasing and convex in beliefs.

The main properties of the equilibrium menu are best understood by decomposing the implications of the arrival of information into three effects. The first component is one of informational value. Each unit sold generates additional value to the firm by facilitating learning. This effect is highest for intermediate values of the market belief. Conversely, as beliefs approach zero or one, the value of information declines, and so do the incentives for additional quantity provision. This effect influences all types in the same way, since the informational content of a unit sold is independent of the buyer purchasing it. However, the effects of new information are not uniform across buyers because of the second and third components.




FIGURE 2: Undiscounted Equilibrium Quantities and Prices: $\theta \sim U[1, 2]$, $\mu_L = 1$, $\mu_H = 8$

The second component is related to efficiency. When positive news arrive, consumers are willing to pay more for each unit, and gains from trade increase. This effect is stronger for high consumer types, who benefit the most from a quality increase.

The third effect is related to rent extraction. Under adverse selection, the differential increase in buyer's valuations tightens the incentive compatibility constraints and increases the informational rents. This raises the cost of providing the appropriate incentives to high-valuation buyers. To understand this, 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 $U(\theta_H) = U(\theta_L) + \mu(\alpha)(\theta_H - \theta_L)q_L$, which is increasing in α . In other words, positive signals generate additional costs to the seller, driving down consumption for low-valuation buyers as beliefs approach one. In principle, the first two effects could suffice to determine decreasing quantity provision for some types, as α goes to one. However, in several cases, the adverse selection effect is necessary in order to overcome the myopic incentives that lead to increasing quantities. The uniform distribution with $\theta_H < 2\theta_L$ (i.e. under full market coverage) is one such case.

The balance of the value of information, efficiency, and rent extraction 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. Figures 3(a) and 3(b) show the construction of the equilibrium quantities from the marginal value of information and the myopic solution for two different buyers.

The equilibrium quantities are given by the vertical sum of the curve $\Sigma(\alpha)V''(\alpha)$ with each straight line $q_m(\alpha, \theta)$. The shape of the resulting quantity allocation depends on the




FIGURE 3: Myopic Quantities and Value of Information: $\theta \sim U[1, 2]$, $\mu_L = 1$, $\mu_H = 8$

buyer’s willingness to pay θ . Finally, notice that the value of information (hence the difference between the equilibrium and the myopic quantities for each type) peaks at a value of α lower than one-half. To understand why this is the case, consider equation (16). The total informational value $Q\Sigma(\alpha)V''(\alpha)$ is equal to the difference between long run average and current flow profits $v(\alpha) - \Pi(\alpha, q)$. The marginal value of information $\Sigma(\alpha)V''(\alpha)$ therefore indicates the contribution of each unit sold to “compensate” for this difference, $(v(\alpha) - \Pi(\alpha, q))/Q$. The numerator of this ratio depends positively on the degree of uncertainty $\alpha(1 - \alpha)$, which is a measure of how much beliefs can be influenced by the signals observed in the current period. At the same time, the total quantity Q is increasing in α , implying that the ratio of the two terms is decreasing at $\alpha = 1/2$. In words, each (additional) unit’s contribution to the learning process depends negatively on the level of sales, which determines the speed of the learning process in the absence of any (additional) experimentation.

4.3 Nonlinear Prices

Combining our results on contract variety and nonmonotonic quantity provision, we can conclude that the arrival of good news extends the range of offered quantities. In particular, when the value of information is initially high, an improvement in beliefs reduces the quantity level offered to the lowest-valuation buyers. This effect corresponds to a strategy based on introductory pricing. When uncertainty is high, even low-valuation buyers are induced to purchase larger quantities through quantity discounts. As the market obtains positive signals, buyers’ valuations increase, but introductory discounts are greatly reduced. As a

consequence, low-valuation buyers reduce their demands. This feature distinguishes the response of the equilibrium menu to the arrival of information from that of the myopic firm's menu. As the market obtains positive signals, the myopic firm increases the quantity level supplied to all buyers. Figure 4 compares the equilibrium menus $(q, \hat{p}(\alpha, q))$ offered by a myopic firm (4(a)) with those offered by a forward-looking firm (4(b)), as described in this section, for several values of α . Figure 4(b) also highlights the response of the lowest available quantity to the arrival of information.

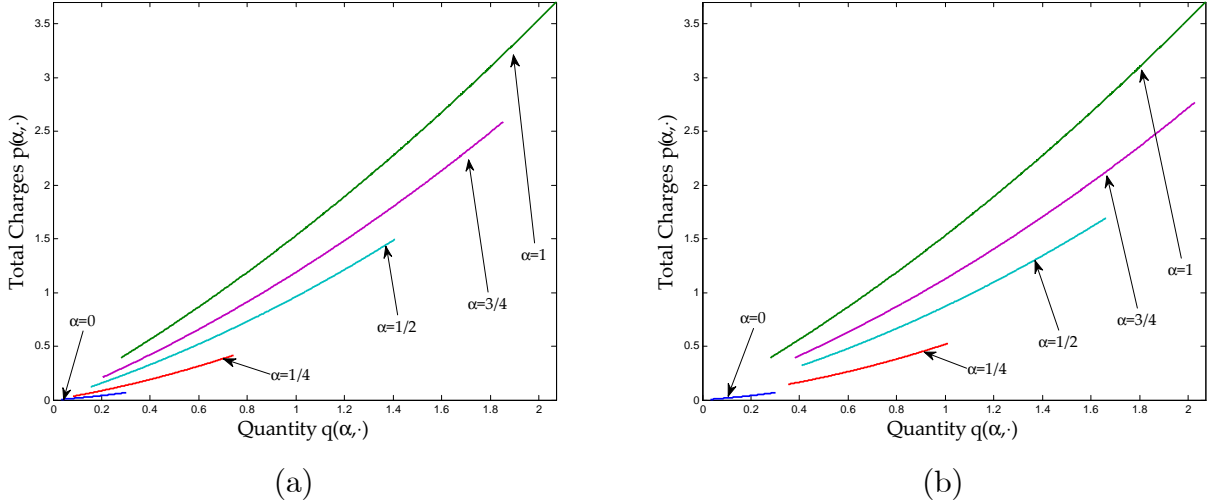


FIGURE 4: Equilibrium Menus: $\theta \sim U[1, 2]$, $\mu_L = 1$, $\mu_H = 8$

The slope of the equilibrium menus corresponds to the marginal prices. When buyers' types are uniformly distributed, marginal prices are given by

$$\hat{p}_q(\alpha, q) = \frac{1}{2} ((q - \Sigma(\alpha) V''(\alpha)) + \mu(\alpha) \theta_H).$$

In the undiscounted uniform case, marginal prices are convex in α . Learning has an intuitive effect on marginal prices. Marginal prices are increasing in α for all quantity levels, provided the difference in quality levels $\mu_H - \mu_L$ is not too high. Conversely, marginal prices are U-shaped in α for all q if learning is relevant enough. Finally, note that the convex costs assumption is responsible for the increasing marginal prices $\hat{p}_{qq}(\alpha, q) > 0$. We use the linear-quadratic setting since it allows for the most straightforward illustration of our results. However, our formulation is essentially equivalent to a model with constant marginal costs and a non separable, quadratic utility function of the form $U(\alpha, \theta, q) = \mu(\alpha) \theta q - q^2/2$. This demand specification, as well as others used in Maskin and Riley (1984), would yield quantity discounts ($\hat{p}_{qq}(\alpha, q) < 0$) for all α, q .

As we have shown, the analysis of the undiscounted problem under full market coverage

delivers explicit solutions that provide insights into the properties of the equilibrium menus. We now extend our findings, by separately relaxing the assumptions of infinite patience and full market coverage.

4.4 Small Positive Discounting

Our first result under small positive discounting uses bounds for the convexity of the value function to establish the concavity of equilibrium quantities in the discounted case. This procedure presents some difficulties, since the second derivative of the value function is unbounded when α goes to zero or one. The only exceptions are given by the myopic profits (since $\Pi_m''(\alpha)$ is a constant), and by the undiscounted profits (since $v''(\alpha) \equiv 0$). In general, as r decreases, the return function $rV(\alpha)$ moves from the quadratic function $\Pi_m(\alpha)$ to the linear function $v(\alpha)$. As in the case of the function $(x(1-x))^k$ for $k \in [1, 2]$, the second derivative is equal to infinity at $x = 0$ and $x = 1$, except for the two limit cases, $k = 1$ and $k = 2$. Our first result extends the concavity property of quantity provision through a careful treatment of the order of limits.

Proposition 7 (Concave Quantities)

For any $\varepsilon \in (0, 1)$, there exists a value of the discount rate r_ε such that, for all $r < r_\varepsilon$, the quantity supply function $q(\alpha, \theta)$ is concave in α for all $\alpha \in [\varepsilon, 1 - \varepsilon]$ and for all $\theta \in \Theta$.

Our second result establishes that all offered quantities are strictly increasing in beliefs when $\alpha = 0$. More importantly, it identifies the minimum patience requirements that allow us to extend the nonmonotonicity of quantity to an arbitrary set of low-valuation buyers. Let $\tilde{\theta}$ be the threshold type defined by (18).

Proposition 8 (Nonmonotonic Quantity Provision)

1. *The quantity $q(\alpha, \theta)$ is increasing in α at $\alpha = 0$ for all r and all θ .*
2. *For every $\theta < \tilde{\theta}$, there exists a value of the discount rate r_θ such that, whenever $r < r_\theta$, $\partial q(1, \theta'; r) / \partial \alpha < 0$ for all $\theta' \in [\theta_L, \theta]$.*

In the next subsection, we relax the full market coverage assumption and analyze the properties of the equilibrium market coverage level.

4.5 Partial Market Coverage

When $\phi(\theta_L) < 0$, it is not always optimal for the monopolist to serve the entire market. We focus on the undiscounted version of the problem and apply the strong long-run average

criterion. For buyers who receive positive quantity provision in equilibrium, the optimal sales level is characterized by first order condition (14):

$$q(\alpha, \theta) = \mu(\alpha) \phi(\theta) + \Sigma(\alpha) V''(\alpha).$$

However, the equilibrium value of information affects the set of buyers receiving positive quantities. In other words, $\Sigma(\alpha) V''(\alpha)$ determines the lowest type served $\theta^*(\alpha)$. After substituting the optimal policy rule as a function of $\Sigma(\alpha) V''(\alpha)$, we can rewrite the firm's problem as follows:

$$v(\alpha) = \int_{\theta^*(\alpha)}^{\theta_H} \frac{1}{2} (\mu(\alpha) \phi(\theta) + \Sigma(\alpha) V''(\alpha))^2 f(\theta) d\theta.$$

The critical type is determined through the equation $q(\alpha, \theta^*(\alpha)) = 0$. In order to obtain a closed form expression for $\Sigma(\alpha) V''(\alpha)$, and hence for $q(\alpha, \theta)$, we assume types are uniformly distributed. We then characterize the equilibrium quantities and market shares in the following proposition.

Proposition 9 (Market Coverage)

The undiscounted equilibrium quantities and market coverage level are given by:

$$\begin{aligned} q(\alpha, \theta) &= (12\mu(\alpha)(\theta_H - \theta_L)v(\alpha))^{1/3} - 2\mu(\alpha)(\theta_H - \theta), \\ 1 - F(\theta^*(\alpha)) &= \frac{\theta_H}{2(\theta_H - \theta_L)} \left(\frac{\alpha\mu_H^2 + (1-\alpha)\mu_L^2}{\mu(\alpha)^2} \right)^{1/3}. \end{aligned}$$

The incentives to experiment lead the firm to serve a larger fraction of types, compared to the myopic solution. These incentives are clearly strongest when the value of information is highest. Market coverage is therefore highest for intermediate values of α , where information is more valuable. However, as in the case of full market coverage, the fraction of buyers served attains a maximum for values of α lower than 1/2.

The case of partial market coverage allows us to clearly show how the arrival of new information is beneficial for some high valuation buyers, but not for others. Figure 5 (on the next page) shows the indirect utility levels for three buyers, as a function of beliefs.

In particular, the lowest valuation buyer shown (θ'') is excluded for high and low enough values of α , while buyers θ and θ' are served for all values of α . However, buyer θ' does not always benefit from the arrival of new (positive) information.




FIGURE 5: Equilibrium Utility Levels: $\theta \sim U[0, 1]$, $\mu_L = 1$, $\mu_H = 8$

5 Intertemporal Patterns

We now consider the dynamics of the equilibrium menus. We derive predictions for the intertemporal evolution of the quantities and prices. The first part of the analysis considers the expected intertemporal patterns from the point of view of market participants. Their beliefs α_t follow the diffusion process described by equation (1). Therefore, by Itô's Lemma, any differentiable function of beliefs $h(\alpha_t, \theta)$, such as prices and quantities, also follows a diffusion process. In particular, the evolution of the process $dh(\alpha_t, \theta)$ is given by

$$dh(\alpha_t, \theta) = \frac{\partial h(\alpha_t, \theta)}{\partial \alpha} d\alpha_t + \frac{1}{2} \frac{\partial^2 h(\alpha_t, \theta)}{(\partial \alpha)^2} (d\alpha_t)^2. \quad (19)$$

Since $\mathbb{E}[d\alpha_t] = 0$, the sign of the drift component of the process $dh(\alpha_t, \theta)$ is determined by the second derivative. In other words, the concavity and convexity properties of any function $h(\alpha_t, \theta)$ may be directly translated into statements about the sign of its unconditional expected change.

Proposition 10 (Unconditional Intertemporal Patterns)

1. For $r = 0$, the quantity $q(\alpha_t, \theta)$ is a supermartingale for all θ .
2. For all r and all θ and θ' , quantity differences $q(\alpha_t, \theta) - q(\alpha_t, \theta')$ are martingales.
3. For all r and all $\theta > \theta'$, total charge differences $p(\alpha_t, \theta) - p(\alpha_t, \theta')$ are submartingales.

Proposition 10 shows that, from the point of view of the agents, supplied quantity levels are expected to decrease over time. Similarly, differences between quantity levels are expected

to shrink. At the same time, differences in the total prices charged to two different buyers are expected to increase. All these findings are consistent with the use of introductory pricing by the firm, which combines lower charges and larger quantities when uncertainty is higher.

While market participants expect beliefs to follow the diffusion process (1), an external observer (*i.e.* the econometrician) knows that the evolution of the process $d\alpha_t$ depends on the true state. Therefore, any empirical prediction about the intertemporal patterns of prices and quantities must be based on the evolution of beliefs conditional on the true quality level. The conditional evolution of beliefs has a non-zero drift component, whose sign depends on the true μ . In particular, for $\mu_i \in \{\mu_L, \mu_H\}$, the general filtering equation (see Liptser and Shiryaev (1977)) is given by

$$d\alpha_t(\mu_i) = \alpha_t(1 - \alpha_t) \frac{\mu_H - \mu_L}{\sigma} \sqrt{Q_t} \left(\frac{\mu_i - \mu(\alpha_t)}{\sigma} \sqrt{Q_t} dt + dz_t \right).$$

Conditional on either state, the drift component of any process $dh(\alpha_t, \theta)$ is no longer uniquely determined by the second partial derivative $\partial^2 h(\alpha_t, \theta) / (\partial \alpha)^2$, but also depends on the first partial derivative $\partial h(\alpha_t, \theta) / \partial \alpha$. In particular, using expression (19), and factoring out common terms, the sign of the drift component of the process $dh(\alpha_t, \theta)$ is determined by the following expressions:

$$\mathbb{E}[dh(\mu_H)] \propto \left(\frac{\partial h(\alpha_t, \theta)}{\partial \alpha} + \frac{\alpha_t}{2} \frac{\partial^2 h(\alpha_t, \theta)}{(\partial \alpha)^2} \right) dt, \quad (20)$$

$$\mathbb{E}[dh(\mu_L)] \propto \left(-\frac{\partial h(\alpha_t, \theta)}{\partial \alpha} + \frac{1 - \alpha_t}{2} \frac{\partial^2 h(\alpha_t, \theta)}{(\partial \alpha)^2} \right) dt. \quad (21)$$

These expressions can be used to derive sufficient conditions under which the expected change in quantities and total charges has an unambiguous sign. In this case, the concavity of the equilibrium quantities does not suffice to conclude that supplied quantities decrease in expectation for all buyers. However, conditional on the bad state μ_L , quantities are expected to decrease for all high-valuation buyers, whose equilibrium quantities are always increasing.

Proposition 11 (Conditional Intertemporal Patterns)

1. *Conditional on the good state ($\mu = \mu_H$):*
 - (a) *for $r = 0$, quantities are expected to decrease whenever $\partial q(\alpha_t, \theta) / \partial \alpha \leq 0$;*
 - (b) *for all r and all $\theta > \theta'$, $q(\alpha_t, \theta) - q(\alpha_t, \theta')$ are submartingales;*
 - (c) *for all r and all $\theta > \theta'$, $p(\alpha_t, \theta) - p(\alpha_t, \theta')$ are submartingales.*
2. *Conditional on the bad state ($\mu = \mu_L$):*
 - (a) *for $r = 0$, quantities are expected to decrease whenever $\partial q(\alpha_t, \theta) / \partial \alpha \geq 0$;*
 - (b) *for all r and all $\theta > \theta'$, $q(\alpha_t, \theta) - q(\alpha_t, \theta')$ are supermartingales.*

If we let $\theta = \theta_H$ and $\theta' = \theta_L$, Proposition 11 suggests that the variety of the offered menu for high quality goods increases over time. Opposite conclusions hold for low quality products.

To summarize, the linear-quadratic model predicts that successful product lines should be characterized by increasing dispersion in prices and in the range of offered quantities. Figure 6(a) shows the results of numerical simulations for the quantity levels offered to two different buyers, conditional on the true quality being high. Figure 6(b) shows the results of numerical simulations for the total charges paid by the same two buyers. As time passes, the quantity supplied to the lower-valuation buyer decreases. However, total charges stay approximately constant, as the firm exploits the consumer's increasing willingness to pay per unit of the product.




FIGURE 6: Simulated Quantities and Total Charges: $\alpha_0 = 1/20$, $\theta \sim U[1, 2]$, $\mu_L = 1$, $\mu_H = 8$

6 Nonlinear Pricing with Competition

The analysis of the monopoly case has allowed us to make precise predictions for the evolution of menus of contracts in markets for experience goods. The monopoly assumption is appropriate in some cases, and the early days of Netflix provide an example. However, often these markets are characterized by imperfect competition, and pricing is strategic. In recent years, for example, Netflix and Blockbuster compete in the market for online DVD rentals. We are therefore motivated to extend our analysis of dynamic menu pricing to a competitive setting. When we introduce strategic pricing in our environment, the role information becomes even more important. The number of units sold by each firm determines the

amount of information obtained by the market about its product. Therefore, from a firm's point of view, increasing its market share contributes towards both revenue maximization and information production. However, a firm might also follow the alternative strategy of reducing market penetration, in order to learn more quickly about its competitors.

The main questions we address in this section are: under which conditions are firms willing to invest in information production about their own product? Are firms willing to invest in information about other products, or do they assign a negative value to learning about their competitors? How do learning incentives affect the equilibrium menus of contracts? Our answers are based on a duopoly model with one product of uncertain quality and one of well-known quality. We measure the two firms' relative strength by comparing the quality of the well-known product and the range of possible qualities for the uncertain product. We find that firms always value learning about their own product, as in the monopoly case. We also find that a firm is willing to invest in learning about its competitor, by reducing its market share, provided it does not hold a particularly strong position.

Consider a duopoly model with menu pricing in a horizontal differentiation framework. Firm 1 is new to the market, and produces a good of uncertain quality $\mu_1 \in \{\mu_L, \mu_H\}$. Firm 2 produces a good of known (safe) quality $\mu_2 = s$. As in our monopoly analysis, consumers value quality uniformly, but have idiosyncratic preferences for the products of each firm. Consumers are indexed by their location $\theta \in [0, 1]$ on a Hotelling line. Firm 1 is located at position 1 and firm 2 at position 0. A consumer's location determines her idiosyncratic taste for each firm's product. For a consumer located at θ , these tastes are given by

$$(\theta_1, \theta_2) = (1 + \theta, 2 - \theta).$$

We assume linear utility, quadratic costs, and uniformly distributed consumers. Given our formulation for (θ_1, θ_2) , the uniform distribution implies that every consumer's virtual valuation for each product is non-negative. The net utility of a buyer type θ purchasing q units from firm j is given by

$$U_j(\theta, \mu_j, q_j, p_j) = \mu_j \cdot \theta_j \cdot q_j - p_j(q_j).$$

Note that each consumer's idiosyncratic valuations for the goods of the two firms are perfectly negatively correlated. In other words, the intensity of a consumer's idiosyncratic taste for her favorite product also determines the intensity of her brand preferences. More specifically, consumers with a higher θ have a higher value for firm 1 and a lower value for firm 2. The incentive compatibility constraints then require firm 1 to provide increasing indirect utility, and firm 2 to provide decreasing indirect utility, as a function of θ .

For each firm $j = 1, 2$, the incentive compatibility constraint is given by

$$U'_j(\theta) = (3 - 2j) \cdot \mu_j q_j(\theta).$$

In particular, if a type $x \in [0, 1]$ is indifferent between the two firms, incentive compatibility implies firm 1 will attract all consumers located at $\theta \geq x$ and firm 2 will serve consumers $\theta < x$. As first shown by Spulber (1989), firms offer the monopoly quantity levels to customers in their market segment, and compete only through the utility offered to the marginal buyer.

In this framework, only the units sold by firm 1 constitute experiments for quality. Firm 2 does not directly control the level of experimentation on each type. Therefore, given market shares $[0, x]$ and $[x, 1]$, firm 2 sets quantity provision at the myopically optimal level. While firm 2 cannot control experimentation on the intensive margin, it can do so on the extensive margin. By letting firm 1 acquire a larger market share, firm 2 can generate more information. Conversely, by pricing more aggressively, firm 2 can increase its own market share, and reduce the overall level of experimentation. Denote by $\Pi_j(\mu_j(\alpha), \theta, q_j)$ the profit level obtained by firm j when selling quantity q_j to buyer θ holding beliefs $\mu_j(\alpha)$ about the product j 's quality (we clearly have $\mu_2(\alpha) \equiv s$). Given a marginal type $x \in [0, 1]$, the two firms' HJB equations may be written as

$$\begin{aligned} rV_1(\alpha) &= \max_{q_1: \Theta \rightarrow \mathbb{R}_+} \left[\int_x^1 \Pi_1(\mu_1(\alpha), \theta, q_1) d\theta + \Sigma(\alpha) V_1''(\alpha) \int_x^1 q_1(\alpha, \theta) d\theta \right], \\ rV_2(\alpha) &= \max_{q_2: \Theta \rightarrow \mathbb{R}_+} \left[\int_0^x \Pi_2(\mu_2(\alpha), \theta, q_2) d\theta + \Sigma(\alpha) V_2''(\alpha) \int_x^1 q_1(\alpha, \theta) d\theta \right]. \end{aligned}$$

Substituting the incentive compatibility constraints and integrating by parts, we obtain formulations for the HJB equations which can be maximized pointwise. The resulting first order conditions determine the equilibrium quantity levels as a function of the value of information:

$$q_1(\alpha, \theta, V_1'') = 2\mu_1(\alpha)\theta + \Sigma(\alpha)V_1''(\alpha) \tag{22}$$

$$q_2(\theta) = 2s(1 - \theta). \tag{23}$$

Equations (22) and (23) determine the optimal quantity provision given market shares. When competing over market shares, firms trade-off the profits made on the marginal buyer ($\Pi_j - u$) with the cost of increasing their market shares, given by the increment in every type's utility $U'_j(x)$ necessary to lure more customers from the competition. Each firm also considers the impact of a marginal change in market shares in terms of information $\Sigma(\alpha)V_j''(\alpha) \cdot q_1(\alpha, x)$. The first order conditions determining the equilibrium indifferent type x and her indirect

utility level u are given by:

$$\Pi_1(\mu_1(\alpha), x, q_1(\alpha, x, V_1'')) - u - (1-x)sq_2(x) + \Sigma(\alpha)V_1''(\alpha)q_1(\alpha, x, V_1'') = 0 \quad (24)$$

$$\Pi_2(\mu_2(\alpha), x, q_2(x)) - u - x\mu(\alpha)q_1(\alpha, x, V_1'') - \Sigma(\alpha)V_2''(\alpha)q_1(\alpha, x, V_1'') = 0 \quad (25)$$

For each value of α , equations (24) and (25) can be solved for the equilibrium market shares and utility level $x(\alpha)$ and $u(\alpha)$ as a function of the marginal values of information $\Sigma(\alpha)V_j''(\alpha)$. The corresponding myopic equilibrium values $x_m(\alpha)$ and $u_m(\alpha)$ are obtained by ignoring the terms $\Sigma(\alpha)V_j''(\alpha)$, and solving the same equations. The analytical expressions for the myopic profits allow us to establish conditions under which both firms assign positive value to information. Analogously to the monopoly case, when beliefs α are high, the risky firm benefits in terms of both profit margins and market shares, so information is valuable. The same logic does not apply in a straightforward way to firm 2. The safe firm might be willing to invest in its own market share, in order to delay learning about its competitor's product. At the same time, when beliefs α are low, firm 2 will also benefit, both in terms of quantities sold to each buyer, and in terms of market shares. We show that the direction of firm 2's learning incentives depends on the relative strength of the two firms. For this reason, let β denote the ratio μ_L/s .

Proposition 12 (Value Function Convexity)

1. $\Pi_{1,m}(\alpha)$ is convex for all β .
2. There exists $\tilde{\beta} > 0$ such that, if $\beta \geq \tilde{\beta}$, then $\Pi_{2,m}(\alpha)$ is convex.
3. If $\Pi_{j,m}(\alpha)$ is convex, then $V_j(\alpha)$ is also convex.

These results confirm the intuition that the risky firm is willing to invest in information generation. Moreover, the safe firm assigns positive value to information when its relative quality is not particularly high. Finally, convexity extends from the myopic profit function to the value function by the same logic as in Theorem 2.

To solve for the Markov equilibrium of this model, we work directly in the undiscounted case. We compute the long run averages $v_j(\alpha)$, as well as the profit levels as a function of both values of information. Once we have done so, the strong long-run average criterion allows us to substitute $rV_j(\alpha)$ with $v_j(\alpha)$ in the two HJB equations, and to solve for the values of information $\Sigma(\alpha)V_j''(\alpha)$. This system must be solved numerically for each α . Figure 7 shows the equilibrium market share of firm 1 as a function the beliefs α .

Just as the incentives for learning imply additional sales in the monopoly model, firm 1's market share lies above the myopic equilibrium level for all values of α . Both firms want to accelerate learning by leaving higher market shares to firm 1. As in the monopoly case, the

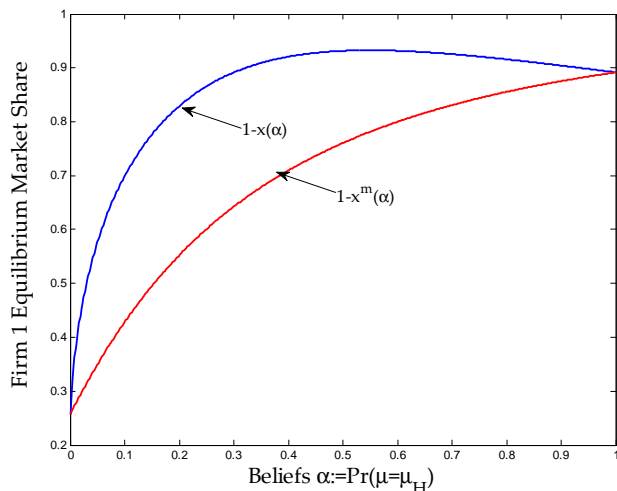


FIGURE 7: Equilibrium Market Shares: $s = 3$, $\mu_L = 2$, $\mu_H = 8$

amount of experimentation is not monotonic in α , and vanishes as beliefs approach zero or one. Furthermore, market shares need not be monotonic in α . When learning about quality is particularly relevant, the dynamic incentives to increase firm 1's sales for intermediate values of α dominate the static forces yielding larger market shares to firm 1 for higher values of α .

7 Concluding Remarks

We have analyzed a firm's dynamic menu pricing strategy in a market characterized by uncertain product quality and heterogeneous buyers. In this environment, the firm can profitably practice second degree price discrimination. Moreover, by adjusting the quantity levels offered to each buyer, the firm can manage the flow of information to the market, and balance information production with short-run revenue maximization. Our model yields tractable closed-form solutions that are informative of the qualitative properties of the optimal contract. It also allows us to predict the patterns of the offered prices and quantities over time. The firm's optimal strategy involves selling larger quantities, even to low-valuation consumers, to generate information when uncertainty is high. As learning occurs, the firm gradually shifts to a more targeted policy that focuses on high-valuation buyers, possibly excluding low-valuation buyers, to extract more surplus.

The model also has clear welfare implications. The additional informational value of each unit sold induces the firm to increase the quantity supplied to each buyer beyond the ideal point of a myopic seller. This counters the distortions induced by adverse selection, which reduce quantities below the efficient level. As a consequence, the firm's incentives to learn lead

to an increase in each buyer's utility and in the overall efficiency of the equilibrium quantity allocation. However, the gradual resolution of the uncertainty is not equally beneficial to all buyers. For example, low-valuation buyers who may be excluded as learning occurs, expect their utility level to decrease over time. Not only them, but also intermediate-valuation buyers, who suffer from more severe distortions when lower types are excluded. This is in sharp contrast with the effects of learning on the highest-valuation buyers. These buyers enjoy larger numbers of units when the product turns out to be of high quality. Moreover, they assign a higher value to each unit. As a consequence, these buyers benefit from information, in expectation terms. These differences in the value of information for different consumers would play an important role in influencing the timing of buyers' choices, in the context of one-time (durable goods) purchases.

The model developed in this paper focuses almost entirely on learning on the firm's side, by abstracting from the effects of individual experience on consumers' beliefs and demand. Though considerably more complicated, a model integrating idiosyncratic and aggregate learning would enable the firm to manage the flow of information to various consumer groups. For example, the firm might offer large quantity discounts to the buyers with the highest valuations, in order to accelerate their learning process, and increase the spread in the distribution of their willingness to pay.

The analysis of strategic environments constitutes a natural extension of this paper. A crucial issue for firms competing in the environment considered here is whether to invest in learning about their own product, and about competitors'. We have provided some intuition for the main trade-offs in a duopoly model. Extending the competitive analysis to richer specifications of brand preferences might provide more insights into the characteristics of menu pricing in imperfectly competitive markets for experience goods.

Another question of interest for future research regards the relative profitability of menu pricing. Several empirical studies have tried to quantify the profit gains due to menu pricing, in comparison to simpler strategies. Our theoretical framework might provide insights into the relationship between firms' choice of pricing scheme and the diffusion of information about the product. Quite surprisingly, evidence from the online DVD and the enterprise software market (more specifically, online hard drive backup services) suggests that firms' strategies follow a pattern that moves away from nonlinear pricing, in favor of linear or flat-rate contracts, as the market gains familiarity with the product.

8 Appendix

8.1 Micro-founded Model for Market Experience

We now provide a formal statistical model for the aggregate market experience and the associated law of motion of market beliefs.⁹ We begin with a finite number of buyers and in discrete time. We are going to suggest a model in which the informativeness of the aggregate market experience is held constant. Informally speaking, the information contributed by the experience of each buyer will become infinitely small, while the market experience will be constant as we increase the number of buyers. In our model, the utility level of a buyer is influenced by two random variables: her willingness-to-pay and her experience with the product. An important feature of the construction will be that, from the point of view of the seller, the randomness due to the willingness-to-pay is going to disappear, but the randomness in the market experience will remain. This is achieved by decreasing the informativeness of each individual buyer's experience proportionally to the size of the market.

Let K be the number of buyers. Each buyer's willingness-to-pay, θ_i , is independently and identically drawn from a distribution $F(\theta_i)$. Each unit j purchased by buyer i generates a normally distributed payoff $\tilde{x}_{ij} \sim N(\mu/K, \sigma^2/K)$. We refer to the realization of \tilde{x}_{ij} , denoted by x_{ij} , as the *experience of buyer i with unit j of the product*. The realized gross utility of a buyer with type θ_i who consumes q_i units with an experience vector $x_i = (x_{i1}, \dots, x_{iq_i})$ is given by

$$U(\theta_i, q_i, x_i) = \gamma(q_i) \cdot \theta_i \cdot \sum_{j=1}^{q_i} x_{ij}.$$

Note that the gross utility is allowed to depend directly on the number of units consumed via the function $\gamma(q)$. The resulting expected utility function

$$\mathbb{E}_{x_i}[U(\theta_i, q_i, x_i)] = \gamma(q_i) \cdot \theta_i \cdot (\mu/K) \cdot q_i,$$

is consistent with the formulation (3) in Section 2.

We assume that the individual experience of each buyer i , x_i , is observed by all market participants, *i.e.* all the buyers and the seller. If each buyer θ_i consumes a quantity level $q(\theta_i)$, the market experience is the sum of the individual experiences $\sum_{i=1}^K \sum_{j=1}^{q(\theta_i)} x_{ij}$. The market experience is now normally distributed with mean $(\mu/K) \sum_{i=1}^K q(\theta_i)$ and variance $(\sigma^2/K) \sum_{i=1}^K q(\theta_i)$. Denote the average number of units purchased by $Q_K = (1/K) \sum_{i=1}^K q(\theta_i)$. An important feature of this construction is that, as we increase the number of buyers (*i.e.* the number of draws, K , from the distribution $F(\theta)$), the realized distribution of willingness-to-pay will coincide with the theoretical distribution, and is thus deterministic.

⁹I thank William Hogan and Richard Zeckhauser for pointing out a mistake in an earlier version.

In particular, the average number of units converges to the expected purchased quantity, $Q_K \rightarrow Q \triangleq \int_{\theta_L}^{\theta_H} q(\theta) f(\theta) d\theta$. The market experience, on the other hand, will remain a true random variable. Thus, in the limit for $K \rightarrow \infty$, the aggregate market experience is normally distributed with mean $Q\mu$ and variance $Q\sigma^2$. As we take the continuous-time limit and use subscripts for time dependence, the flow of new information follows a Brownian motion with drift $Q_t\mu$ and variance $Q_t\sigma^2$:

$$d\tilde{x}_t = Q_t\mu dt + \sqrt{Q_t}\sigma dz_t.$$

With this structure for the information flow, one can use the filtering equations (see Theorem 9.1 in Liptser and Shiryaev (1977)) to derive the evolution of beliefs α_t :

$$d\alpha_t = \alpha_t(1 - \alpha_t) \frac{\mu_H - \mu_L}{\sigma} \sqrt{Q_t} dz_t.$$

8.2 Proofs of Propositions

Proof of Proposition 1. (1.) By the implicit function theorem and assumption 2, the first partial derivatives of the myopic solution are given by

$$\begin{aligned} \frac{\partial q_m(\alpha, \theta)}{\partial \alpha} &= -\frac{(\mu_H - \mu_L) \phi(\theta) u'(q)}{\mu(\alpha) \phi(\theta) u''(q) - c''(q)} > 0, \\ \frac{\partial q_m(\alpha, \theta)}{\partial \theta} &= -\frac{\mu(\alpha) \phi'(\theta) u'(q)}{\mu(\alpha) \phi(\theta) u''(q) - c''(q)} > 0. \end{aligned}$$

(2.) Apply the envelope theorem and use part (1.) to obtain the following expressions for the derivatives of $\Pi_m(\alpha)$:

$$\begin{aligned} \Pi'_m(\alpha) &= \mathbb{E}_\theta [(\mu_H - \mu_L) \phi(\theta) u(q_m)] > 0, \\ \Pi''_m(\alpha) &= \mathbb{E}_\theta \left[(\mu_H - \mu_L) \phi(\theta) u'(q_m) \frac{\partial q_m(\alpha, \theta)}{\partial \alpha} \right] > 0, \end{aligned}$$

which ends the proof. ■

The next lemma shows how to reformulate the HJB equation (11) as a second order differential equation.

Lemma 1 (Boundary-Value Problem)

Let $\Sigma(\alpha)$ be defined by (2) and let $Q = \int_{\Theta} q(\theta) dF(\theta)$. For all $\alpha \in (0, 1)$, the HJB equation (11) may be written as the following boundary value problem:

$$V''(\alpha) = \min_{q: \Theta \rightarrow \mathbb{R}_+} \frac{rV(\alpha) - \Pi(\alpha, q)}{\Sigma(\alpha) Q}, \quad (26)$$

with boundary conditions

$$rV(0) = \Pi_m(0), \quad (27)$$

$$rV(1) = \Pi_m(1). \quad (28)$$

Proof. The ratio in (26) is well defined for all $\alpha \in (0, 1)$, since Q must be positive for any q solving the HJB equation (11). Suppose instead that $Q = 0$. Then we know $q \neq q_m$, since $q_m(\alpha, \theta) > 0$ for all α and θ . Therefore, $\Pi(\alpha, q) < \Pi_m(\alpha)$. However, in the HJB equation (11), $Q = 0$ would imply $rV(\alpha) = 0$, which is impossible since we know the firm can guarantee itself $\Pi_m(\alpha)$ in every period by setting $q = q_m$. It follows that, if a function $V(\alpha)$ solves the HJB equation (11), then it must also solve the differential equation (26). ■

We now state an existence theorem for boundary value problems due to Bernfeld and Lakshmikantham (1974), which we then use to prove Theorem 1. This result requires the concept of supersolution and subsolution and the introduction of a regularity condition.

Consider a second order differential equation of the form

$$V'' = G(\alpha, V, V') \quad (29)$$

on an open interval $J_\varepsilon = (\varepsilon, 1 - \varepsilon)$ with $\varepsilon \geq 0$. Let V_L and V_H be functions with continuous second derivatives on J . The function V_L is called a subsolution of (29) if $V_L'' \geq G(\alpha, V_L, V_L')$ on J . Similarly, a function V_H is a supersolution if $V_H'' \leq G(\alpha, V_H, V_H')$ on J . If these inequalities are strict, these functions are called strict sub- and supersolutions. Fix two functions V_H and V_L such that $V_L \leq V_H$ on \bar{J} . The function $G(\alpha, V, V')$ is said to be regular with respect to V_H and V_L if it is continuous on $S_\varepsilon \{(\alpha, V_0, V_1) \in J_\varepsilon \times \mathbb{R} \times \mathbb{R} : V_L(\alpha) \leq V_0 \leq V_H(\alpha)\}$ and there is a constant $C(\varepsilon)$ such that $|G(\alpha, V, V_1)| \leq C(\varepsilon)(1 + |V_1|^2)$ on S_ε .

We can adapt Theorem 1.5.1 in Bernfeld and Lakshmikantham (1974) to our framework, to show existence of a solution.

Lemma 2 (Existence and Uniqueness)

Consider an interval $J_\varepsilon \triangleq (\varepsilon, 1 - \varepsilon)$. Suppose V_L is a subsolution and V_H a supersolution of (29) on J_ε , and $V_L \leq V_H$. Suppose further that G is regular with respect to V_L and V_H on \bar{J}_ε . Given any pair of boundary conditions $V(\varepsilon) \in [V_L(\varepsilon), V_H(\varepsilon)]$ and $V(1 - \varepsilon) \in [V_L(1 - \varepsilon), V_H(1 - \varepsilon)]$, (29) has a C^2 solution on J_ε which satisfies the boundary conditions. Moreover, for all $\alpha \in \bar{J}_\varepsilon$, $V_L(\alpha) \leq V(\alpha) \leq V_H(\alpha)$. If V_L is a strict subsolution, $V > V_L$ and if V_H is a strict supersolution $V < V_H$ on J_ε . Moreover, for all $\alpha \in J_\varepsilon$, $|V'(\alpha)| < N$, where N only depends on $C(\varepsilon)$ and on the functions V_L and V_H .

We also adapt Corollary 1.5.1 from Bernfeld and Lakshmikantham (1974) to show convergence properties of our solution.

Lemma 3 (Uniform Convergence)

Under the assumptions of Lemma 2, any infinite sequence of solutions of (29), with $V_L(\alpha) \leq V(\alpha) \leq V_H(\alpha)$ on J_ε has a uniformly convergent subsequence converging to a solution of (29) on J_ε .

We can now use these results and follow the steps of Faingold and Sannikov (2007) to prove existence and uniqueness of a solution.

Claim 1 *The myopic profit function $\Pi_m(\alpha)/r$ is a strict subsolution of (26) on $(0, 1)$.*

Proof. By Proposition 1, $\Pi_m''(\alpha) > 0$ and $\min_q ((\Pi_m(\alpha) - \Pi(\alpha, q)) / \Sigma(\alpha) Q) = 0$. ■

Claim 2 *The long run payoff $v(\alpha)/r$ is a strict supersolution of (26) on $(0, 1)$.*

Proof. We know that $v(\alpha)$ is linear by definition, while $\min_q ((v(\alpha) - \Pi(\alpha, q)) / \Sigma(\alpha) Q) > 0$. In fact, $\max_q \{\Pi(\alpha, q)\} \triangleq \Pi_m(\alpha)$, is a strictly convex function. Therefore, $\Pi_m(\alpha) < v(\alpha)$ on $(0, 1)$, and $v(\alpha) - \Pi(\alpha, q) > 0$ for all $\alpha \in (0, 1)$ and all functions q . ■

Claim 3 *Fix an interval $J_\varepsilon = (\varepsilon, 1 - \varepsilon)$. The boundary value problem (26) is regular with respect to Π_m and v on J_ε .*

Proof. It suffices to show that there exists a constant $C > 0$ such that, for all $(\alpha, V) \in J_\varepsilon \times \mathbb{R}_+$ with $rV \in [\Pi_m(\alpha), v(\alpha)]$, the following obtains:

$$\min_q \frac{rV - \Pi(\alpha, q)}{\Sigma(\alpha) Q} \leq C. \tag{30}$$

We know that this ratio is always positive and that the first term in the numerator is bounded from above by $v(\alpha)$. Furthermore, we can show that Q is bounded from below by Q_m . Suppose in fact that

$$\tilde{q} = \arg \min_q ((rV - \Pi(\alpha, q)) / \Sigma(\alpha) Q),$$

and that $\tilde{Q} < Q_m$. Then we would have

$$(rV - \Pi(\alpha, \tilde{q})) / \Sigma(\alpha) \tilde{Q} < (rV - \Pi(\alpha, q_m)) / \Sigma(\alpha) Q_m,$$

which yields a contradiction. In fact, $\tilde{Q} < Q_m$ implies the right hand side's denominator is larger than the left hand side's, while $\Pi(\alpha, q_m) = \Pi_m(\alpha) > \Pi(\alpha, \tilde{q})$ implies the numerator of the right hand side is smaller than the left hand side's. Moreover, if the solution to (30) is different from q_m , then it must achieve a lower value than q_m does. We can then define the uniform bound as

$$C(\varepsilon) = \max_{\alpha \in J_\varepsilon} \left\{ \frac{v(\alpha) - \Pi_m(\alpha)}{\Sigma(\alpha) Q_m(\alpha)} \right\}, \quad (31)$$

which ends the proof. ■

Proof of Theorem 1. (1.) We know the HJB is equivalent to the boundary value problem (26). Furthermore, this problem satisfies all conditions of Lemma 2. Therefore, for all $\varepsilon > 0$, the boundary value problem (26) admits a C^2 solution on $[\varepsilon, 1 - \varepsilon]$ with boundary conditions $rV(\varepsilon) \in [\Pi_m(\varepsilon), v(\varepsilon)]$ and $rV(1 - \varepsilon) \in [\Pi_m(1 - \varepsilon), v(1 - \varepsilon)]$.

Now let $\varepsilon = 1/n$ and fix the closed interval $\bar{J}_n \triangleq [1/n, 1 - 1/n]$. Similarly, let $s \geq n$ and consider a solution $V_s(\alpha)$ to (26) on the interval $[1/s, 1 - 1/s]$. Define a sequence of functions V_s^n , where for each $s_j > n$, $V_{s_j}^n(\alpha)$ is the restriction of $V_{s_j}(\alpha)$ to \bar{J}_n . By Lemma 3, for each n , the sequence V_s^n has a converging subsequence. By a standard diagonalization argument, there exist a convergent subsequence (which we define as V_n) converging pointwise to a function defined on $V : (0, 1) \rightarrow \mathbb{R}$. By Lemma 2, $|V_n'|$ is uniformly bounded, hence on any closed subinterval $\bar{J}_n \subset (0, 1)$, $V_n \rightarrow V$ uniformly on I . Moreover, for a given interval \bar{J}_n , the bound $C(n)$ defined in (31) constitutes a uniform bound on $|V_n''|$. Therefore, V_n' is locally Lipschitz, hence it converges uniformly to V' on any closed subinterval $\bar{J}_n \subset (0, 1)$.

The solution to the boundary value problem (26) is unique. Suppose instead there were two solutions $V_1(\alpha)$ and $V_2(\alpha)$ to (26), with $V_1(\alpha) \neq V_2(\alpha)$. Without loss of generality, suppose $V_2 > V_1$ for some α . Define $G(\alpha, V, V') \triangleq \min_q ((rV - \Pi(\alpha, q)) / \Sigma(\alpha) Q)$. The function G is strictly increasing in V by the envelope theorem. Since the boundary conditions are identical, the function $V_2 - V_1$ attains a local maximum on $(0, 1)$ with $V_2 > V_1$. At the maximum, $V_2'' - V_1'' < 0$; therefore, the HJB equations imply $G(\alpha, V_1, V') > G(\alpha, V_2, V')$ which contradicts $V_1 < V_2$.

(2.) Under assumption 2, the pointwise maximization of (12) admits a unique solution. We know from part (1.) that a solution $V(\alpha)$ exists. Therefore $q(\alpha, \theta)$ is the only policy attaining it. We can then apply the implicit function theorem to obtain the following expressions for the first partial derivatives:

$$\begin{aligned} \frac{\partial q(\alpha, \theta)}{\partial \alpha} &= - \frac{(\mu_H - \mu_L) \phi(\theta) u'(q) + (d/d\alpha) (\Sigma(\alpha) V''(\alpha))}{\mu(\alpha) \phi(\theta) u''(q) - c''(q)}, \\ \frac{\partial q(\alpha, \theta)}{\partial \theta} &= - \frac{\mu(\alpha) \phi'(\theta) u'(q)}{\mu(\alpha) \phi(\theta) u''(q) - c''(q)}. \end{aligned}$$

By assumption (b), these ratios are well defined. Moreover, formulation (26) and the envelope theorem imply that $(d/d\alpha)(\Sigma(\alpha)V''(\alpha))$ is equal to $Q^{-1}(rV'(\alpha) - (\mu_H - \mu_L)\mathbb{E}_\theta[\phi(\theta) \cdot u(q)])$, and therefore it is continuous in α .

(3.) We verify three necessary conditions for the application of a verification theorem. First, by part (1.), there exists a \mathcal{C}^2 solution $V(\alpha)$ to the HJB equation. Second, the solution to the HJB equation delivers bounded expected profits for all α (since $V(\alpha)$ is bounded by $v(\alpha)/r$). It follows that $\limsup_{t \rightarrow \infty} e^{-rt}\mathbb{E}_\alpha(V(\alpha_t)) = 0$. Third, from part (2.), there exists a \mathcal{C}^1 policy $q : [0, 1] \times \Theta \rightarrow \mathbb{R}_+$ that maximizes the right-hand side of the HJB equation (11). We can therefore apply Theorem III.9.1 in Fleming and Soner (2006) and conclude that $V(\alpha)$ achieves the maximum of (10). ■

Proof of Theorem 2. From the HJB equation (11), it follows that $rV(0) = \Pi_m(0)$ and $rV(1) = \Pi_m(1)$. Consequently, if there exists a point α for which $\Pi_m(\alpha) > rV(\alpha)$, then there must also exist a point $\hat{\alpha}$ where that the function $\Pi_m - rV$ attains a local maximum. Therefore, $\Pi_m''(\hat{\alpha}) - rV''(\hat{\alpha}) < 0$. By Proposition 1, this implies rV is convex at $\hat{\alpha}$. However, whenever $V''(\alpha)$ is positive, the HJB equation implies $rV(\alpha) > \max_q [\Pi(\alpha, q)] \triangleq \Pi_m(\alpha)$, contradicting $\Pi_m(\alpha) > rV(\alpha)$. We therefore know that $rV(\alpha) \geq \Pi_m(\alpha)$ for all α . A similar argument rules out the case $V''(\alpha) < 0$ and $rV(\alpha) \geq \Pi_m(\alpha)$ for any α . It follows that we must have $V''(\alpha) \geq 0$ for all α . ■

Proof of Proposition 2. (1.) Define the return function $W_r(\alpha) \triangleq rV(\alpha)$, and the function $G(\alpha, W, W') \triangleq \min_q (r\sigma^2(W - \Pi(\alpha, q)) / (\mu_H - \mu_L)^2 (\alpha(1 - \alpha))^2 Q)$. The boundary conditions for (26) are given by $W(\alpha) = \Pi_m(\alpha)$ for $\alpha \in \{0, 1\}$. To prove the result, let $r_2 > r_1$ and suppose that for some α , $W_{r_2}(\alpha) > W_{r_1}(\alpha)$. Then $W_{r_2} - W_{r_1}$ must attain a local maximum. At the maximum point, we then have $W_{r_2}'' - W_{r_1}'' < 0$. The formulation (26) of the HJB equation then implies $G(\alpha, W_{r_1}, W_{r_1}') > G(\alpha, W_{r_2}, W_{r_2}')$, contradicting $r_2 > r_1$ and $W_{r_2}(\alpha) > W_{r_1}(\alpha)$. Since σ and r both enter (26) multiplicatively, an identical argument shows that $W(\alpha)$, and hence $V(\alpha)$, depend negatively on σ .

(2.) Holding $\mu(\alpha_0)$ constant while increasing $(\mu_H - \mu_L)$ induces a mean-preserving spread in the process μ_t . Since the profit function is linear in μ , the value function $\Pi^*(\alpha_0)$ increases, and so does the return function $W(\alpha_0; r)$. Since $\Sigma(\alpha)V''(\alpha)$ is related to $W_r(\alpha)$ by equation (26), a straightforward application of the envelope theorem delivers that the value of information depends positively on the value of the problem, and hence on the return function $W_r(\alpha)$, and on the difference $(\mu_H - \mu_L)$. ■

Proof of Proposition 3. (1.) Let $\Lambda(\alpha) = \Sigma(\alpha)V''(\alpha)$. From first order condition (13)

and the implicit function theorem, we have

$$\frac{\partial q(\alpha, \theta, \Lambda)}{\partial \Lambda} = -\frac{1}{\mu(\alpha) \phi(\theta) u''(q) - c''(q)} > 0.$$

Since the value of information $\Lambda(\alpha) \equiv 0$ in the myopic case and $\Lambda(\alpha) > 0$ in the forward-looking case, quantity is higher in the latter setting.

(2.) Similarly, quantity is increasing in the value of information $\Lambda(\alpha)$ for all α and θ . ■

Proof of Proposition 4. (1.) Since $\Sigma(\alpha) V''(\alpha) \geq 0$ and $\phi(\theta)$ is increasing, for any level q offered at α both by the myopic and the forward looking firm, the corresponding marginal price $\hat{p}_q(\alpha, q)$ is lower in the latter case.

(2.) Since $\phi(\theta)$ is increasing, the higher the value of information, the lower the marginal prices. ■

Proof of Proposition 5. (1.) Directly substituting $\Sigma(\alpha) V''(\alpha) = 0$ in (14) gives us the expression for $q_m(\alpha, \theta)$. From first order condition (14), and using constraint (4), we obtain

$$\begin{aligned} p(\alpha, \theta) &= \mu(\alpha) \theta q(\alpha, \theta) - U(\alpha, \theta) \\ &= \mu(\alpha) \left(\theta (\mu(\alpha) \phi(\theta) + \Sigma(\alpha) V''(\alpha)) - \int_{\theta_L}^{\theta} (\mu(\alpha) \phi(s) + \Sigma(\alpha) V''(\alpha)) ds \right) \\ &= \mu^2(\alpha) \left(\theta \phi(\theta) + \int_{\theta_L}^{\theta} \phi(s) ds \right) + \mu(\alpha) \theta_L \Sigma(\alpha) V''(\alpha). \end{aligned} \quad (32)$$

(2.) Substituting first order condition (14) into the objective function, we obtain:

$$\begin{aligned} rV(\alpha) &= \int_{\theta_L}^{\theta_H} \frac{1}{2} (\mu(\alpha) \phi(\theta) + \Sigma(\alpha) V''(\alpha))^2 f(\theta) d\theta \\ &= \int_{\theta_L}^{\theta_H} \frac{1}{2} (\mu(\alpha) \phi(\theta))^2 f(\theta) d\theta + \mu(\alpha) \mathbb{E}_{\theta}[\phi] \Sigma(\alpha) V''(\alpha) + \frac{1}{2} (\Sigma(\alpha) V''(\alpha))^2 \end{aligned} \quad (33)$$

The first term in (33) is exactly the expression for the firm's myopic profits $\Pi_m(\alpha)$ in this context. We can then solve explicitly for $\Sigma(\alpha) V''(\alpha)$ and obtain

$$\Sigma(\alpha) V''(\alpha) = -\mu(\alpha) \mathbb{E}_{\theta}[\phi] + \sqrt{(\mu(\alpha) \mathbb{E}_{\theta}[\phi])^2 + 2(rV(\alpha) - \Pi_m(\alpha))},$$

which ends the proof. ■

Proof of Proposition 6. (1.) From condition (14), the difference in the quantity levels supplied to types θ and θ' is equal to $\mu(\alpha) (\phi(\theta) - \phi(\theta'))$, hence it is linear in α .

(2.) Differences between total charges are given by $p(\alpha, \theta) - p(\alpha, \theta')$ which simplifies to $\mu^2(\alpha) \int_{\theta'}^{\theta} s \phi'(s) ds$. Hence, these differences are positive and convex in α for all $\theta > \theta'$. ■

Proof of Theorem 3. (1.) The first term in expression (17) is linear in α . The term inside the square root is concave, since its second derivative with respect to α is given by

$$2(\mathbb{E}_{\theta}[\phi](\mu_H - \mu_L))^2 - 2(\text{Var}[\phi] + \mathbb{E}_{\theta}[\phi]^2)(\mu_H - \mu_L)^2 = -2(\mu_H - \mu_L)^2 \text{Var}[\phi].$$

Therefore, $q(\cdot, \theta)$ is a concave function of α .

(2.) Using the concavity of $q(\cdot, \theta)$, and the fact that $\partial^2 q(\alpha, \theta) / \partial \alpha \partial \theta = (\mu_H - \mu_L) \phi'(\theta) > 0$, we can identify the critical type $\tilde{\theta}$ receiving nonmonotonic quantity provision by setting $\partial q(1, \tilde{\theta}) / \partial \alpha = 0$.

$$\begin{aligned} \frac{\partial q(1, \theta)}{\partial \alpha} &\propto \phi - \mathbb{E}_{\theta}[\phi] + \frac{2\mu_H \mathbb{E}_{\theta}[\phi]^2 - (\text{Var}[\phi] + \mathbb{E}_{\theta}[\phi]^2)(\mu_H - \mu_L)}{2\mu_H \mathbb{E}_{\theta}[\phi]} \\ &\propto \frac{2\mu_H \mathbb{E}_{\theta}[\phi](\phi - \mathbb{E}_{\theta}[\phi]) + 2\mu_H (\mathbb{E}_{\theta}[\phi])^2 - (\text{Var}[\phi] + \mathbb{E}_{\theta}[\phi]^2)(\mu_H - \mu_L)}{2\mu_H \mathbb{E}_{\theta}[\phi]}. \end{aligned}$$

The threshold $\tilde{\theta}$ is defined by the following equation:

$$\phi(\tilde{\theta}) = \frac{\text{Var}[\phi] + \mathbb{E}_{\theta}[\phi]^2}{2\mathbb{E}_{\theta}[\phi]} \frac{\mu_H - \mu_L}{\mu_H},$$

which completes the proof. ■

Proof of Proposition 7. Consider first order condition (14) for the equilibrium quantity function $q(\alpha, \theta)$. Parametrize the solution $q(\alpha, \theta)$ and the value function $V(\alpha)$ by the discount rate r . The first derivative with respect to α is given by

$$\frac{\partial q(\alpha, \theta; r)}{\partial \alpha} = (\mu_H - \mu_L)(\phi(\theta) - \mathbb{E}_{\theta}[\phi]) + \frac{rV'(\alpha; r) - (\mu_H - \mu_L)\mu(\alpha)\text{Var}[\phi]}{\sqrt{2rV(\alpha; r) - \mu(\alpha)^2\text{Var}[\phi]}}. \quad (34)$$

The second derivative is given by

$$\frac{\partial^2 q(\alpha, \theta; r)}{(\partial \alpha)^2} = \frac{rV''(\alpha; r) - (\mu_H - \mu_L)^2 \text{Var}[\phi] - (rV'(\alpha; r) - (\mu_H - \mu_L)\mu(\alpha)\text{Var}[\phi])^2}{\sqrt{2rV(\alpha; r) - \mu(\alpha)^2\text{Var}[\phi]}}. \quad (35)$$

Now consider an interval $[\varepsilon, 1 - \varepsilon]$ with $\varepsilon > 0$. By equation (30), the second derivative of the value function $V''(\alpha; r)$ is uniformly bounded from above by $C(\varepsilon)$ for all r . The bound $C(\varepsilon)$ is defined in equation (31). From expression (35), it is clear that $rV''(\alpha; r) -$

$(\mu_H - \mu_L)^2 \text{Var}[\phi] \leq 0$ ensures that $\partial^2 q(\alpha, \theta; r) / (\partial \alpha)^2 < 0$. Therefore, if the discount rate r is lower than the threshold $r_\varepsilon \triangleq (\mu_H - \mu_L)^2 \text{Var}[\phi] / C(\varepsilon)$, then quantity provision $q(\alpha, \theta; r)$ is concave in α over the interval $[\varepsilon, 1 - \varepsilon]$. Furthermore, since the second derivative $\partial^2 q(\alpha, \theta; r) / (\partial \alpha)^2$ does not depend on the buyer's type, the result holds for all θ . ■

Proof of Proposition 8. (1.) Consider again the derivative $\partial q(\alpha, \theta; r) / \partial \alpha$, given in equation (34). Evaluate expression (34) at $\alpha = 0$. Since we know that $rV(\alpha; r) \geq \Pi_m(\alpha)$ for all α and for all r , we can conclude that $rV'(0; r) \geq \Pi'_m(0)$. Using the fact that $rV(0; r) = \Pi_m(0)$ for all r , and that $\Pi'_m(\alpha) = (\mu_H - \mu_L) \mu(\alpha) \mathbb{E}_\theta[\phi^2]$, we obtain the following expressions:

$$\begin{aligned} \frac{\partial q(0, \theta; r)}{\partial \alpha} &= (\mu_H - \mu_L) (\phi(\theta) - \mathbb{E}_\theta[\phi]) + \frac{rV'(0; r) - (\mu_H - \mu_L) \mu_L \text{Var}[\phi]}{\sqrt{\mu_L^2 (\mathbb{E}_\theta[\phi^2] - \text{Var}[\phi])}} \\ &\geq (\mu_H - \mu_L) (\phi(\theta) - \mathbb{E}_\theta[\phi]) + (\mu_H - \mu_L) \sqrt{\mathbb{E}_\theta[\phi^2] - \text{Var}[\phi]} \\ &= (\mu_H - \mu_L) \phi(\theta) > 0. \end{aligned}$$

Therefore, quantity provision is increasing in α around $\alpha = 0$ for all types θ with positive virtual valuation $\phi(\theta)$.

(2.) Evaluate expression (34) at $\alpha = 1$ and let $W_r(\alpha) = rV(\alpha; r)$. We can then write $\partial q(\alpha, \theta; r) / \partial \alpha$ as

$$\frac{\partial q(1, \theta; r)}{\partial \alpha} = (\mu_H - \mu_L) (\phi(\theta) - \mathbb{E}_\theta[\phi]) + \frac{W'_r(1) - (\mu_H - \mu_L) \mu_H \text{Var}[\phi]}{\mu_H \mathbb{E}_\theta[\phi]}. \quad (36)$$

We know the derivative $W'_r(1)$ is increasing in r , since $W_r(\alpha)$ is convex in α and decreasing in r for all α , and, at $\alpha = 1$, we have $W_r(1) = \Pi_m(1)$ for all r . It follows the right-hand side of (36) is increasing in r . The right-hand side of (36) is also increasing in θ , since it depends positively on $\phi(\theta)$. In the undiscounted case, we have $W'_r(1) = v'(1)$. When $r = 0$, we can identify a threshold type $\tilde{\theta}$ that solves $\partial q(1, \tilde{\theta}; 0) / \partial \alpha = 0$. Moreover, since $W_\alpha(1; r)$ is increasing in r , for each ε we can find a discount rate r_ε such that $|W'_r(1) - v'(1)| < \varepsilon$ for all $r < r_\varepsilon$. Since the right-hand side of (36) is increasing in θ , for any θ' lower than the undiscounted threshold $\tilde{\theta}$, we can find a value for the discount rate r low enough so that θ' solves $\partial q(1, \theta'; r) / \partial \alpha = 0$. For all $r < r_{\theta'}$, we then obtain decreasing quantities $q(\alpha, \theta)$ at $\alpha = 1$ for all $\theta \in [\theta_L, \theta']$. ■

Proof of Proposition 9. The firm's HJB equation is given by

$$v(\alpha) = \int_{\theta^*(\alpha)}^{\theta_H} \frac{1}{2} (\mu(\alpha) \phi(\theta) + \Sigma(\alpha) V''(\alpha))^2 f(\theta) d\theta,$$

which may also be written as

$$v(\alpha) = \frac{1}{6\mu(\alpha)} \int_{\theta^*(\alpha)}^{\theta_H} \frac{d}{d\theta} (\mu(\alpha) \phi(\theta) + \Sigma(\alpha) V''(\alpha))^3 \frac{f(\theta)}{\phi'(\theta)} d\theta.$$

We assume that types θ are uniformly distributed on $[\theta_L, \theta_H]$. Under the uniform distribution $f(\theta)/\phi'(\theta)$ is a constant equal to $(2(\theta_H - \theta_L))^{-1}$. We can then integrate out the previous expression, solve the equation

$$v(\alpha) = \frac{(\mu(\alpha) \theta_H + \Sigma(\alpha) V''(\alpha))^3}{12\mu(\alpha) (\theta_H - \theta_L)},$$

for the value of information, and obtain the result. ■

Proof of Proposition 10. Since the drift component of the process $dh(\alpha_t, \theta)$ is given by

$$\mathbb{E}[dh(\alpha_t, \theta)] = \frac{\partial^2 h(\alpha_t, \theta)}{(\partial \alpha)^2} \Sigma(\alpha_t) Q_t dt, \quad (37)$$

the result follows directly from equation (37), Proposition 6 and Theorem 3. ■

Proof of Proposition 11. This result follows directly from equations (20) and (21), from Proposition 6, and from Theorem 3. ■

Proof of Proposition 12. (1.) Imposing $\Sigma(\alpha) V_j''(\alpha) \equiv 0$ in conditions (24) and (25), one obtains the following expressions for the myopic equilibrium market shares and utility levels

$$\begin{aligned} x_m(\alpha) &= \frac{4s^2}{\mu_1(\alpha)^2 + 3s^2 + \sqrt{(\mu_1(\alpha)^2 + 3s^2)^2 + 8s^2(\mu_1(\alpha)^2 - s^2)}}, \\ u_m(\alpha) &= 2(1 - x_m(\alpha))s^2 - 2(x_m(\alpha)\mu_1(\alpha))^2, \end{aligned}$$

The resulting myopic profit levels are given by

$$\begin{aligned} \Pi_{1,m}(\alpha) &= \int_{x_m(\alpha)}^1 \left(\frac{1}{2} (2\mu_1(\alpha)\theta)^2 - u_m(\alpha) \right) d\theta, \\ \Pi_{2,m}(\alpha) &= \int_0^{x_m(\alpha)} \left(\frac{1}{2} (2s(1-\theta))^2 - u_m(\alpha) \right) d\theta. \end{aligned}$$

Defining $k(\alpha) \triangleq \mu_1(\alpha)/s$, one can show that the sign of the second derivative of the profit functions only depends on k . This allows us to prove that $\Pi''_{1,m}(\alpha) > 0$ for all α and all k .

(2.) Similarly, $\Pi''_{2,m}(\alpha) > 0$ whenever $k(\alpha) \geq \tilde{\beta}$, with $\tilde{\beta} > 0$. Moreover, $\beta = \min_{\alpha} k(\alpha) = \mu_L/s$. Therefore, if $\beta > \tilde{\beta}$, then $\Pi_{2,m}(\alpha)$ is convex.

(3.) The proof that convexity of $\Pi_{j,m}(\alpha)$ implies convexity of $V_i(\alpha)$ follows the same steps as the proof of Theorem 2, for each firm j separately. ■

References

- ACKERBERG, D. (2003): “Advertising, learning, and consumer choice in experience good markets: an empirical examination,” *International Economic Review*, 44(3), 1007–1040.
- AKÇURA, M., F. GÖNUL, AND E. PETROVA (2004): “Consumer Learning and Brand Valuation: An Application on Over-the-Counter Drugs,” *Marketing Science*, 23(1), 156–169.
- BERGEMANN, D., AND J. VÄLIMÄKI (1997): “Market Diffusion with Two-Sided Learning,” *RAND Journal of Economics*, 28(4), 773–795.
- (2002): “Entry and Vertical Differentiation,” *Journal of Economic Theory*, 106(1), 91–125.
- (2006): “Dynamic Pricing of New Experience Goods,” *Journal of Political Economy*, 114(4), 713–743.
- BERNFELD, S., AND V. LAKSHMIKANTHAM (1974): *An Introduction to Nonlinear Boundary Value Problems*. Academic Press.
- BOLTON, P., AND C. HARRIS (1999): “Strategic Experimentation,” *Econometrica*, 67(2), 349–374.
- BRADEN, D., AND S. OREN (1994): “Nonlinear Pricing to Produce Information,” *Marketing Science*, 13(3), 310–326.
- CRAWFORD, G., AND M. SHUM (2005): “Uncertainty and Learning in Pharmaceutical Demand,” *Econometrica*, 73(4), 1137–1173.
- 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.
- ERDEM, T., AND M. KEANE (1996): “Decision-making Under Uncertainty: Capturing Dynamic Brand Choice Processes in Turbulent Consumer Goods Markets,” *Marketing Science*, 15, 1–20.
- FAINGOLD, E., AND Y. SANNIKOV (2007): “Equilibrium Degeneracy and Reputation Effects in Continuous Time Games,” *Cowles Foundation Discussion Paper 1624*.
- FLEMING, W., AND H. SONER (2006): *Controlled Markov Processes and Viscosity Solutions*. Springer.
- GOETTLER, R., AND K. CLAY (2006): “Tariff Choice with Consumer Learning: Sorting-Induced Biases and Illusive Surplus,” *Working Paper, Carnegie Mellon University*.

- GOWRISANKARAN, G., AND M. RYSMAN (2007): “Dynamics of Consumer Demand for New Durable Consumer Goods,” *Working paper, Boston University*.
- HITSCH, G. (2006): “An Empirical Model of Optimal Dynamic Product Launch and Exit Under Demand Uncertainty,” *Marketing Science*, 25(1), 25.
- ISRAEL, M. (2005): “Services as Experience Goods: An Empirical Examination of Consumer Learning in Automobile Insurance,” *American Economic Review*, 95(5), 1444–1463.
- KELLER, G., AND S. RADY (1999): “Optimal Experimentation in a Changing Environment,” *Review of Economic Studies*, 66(3), 475–507.
- LARKIN, I. (2008): “Bargains-then-ripoffs: Innovation, Pricing and Lock-in in Enterprise Software,” *Proceedings of the Academy of Management*.
- LEWIS, T., AND H. YILDIRIM (2002): “Learning by Doing and Dynamic Regulation,” *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,” *RAND Journal of Economics*, 15(2), 171–196.
- MOSCARINI, G., AND L. SMITH (2001): “The Optimal Level of Experimentation,” *Econometrica*, 69(6), 1629–1644.
- 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,” *Economic Journal*, 38(152), 543–559.
- SONG, I., AND P. CHINTAGUNTA (2003): “A Micromodel of New Product Adoption with Heterogeneous and Forward-Looking Consumers: Application to the Digital Camera Category,” *Quantitative Marketing and Economics*, 1(4), 371–407.
- SPULBER, D. (1989): “Product Variety and Competitive Discounts,” *Journal of Economic Theory*, 48(2), 510–525.
- VILLAS-BOAS, J. (2004): “Consumer Learning, Brand Loyalty, and Competition,” *Marketing Science*, 23(1), 134–145.
- (2006): “Dynamic Competition with Experience Goods,” *Journal of Economics & Management Strategy*, 15(1), 37–66.