

MIT Open Access Articles

Menu Pricing and Learning

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Bonatti, Alessandro 2011. "Menu Pricing and Learning." American Economic Journal: Microeconomics, 3(3): 124–63.

As Published: <http://dx.doi.org/10.1257/mic.3.3.124>

Publisher: American Economic Association

Persistent URL: <http://hdl.handle.net/1721.1/65904>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike 3.0



Menu Pricing and Learning

By ALESSANDRO BONATTI*

We analyze the design of dynamic menus to sell experience goods. The quality of the product is initially unknown, and the total quantity sold in each period determines the amount of information in the market. We characterize the optimum menu as a function of consumers' beliefs, and the dynamic adjustments resulting from the diffusion of information. The firm faces a dynamic trade-off between gains from trade, information production, and information rents. It initially charges lower prices, sacrificing short-term revenue to increase sales. As more information is revealed, prices increase, and low-valuation buyers are excluded, even when the product's quality is high.

JEL: D42, D82, D83, L12.

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

Learning plays a crucial role in many markets and other strategic environments. In particular, in markets for new products and services, sellers face uncertainty over the product's fit to 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.

In these markets, heterogeneity in consumers' willingness to pay for the product creates the opportunity for firms to profitably adopt price discrimination techniques, such as menu pricing. In addition, information about a product's performance is widely and publicly accessible through an increasing number of channels.¹ The availability of such aggregate information in a dynamic environment enables firms to modify their menu prices on the basis of the opinion of their customers. This aspect is particularly relevant in markets for experience goods, because the diffusion of information is endogenous to the behavior of market participants: consumers' purchasing decisions and firms' pricing strategies determine the level of sales, and hence the amount of information conveyed to the market. In this scenario, a forward-looking firm must screen consumers in order to maximize revenues, while taking into account the informational value of sales. By selling additional units of the product (for example, by offering introductory discounts), the

* MIT Sloan School of Management, 100 Main Street, Cambridge MA 02142, bonatti@mit.edu. This paper is based on Chapter 1 of my doctoral dissertation at Yale University. I am indebted to Dirk Bergemann, Johannes Hörner, and Ben Polak for their invaluable help and encouragement throughout this project. For many helpful discussions, I also wish to thank Rossella Argenziano, Eduardo Faingold, Dino Gerardi, Michael Grubb, William Hogan, Marco Ottaviani, 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.

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

firm accelerates the buyers' learning process, thereby trading off (a) the long-run profits that accrue due to the diffusion of information against (b) the maximization of current revenue.

In this paper, we address the issue of designing dynamic menus to sell experience goods. We characterize the evolution of menu prices as information about product quality is gradually revealed, and examine the interaction of the screening and learning problems. We develop a dynamic 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 price differentially. The quality of the product is unknown initially; 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. As a result, the firm can control the information flow to the market by adjusting the level of sales. Learning about the product occurs faster as more units are sold; hence, the firm might 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 quantity of the product that is supplied to each buyer is determined by the combination of three components. The first of these components is the generation of information. Learning occurs through consumption, and each unit sold provides additional information. Thus the firm wants to sell additional units to gain more information when uncertainty about quality is high and beliefs are more responsive to news. The second component is related to efficiency. As consumers grow more optimistic about the quality of the product, their willingness to pay increases, thereby creating the opportunity for the firm to realize larger gains from trade. Therefore, the firm offers larger quantities in this case. The third component is adverse selection. 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 fewer units to buyers who have a lower willingness to pay.

The firm pursues the dual objectives of generating information and screening consumers simultaneously. However, the balance between the two goals 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. As more information is gained, the firm gradually adopts 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 high. In greater detail, as all consumers become more optimistic about the quality of the product, 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 combination of the learning and screening goals has three main effects: (i) the quantity offered to a low-valuation buyer need not be a monotonic function of her posterior beliefs about the product's quality; (ii) successful products are

characterized by a greater price dispersion and a wider variety of available quantities; (iii) for successful products, the firm expands the range of offered quantities through the addition of new options, at both the top and at the bottom of the menu.

In the model, learning occurs on the basis of aggregate information. More precisely, we assume that each consumer's action (quantity choice) and payoff (experienced quality) 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 study for two reasons. First, in large markets, consumers realize that others' experience is also indicative of the underlying quality of the product and take public information into account. More importantly, the study presented herein is interested in modeling the firm's optimal response to variations in demand that arise from the arrival of new information. As such, it focuses only on information that the firm can use in order to determine its strategies. 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 study, 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 consider only the market's observable aggregate experience.

The model is well suited to analyzing several different markets. The market for enterprise software provides an interesting 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, and the renting of more licenses enables the firm to benefit more from a higher quality product. This is so because, in this market, each employee using the software constitutes an experiment for product quality, so that the number of seats may be tied directly to the rate at which information arrives. Moreover, the rental contracts and their corresponding prices can easily be adjusted. Finally, network externalities between firms are not a significant issue in enterprise software, because it is designed for internal use; hence, the private values framework is realistic.²

As an alternate example, consider the market for online DVD rentals. Companies such as Netflix offer membership plans that charge a fixed monthly fee and specify the number of movies a consumer 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 (suggesting new titles based on each buyer's ratings of other films) is a common component in determining the overall quality of the service.³ With this interpretation, each movie rented constitutes an informative experiment about the product's quality. It is reasonable to assume that customers with a 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 choice of plan made by the consumer can easily be adjusted. Finally,

²See Ian Larkin (2008) for a detailed analysis of some frequently used contractual arrangements in this industry.

³The recommender system is considered the main determinant of product quality by Netflix executives. See for example "If You Liked This, Sure to Love That," *The New York Times*, November 23, 2008.

Netflix subscribers exchange information about their experience through a surprisingly large number of channels.⁴ This means that information about the overall performance of the service circulates very rapidly.

Netflix launched their 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.




FIGURE 1. NETFLIX DVD RENTAL PLANS 2002-2005

In 2002, the Netflix menu offer consisted of two plans, which allowed 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.⁵ In 2003 and 2004, Netflix modified its offer of plans to a four-item menu, while raising unit prices across the product line. It added several more options in 2005, while at the same time reducing all prices slightly, possibly due to competitive pressures from Blockbuster. 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 to eight in 2005. Finally, the lowest quantity offered decreased from two rentals at a time in 2002 to one in 2005.⁶

This study enriches the literature on screening by extending nonlinear pricing techniques beyond the canonical, static environment to a model in which information is revealed over time. It therefore builds upon the classic studies in price discrimination, such

⁴For example, <http://www.hackingnetflix.com> and <http://blog.netflix.com> are two of the most popular blogs among Netflix customers.

⁵The 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>.

⁶We here ignore plans that impose a limit on the number of monthly rentals. If we were to include them, the lowest quantity would be given by one DVD at a time, up to four per month.

as Michael Mussa and Sherwin Rosen (1978) and Eric Maskin and John Riley (1984). At the same time, it is tightly connected to continuous time models in which the flow of information is controlled by one or more of the agents. These papers include the works of Patrick Bolton and Christopher Harris (1999) and Godfrey Keller, Sven Rady and Martin Cripps (2005) on strategic experimentation, Godfrey Keller and Sven Rady (1999) on experimentation by a monopolist in a changing environment, Giuseppe Moscarini and Lones Smith (2001) on the optimal level of experimentation, and Eduardo Faingold and Yuliy Sannikov (2011) on reputation in continuous time. In particular, we use the method of Keller and Rady (1999) to show the existence and uniqueness of a solution to the firm's problem.

Our analysis also complements several models of introductory and dynamic pricing under uncertainty about product quality. The main work in this area is due to Dirk Bergemann and Juuso Välimäki (1997), Dirk Bergemann and Juuso Välimäki (2002), Dirk Bergemann and Juuso Välimäki (2006), J. Miguel Villas-Boas (2004), and J. Miguel Villas-Boas (2006). In particular, Bergemann and Välimäki (1997) and Bergemann and Välimäki (2002) analyze a duopoly model of price competition where market participants are uncertain about the degree of horizontal or vertical differentiation of the two firm's products, while Bergemann and Välimäki (2006) consider dynamic monopoly pricing in a private values environment. We discuss these papers at length in Section V. Our paper is also related to the dynamic pricing models in Subir Bose, Gerhard Orosel, Marco Ottaviani and Lise Vesterlund (2006) and Subir Bose, Gerhard Orosel, Marco Ottaviani and Lise Vesterlund (2008), in which buyers take actions sequentially, based on the history of previous purchases, prices, as well as their private information about a common value component. In contrast, in our model, each buyer's willingness to pay is determined by her own (and others') past experience with the product, and her private information concerns an idiosyncratic component.

The problem of generating information through sales was first studied, in the context of a screening model, by David J. Braden and Shmuel S. Oren (1994), who introduce uncertainty over the distribution of buyers' willingness to pay. In their model, one buyer arrives in each period, and her choice from the firm's menu provides information about the true distribution of types. Braden and Oren (1994) and our paper share the conclusion that excluding types early on reduces the amount of information generated and is therefore suboptimal. However, in Braden and Oren (1994) information is only obtained by avoiding bunching and exclusion of types. The learning problem is therefore separate from profit maximization, because learning considerations do not affect the quantity levels offered to each buyer.

Generating (public) information proves to be beneficial in our context. This is indeed similar to the findings of Marco Ottaviani and Andrea Prat (2001). However, our result is based on the convexity of the firm's profits as a function of the unknown product quality, as opposed to the effect of an affiliated public signal on buyer's information rents.

Finally, the techniques used in this study also relate to the models developed by Tracy R. Lewis and Huseyin Yildirim (2002) and by Jan Boone and Joel Shapiro (2008). In particular, in the dynamic regulation of Lewis and Yildirim (2002), 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. In another closely related contribution, Dennis Gärtner (2010) analyzes a two-period regulation model in which the rates of learning-by-doing of different types affect the dynamics of output distortions.

In addition to this theoretical body of work, recent empirical literature has attempted 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 Daniel A. Akerberg (2003), M. Tolga Akçura, Füsün F. Gonul and Elina Petrova (2004), Gregory S. Crawford and Matthew Shum (2005), Tulin Erdem and Michael P. Keane (1996), Ronald L. Göttinger and Karen Clay (2011), Gautam Gowrisankaran and Marc Rysman (2007), and Mark Israel (2005). From a different perspective, Günter J. Hitsch (2006) and Inseong Song and Pradeep K. 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. The study reported herein 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.

I. The Model

A. 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 of the firm's product depends on both 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 continuously differentiable distribution $F(\theta)$.

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 that represents the quality of the match between the product and the needs of the market. 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, product quality and personal taste interact multiplicatively. Hence, buyers with a higher willingness to pay benefit more from a higher quality product. We assume that each buyer makes a purchase decision in every period, and that she can freely switch between purchasing different quantities. 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 unknown initially 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 adjusted flexibly. In a direct mechanism, the firm's strategy is a pair of piecewise differentiable functions $q_t : \Theta \rightarrow \mathbb{R}_+$ and $p_t : \Theta \rightarrow \mathbb{R}_+$ in each period. These functions determine the quantity and the total charges assigned to each buyer θ . Suppose each buyer purchases quantity $q_t(\theta)$ and pays total charges of $p_t(\theta)$. The firm then obtains flow profits of

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

The social gains from trade that are realized by selling quantity q to type θ , when the product quality is μ , are given by $\mu\theta u(q) - c(q)$. We assume it is always efficient to sell a positive quantity level to every buyer.

ASSUMPTION 2 (Social Gains from Trade):

For all μ and θ , $\mu\theta u'(\theta) - c'(\theta) > 0$.

We now define the virtual valuation of buyer θ as

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

Under assumption 1, virtual valuations are increasing in θ . We then consider the virtual surplus, $\mu\phi(\theta)u(q) - c(q)$, and we introduce the following assumption.

ASSUMPTION 3 (Concave Virtual Surplus):

1) The virtual surplus is strictly concave in q for all μ and θ .

2) For all μ and θ , $\lim_{q \rightarrow \infty} [\mu \phi(\theta) u'(q) - c'(q)] = -\infty$.

We use this assumption in our dynamic analysis to ensure that the optimal quantity is bounded, and can be characterized by a first order condition whenever strictly positive. It is satisfied, for example, by the model in Mussa and Rosen (1978). An *alternative* assumption, that does not require concavity of the virtual surplus, is the following.

ASSUMPTION 4 (Constant Marginal Cost):

- 1) $c(q) = cq$, with $c > 0$.
- 2) $u(q)$ is strictly concave, and $\lim_{q \rightarrow \infty} u'(q) = 0$.

This is the case, for example, in the model of Maskin and Riley (1984). This assumption ensures that the monopolist does not want to serve types with a negative virtual valuation, and that the first order conditions are sufficient whenever quantity provision is positive. Our results hold under either of these assumptions.

In this sense, our model can accommodate (among others) both the Mussa and Rosen (1978) specification with linear utility and convex cost, and the Maskin and Riley (1984) formulation with concave utility and constant marginal cost.

B. Information and Learning

Information about product quality may only be obtained through consumption. We now provide a formal treatment of the aggregate market experience and the associated law of motion of beliefs. In particular, we adapt the model in Bergemann and Välimäki (1997) to allow for multi-unit demand.

We begin with a finite number of buyers and discrete time. We are going to suggest a model in which the informativeness of the aggregate market experience is held constant as the number of buyers increases. In other words, each additional buyer does not lead to a larger, more informative market. Instead, we interpret a larger number of buyers as a more fragmented consumer population, in which each individual buyer purchases units of a smaller size. Formally, 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 signal $\tilde{x}_{ij} \sim N(\mu/K, \sigma^2/K)$. 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.

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 $\sum_{i=1}^K \sum_{j=1}^{q(\theta_i)} x_{ij}$. Denote the average number of units purchased by $Q_K = (1/K) \sum_{i=1}^K q(\theta_i)$. The market experience is now normally distributed with mean $Q_K \mu$ and variance $Q_K \sigma^2$.

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. In particular, we obtain that 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 $Q_t\sigma^2$,

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

With this structure for the information flow, one can use the filtering equations⁷ to derive the evolution of the posterior beliefs α_t :

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

In our model, the information contributed by the experience of each buyer becomes infinitely small. An alternative model would keep the informativeness of each individual signal constant. As a result, the law of large numbers would imply that the quality of the product is learned instantaneously. In contrast, by holding the aggregate market experience constant as we increase the number of buyers, we intend to capture the relevant features of new online services with a large diffusion, such as Netflix. In particular, at each instant, the quality of the product is not perfectly revealed, and the individual's experience is of negligible importance, relative to the entire market's. Finally, we point out that several other signal structures would also preserve the imperfect informativeness of the market experience. For example, we could consider common quality shocks to the individual experiences, under the assumption that larger production levels reduce the variance of these shocks.

To summarize, information is imperfect but symmetric at all points in time. The posterior beliefs follow a martingale and as a result the process has a zero drift. The level of aggregate sales determines the total number of experiments with the product, and hence the rate at which the firm and the consumers learn about its quality.

We now define the following function:

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

This function captures the marginal contribution of each unit sold to the variance of the belief process $(d\alpha_t)^2 = 2Q_t \Sigma(\alpha_t) dt$. The variance is increasing in the degree of dispersion $\alpha_t (1 - \alpha_t)$ and in the signal-to-noise ratio $(\mu_H - \mu_L)/\sigma$. Posterior beliefs evolve

⁷ See Theorem 9.1 in Robert Liptser and Albert Shiryaev (1977). Equivalently, under our two-point prior assumption, one can use Bayes' rule to compute the posterior belief $\alpha_{t+\Delta t}$, and take the limit for $\Delta t \rightarrow 0$. See, for example, the steps in Patrick Bolton and Christopher Harris (2000).

more quickly when current uncertainty is high and when signals are precise. Finally, we stress that the changes in beliefs are determined endogenously, because the total quantity sold Q_t depends on the firm's pricing and on the consumers' purchasing decisions. In other words, the firm can control the rate of information flow to the market by adjusting the level of sales.

II. Equilibrium Analysis

As a first step in the equilibrium analysis, we characterize the incentive-compatible menus of contracts. 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 her beliefs α_t and the firm's menu offer. Because quality μ may only take one of two values, the current posterior belief α_t is a sufficient statistic for the firm's problem at every point in time. Therefore, we denote by $(q(\alpha_t, \theta), p(\alpha_t, \theta))$ the menu offered by the firm when the posterior beliefs are given by α_t . We also denote by $U(\alpha_t, \theta, \theta')$ the expected utility of a buyer with willingness to pay θ who purchases the item $(q(\alpha_t, \theta'), p(\alpha_t, \theta'))$ intended for a buyer of type θ' :

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

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. By standard arguments, these are equivalent to:

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

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

Equation (4) shows that the firm must concede higher information rents when beliefs become more optimistic. This effect is due to the complementarity between product quality and buyers' willingness to pay. Buyers' valuations depend positively on the posterior beliefs α_t , hence positive news allow the firm to charge higher prices. However, as α_t increases, the difference between any two buyers' willingness to pay also increases, thereby creating stronger incentives to misreport one's type. This means that for high values of α_t , the incentive compatibility constraints are more difficult to satisfy.

Finally, the buyers' participation constraints are given by

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

A. Myopic Benchmark

Consider the problem of 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). Following the standard procedure for one-dimensional screening problems, we substitute 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 the posterior probability α_t and quantities $q(\alpha_t, \theta)$. Assumption 1 ensures that constraint (5) holds in equilibrium. The firm's flow profits are given by

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

where $\phi(\theta)$ denotes the virtual valuation, 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 myopic solution is obtained by maximizing (7) pointwise. The first-order condition for the provision of quantity is given by

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

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. The expected product quality $\mu(\alpha_t)$ acts as a 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) When positive, the myopic quantity $q_m(\alpha_t, \theta)$ is strictly increasing in α_t and θ .
- 2) The myopic profit function $\Pi_m(\alpha_t)$ is strictly increasing and strictly convex in α_t .

The convexity of the myopic profit function has implications for the firm's incentives to learn about the quality of its product. This result is quite intuitive. More optimistic beliefs improve every buyer's willingness to pay, and the firm can charge higher unit prices. Moreover, the firm finds it profitable to sell a larger number of units. As a result, the myopic profit function increases more than linearly with the posterior beliefs α_t .

Therefore, a myopic 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 expected payoff of this lottery as the complete information average profit,

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

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

B. Dynamic Solution

In order to design the dynamically optimal menu prices, we consider a sequence of quantity supply functions $q_t : \Theta \rightarrow \mathbb{R}_+$. The incentive compatibility and participation constraints (4)–(6) uniquely determine the corresponding sequence of total charges $p_t : \Theta \rightarrow \mathbb{R}_+$. Therefore, we can express the (forward-looking) firm's objective function as

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

Our first result is instrumental to determining whether the firm assigns a positive value to information.

THEOREM 1 (Convexity of the Value Function):

The value function V^ is continuous and convex.*

The intuition for Theorem 1 is straightforward. For a fixed quantity supply function, profits are linear in α_t . Clearly, the firm can improve on these linear profits by reacting to information. The main implication of Theorem 1 is that the forward-looking firm is willing to give up some revenue in the short run (i.e. to depart from $\Pi_m(\alpha_t)$), in exchange for more information generated through sales. The evolution of the posterior beliefs α_t is controlled by the law of motion (1). Using the law of motion for beliefs and Itô's Lemma, we can write the Hamilton-Jacobi-Bellman (HJB) equation for the firm's problem as

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

The firm's value function differs from the myopic profit function only through the term $Q\Sigma(\alpha_t) V''(\alpha_t)$, which is positive by Theorem 1, and proportional to the total quantity sold Q . Remember that each unit sold provides an informative signal whose effect on the posterior beliefs depends on the variance $\Sigma(\alpha_t)$. Therefore, the term $Q\Sigma(\alpha_t)$ can be interpreted as the amount of information generated through sales. The term $V''(\alpha_t)$ represents the marginal value of information. As such, it determines the firm's incentives to increase the speed at which customers learn about the product's quality. Note that

information has no value when $\alpha_t = 0$ and $\alpha_t = 1$, because beliefs no longer change in those cases.

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

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

We now prove the existence of a solution to this 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 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.

THEOREM 2 (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 $V^*(\alpha)$ of (10).*

The proof of (1.) and (2.) adapts the method in Keller and Rady (1999), which is based on super- and subsolutions to a two-point boundary-value problem and deals with the singularities of the differential equation at both ends of the unit interval.⁸ The proof of (3.) uses a standard verification theorem.

We now derive some elementary properties of the policy function. 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) *Fix an α , and consider all pairs (μ_L, μ_H) such that $\alpha\mu_H + (1-\alpha)\mu_L = \mu$ for some $\mu > 0$. Then $rV(\alpha)$ and $\Sigma(\alpha) V''(\alpha)$ are increasing in the difference $\mu_H - \mu_L$.*

⁸A similar method is also used, in the context of a continuous-time reputation model, by Faingold and Sannikov (2011).

As expected, the precision of the individual signals and the firm's patience level increase the value of information. Proposition (2) also shows that the returns to experimentation increase in the relevance of the learning process, as measured by the difference in the possible quality levels. As we discuss in the next section, higher returns to experimentation induce the monopolist to increase the quantity sold.

C. Properties of the Equilibrium Menus

Pointwise maximization of the firm's objective (the right-hand side of equation (12)) yields an intuitive expression for the optimal quantity provision. In particular, the equilibrium quantities $q(\alpha_t, \theta)$ are given by the solution to the first-order condition

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

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, because 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 (Equilibrium Quantities):

- 1) Quantities $q(\alpha, \theta)$ are everywhere higher than in the myopic solution ($q_m(\alpha, \theta)$).
- 2) Quantities $q(\alpha, \theta)$ are weakly increasing in the value of information $\Sigma(\alpha) V''(\alpha)$, and strictly increasing if $q(\alpha, \theta) > 0$.

Proposition 3 shows that the firm induces market experimentation by selling quantities in excess of the myopic optima for all α and θ . This leads *inter alia* 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 degree of patience, and in the precision of the signals. However, the value of information $\Sigma(\alpha) V''(\alpha)$, as well as quantities and levels of market coverage, are typically not monotonic in the posterior beliefs α . In particular, the firm has no incentive to experiment when beliefs are degenerate and $\alpha \in \{0, 1\}$.

The quantities $q(\alpha, \theta)$ in the direct mechanism can be linked to the actual price-quantity menus offered by the firm in an indirect mechanism. We do so through a non-linear price function $\hat{p}(\alpha, q)$. This function defines the total amount charged by the firm for q units of the product, when the posterior beliefs are given by α . Consumers maximize their utility given the firm's current menu offer. This allows us to characterize the marginal prices charged on each unit via the buyer's first-order condition

$$(14) \quad \mu(\alpha) \theta(\alpha, q) u'(q) - \hat{p}_q(\alpha, q) = 0.$$

In equation (14), $\theta(\alpha, q)$ denotes the buyer who purchases quantity q in equilibrium. 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 $\Sigma(\alpha) V''(\alpha)$.*

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 further 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 levels of quality of the product, but it is not monotonic in consumers' posterior beliefs about quality.

III. Linear-Quadratic Model

We now specify our model to the Mussa and Rosen (1978) functional form assumptions of linear utility ($u(q) = q$) and quadratic costs ($c(q) = q^2/2$). These assumptions allow us to identify separately the role of the value of information in determining the changes in the equilibrium menus as a function of beliefs.⁹ In particular, the first-order condition (13) now provides an explicit expression for the provision of quantity

$$(15) \quad q(\alpha, \theta) = \max \{ \mu(\alpha) \phi(\theta) + \Sigma(\alpha) V''(\alpha), 0 \}.$$

In this section, we first characterize the solution for a setting in which all buyers participate, and the firm has a positive discount rate r . We turn to the undiscounted limit to describe the effects of information in finer 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.

⁹These functional form assumptions allow for the most straightforward illustration of our results. However, our analysis can be easily adapted to the case of constant marginal costs and quadratic utility, as in Maskin and Riley (1984).

A. Full Market Coverage and Positive Discounting

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

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

In addition, it is immediate to show that the myopic profits Π_m are a quadratic function of α . The following proposition relates the equilibrium menus to the myopic benchmark.

PROPOSITION 5 (Quantities and Prices):

1) *The equilibrium quantities and prices are given by*

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

2) *The marginal value of information is given by*

$$(16) \quad \Sigma(\alpha) V''(\alpha) = -\mu(\alpha) \mathbb{E}_{\theta}[\phi] + \sqrt{(\mu(\alpha) \mathbb{E}_{\theta}[\phi])^2 + 2(rV(\alpha) - \Pi_m(\alpha))}.$$

A few remarks are in order. First, each type receives $\Sigma(\alpha) V''(\alpha)$ units over and above the myopic quantity supply. These additional units constitute the marginal level of experimentation by the firm, which is constant across buyers θ . However, prices only exceed the corresponding myopic level $p_m(\alpha, \theta)$ 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 price that is low enough to convince them to purchase more. Second, the number of additional units $\Sigma(\alpha) V''(\alpha)$ need not increase monotonically in the posterior beliefs α . Third, the marginal value of information does not depend solely on the difference $rV(\alpha) - \Pi_m(\alpha)$, but also directly on the current level of demand, which is captured by $\mu(\alpha) \mathbb{E}_{\theta}[\phi]$. This term is equal to the total quantity sold by the myopic seller, which can be viewed as the “default” amount of experimentation. In other words, the interaction of the monopolist's screening and learning goals yields optimal quantities that depend on the speed of the learning process in the absence of any additional investment in information production.

A further implication of Proposition 5 is that the effects of new information on the supplied quantities depend on the consumer's willingness to pay θ . Combining the first result of Proposition 5 with equation (16), the equilibrium quantity levels may be written

as

$$(17) \quad q(\alpha, \theta) = \mu(\alpha) (\phi(\theta) - \mathbb{E}_\theta[\phi]) + \sqrt{(\mu(\alpha) \mathbb{E}_\theta[\phi])^2 + 2(rV(\alpha) - \Pi_m(\alpha))}.$$

Consequently, types with a virtual valuation above the average $\mathbb{E}_\theta[\phi]$ benefit more from an increase in the posterior α than those with below-average virtual valuations. At the same time, differences between the quantities 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.

PROPOSITION 6 (Contract Variety):

- 1) For all $\theta > \theta'$, differences in quantity $q(\alpha, \theta) - q(\alpha, \theta')$ 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 .

This result is informative of the dynamics of the variety of the equilibrium menu. If we let $\theta = \theta_H$ and $\theta' = \theta_L$, we obtain that increases in the posterior 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 characterize explicitly the behavior of the equilibrium menus as a function of α . This requires solving the differential equation (16) for the firm's value function. Unfortunately, this differential equation is a second-order, nonlinear problem that does not have an analytical solution. However, we are able to obtain closed-form solutions by analyzing the undiscounted version of the firm's problem.

B. No Discounting

For the analysis of the undiscounted version of the problem, we adopt the *strong long-run average* criterion.¹⁰ 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.

With reference to our model, the strong long-run average criterion may be summarized as follows. By the martingale convergence theorem, beliefs converge to either μ_L or μ_H . In the limit for $r \rightarrow 0$, the return function $rV(\alpha)$ converges to the complete information average payoff $v(\alpha)$ defined in (9). However, many policy functions attain the long-run

¹⁰This criterion was pioneered by Frank P. Ramsey (1928). In more recent work, Prajit K. Dutta (1991) discusses the relationship between the strong long-run average and other criteria for undiscounted optimization in a discrete-time framework.

average value $v(\alpha)$, independently of their finite time properties. Dutta (1991) considers the undiscounted stream of payoffs, net of their long run averages,

$$(18) \quad V(\alpha_0) \triangleq \sup_{q_t: \Theta \rightarrow \mathbb{R}_+} \mathbb{E}_{\alpha} \left[\int_0^{\infty} (\Pi(\alpha_t, q_t) - v(\alpha_t)) d\alpha_t \mid \alpha_0 \right].$$

Dutta (1991) proves that the policy q_t maximizing (18) represents the limit for $r \rightarrow 0$ of the policy functions that maximize 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 such criteria as the limit of the means. We can therefore write the undiscounted analog of the HJB equation (11) as

$$(19) \quad v(\alpha) = \sup_{q: \Theta \rightarrow \mathbb{R}_+} [\Pi(\alpha, q) + Q \Sigma(\alpha) V''(\alpha)],$$

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 shows that the value of information increases as the firm's discount rate decreases. In the linear-quadratic, undiscounted case, we can solve for this value in closed form, and express the equilibrium quantities as

$$(20) \quad q(\alpha, \theta) = \mu(\alpha) (\phi(\theta) - \mathbb{E}_{\theta}[\phi]) + \sqrt{(\mu(\alpha) \mathbb{E}_{\theta}[\phi])^2 + (\text{Var}[\phi] + \mathbb{E}_{\theta}[\phi]^2) \alpha (1 - \alpha) (\mu_H - \mu_L)^2}.$$

This expression is obtained by substituting (17) into equation (19), and then solving for $\Sigma(\alpha) V''(\alpha)$ directly. The key properties of the equilibrium quantity supply are given in the next Theorem.

THEOREM 3 (Undiscounted Equilibrium Quantities):

- 1) *The equilibrium quantities $q(\alpha, \theta)$ are 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 qualitative implications for the evolution of equilibrium quantities. Contrary to the myopic case, a set of types $[\theta_L, \tilde{\theta}]$ does not always receive greater quantities as the posterior beliefs increase. The threshold type $\tilde{\theta}$ identified in Theorem 3 satisfies the following equation:

$$(21) \quad \phi(\tilde{\theta}) = \frac{\text{Var}[\phi] + \mathbb{E}_{\theta}[\phi]^2}{2\mathbb{E}_{\theta}[\phi]} \frac{\mu_H - \mu_L}{\mu_H}.$$

Therefore, the fraction of types who receive nonmonotonic quantities is increasing in (a) the relative difference between the two quality levels $(\mu_H - \mu_L) / \mu_H$ and (b) the dispersion of buyers' valuations $\text{Var}[\phi(\theta)]$. The latter result follows from the fact that the

firm's equilibrium profit *on each type* θ , given by $p(\alpha, \theta) - c(q(\alpha, \theta))$, is convex in θ . Therefore, an increase in the spread of the distribution $F(\theta)$ improves the firm's profits, thereby making the learning process more significant. The concavity of equilibrium quantities suggests that experimentation is greater when beliefs about the quality of the product are intermediate. Figure 2(a) confirms this intuition. In this figure, we show the quantities supplied to three different buyers $\theta < \theta' < \theta''$ as a function of α . The lowest type has a zero virtual valuation and would never be served in the myopic case. Figure 2(b) illustrates the equilibrium total charges. Consistent with the result from Proposition 6, the differences between the total charges paid by different buyers are increasing and convex in α .




FIGURE 2. QUANTITIES AND TOTAL CHARGES: $\theta \sim U[1, 2]$, $\mu_L = 1$, $\mu_H = 8$, $\theta = 1$, $\theta' = 31/24$, $\theta'' = 2$.

The main properties of the equilibrium menu are best understood by decomposing the implications of the arrival of information into three effects. The first effect is related to information value. Each unit sold generates additional value to the firm by facilitating learning. This effect is strongest when beliefs are intermediate, and uncertainty is highest. Conversely, as beliefs approach zero or one, the value of information declines, and so do the incentives to provide greater quantities. This effect influences all types in the same way, because the informational content of a unit that is sold is independent of the buyer who purchases it.

The second effect is related to efficiency. When positive news arrive, consumers are willing to pay more for each unit; hence, 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 adverse selection. The differential increase in buyer's valuations tightens the incentive compatibility constraints and increases the information rents. This raises the cost of screening consumers. To understand this, remember that in a two-type static model, the information rent of the high type is equal to $U(\theta_H) =$

$U(\theta_L) + (\theta_H - \theta_L)q_L$. 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 news generates an additional cost to the seller, thereby driving down consumption for low-valuation buyers as beliefs approach one.

The combined effects of information value, efficiency, and adverse selection determine a set of types for which the provision of quantity 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 for two different buyers.

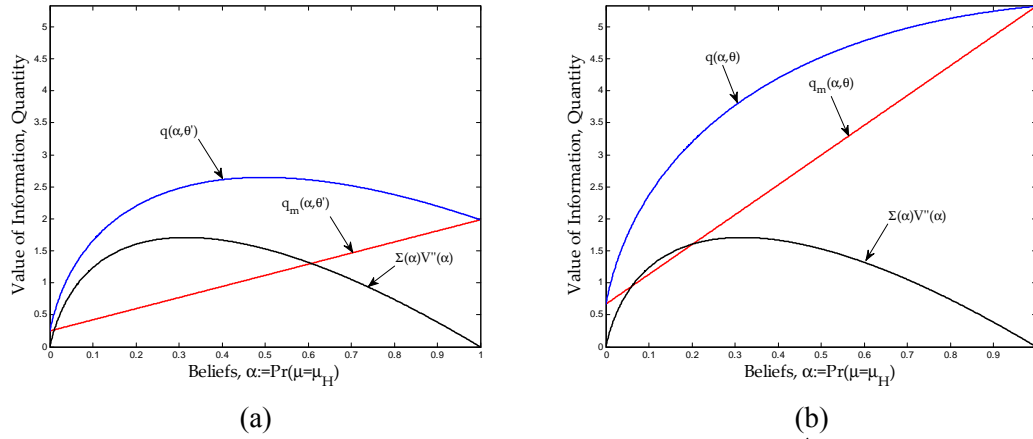



FIGURE 3. QUANTITIES AND VALUE OF INFORMATION: $\theta \sim U[1, 2]$, $\mu_L = 1$, $\mu_H = 8$, $\theta' = 9/8$, $\theta = 11/8$.

The equilibrium quantities are given by the vertical sum of the marginal value of information $\Sigma(\alpha)V''(\alpha)$ with the myopic solution $q_m(\alpha, \theta)$. A peculiar result of this model is that the value of information (and hence the difference between the equilibrium and the myopic quantities) peaks at a value of α lower than one-half. To understand why this is the case, consider equation (19). The total information 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 this difference. The value of $v(\alpha) - \Pi(\alpha, q)$ depends positively on the degree of uncertainty $\alpha(1 - \alpha)$, which is a measure of how much posterior beliefs can be influenced by the signals observed in the current period. At the same time, the total quantity Q is increasing in α , which implies that the ratio $(v(\alpha) - \Pi(\alpha, q))/Q$ is decreasing at $\alpha = 1/2$. In other words, since the myopic firm's total sales are increasing in α , learning will occur faster when beliefs are high, even in the absence of any (additional) experimentation. This lowers the gap between the full information and the incomplete information profits, and reduces the information value of each (additional) unit sold.

C. Nonlinear Prices

Our result on the nonmonotonic provision of quantity can be related to 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 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 distributed uniformly, marginal prices are given by

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

In the undiscounted case, learning has an intuitive effect on marginal prices. Marginal prices are increasing in α for each quantity, provided the difference in quality levels $\mu_H - \mu_L$ is not too high. Conversely, for high values of $\mu_H - \mu_L$, marginal prices are U-shaped in α for all 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.

D. Small Positive Discounting

Many results obtained in the undiscounted limit extend to the case of a positive discount rate. In particular, we can use bounds for the convexity of the value function to establish the concavity of equilibrium quantities under small positive discounting. This procedure presents some difficulties, because 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 (because $\Pi_m''(\alpha)$ is a constant), and by the undiscounted profits (because $v''(\alpha) \equiv 0$). Our first result extends the concavity property of the provision of quantity 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 increasing in the posterior beliefs when $\alpha = 0$. More importantly, it identifies the minimum degrees of patience required to extend the nonmonotonic quantities result to an arbitrary set of low-valuation buyers. For this purpose, let $\tilde{\theta}$ be the threshold type defined by (21).

PROPOSITION 8 (Nonmonotonic Quantities):

- 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 $\theta' \in [\theta_L, \theta]$.*

E. Partial Market Coverage

When the distribution of types is such that $\phi(\theta_L) < 0$, it is not optimal for the monopolist to serve the entire market for all values of the posterior beliefs. In what follows, we focus on the undiscounted version of the problem and apply the strong long-run average criterion. For buyers who are offered positive quantities in equilibrium, the optimal sales level is characterized by the first-order condition (15):

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

However, the equilibrium value of information also affects the set of buyers who receive positive quantities. In other words, $\Sigma(\alpha) V''(\alpha)$ determines the lowest type served,

which we denote by $\theta^*(\alpha)$. 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 that types are distributed uniformly. We then obtain the following characterization of the equilibrium quantities and market coverage levels.

PROPOSITION 9 (Market Coverage):

Assume types θ are uniformly distributed on $[\theta_L, \theta_H]$. The undiscounted equilibrium level of market coverage and equilibrium quantities are given by

$$\begin{aligned} \theta^*(\alpha) &= \theta_H - \frac{\theta_H}{2} \left(\frac{\alpha \mu_H^2 + (1 - \alpha) \mu_L^2}{\mu(\alpha)^2} \right)^{1/3}, \\ q(\alpha, \theta) &= \max\{(12\mu(\alpha)(\theta_H - \theta_L)v(\alpha))^{1/3} - 2\mu(\alpha)(\theta_H - \theta), 0\}. \end{aligned}$$

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 greatest. 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 marginal value of information (and hence the fraction of buyers who are served) attains a maximum when α is lower than $1/2$.

The case of partial market coverage allows us to show clearly how the arrival of new information benefits some high valuation buyers, but not others. Figure 5 shows the indirect utility levels for three buyers, as a function of α .

In particular, the lowest valuation buyer shown (θ) is excluded for some high and low 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.

IV. Intertemporal Patterns

We are now interested in deriving predictions for the intertemporal evolution of the equilibrium menus. We first consider the point of view of participants in the market. Their posterior beliefs α_t follow the diffusion process described by equation (1). Therefore, by Itô's Lemma, any twice differentiable function $h(\alpha_t, \theta)$, such as prices and quantities, also follows a diffusion process. In particular, the law of motion of $h(\alpha_t, \theta)$ is given by

$$(22) \quad dh(\alpha_t, \theta) = \frac{\partial h(\alpha_t, \theta)}{\partial \alpha} d\alpha_t + \frac{1}{2} \frac{\partial^2 h(\alpha_t, \theta)}{(\partial \alpha)^2} (d\alpha_t)^2.$$




FIGURE 5. EQUILIBRIUM UTILITY LEVELS: $\theta \sim U[0, 1]$, $\mu_L = 1$, $\mu_H = 8$, $\theta = 4/9$, $\theta' = 5/9$, $\theta'' = 2/3$.

Given that $\mathbb{E}[da_t] = 0$, the sign of the drift component of the process $dh(\alpha_t, \theta)$ is determined by the second partial derivative $\partial^2 h(\alpha_t, \theta) / (\partial \alpha)^2$. In other words, the concavity and convexity properties of any function $h(\alpha_t, \theta)$ may be translated directly into statements about the sign of its expected changes.

Throughout this section, we maintain the linear-quadratic functional form assumptions.

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, quantities are expected to decrease over time. Conversely, differences between quantities offered to different buyers are expected to remain constant over time. Finally, differences in the total prices charged to 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.

The posterior beliefs α_t of market participants follow the diffusion process (1). However, from the point of view of an external observer (*i.e.* the econometrician), the evolution of the process da_t depends on the true underlying quality level. Therefore, any empirical prediction about the intertemporal patterns of prices and quantities must be

based on the *conditional* law of motion of beliefs. The conditional changes in beliefs have a non-zero drift component, whose sign depends on the true μ . In particular, for $\mu \in \{\mu_L, \mu_H\}$, the general filtering equation (see Liptser and Shiryaev (1977)) is given by

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

The drift component of a $dh(\alpha_t, \theta)$ 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 (22), and factoring out common terms, the sign of the drift component of the process $dh(\alpha_t, \theta)$ is determined by the following expressions:

$$(23) \quad \mathbb{E}[dh(\mu_H)] \propto \left(\frac{\partial h(\alpha_t, \theta)}{\partial \alpha} + \frac{\alpha_t}{2} \frac{\partial^2 h(\alpha_t, \theta)}{(\partial \alpha)^2} \right) dt,$$

$$(24) \quad \mathbb{E}[dh(\mu_L)] \propto \left(-\frac{\partial h(\alpha_t, \theta)}{\partial \alpha} + \frac{1 - \alpha_t}{2} \frac{\partial^2 h(\alpha_t, \theta)}{(\partial \alpha)^2} \right) dt.$$

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 is no longer sufficient to conclude that supplied quantities decrease in expectation for all buyers. However, conditional on the bad state μ_L , quantities are expected to decrease over time for all high-valuation buyers, since their equilibrium quantities are increasing in α .

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 products increases over time. Opposite conclusions hold for low-quality products.

To summarize, our 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 quantities offered to two different buyers, with a prior belief $\alpha_0 = 1/20$, and assuming that the actual quality is 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. QUANTITIES AND TOTAL CHARGES: $\theta \sim U[1, 2]$, $\mu_L = 1$, $\mu_H = 5$, $\theta' = 6/5$, $\theta = 8/5$.

V. Discussion

We now discuss the relationship between our results and other dynamic pricing models. The main questions of interest are: (i) Which results are due to the combination of learning and price discrimination? (ii) What is the role of the multiplicative interaction between consumers' tastes and product quality?

A. Single-Price Benchmarks

The papers in the literature that best serve as single-price benchmarks for the present work are (monopoly versions of) the models in Bergemann and Välimäki (1997) and Bergemann and Välimäki (2002). In these papers, the utility level of a buyer is influenced by two random variables: her willingness-to-pay and her experience with the product. The main difference with the work of Bergemann and Välimäki is that we allow consumers to have multi-unit demands, and the firm to price discriminate.

In Bergemann and Välimäki (1997) and Bergemann and Välimäki (2002), the firm charges lower prices, relative to the myopic solution. Furthermore, when the value of information is sufficiently high, the equilibrium prices can increase following both good and bad news about the quality of the product. Our model shares the same intuition for the positive value of information, and hence for introductory pricing. The novelty of our framework is in the set of instruments available to the firm, namely the ability of the

monopolist to choose both the price and the quantity offered to each buyer. In our setting, the diffusion of information impacts buyers differentially. In particular, good news can benefit high-valuation types and hurt low-valuation buyers.

With a slight change in interpretation, we can view q as a one-dimensional product characteristic (e.g. “quality”), and μ as the match value of the product’s features with the consumers’ tastes.¹¹ In such a model, the firm offers different versions of the product as a function of the market’s beliefs about the value of the match. Furthermore, the equilibrium product line variety does not respond to good and bad news symmetrically. Indeed, bad news lead to a contraction of the firm’s menu, while good news lead to an increase in product line variety.

It can also be useful to contrast our framework with the idiosyncratic learning model in Bergemann and Välimäki (2006). This paper examines dynamic pricing of experience goods when buyers are ex-ante identical and learn their true value through consumption. As a result of the private values environment, and in contrast to our model, the equilibrium price patterns are deterministic. Bergemann and Välimäki (2006) show that the equilibrium prices can be either increasing or decreasing over time. In particular, in mass markets, the firm serves informed buyers with progressively lower valuations, as the size of the uninformed consumer population decreases. This causes prices to decrease. However, the reason for decreasing prices is related to the firm moving along the demand curve. In our model, prices are stochastic and decline following negative signals about a common value component. Finally, when the price in Bergemann and Välimäki (2006) is decreasing, it always lies above the static monopoly price, which further highlights the different role of experimentation in the two models.

B. Product Quality and Idiosyncratic Tastes

In our model, the effects of information on the quantities offered by the firm depend on the interaction between consumers’ willingness to pay and product quality. We have assumed a multiplicative interaction, but depending on the application, different demand specifications may be more appropriate. A plausible alternative specification for each buyer’s complete information utility is an additive one, such as $U = (\mu + \theta) u(q)$. In this case, product quality shifts the distribution of consumers’ willingness to pay. Under full market coverage, changes in beliefs modify the quantity sold to each buyer in the same direction. This is in contrast with our finding in Section III.A, in which the amount of experimentation is constant across buyers, but information may increase one buyer’s consumption level, and decrease another’s. Nevertheless, under a linear specification, the firm still adopts introductory pricing, and serves more buyers than in the myopic solution. An even simpler demand function would be $U = \mu + \theta u(q)$. This is equivalent to shifting the buyers’ participation constraint. If μ is allowed to take negative values, the firm solves a standard optimal stopping problem, in order to determine for which beliefs α it should quit the market. When in the market, the firm sells nonmonotonic quantity levels to all buyers. This occurs because product quality and the number of

¹¹ In this case, we would assume that higher quality items generate more information about the value of the match.

units purchased by each buyer do not interact, and only the learning effect is present. Therefore, the diffusion of information affects all buyers in the same qualitative way.

The most interesting alternative formulation is perhaps one in which the consumers' tastes are closer together when the product is of high quality. This could be the case when user-friendliness or other characteristics make a high quality product more easily accessible by many users.¹² Indeed, consider the utility specification $U = (\beta\mu + \theta/\mu)u(q)$, with $\beta > 0$. For simplicity, we focus on the linear-quadratic model under full market coverage. The dynamic optimal quantity levels are given by

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

For β large enough, both the myopic quantity provision and the dynamic quantity levels are increasing in α for all types. In this model, the optimal myopic quantity is convex in α , and hence buyers expect quantity levels to increase over time when the discount rate is sufficiently high. However, for a low discount rate (or a large enough difference $\mu_H - \mu_L$), the optimal dynamic quantity provision is again a supermartingale, and hence expected to decrease over time. This again highlights the value of information, and emphasizes how the results on introductory pricing do not rely on the multiplicative specification of buyers' preferences. Finally, in contrast to the results in Proposition 6, differences in the quantity levels provided to any two types are decreasing in α . This result is in line with the findings of Gärtner (2010).

VI. Concluding Remarks

We have analyzed the dynamic menu pricing strategy of a new firm, when the quality of its product is initially unknown. Buyers assess the quality of the product uniformly, but have different willingness to pay, which makes it profitable to practice second-degree price discrimination. By adjusting the quantities offered to each buyer, the firm can manage the flow of information to the market, and balance the diffusion of information with the maximization of short-run revenue. The model yields tractable closed-form solutions that enable us to predict the intertemporal patterns of the equilibrium prices and quantities. It also has clear welfare implications, and extends quite naturally to the analysis of competitive environments. We now provide some remarks on these two issues.

Welfare Analysis: The information 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 effect counters the downward distortions induced by adverse selection. As a consequence, experimentation by the monopolist increases each buyer's utility level, as well as the efficiency of the allocation. However, the gradual resolution of the uncertainty is not equally beneficial to all buyers. Low-valuation buyers (who may be excluded as learning occurs) expect their utility level to decrease over time. This is also the case for intermediate-valuation buyers, who face higher prices once low-valuation buyers have been excluded.

¹²I thank an anonymous referee for suggesting this interpretation.

When posterior beliefs become more optimistic, high-valuation buyers consume larger quantities, and assign a higher value to each unit. Their indirect utility is therefore a convex function of the posterior beliefs and consequently these buyers benefit from the diffusion of information.

The aggregate quantity sold in equilibrium is inefficiently low. This is due both to a (static) adverse selection effect and to lower (dynamic) incentives to experiment. In each period, the social planner would not impose downward quantity distortions. Rather, the planner would shift rents from the firm to the consumers, and achieve incentive compatibility by lowering prices. This shift is welfare-improving. Due to the fact that quality and quantity are complements in the buyers' utility function, these gains in efficiency are amplified when the product is of high quality. Compared to a monopolist, the social planner assigns a larger value to information, and hence sells even larger quantities in order to experiment more.

Dynamic Competition: The assumption of a monopoly environment is appropriate in some cases. An example is the early days of Netflix. However, markets for experience goods are often characterized by imperfect competition, and pricing is strategic. We are therefore motivated to extend our analysis of dynamic menu pricing to a competitive setting. For example, consider a model in which a new entrant faces a single “safe” incumbent. We assume the two products are horizontally differentiated, or, in other words, that consumers have idiosyncratic preferences for the products of each firm. In this environment, the role of information becomes even more important. A crucial issue for both firms is whether to invest in learning about the entrant's product. In particular, the entrant can affect the speed of information diffusion on both the intensive margin, through the number of units sold to each buyer, and the extensive margin, by controlling market shares. The incumbent, who is selling a product of known quality, can only affect learning on the extensive margin: pricing aggressively reduces learning about the entrant, while accommodating the new firm accelerates it. We find that the entrant is always willing to invest in acquiring information, while the incumbent regards acquiring information as beneficial only if it believes the relative quality of its product is not very high. Furthermore, experimentation drives the entrant's market share above its myopic equilibrium level, for all values of the posterior beliefs. As in the monopoly case, the amount of experimentation is nonmonotonic, and the entrant's market share is largest when uncertainty about the quality of its product is high.

Relaxing the symmetric learning assumption and extending the competitive analysis to richer specifications of brand preferences are two directions for future research that should provide more insights into the dynamics of (competitive) menu pricing in markets for experience goods.

APPENDIX

PROOF OF PROPOSITION 1:

(1.) By the implicit function theorem, whenever (8) admits a positive solution, the partial derivatives of the myopic supply function are given by

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

When $q_m(\alpha, \theta) > 0$, both these expressions are positive under either assumption 3 or 4. Note that whenever $\phi(\theta) \leq 0$, the myopic quantity is zero.

(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 proof of the next theorem adapts the one in Godfrey Keller and Sven Rady (1997) to the case of nonlinear pricing.

PROOF OF THEOREM 1:

Fix a quantity supply function $q : \Theta \rightarrow \mathbb{R}_+$. Definition (7) shows that $\Pi(\alpha, q)$ is linear in α . Therefore, we can write the expected discounted stream of profits as

$$\begin{aligned}V(\alpha, q) &= \alpha \mathbb{E}_{\mu=\mu_H} \left[\int_0^\infty e^{-rt} \Pi(\alpha_t, q) d\alpha_t \mid \alpha_0 = \alpha \right] \\ &\quad + (1 - \alpha) \mathbb{E}_{\mu=\mu_L} \left[\int_0^\infty e^{-rt} \Pi(\alpha_t, q) d\alpha_t \mid \alpha_0 = \alpha \right].\end{aligned}$$

Now let $\alpha = \eta \alpha_1 + (1 - \eta) \alpha_2$, with $\eta \in [0, 1]$. We then have

$$\begin{aligned}V(\alpha, q) &= \eta V(\alpha_1, q) + (1 - \eta) V(\alpha_2, q) \\ &\leq \eta V^*(\alpha_1) + (1 - \eta) V^*(\alpha_2),\end{aligned}$$

by definition of the value function V^* . Taking the supremum of the left-hand side with respect to q establishes the convexity of V^* . Therefore, to establish continuity, we only need to check at $\alpha = 0$ and $\alpha = 1$. Suppose V^* were not continuous at $\alpha = 0$. Because of convexity, this implies $\lim_{\alpha \rightarrow 0^+} V^*(\alpha) < V^*(0)$, which in turn means there exists a policy q such that $V(0, q) > \lim_{\alpha \rightarrow 0^+} V^*(\alpha)$. But the strict inequality would continue

to hold in a neighborhood of $\alpha = 0$, contradicting the definition of V^* . An identical argument can be used to show continuity at $\alpha = 1$.

The next lemma follows the steps in Keller and Rady (1997), and shows that the HJB equation (11) may be reformulated as a boundary value problem.

LEMMA 1 (Boundary-Value Problem):

Let $\Sigma(\alpha)$ be defined by (2) and let $Q = \int_{\Theta} q(\theta) dF(\theta)$. The HJB equation (11) may be reformulated as

$$(A1) \quad V''(\alpha) = \min_{q: \Theta \rightarrow \mathbb{R}_+} \frac{rV(\alpha) - \Pi(\alpha, q)}{\Sigma(\alpha)Q},$$

with boundary conditions

$$(A2) \quad rV(0) = \Pi_m(0)$$

$$(A3) \quad rV(1) = \Pi_m(1),$$

for all (α, V, V'') with $V'' \geq 0$ and (α, V) lying in the set

$$\mathcal{A} = \{(\alpha, V) \in (0, 1) \times \mathbb{R} : rV \geq \Pi_m\}.$$

PROOF:

Denote $V''(\alpha)$ by s and $rV(\alpha)$ by v . Then define the function

$$B(\alpha, v, s, q) \triangleq \Sigma(\alpha)Qs + \Pi(\alpha, q) - v,$$

and

$$B^*(\alpha, v, s) = \max_{q: \Theta \rightarrow \mathbb{R}_+} B(\alpha, v, s, q).$$

For all (α, v, s) with $s \geq 0$ and $(\alpha, V) \in \mathcal{A}$, the HJB equation (11) can be written as $\max_{q: \Theta \rightarrow \mathbb{R}_+} B(\alpha, v, s, q) = 0$, or $B^*(\alpha, v, s) = 0$. Assumption 2 implies that $Q > 0$. Consequently $B^*(\alpha, v, s) = 0$ if and only if

$$\max_{q: \Theta \rightarrow \mathbb{R}_+} \frac{B(\alpha, v, s, q)}{Q} = 0,$$

which in turn is equivalent to

$$\Sigma(\alpha)s = \min_{q: \Theta \rightarrow \mathbb{R}_+} \frac{v - \Pi(\alpha, q)}{Q}.$$

Therefore, any function that solves (A1), subject to boundary conditions (A2) and (A3), must also solve the HJB equation (11).

We now state an existence theorem for boundary value problems due to Stephen R. Bernfeld and V. Lakshmikantham (1974), which we then use to prove Theorem 2. 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

$$(A4) \quad V'' = G(\alpha, V, V'),$$

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 (A4) 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 (A4) 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 $V(1 - \varepsilon) \in [V_L(1 - \varepsilon), V_H(1 - \varepsilon)]$, (A4) 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 the convergence properties of our solution.

LEMMA 3 (Uniform Convergence):

Under the assumptions of Lemma 2, any infinite sequence of solutions of (A4), with $V_L(\alpha) \leq V(\alpha) \leq V_H(\alpha)$ on J_ε has a uniformly convergent subsequence converging to a solution of (A4) on J_ε .

We can now use these results to prove existence and uniqueness of a solution.

CLAIM 1: *The myopic profit function $\Pi_m(\alpha)/r$ is a strict subsolution of (A1) 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 (A1) 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 (A1) 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:

$$(A5) \quad \min_q \frac{rV - \Pi(\alpha, q)}{\Sigma(\alpha) Q} \leq C.$$

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 (A5) is different from q_m , then it must achieve a lower value than q_m does. We can then define the uniform bound as

$$(A6) \quad C(\varepsilon) = \max_{\alpha \in J_\varepsilon} \left\{ \frac{v(\alpha) - \Pi_m(\alpha)}{\Sigma(\alpha) Q_m(\alpha)} \right\},$$

which ends the proof.

PROOF OF THEOREM 2:

(1.) We know the HJB is equivalent to the boundary value problem (A1). Furthermore, this problem satisfies all conditions of Lemma 2. Therefore, for all $\varepsilon > 0$, the boundary value problem (A1) 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 (A1) on the interval $[1/s, 1 - 1/s]$. Define

the function V_s^n as the restriction of V_s 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 exists a convergent subsequence (which we define as V_n) converging pointwise to a function $V : (0, 1) \rightarrow \mathbb{R}$. By Lemma 2, $|V_n'|$ is uniformly bounded, hence $V_n \rightarrow V$ uniformly on any closed subinterval of $(0, 1)$. Moreover, the constant $C(1/n)$ defined in (A6) yields a uniform bound for $|V_n''|$ on \bar{J}_n . Therefore, V_n' is locally Lipschitz, hence converges uniformly to V' on any closed subinterval of $(0, 1)$. Finally, a standard continuity argument shows that the limit function V actually solves (A1).

The solution to the boundary value problem (A1) is unique. Suppose instead there were two solutions $V_1(\alpha)$ and $V_2(\alpha)$ to (A1), with $V_1(\alpha) \neq V_2(\alpha)$. Without loss of generality, suppose $V_2 > V_1$ for some α , and define the function $G(\alpha, V, V') \triangleq \min_q ((rV - \Pi(\alpha, q)) / \Sigma(\alpha) Q)$. 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, $V_2'' - V_1'' \leq 0$; therefore, the HJB equations imply $G(\alpha, V_1, V_1') \geq G(\alpha, V_2, V_2')$ which contradicts $V_1 < V_2$.

(2.) Under either assumption 3 or 4, 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:

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

Under either assumption 3 or 4, these ratios are well defined whenever (13) admits a positive solution.

Formulation (A1) 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 conditions for the application of a verification theorem. First, by part (1.), there exists a 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 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 Wendell H. Fleming and H. Mete Soner (2006) and conclude that $V(\alpha)$ achieves the maximum of (10).

PROOF OF PROPOSITION 2:

(1.) Define the return function $W(\alpha) \triangleq rV(\alpha)$, 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 (A1) 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 $W''_{r_2} - W''_{r_1} \leq 0$. The formulation (A1) of the HJB equation then implies $G(\alpha, W_{r_1}, W'_{r_1}) \geq G(\alpha, W_{r_2}, W'_{r_2})$, contradicting $r_2 > r_1$ and $W_{r_2}(\alpha) > W_{r_1}(\alpha)$. Since σ and r both enter (A1) 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 $V^*(\alpha_0)$ increases, and so does the return function $W(\alpha_0; r)$. Since $\Sigma(\alpha) V''(\alpha)$ is related to $W_r(\alpha)$ by equation (A1), 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

$$(A7) \quad \frac{\partial q(\alpha, \theta, \Lambda)}{\partial \Lambda} = -\frac{1}{\mu(\alpha) \phi(\theta) u''(q) - c''(q)} \geq 0,$$

with a strict inequality if $q(\alpha, \theta, \Lambda) > 0$. This is clear under assumption 3 (concave virtual surplus). Under assumption 4, it is immediate to show that we must have $c > \Lambda$ for all α , and therefore $\phi(\theta) \leq 0$ implies the optimal quantity is zero. It follows that the denominator in (A7) is strictly negative whenever quantity is strictly positive. Finally, because the value of information is identically equal to zero in the myopic case and it is given by $\Lambda(\alpha) > 0$ 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 quantity q offered 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 (15) gives us the expression for $q_m(\alpha, \theta)$. From first order condition (15), and using constraint (4), we obtain

$$\begin{aligned} p(\alpha, \theta) &= \mu(\alpha) \theta q(\alpha, \theta) - U(\alpha, \theta) \\ &= \mu(\alpha) \left(\theta (\mu(\alpha) \phi(\theta) + \Sigma(\alpha) V''(\alpha)) - \int_{\theta_L}^{\theta} (\mu(\alpha) \phi(s) + \Sigma(\alpha) V''(\alpha)) ds \right) \\ (A8) &= \mu^2(\alpha) \left(\theta \phi(\theta) + \int_{\theta_L}^{\theta} \phi(s) ds \right) + \mu(\alpha) \theta_L \Sigma(\alpha) V''(\alpha). \end{aligned}$$

(2.) Substituting first order condition (15) into the objective function, we obtain:

$$\begin{aligned} rV(\alpha) &= \int_{\theta_L}^{\theta_H} \frac{1}{2} (\mu(\alpha) \phi(\theta) + \Sigma(\alpha) V''(\alpha))^2 f(\theta) d\theta \\ (A9) \quad &\int_{\theta_L}^{\theta_H} \frac{1}{2} (\mu(\alpha) \phi(\theta))^2 f(\theta) d\theta + \mu(\alpha) \mathbb{E}_\theta[\phi] \Sigma(\alpha) V''(\alpha) + \frac{1}{2} (\Sigma(\alpha) V''(\alpha))^2. \end{aligned}$$

The first term in (A9) 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 (15), 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 (20) 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 (15) 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

$$(A10) \quad \frac{\partial q(\alpha, \theta; r)}{\partial \alpha} = (\mu_H - \mu_L)(\phi(\theta) - \mathbb{E}_\theta[\phi]) + \frac{rV'(\alpha; r) - (\mu_H - \mu_L)\mu(\alpha)\text{Var}[\phi]}{\sqrt{2rV(\alpha; r) - \mu(\alpha)^2\text{Var}[\phi]}}.$$

The second derivative is given by

$$(A11) \quad \frac{\partial^2 q(\alpha, \theta; r)}{(\partial \alpha)^2} = \frac{rV''(\alpha; r) - (\mu_H - \mu_L)^2\text{Var}[\phi] - (rV'(\alpha; r) - (\mu_H - \mu_L)\mu(\alpha)\text{Var}[\phi])^2}{\sqrt{2rV(\alpha; r) - \mu(\alpha)^2\text{Var}[\phi]}}.$$

Now consider an interval $[\varepsilon, 1 - \varepsilon]$ with $\varepsilon > 0$. We know that the second derivative of the value function $V''(\alpha; r)$ is uniformly bounded from above for all r . The bound $C(\varepsilon)$ is defined in equation (A6). From expression (A11), we know that if $rV''(\alpha; r) - (\mu_H - \mu_L)^2\text{Var}[\phi] \leq 0$, then $\partial^2 q(\alpha, \theta; r)/(\partial \alpha)^2 < 0$. 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 (A10). Evaluate expression (A10) 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 (A10) at $\alpha = 1$ and let $W_r(\alpha) = rV(\alpha; r)$. We can then write $\partial q(\alpha, \theta; r)/\partial \alpha$ as

$$(A12) \quad \frac{\partial q(1, \theta; r)}{\partial \alpha} = (\mu_H - \mu_L)(\phi(\theta) - \mathbb{E}_\theta[\phi]) + \frac{W'_r(1) - (\mu_H - \mu_L)\mu_H\text{Var}[\phi]}{\mu_H\mathbb{E}_\theta[\phi]}.$$

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 (A12) is increasing in r . The right-hand side of (A12) 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 $W'_r(1)$ 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 (A12) 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 $\Sigma(\alpha) V''(\alpha)$, and obtain the result in the text.

PROOF OF PROPOSITION 10:

The drift component of the process $dh(\alpha_t, \theta)$ is given by

$$(A13) \quad \mathbb{E}[dh(\alpha_t, \theta)] = \frac{\partial^2 h(\alpha_t, \theta)}{(\partial \alpha)^2} \Sigma(\alpha_t) Q_t dt.$$

The result then follows directly from equation (A13), Proposition 6 and Theorem 3.

PROOF OF PROPOSITION 11:

This result follows directly from equations (23) and (24), from Proposition 6, and from Theorem 3.

*

REFERENCES

- Ackerberg, Daniel A.** 2003. "Advertising, learning, and consumer choice in experience good markets: an empirical examination." *International Economic Review*, 44(3): 1007–1040.
- Akçura, M. Tolga, Füsün F. Gonul, and Elina Petrova.** 2004. "Consumer Learning and Brand Valuation: An Application on Over-the-Counter Drugs." *Marketing Science*, 23(1): 156–169.
- Bergemann, Dirk, and Juuso Välimäki.** 1997. "Market Diffusion with Two-Sided Learning." *RAND Journal of Economics*, 28(4): 773–795.
- Bergemann, Dirk, and Juuso Välimäki.** 2002. "Entry and Vertical Differentiation." *Journal of Economic Theory*, 106(1): 91–125.
- Bergemann, Dirk, and Juuso Välimäki.** 2006. "Dynamic Pricing of New Experience Goods." *Journal of Political Economy*, 114(4): 713–743.
- Bernfeld, Stephen R., and V. Lakshmikantham.** 1974. *An Introduction to Nonlinear Boundary Value Problems*. Academic Press.
- Bolton, Patrick, and Christopher Harris.** 1999. "Strategic Experimentation." *Econometrica*, 67(2): 349–374.
- Bolton, Patrick, and Christopher Harris.** 2000. "Strategic Experimentation: the Undiscounted Case." *Incentives and Organization: Essays in Honour of James Mirrlees*.
- Boone, Jan, and Joel Shapiro.** 2008. "Selling to Consumers with Endogenous Types." *Working Paper, Tilburg University and Universitat Pompeu Fabra*.
- Bose, Subir, Gerhard Orosel, Marco Ottaviani, and Lise Vesterlund.** 2006. "Dynamic monopoly pricing and herding." *Rand Journal of Economics*, 37(4): 910–928.
- Bose, Subir, Gerhard Orosel, Marco Ottaviani, and Lise Vesterlund.** 2008. "Monopoly pricing in the binary herding model." *Economic Theory*, 37(2): 203–241.
- Braden, David J., and Shmuel S. Oren.** 1994. "Nonlinear Pricing to Produce Information." *Marketing Science*, 13(3): 310–326.
- Crawford, Gregory S., and Matthew Shum.** 2005. "Uncertainty and Learning in Pharmaceutical Demand." *Econometrica*, 73(4): 1137–1173.
- Dutta, Prajit 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, Tulin, and Michael P. Keane.** 1996. "Decision-making Under Uncertainty: Capturing Dynamic Brand Choice Processes in Turbulent Consumer Goods Markets." *Marketing Science*, 15: 1–20.
- Faingold, Eduardo, and Yuliy Sannikov.** 2011. "Reputation in Continuous-Time Games." *Econometrica*, forthcoming.
- Fleming, Wendell H., and H. Mete Soner.** 2006. *Controlled Markov Processes and Viscosity Solutions*. Springer.
- Gärtner, Dennis.** 2010. "Monopolistic screening under learning by doing." *Rand Journal of Economics*, 41(3): 574–597.
- Göttler, Ronald L., and Karen Clay.** 2011. "Tariff Choice with Consumer Learning and Switching Costs." *Journal of Marketing Research*, forthcoming.
- Gowrisankaran, Gautam, and Marc Rysman.** 2007. "Dynamics of Consumer Demand for New Durable Consumer Goods." *Working paper, Boston University*.
- Hitsch, Günter J.** 2006. "An Empirical Model of Optimal Dynamic Product Launch and Exit Under Demand Uncertainty." *Marketing Science*, 25(1): 25.
- Israel, Mark.** 2005. "Services as Experience Goods: An Empirical Examination of Consumer Learning in Automobile Insurance." *American Economic Review*, 95(5): 1444–1463.
- Keller, Godfrey, and Sven Rady.** 1997. "Optimal Experimentation in a Changing Environment." *STICERD-Theoretical Economics Paper Series*.
- Keller, Godfrey, and Sven Rady.** 1999. "Optimal Experimentation in a Changing Environment." *Review of Economic Studies*, 66(3): 475–507.
- Keller, Godfrey, Sven Rady, and Martin Cripps.** 2005. "Strategic experimentation with exponential bandits." *Econometrica*, 39–68.
- Larkin, Ian.** 2008. "Bargains-then-ripoffs: Innovation, Pricing and Lock-in in Enterprise Software." *Proceedings of the Academy of Management*.
- Lewis, Tracy R., and Huseyin Yildirim.** 2002. "Learning by Doing and Dynamic Regulation." *RAND Journal of Economics*, 33(1): 22–36.
- Liptser, Robert, and Albert Shiryaev.** 1977. *Statistics of Random Processes, I*. Springer-Verlag, New York.
- Maskin, Eric, and John Riley.** 1984. "Monopoly with Incomplete Information." *RAND Journal of Economics*, 15(2): 171–196.
- Moscarini, Giuseppe, and Lones Smith.** 2001. "The Optimal Level of Experimentation." *Econometrica*, 69(6): 1629–1644.

- Mussa, Michael, and Sherwin Rosen.** 1978. "Monopoly and Product Quality." *Journal of Economic Theory*, 18(2): 301–317.
- Ottaviani, Marco, and Andrea Prat.** 2001. "The value of public information in monopoly." *Econometrica*, 69(6): 1673–1683.
- Ramsey, Frank P.** 1928. "A Mathematical Theory of Saving." *Economic Journal*, 38(152): 543–559.
- Song, Inseong, and Pradeep K. 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.
- Villas-Boas, J. Miguel.** 2004. "Consumer Learning, Brand Loyalty, and Competition." *Marketing Science*, 23(1): 134–145.
- Villas-Boas, J. Miguel.** 2006. "Dynamic Competition with Experience Goods." *Journal of Economics & Management Strategy*, 15(1): 37–66.