

Dynamic Pricing and Product Differentiation with Cost Uncertainty and Learning

N. Bora Keskin, John R. Birge

Dynamic Pricing and Product Differentiation with Cost Uncertainty and Learning

(Authors' names blinded for peer review)

Motivated by applications in the health insurance industry, we consider a seller who designs and sells a set of vertically differentiated products to a population of quality-sensitive customers. The seller's business environment entails an uncertainty about production costs. We characterize the seller's optimal price-quality schedule in the cases of: (a) static cost uncertainty, and (b) dynamic learning about cost uncertainty through noisy observations on an underlying cost curve. We prove that the seller's optimal quality allocations in (a) and (b) stand in stark contrast: While a seller facing static cost uncertainty *degrades* the quality in its product offering, a dynamically learning seller *improves* the quality of its products to accelerate information accumulation. In the case of dynamic learning, we prove that the seller exercises the most extreme experimentation on less quality-sensitive customers. We also extend our results to the cases of commonly used regulations in health insurance industry, and show how such regulations moderate the interplay between uncertainty and learning.

Key words: Dynamic programming, Bayesian learning, pricing, self-selection, exploration-exploitation.

1. Introduction

Pricing and product differentiation decisions have dual goals: First, they can help firms design products for the varying needs of potential customers, and increase profits by implementing a menu of products from which customers with some unobservable characteristics can self-select. Secondly, these decisions can also serve as tools to learn about the potential uncertainties present in a firm's business environment, thereby enabling the firm to adaptively improve its product design and increase future profits. Some of the key questions in this context are the following: If a firm's business environment entails uncertainties that can influence production costs, how do these uncertainties affect the firm's optimal pricing and product differentiation decisions? What are the benefits and consequences of using these decisions as learning tools? How should a firm design its product offering to learn optimally?

This paper sheds light on the aforementioned questions by introducing a dynamic pricing and production differentiation problem with cost uncertainty and learning. In our problem formulation, we consider a firm that faces an uncertainty about its production costs. There is a *cost curve* that gives the cost of production as a function of a certain quality that makes the product more desirable to customers. We model the firm's uncertainty about costs as a Bernoulli belief distribution on two possible cost curves; that is, the firm's production costs may turn out to be high or low due to the unresolved uncertainty in the firm's business environment. To measure the effect of this cost uncertainty, we use as a benchmark the case of perfect cost information, which corresponds to the

expected profit performance of a clairvoyant who knows the underlying cost curve. In this setting, we characterize the firm's optimal pricing and production differentiation policy and show that the presence of cost uncertainty leads to a quality degradation in the firm's optimal policy, relative to the expected quality allocation of the clairvoyant. In stark contrast, if the firm can make sequential observations on the cost curve, and dynamically update its product offering, we show that the firm optimally improves the qualities in its product offering to accelerate learning. Assuming that the population of customers are heterogeneous in terms of their sensitivity to quality, we examine the structure of the firm's optimal learning policy and identify that the firm exercises more experimentation on less quality-sensitive customers. Furthermore, we generalize our analysis to the cases of commonly used regulations that subsidize quality in practice, namely a free outside option with positive but low quality and price subsidies.

Vertically differentiated products are present in many industries. A notable example is the health insurance industry where insurance contracts are designed and priced to satisfy significantly varying healthcare needs of insurance customers. Thus, insurance companies can sell a range of products tailored to their customers' differing preferences. However, when there is a new regulation or a technological innovation, an insurance company's cost of providing the same service can shift upward or downward, creating a cost uncertainty from the company's perspective. In a recent study, Kowalski (2014) reports that the introduction of the Affordable Care Act (ACA) in the United States has led to increased average costs of health insurance companies in some states, but reduced average costs in some others. According to Kowalski (2014), some of this effect can be explained by technological issues such as the glitches in the enrollment system and political issues such as the degree of cooperation between the federal and state governments. Moreover, the introduction of such regulations enlarges the insurance companies' customer base to include individuals with little or no understanding of how to use health insurance (Goodnough 2014), creating further uncertainty about the potential cost of providing service. Motivated by this application, our paper focuses on the possible theoretical implications of cost uncertainties on a firm's pricing and product design decisions. To that end, we develop a stylized framework of dynamic pricing and product differentiation with cost uncertainty, with the hope of gaining insights into more realistic product design problems.

Summary of main contributions. This paper makes four contributions to the literature on learning and earning. First, we develop and analyze a dynamic learning model that extends the research literature on model uncertainty and dynamic learning to the case of self-selection mechanisms. As will be discussed below, the vast majority of the related research has focused on either the dynamic control of posted prices or the optimal timing of new technology adoption. Our paper shows how

product design decisions such as product differentiation and pricing can be jointly used a learning tool, and we hope that this work will lead to further research on dynamic product design with cost uncertainty. Our analysis of the dynamic learning model provides key tools for further studying dynamic nonlinear pricing problems with model uncertainty and learning. In our analysis, we define a *conditional value function* that is adapted to the quality choices of customers and then characterize the optimal pricing and quality allocation decisions based on the conditional value of learning. Second, our characterization of the optimal product offering identifies two starkly contrasting effects in the cases of: (a) static cost uncertainty, and (b) dynamic learning about cost uncertainty. While the presence of cost uncertainty makes a firm *degrade* the quality in its product offering, an opportunity to learn about this cost uncertainty makes the firm *improve* the quality of its products. It is perhaps worth noting that our quality degradation result, which is driven by uncertainty, is distinct from the standard quality degradation results in nonlinear pricing theory, which are usually caused by monopolistic distortion. Third, we provide key structural insights into the design of optimal experimentation for dynamic learning about cost uncertainty. Our analysis reveals that a firm projects the value of learning onto its product offering by improving the quality uniformly for all products and implementing the strongest quality improvement for the products that target less quality-sensitive customers. We believe that this structure of our optimal policy can provide a practical guideline for designing products in presence of cost uncertainty. Finally, we extend our analysis to the case of commonly used regulations in the health insurance industry, shedding some light on how these regulations would moderate the interplay between static uncertainty and dynamic learning, and how a firm should design its product offering under such regulations.

Related literature. Our work is related to three streams of research, namely (i) the economics literature on nonlinear pricing and self-selection, (ii) the decision analysis literature on dynamic programming with Bayesian learning, and (iii) the operations research and management science (OR/MS) and economics literature on dynamic pricing with model uncertainty and learning.

With regard to (i), there is a rich economics literature on vertical product differentiation and self-selection. The origins of this literature can be traced back to Mussa and Rosen (1978) and Maskin and Riley (1984), who have formulated and analyzed the problem of designing and pricing a set of products to maximize profits earned from a population of quality-sensitive customers. A key result derived in this framework is that monopolistic pricing reduces the quality in a seller's product offering, relative to perfectly competitive pricing. Besanko, Donnenfeld and White (1987, 1988) have studied how this quality distortion due to monopoly power can be remedied by implementing regulations such as a price ceiling or minimum quality standards, which are motivated by the regulations in the

cable television industry in the United States. Another paper that studies the monopoly distortion with regulations is due to Jullien (2000), who has investigated sufficient conditions for customers' participation in the presence of price cap regulations. Rochet and Stole (2002) have extended this literature by introducing randomness to the customers' valuations for outside options, showing that such random outside options can eliminate some of the monopoly distortion effect. More recently, Bergemann, Shen, Xu and Yeh (2012) have studied a nonlinear pricing problem with a finite menu constraint, characterizing the loss due to the limitation to a finite menu of products. To the best of our knowledge, the aforementioned economics literature on nonlinear pricing and product differentiation has not focused on cost model uncertainty and dynamic learning. Our work contributes to this literature by introducing and analyzing the optimal dynamic control of a nonlinear price-quality schedule in the presence of cost model uncertainty and learning. In particular, we show that the presence of a static cost uncertainty could result in further quality reduction in a monopolist's product offering and that a dynamically learning monopolist can improve quality significantly to increase future gains.

With regard to (ii), a main focus of attention has been the analysis of technology adoption problems with uncertainty about benefits and costs of adoption. In this context, McCardle (1985) and Lippman and McCardle (1987) have studied optimal stopping problems in which a firm explores the profitability of a technology to decide whether to adopt the technology or not and have characterized the optimal technology adoption policies in their settings. Ulu and Smith (2009) have analyzed a more general model of technology adoption involving non-stationary technologies and fairly general signal processes and have derived key comparative statics results regarding the optimal policy. In a more recent study, Smith and Ulu (2012) have considered the case of repeated technology adoption decisions with time-varying costs and benefits, showing that the structure of the optimal policy for repeated technology adoptions differs significantly from the one for single technology adoption. Apart from these studies, Mazzola and McCardle (1996) have focused on a production planning problem where a firm has recently adopted a new technology and faces an uncertainty about its learning curve and have showed that in the presence of learning-curve uncertainty the firm's optimal production does not necessarily increase with cumulative production. Like the majority of the above studies, our paper employs a Bayesian decision model based on dynamic programming. But, the problem we study is a dynamic control problem rather than an optimal stopping problem, which is used in many of the studies mentioned above (see, e.g., McCardle 1985, Lippman and McCardle 1987, Ulu and Smith 2009). More importantly, the decisions in our framework, namely nonlinear price-quality schedules of the firm, are continuous mappings and hence have very high dimensions compared to the decision variables in the aforementioned decision analysis studies. To analyze this framework, we introduce and study a

conditional value function that is adapted to the latest customer choice observed by the firm, and provide an analysis of how to use high dimensional price-quality schedules as a dynamic learning tool.

With regard to (iii), there has been substantial effort to study the tradeoff between learning and earning in dynamic pricing applications. In the economics literature, Keller and Rady (1999) have studied a continuous-time stochastic control problem in which a firm is uncertain about the time-varying demand curve for its product and learns about the unknown demand curve by experimenting with sales quantities, and have characterized the optimal policy in their setting. Following this work, Bergemann and Välimäki (2000) and Keller and Rady (2003) have formulated and analyzed duopolistic competition models for dynamic pricing. Bergemann and Välimäki (2000) have considered a market with an incumbent seller and an entrant, assuming that buyers and sellers are uncertain about the value of the entrant's product and receive publicly observable signals about the unknown product value. In this setting, they have analyzed the Markov perfect equilibrium that describes the pricing strategies of the sellers. On the other hand, Keller and Rady (2003) have considered two sellers who are uncertain about the parameters of the demand curves for their substitutable products, and have characterized the sellers' pricing strategies in a Markov perfect equilibrium. There are several distinguishing features that differentiate our work from these studies. First, as explained in the preceding paragraph, our model involves the dynamic control of a price-quality schedule. To be more precise, the seller in our problem dynamically designs a continuum of price-quality pairs throughout the selling horizon, which separates our work from standard dynamic pricing problems in terms of modeling and analysis. Second, unlike the above studies, we investigate the impact of cost uncertainty on the optimal policy, and identify the contrasting effects of static cost uncertainty and dynamic learning about cost uncertainty. This analysis provides unique insights about how to design experiments with a price-quality schedule with the aim of resolving an underlying cost uncertainty. Third, we use discrete-time dynamic programming rather than continuous-time stochastic control, which makes our mathematical tools and proof techniques distinct. On top of the related economics literature discussed above, there has been recent interest in the OR/MS literature on dynamic pricing problems with demand learning. An early study in this research stream is due to Aviv and Pazgal (2005), who analyze a dynamic pricing problem in which a firm sells a single product while facing Bayesian uncertainty about the demand curve for its product. More recently, Araman and Caldentey (2009), Farias and van Roy (2010), and Harrison, Keskin and Zeevi (2012) have further investigated dynamic pricing with Bayesian demand learning and characterized well-performing policies that balance the tradeoff between learning and earning. A common feature of these studies is that they focus on the pricing of a single product, rather than multiple products. While the extension from single-product

pricing to multi-product pricing has been recently studied by Keskin and Zeevi (2014) and den Boer (2014), these papers, too, consider a pre-determined set of products. In summary, the aforementioned OR/MS research on dynamic pricing and learning has essentially focused on using dynamic posted-price mechanisms as experimentation tools, either for a single product or for an exogenously given set of products. Unlike this research stream, we study the possibility of jointly using pricing and product differentiation decisions as learning tools and explore how learning affects the design of self-selection mechanisms. Moreover, the vast majority of the recent OR/MS work on dynamic pricing and learning have concentrated on near-optimal learning policies, whereas we investigate in this paper the structure of exactly optimal learning policies for dynamic pricing and product differentiation.

Organization of the paper. This paper is organized as follows. Section 2 describes the basic problem formulation and analyzes the case of static cost uncertainty. In Section 3, we extend our basic problem formulation and analysis to the case of dynamic learning. Section 4 investigates how regulations moderate the effects of static cost uncertainty and dynamic learning. Finally, Section 5 concludes with possible extensions of this study. Proofs of key results are deferred to appendices.

2. Problem Formulation and Analysis of Static Cost Uncertainty

Consider a firm that sells a product that can be produced at different levels of quality. The product's quality q takes values in $\mathcal{Q} = [0, \infty)$. The firm can produce a range of differentiated qualities in $\mathcal{Q}$ and can charge a different price for every distinct quality level in its product offering. If a customer purchases a product with quality q at price p , then the customer derives a net utility of

$$U(\theta, q, p) = \theta q - p, \quad (2.1)$$

where θ denotes the customer's sensitivity to quality. There is a population of potential customers whose quality sensitivities are unobservable and distributed according to a density function $f(\theta)$ on $\Theta = [\theta_{\min}, \theta_{\max}]$. We assume that $f(\theta)$ has increasing failure rate; i.e., $f(\theta)/\bar{F}(\theta)$ increases in θ , where $\bar{F}(\theta) = \int_{\theta}^{\theta_{\max}} f(\xi) d\xi$. (An important example that satisfies this property is the uniform distribution; other noteworthy examples are Gaussian and exponential distributions.) The expected cost of producing and selling one unit of the product at quality level q is $c(q)$. Due to technological innovations and operational regulations that can potentially happen after the time of sales, the firm is uncertain about the cost curve $c(\cdot)$ at the time of sales. We model this uncertainty via a Bernoulli prior distribution on two possible cost curves. Prior to sales, nature chooses either $c_0(\cdot)$ or $c_1(\cdot)$ as the underlying cost curve, where $c_i(q) = a_i q^2$ for $i = 0, 1$, and $0 < a_0 < a_1 < \infty$. Without observing the underlying cost curve, the firm initially believes that

$$c(\cdot) = \begin{cases} c_1(\cdot) & \text{with probability } b \\ c_0(\cdot) & \text{with probability } 1 - b, \end{cases} \quad (2.2)$$

where $b \in [0, 1]$. We hereafter refer to b as the firm's *prior belief*. To serve its customers, the firm offers the product with quality q at a price $P(q)$ for all $q \in \mathcal{Q}$. As will be explained below, the firm can implement a selling mechanism $\mathcal{S} = \{(p(\theta), q(\theta)) : \theta \in \Theta\}$ that assigns a price-quality pair $(p(\theta), q(\theta))$ to each quality sensitivity parameter θ , ensuring that every customer selects the price-quality pair assigned to its quality sensitivity. For brevity, we hereafter refer to a pair of pricing and quality allocation functions, $p : \Theta \rightarrow \mathbb{R}_+$ and $q : \Theta \rightarrow \mathcal{Q}$ respectively, as a *product offering*. Therefore, the firm's objective is to choose a product offering to maximize the expected profit from all potential customers; that is, the firm solves the following problem:

$$V(b) = \max_{p(\cdot), q(\cdot)} \left\{ \int_{\theta_{\min}}^{\theta_{\max}} (p(\theta) - C(b, q(\theta))) f(\theta) d\theta \right\} \quad (2.3a)$$

$$\text{s.t. } \theta q(\theta) - p(\theta) \geq \theta q(\tilde{\theta}) - p(\tilde{\theta}) \quad \text{for all } \theta, \tilde{\theta} \in \Theta, \quad (2.3b)$$

$$\theta q(\theta) - p(\theta) \geq 0 \quad \text{for all } \theta \in \Theta, \quad (2.3c)$$

$$q(\theta) \geq 0 \quad \text{for all } \theta \in \Theta, \quad (2.3d)$$

where $C(b, q) = bc_1(q) + (1 - b)c_0(q)$ is the expected cost of production when the firm's prior belief is b and the quality level is q . The constraints (2.3b) are usually referred to as *incentive compatibility* (IC) constraints, which ensure that a customer with quality sensitivity parameter θ prefers $(p(\theta), q(\theta))$ to any other price-quality pair in the firm's product offering. The constraints (2.3c) are referred to as *individual rationality* (IR) constraints, which ensure that a customer with quality sensitivity parameter θ prefers $(p(\theta), q(\theta))$ to an outside option that offers a net utility of zero. The optimal objective value of the problem in (2.3), which is denoted by $V(b)$, is the expected maximized profit as a function of the firm's initial belief b .

Our first goal is to express the optimization problem in (2.3) in a simpler form. To that end, we start with characterizing a set of conditions that describe the constraints of this problem. By standard results in nonlinear pricing theory (see, e.g., Mussa and Rosen 1978), we have the following proposition.

PROPOSITION 1. (incentive compatibility and individual rationality)

- (i) *The IC constraints (2.3b) hold if and only if $q(\cdot)$ is increasing and differentiable almost everywhere, and $p'(\theta) = \theta q'(\theta)$ for $\theta \in \Theta$, except on a set of measure zero.*
- (ii) *Assume that the IC constraints (2.3b) hold, and that $\theta_{\min} q(\theta_{\min}) - p(\theta_{\min}) = 0$. Then the IR constraints (2.3c) also hold.*

Proposition 1(i) transforms the constraints of (2.3) into a local differential relationship between pricing and quality allocation functions. In particular, it prescribes a particular marginal price increase

for an additional unit of quality. Using Proposition 1, we deduce that the pricing function $p(\cdot)$ must satisfy the following to induce the IC and IR constraints:

$$\begin{aligned} p(\theta) &= p(\theta_{\min}) + \int_{\theta_{\min}}^{\theta} p'(\xi) d\xi \stackrel{(a)}{=} p(\theta_{\min}) + \int_{\theta_{\min}}^{\theta} \xi q'(\xi) d\xi \\ &\stackrel{(b)}{=} p(\theta_{\min}) + \theta q(\theta) - \theta_{\min} q(\theta_{\min}) - \int_{\theta_{\min}}^{\theta} q(\xi) d\xi \\ &\stackrel{(c)}{=} \theta q(\theta) - \int_{\theta_{\min}}^{\theta} q(\xi) d\xi \end{aligned} \quad (2.4)$$

for all $\theta \in \Theta$, where: (a) follows by Proposition 1(i), (b) follows by integration by parts, and (c) follows by Proposition 1(ii). This simplifies the optimization problem in (2.3) as follows:

$$V(b) = \max_{q(\cdot) \in \mathcal{A}^*} \left\{ \int_{\theta_{\min}}^{\theta_{\max}} \left(\theta q(\theta) - \int_{\theta_{\min}}^{\theta} q(\xi) d\xi - C(b, q(\theta)) \right) f(\theta) d\theta \right\}, \quad (2.5)$$

where $\mathcal{A}^*$ is the set of functions $q: \Theta \rightarrow \mathbb{R}_+$ that are increasing and differentiable almost everywhere. Using integration by parts, and the fact that $F(\theta_{\max}) = \int_{\Theta} f(\theta) d\theta = 1$, we note that the second term in the integrand above is equal to $-\overline{F}(\theta)q(\theta)/f(\theta)$, because

$$\int_{\theta_{\min}}^{\theta_{\max}} \int_{\theta_{\min}}^{\theta} q(\xi) f(\theta) d\xi d\theta = F(\theta_{\max}) \int_{\theta_{\min}}^{\theta_{\max}} q(\xi) d\xi - \int_{\theta_{\min}}^{\theta_{\max}} F(\theta) q(\theta) d\theta = \int_{\theta_{\min}}^{\theta_{\max}} \overline{F}(\theta) q(\theta) d\theta. \quad (2.6)$$

Consequently, we can state the firm's optimization problem as

$$V(b) = \max_{q(\cdot) \in \mathcal{A}^*} \left\{ \int_{\theta_{\min}}^{\theta_{\max}} \left(\theta q(\theta) - \frac{\overline{F}(\theta)}{f(\theta)} q(\theta) - C(b, q(\theta)) \right) f(\theta) d\theta \right\}. \quad (2.7)$$

Let $q^*(b, \theta)$ be optimal quality allocation for the problem in (2.7). In our next result, we present how $q^*(b, \theta)$ is influenced by the firm's belief b .

PROPOSITION 2. (optimal quality allocation) *In any open subset of $[0, 1] \times \Theta$ on which the optimal quality allocation $q^*(b, \theta)$ is positive, $q^*(b, \theta)$ is strictly decreasing and strictly convex in b .*

In Proposition 2, we describe how the firm's optimal quality allocation reacts to cost uncertainty. As the firm's belief b in the higher cost curve $c_1(\cdot)$ increases, the firm's expected cost of production increases. In response to this increase in expected costs, the firm optimally reduces its quality offering uniformly for all customers, which makes $q^*(b, \theta)$ decrease in b . A more interesting point is that the quality reduction occurs at a decreasing rate. The main reason behind this result is that, when the quality allocation $q^*(b, \theta)$ is an interior solution, it decreases in b in inverse proportion to the marginal cost of production, $\partial C(b, q)/\partial q$. Because the relative changes in the marginal cost becomes smaller at larger values of b , the firm chooses to reduce quality at a decreasing rate, which makes $q^*(b, \theta)$ convex in b .

Our next goal is to assess the inefficiency due to the uncertainty in the cost structure. As a performance benchmark to measure the value of cost information, let us consider the expected profit of a clairvoyant who knows the underlying cost curve. Depending on the cost curve, which could be $c_0(\cdot)$

or $c_1(\cdot)$, the clairvoyant would solve (2.3) with $b = 0$ or $b = 1$. The resulting optimal objective value, which equals $V(i)$ and is hereafter called the *clairvoyant profit*, is the maximized profit under the cost curve $c_i(\cdot)$ for $i = 0, 1$. We define the firm's loss due to uncertainty as the difference between the expected clairvoyant profit and the firm's optimal profit, which is given by

$$\Delta(b) := bV(1) + (1 - b)V(0) - V(b) \quad \text{for all } b \in [0, 1]. \quad (2.8)$$

To characterize the impact of uncertainty on the firm's quality allocation, let us denote by $\delta(b, \theta)$ the deviation of $q^*(b, \theta)$ from the expected clairvoyant quality allocation; i.e.,

$$\delta(b, \theta) := bq^*(1, \theta) + (1 - b)q^*(0, \theta) - q^*(b, \theta) \quad \text{for all } b \in [0, 1] \text{ and } \theta \in \Theta. \quad (2.9)$$

Note that positive (negative) values of $\delta(b, \theta)$ correspond to quality decreases (increases), relative to the case of perfect cost information. In our next result, we derive a lower bound on the loss due to uncertainty.

PROPOSITION 3. (firm's loss due to uncertainty) *$\Delta(b)$ is concave, and there exists a finite positive constant K such that $\Delta(b) \geq Kb(1 - b) \geq 0$ for all $b \in [0, 1]$.*

Proposition 3 characterizes how the uncertainty about costs can affect the firm's profit performance. It shows that the presence of cost uncertainty makes the firm incur a non-negative loss relative to the clairvoyant and that the amount of this loss is greater than a function that increases with the variance of the firm's belief distribution, namely $b(1 - b)$. Figure 1 displays the loss function $\Delta(\cdot)$ in a numerical example to demonstrate this result.




Figure 1 Firm's loss due to uncertainty. The loss $\Delta(b)$ attains larger values for intermediate values of the firm's belief b . The consumers' quality sensitivity parameters are uniformly distributed between 0 and 1. The cost parameters are $a_0 = 0.1$ and $a_1 = 1$.

In the last result of this section, we show the extent of quality degradation due to cost uncertainty.

THEOREM 1. (quality degradation due to uncertainty) *There exists an increasing function $g: \Theta \rightarrow \mathbb{R}_+$ such that $g(\theta_{\max}) > 0$ and $\delta(b, \theta) \geq g(\theta)b(1 - b) \geq 0$ for all $b \in [0, 1]$ and $\theta \in \Theta$.*

Theorem 1 shows that the firm reacts to the cost uncertainty by choosing positive values for $\delta(b, \theta)$ to degrade the quality allocation in its product offering. We note that this effect is different than the standard quality degradation results in nonlinear pricing theory: As shown by Mussa and Rosen

(1978), the firm is incentivized to degrade quality for customers whose quality sensitivity parameters are less than $\theta_{\max}$; we already see this effect in both the optimal and the clairvoyant quality allocations. On top of this effect, Theorem 1 shows that the quality allocation in the case of cost uncertainty is strictly lower than the expected quality allocation under perfect cost information, which creates another distortion in the dimension of beliefs. The magnitude of this distortion increases with: (a) the customers' quality sensitivity, θ ; and (b) the variance of the firm's belief distribution, $b(1 - b)$. In stark contrast to the standard quality degradation result of Mussa and Rosen (1978), which prescribes no quality distortion for top consumers, Theorem 1 implies that the quality degradation due to cost uncertainty affects the quality-sensitive consumers the most, thereby creating a substantial quality distortion for top consumers.

Figure 2 displays an optimal quality allocation example, in which the customers' quality sensitivity parameters are distributed uniformly on the unit interval. The two quality degradation effects mentioned above are shown as arrows (A) and (B). The monopoly distortion effect in (B) vanishes at the upper boundary of $\Theta = [0, 1]$, whereas the uncertainty distortion effect in (A) attains its maximum value at the upper boundary of $\Theta = [0, 1]$, leading to significant quality degradation at intermediate belief levels that are away from 0 or 1.

3. Dynamic Learning about Cost Uncertainty

In this section, we extend the original problem formulation in the preceding section to the case where the firm dynamically learns about the underlying cost curve. Suppose that the firm now sells its products over T sales periods, where each sales period entails a learning opportunity about costs. In every period $t = 1, 2, \dots, T$, the firm first chooses a product offering $\mathcal{S}_t = \{(p_t(\theta), q_t(\theta)) : \theta \in \Theta\}$, which is a set of price-quality pairs to be offered to its customers. After that, a customer whose quality sensitivity θ_t is randomly drawn from the density $f(\cdot)$ arrives and selects a price-quality pair from the firm's product offering $\mathcal{S}_t$ (or the outside option of zero net utility). Let $Q_t = q_t(\theta_t)$ denote the quality choice of the customer arriving in period t . Following a sale, the firm incurs the following cost of production:

$$C_t = c(Q_t) + \epsilon_t \quad \text{for } t = 1, 2, \dots, \quad (3.1)$$

where $c(\cdot)$ is underlying cost curve which is unknown to the firm, and $\epsilon_t \stackrel{\text{iid}}{\sim} \mathcal{N}(0, \sigma^2)$ are idiosyncratic cost shocks of consumers.

To formally define dynamic pricing-and-quality-allocation policies in this framework, let us first recall the firm's prior belief, which was originally expressed in (2.2):

$$c(\cdot) = \begin{cases} c_1(\cdot) & \text{with probability } b_1 \\ c_0(\cdot) & \text{with probability } 1 - b_1, \end{cases} \quad (3.2)$$




Figure 2 Optimal quality allocation and distortions. The shaded surface displays the optimal quality allocation as function of: (i) the firm's belief $b \in [0, 1]$, and (ii) the consumers' quality sensitivity parameter θ , which is assumed to be uniformly distributed between 0 and 1. The cost parameters are $a_0 = 0.1$ and $a_1 = 1$. For positive values of $q^*(b, \theta)$, we observe that $q^*(b, \theta)$ decreases convexly in b . The dashed line on the left shows the cross section of the quality allocation in the case of perfect cost information, and the arrow (A) shows the firm's quality degradation due to cost uncertainty. On the other hand, the dashed line on the right shows the cross section of the quality allocation in the case of perfect competition (or equivalently, a benevolent quality allocation that collects zero profit from all customers who purchase a positive quality), and the arrow (B) shows the firm's quality degradation to capture consumer surplus.

where $b_1 \in [0, 1]$. (The subscript in the notation b_1 is meant to emphasize that this is the firm's belief in period 1.) In this setting, we define an *admissible policy* as a sequence of non-anticipating functions $\pi = \{\pi_t, t = 1, 2, \dots\}$, where π_t is measurable mapping from $[0, 1]^t$ to $\mathcal{A}$, and $\mathcal{A}$ is the set of product offerings $\mathcal{S} = \{(p(\theta), q(\theta)) : \theta \in \Theta\}$ that satisfy (2.3b), (2.3c), and (2.3d). Given a policy π , we define the sequences of posterior beliefs $\{b_t \in [0, 1], t = 1, 2, \dots\}$ and product offerings $\{\mathcal{S}_t \in \mathcal{A}, t = 1, 2, \dots\}$ with the following recursive relations. In every period t , the firm first chooses $\mathcal{S}_t = \pi_t(b_1, \dots, b_t)$. After that, the firm observes the arriving customer's quality choice Q_t , incurs the cost C_t , and then computes b_{t+1} via Bayes rule:

$$b_{t+1} = \frac{b_t L_1(C_t, Q_t)}{b_t L_1(C_t, Q_t) + (1 - b_t) L_0(C_t, Q_t)}, \quad (3.3)$$

where: $L_i(c, q) = \frac{1}{\sigma} \phi\left(\frac{c - c_i(q)}{\sigma}\right)$ for $i = 0, 1$, and $\phi(\cdot)$ is the standard Gaussian density. In (3.3), $L_i(C_t, Q_t)$ denotes the likelihood of incurring the cost C_t to produce a product with quality Q_t if the expected cost curve is $c_i(\cdot)$. For notational brevity, we re-express the preceding application of Bayes rule in the following more compact form:

$$b_{t+1} = \beta(b_t, C_t, Q_t), \quad (3.4)$$

where $\beta(b, c, q) = bL_1(c, q) / (bL_1(c, q) + (1 - b)L_0(c, q))$. As a consequence of the preceding recursive relations, every policy π induces a probability measure $\mathbb{P}^\pi\{\cdot\}$ on the sample space of posterior beliefs

$[0, 1]^T$. To construct this probability measure, let us define the following transition density for beliefs:

$$\begin{aligned} \psi_S(x, y) = & x \int_{\Theta} \int_{\{\gamma \in \mathbb{R} : \beta(x, \gamma, q(\theta)) = y\}} L_1(\gamma, q(\theta)) f(\theta) d\gamma d\theta \\ & + (1 - x) \int_{\Theta} \int_{\{\gamma \in \mathbb{R} : \beta(x, \gamma, q(\theta)) = y\}} L_0(\gamma, q(\theta)) f(\theta) d\gamma d\theta \end{aligned} \quad (3.5)$$

for all $x, y \in [0, 1]$ and $S \in \mathcal{A}$, where $q(\cdot)$ is the quality allocation function in the product offering S . Because the transition density $\psi_S(x, y)$ depends only on the quality allocation $q(\cdot)$ in S , we will also use the equivalent notation $\psi_{q(\cdot)}(x, y) = \psi_S(x, y)$ when necessary. Thus, given the value of $b_1 = b$, $\mathbb{P}^\pi\{\cdot\}$ is defined by the following relations: $\mathbb{P}^\pi\{b_1 = b\} = 1$, and

$$\mathbb{P}^\pi\{b_{t+1} \in dy \mid b_1, \dots, b_t\} = \psi_S(b_t, y) dy \quad (3.6)$$

for all $t = 1, 2, \dots, T$, and $y \in [0, 1]$.

The firm's objective is to choose an admissible policy to maximize its expected profit over T periods. For brevity, let $r(b, S)$ be the firm's expected single-period profit, expressed as a function of the firm's belief b and product offering $S = \{(p(\theta), q(\theta)) : \theta \in \Theta\} \in \mathcal{A}$; i.e.,

$$r(b, S) := \int_{\Theta} (p(\theta) - C(b, q(\theta))) f(\theta) d\theta \quad (3.7)$$

for $b \in [0, 1]$ and $S \in \mathcal{A}$. To maximize the expected over T periods, the firm solves

$$\max_{\pi \in \Pi} \mathbb{E}^\pi \left\{ \sum_{t=1}^T r(b_t, S_t) \mid b_1 = b \right\}, \quad (3.8)$$

where $b \in [0, 1]$, Π is the set of all admissible policies, and $\mathbb{E}^\pi\{\cdot\}$ is the expectation operator associated with $\mathbb{P}^\pi\{\cdot\}$. We hereafter refer to the problem in (3.8) as the *dynamic learning problem*. Our next result constructs the optimal value function for the dynamic learning problem and the associated optimal policy.

PROPOSITION 4. (optimal value function and optimal policy)

(i) **(existence)** *There exists a unique sequence of functions $\{V_n(\cdot), n = 0, 1, \dots, T\}$ that satisfy*

$$V_n(x) = \max_{S \in \mathcal{A}} \left\{ r(x, S) + \int_0^1 V_{n-1}(y) \psi_S(x, y) dy \right\} \quad \text{for } x \in [0, 1] \text{ and } n = 1, \dots, T, \quad (3.9a)$$

$$V_0(x) = 0 \quad \text{for } x \in [0, 1]. \quad (3.9b)$$

(ii) **(verification)** *$V_T(\cdot)$ satisfies $V_T(x) \geq \max_{\pi \in \Pi} \mathbb{E}^\pi \left\{ \sum_{t=1}^T r(b_t, S_t) \mid b_1 = x \right\}$ for $x \in [0, 1]$.*

(iii) **(attainment)** *Let π^* be an admissible policy satisfying $\pi_t^*(b_1, \dots, b_t) = \varphi_{T-t+1}(b_t)$, where*

$$\varphi_n(x) := \arg \max_{S \in \mathcal{A}} \left\{ r(x, S) + \int_0^1 V_{n-1}(y) \psi_S(x, y) dy \right\} \quad (3.10)$$

for $x \in [0, 1]$ and $n = 1, \dots, T$. Then,

$$\mathbb{E}^{\pi^*} \left\{ \sum_{t=1}^T r(b_t, S_t^*) \mid b_1 = x \right\} = V_T(x) \quad \text{for } x \in [0, 1], \quad (3.11)$$

where $\{S_t^, t = 1, 2, \dots\}$ is the sequence of product offerings under policy π^* .*

Proposition 4 shows that the optimal value function for the firm's dynamic learning problem is characterized by the recursive relation in (3.9a), which is usually referred to as the *Bellman equation*, and the boundary condition in (3.9b). More importantly, this result verifies that the optimal value is attained by a Markovian policy that uses only the firm's most recent belief to determine the product offering in a given period. We will now analyze the Bellman equation in detail to gain further insights into the structure of the optimal value function and the optimal policy. By Proposition 1, we know that the quality allocation function $q_t(\cdot)$ must be increasing and differentiable almost everywhere to induce IC and IR in period t . Moreover, as argued in the derivation of (2.4) in the preceding section, the pricing function in period t , namely $p_t(\cdot)$, satisfies the following relation to induce IC and IR:

$$p_t(\theta) = \theta q_t(\theta) - \int_{\theta_{\min}}^{\theta} q_t(\xi) d\xi \quad \text{for all } \theta \in \Theta. \quad (3.12)$$

Given the above structure for the pricing function $p_t(\cdot)$, we note that the product offering in period t , namely $\mathcal{S}_t$, is completely characterized by its quality allocation function $q_t(\cdot)$. Consequently, we can re-express the Bellman equation in (3.9a) as follows:

$$V_n(b) = \max_{q(\cdot) \in \mathcal{A}^*} \left\{ \int_{\theta_{\min}}^{\theta_{\max}} \left(\theta q(\theta) - \int_{\theta_{\min}}^{\theta} q(\xi) d\xi - C(x, q(\theta)) \right) f(\theta) d\theta + \int_0^1 V_{n-1}(y) \psi_{q(\cdot)}(b, y) dy \right\}$$

for all $b \in [0, 1]$ and $n = 1, \dots, T$, where $\mathcal{A}^*$ is the set of functions $q: \Theta \rightarrow \mathbb{R}_+$ that are increasing and differentiable almost everywhere. Recalling the definition of the transition density for beliefs in (3.5), we further obtain

$$\begin{aligned} V_n(b) = \max_{q(\cdot) \in \mathcal{A}^*} \left\{ \int_{\theta_{\min}}^{\theta_{\max}} \left(\theta q(\theta) - \int_{\theta_{\min}}^{\theta} q(\xi) d\xi - C(b, q(\theta)) \right) f(\theta) d\theta \right. \\ \left. + b \int_{\theta_{\min}}^{\theta_{\max}} \left(\int_{-\infty}^{\infty} V_{n-1}(\beta(b, \gamma, q(\theta))) L_1(\gamma, q(\theta)) d\gamma \right) f(\theta) d\theta \right. \\ \left. + (1-b) \int_{\theta_{\min}}^{\theta_{\max}} \left(\int_{-\infty}^{\infty} V_{n-1}(\beta(b, \gamma, q(\theta))) L_0(\gamma, q(\theta)) d\gamma \right) f(\theta) d\theta \right\} \quad (3.13) \end{aligned}$$

for all $b \in [0, 1]$ and $n = 1, \dots, T$. Noting that all of the three integrals on the right hand side of (3.13) have the common domain of $[\theta_{\min}, \theta_{\max}]$, we can simplify (3.13) by introducing the following *conditional value function*:

$$\mathcal{V}_n(b, q) := b \int_{-\infty}^{\infty} V_n(\beta(b, \gamma, q)) L_1(\gamma, q) d\gamma + (1-b) \int_{-\infty}^{\infty} V_n(\beta(b, \gamma, q)) L_0(\gamma, q) d\gamma \quad (3.14)$$

for all $b \in [0, 1]$, $q \in \mathcal{Q}$, and $n = 0, 1, \dots, T$. The value of $\mathcal{V}_n(b, q)$ corresponds to the firm's optimal profit-to-go in the remaining n periods, conditional on the firm selling a product of quality q . Using this definition of the conditional value function, we express (3.13) as

$$V_n(b) = \max_{q(\cdot) \in \mathcal{A}^*} \left\{ \int_{\theta_{\min}}^{\theta_{\max}} \left(\theta q(\theta) - \int_{\theta_{\min}}^{\theta} q(\xi) d\xi - C(b, q(\theta)) + \mathcal{V}_{n-1}(b, q(\theta)) \right) f(\theta) d\theta \right\} \quad (3.15)$$

for all $b \in [0, 1]$ and $n = 1, \dots, T$. As argued in the derivation of (2.6), we use integration by parts to deduce that $\int_{\theta_{\min}}^{\theta} q(\xi) d\xi = \bar{F}(\theta)q(\theta)/f(\theta)$, which further simplifies (3.15) to the following:

$$V_n(b) = \max_{q(\cdot) \in \mathcal{A}^*} \left\{ \int_{\theta_{\min}}^{\theta_{\max}} \left(\theta q(\theta) - \frac{\bar{F}(\theta)}{f(\theta)} q(\theta) - C(b, q(\theta)) + \mathcal{V}_{n-1}(b, q(\theta)) \right) f(\theta) d\theta \right\} \quad (3.16)$$

for all $b \in [0, 1]$ and $n = 1, \dots, T$. Let $q_n^*(b, \theta)$ be optimal quality allocation for the problem in (3.16); i.e.,

$$q_n^*(b, \cdot) := \arg \max_{q(\cdot) \in \mathcal{A}^*} \left\{ \int_{\theta_{\min}}^{\theta_{\max}} \left(\theta q(\theta) - \frac{\bar{F}(\theta)}{f(\theta)} q(\theta) - C(b, q(\theta)) + \mathcal{V}_{n-1}(b, q(\theta)) \right) f(\theta) d\theta \right\} \quad (3.17)$$

for all $b \in [0, 1]$ and $n = 1, \dots, T$.

As in Section 2, we will measure the inefficiency due to cost uncertainty by comparing the firm's performance with the performance of a clairvoyant who knows the underlying cost structure. To that end, let us first generalize the definitions of the firm's loss and quality degradation functions, which were originally expressed in (2.8) and (2.9) respectively. Define

$$\Delta_n(b) := bV_n(1) + (1-b)V_n(0) - V_n(b) \quad \text{for all } b \in [0, 1] \text{ and } n = 1, \dots, T, \quad (3.18)$$

$$\delta_n(b, \theta) := bq_n^*(1, \theta) + (1-b)q_n^*(0, \theta) - q_n^*(b, \theta) \quad \text{for all } b \in [0, 1], \theta \in \Theta, \text{ and } n = 1, \dots, T. \quad (3.19)$$

Note that the functions $V_1(b)$, $\Delta_1(b)$, $q_1^*(b, \theta)$, and $\delta_1(b, \theta)$ in the above definitions correspond to $V(b)$, $\Delta(b)$, $q^*(b, \theta)$, and $\delta(b, \theta)$, respectively, analyzed in Section 2. For the definition of quality degradation in (3.19), readers are reminded that positive (negative) values of $\delta_n(b, \theta)$ correspond to quality decreases (increases) relative to the expected quality allocation with perfect cost information.

In the next proposition, we show that the firm's n -period loss due to cost uncertainty, namely $\Delta_n(b)$, is increasing in n .

PROPOSITION 5. (monotonicity of firm's loss due to uncertainty) $\Delta_{n+1}(b) \geq \Delta_n(b)$ for all $b \in [0, 1]$ and $n = 1, \dots, T-1$.

Using Proposition 5, we obtain the following result:

PROPOSITION 6. (firm's loss due to uncertainty) $\Delta_n(b)$ is concave, and there exists a finite positive constant K such that $\Delta_n(b) \geq Kb(1-b) \geq 0$ for all $b \in [0, 1]$ and $n = 1, \dots, T$.

Propositions 5 and 6 extend Proposition 3 in Section 2 by establishing that the firm's loss in the n -period dynamic learning problem increases with n and attains larger values for intermediate values of the firm's belief b . Figure 3 shows this behavior of $\Delta_n(b)$.

The preceding results in this section show that the profit implications of static cost uncertainty, which was analyzed in Section 2, also hold in the case of dynamic learning. But, the opportunity of dynamic learning has a contrasting effect on the firm's optimal quality allocation. The quality distortion due to uncertainty, which was established in Theorem 1, is mitigated in the dynamic learning problem. To show the extent of this mitigation we present the next two results:




Figure 3 Firm's loss due to uncertainty in the dynamic learning problem. The loss $\Delta_n(b)$ is increasing in n , and is concave in b . The consumers' quality sensitivity parameters are uniformly distributed between 0 and 1. The cost parameters are $a_0 = 0.1$ and $a_1 = 1$.

THEOREM 2. (uniform quality improvement for learning) Let $b \in [0, 1]$ and $n \geq 2$. Then, $\delta_n(b, \theta) \leq \delta_1(b, \theta)$ for all $\theta \in \Theta$.

Theorem 2 states that, in the dynamic learning problem, the quality distortion due to uncertainty is mitigated for *all* types of customers, implying that the opportunity of learning about the cost structure improves quality uniformly for all potential customers. The main reason behind this result is that selling the product at a higher quality provides a clearer signal about the underlying cost curve; hence the firm can use quality improvements as a tool to accelerate its learning. Consequently, the firm chooses to increase the quality of its product offering to eliminate future losses due to cost uncertainty. A key feature of Theorem 2 is that the aforementioned quality improvement is offered to every customer, establishing that learning is valuable even when the firm's expected profit is conditioned upon the quality sensitivity of an arriving customer. Figure 4 displays the optimal quality allocation function $q_n^*(b, \theta)$ in a numerical example of the dynamic learning problem.

In our next result we prove that, in the dynamic learning problem, the quality degradation result in Theorem 1 is entirely reversed for a subset of customers.

THEOREM 3. (strong quality improvement for learning) Let $b \in (0, 1)$, $n \geq 2$, and assume that: (i) $\theta_{\min} \leq 1/f(\theta_{\min})$, and (ii) $\sigma \leq \sigma_0$ where σ_0 is a finite positive constant. Then, there exists an open set $E_{b,n} \subseteq \Theta$ of strictly positive measure such that $\delta_n(b, \theta) < 0$ for all $\theta \in E_{b,n}$.

While the uniform quality improvement result in Theorem 2 prescribes higher quality levels for all types of customers, such quality improvements cannot be arbitrarily large because a sufficiently large quality improvement for all customers would lead to a net loss almost surely; i.e., the firm's n -period profit would be negative under both of the possible cost curves $c_0(\cdot)$ and $c_1(\cdot)$. It is therefore important to understand if there exist customers who should be offered quality levels higher than the expected clairvoyant quality allocation. By Theorem 1 we know that there exist no such customers in the case




Figure 4 Optimal quality allocation in the dynamic learning problem. The shaded surface shows the optimal quality allocation in the beginning of $n = 20$ sales opportunities, where the problem parameters are the same as in the example in Figure 2. The arrow (A') shows the firm's quality improvement relative to the case of static cost uncertainty displayed in Figure 2.

of static cost uncertainty, and Theorem 3 provides conditions for the existence of such customers in the case of dynamic learning. We will now provide an interpretation for these conditions and explain why we need them for the strong quality improvement result in Theorem 3. First, condition (i) states that there exists a customer with a sufficiently small quality sensitivity parameter at the bottom end of $\Theta = [\theta_{\min}, \theta_{\max}]$, which means that the firm does not necessarily serve all customers with a strictly positive quality level. For example, a simple case in which condition (i) holds is where $\theta_{\min} = 0$. In the absence of condition (i), e.g., if $\theta_{\min}$ is positive and arbitrarily close to $\theta_{\max}$, even a short-sighted quality allocation that makes no attempt to learn could offer very high quality levels to all customers, which would eliminate the firm's incentive to further increase quality for learning. Second, condition (ii) states the standard deviation of the customers' cost shocks, namely σ , is bounded above by a constant σ_0 . Note that, if σ is arbitrarily large, then the information that can be gleaned in a sales period becomes very small. Again, this would essentially eliminate the firm's incentive to improve quality for learning. This is why we use conditions (i) and (ii) to prove the existence of customers for which the firm substantially increases quality and entirely reverses the quality distortion due to uncertainty. Figure 5 displays an example for the subset of customers who would experience the strong quality improvement in the dynamic learning problem. Note that, among the customers who would make a purchase, the less quality-sensitive ones experience stronger quality improvement. This feature of the firm's optimal policy can be explained by viewing quality improvement as a tool for learning. Learning from less quality-sensitive customers is relatively more profitable because the marginal cost of production is smaller for these customers.




Figure 5 Strong quality improvement for learning. Panels (a) and (b) display cross sections of quality allocation functions at $b = 0.1$ in the cases of (a) static cost uncertainty, and (b) dynamic learning with $n = 20$ sales opportunities, respectively. The dashed curves show the expected clairvoyant quality allocation $q^\circ(b, \theta) = bq_n^*(1, \theta) + (1 - b)q_n^*(0, \theta) = bq_1^*(1, \theta) + (1 - b)q_1^*(0, \theta)$, and the bold curves show the firm's optimal quality allocation. As shown in panel (b), in the case of dynamic learning, there exists a subset of customer types, namely $E_{b,n}$, for which the firm's optimal quality allocation exceeds the expected clairvoyant quality allocation. Panel (a) shows there is no such subset of customers in the case of static uncertainty. (The problem parameters are the same as in the example in Figure 2.)

4. Implications of Subsidizing Quality

This section investigates how regulations that subsidize quality influence the interplay between static cost uncertainty and dynamic learning. In our analysis, we will consider two families of commonly used regulations that subsidize quality in the health insurance industry: (a) providing a free low-quality outside option to all potential customers, and (b) subsidizing the prices that customers pay for the product. The implementation of (a) can vary significantly by country. According to 42 U.S. Code §1395dd, emergency departments of the hospitals in the United States are legally obliged to provide examination and (if necessary) stabilizing treatment for any patient visiting the emergency department, regardless of the patient's ability to pay. On the other hand, the National Health Service (NHS) system in England provides free healthcare to all citizens and legal immigrants at the point of use. The extent of the free healthcare naturally influences the design of products in private insurance sector, which motivates the extension of our preceding analysis to the case of a free outside option with positive quality. With regard to (b), a prominent example is the Affordable Care Act (Pub. L. 111-148, 124 Stat. 119) in the United States, which provides subsidies for the insurance premiums of individuals whose household income is near the poverty line, with gradually increasing subsidies for those closer to the poverty line. Such premium subsidies have a direct impact on the pricing decisions of insurers, bringing forth the question of whether our findings on cost uncertainty and learning in the preceding sections extend to the case of subsidized prices. Motivated by these applications in health insurance, we will now study the combined effect of cost uncertainty and quality-subsidizing regulations in our framework. For

both (a) and (b), we will begin with verifying the implementability of the regulation, which includes deriving conditions for incentive compatibility and individual rationality in product design, as well as investigating whether the regulation in question can be implemented as a welfare-increasing regulation. That is, for both (a) and (b), we will show the existence of quality-subsidizing regulations that can increase social welfare, relative to the case of no regulation. After that, we will generalize our results on cost uncertainty and learning in Sections 2 and 3 to the cases of quality-subsidizing regulations.

4.1. Free Low-quality Outside Option

Suppose that every potential customer has the outside option of obtaining a product with quality $q_0 > 0$ at zero price. Let us first consider how the presence of this outside option will affect the firm's pricing and quality allocation decisions. As in our analysis of the original problem in Section 2, the firm's optimal product offering $\mathcal{S} = \{(p(\theta), q(\theta)) : \theta \in \Theta\}$ prescribes a threshold $\theta_c \in [\theta_{\min}, \theta_{\max})$ such that the customers whose quality sensitivities are greater than θ_c will be served with a positive quality allocation. But this time, the outside option is more attractive; hence θ_c could be higher relative to the original problem. (Note that, in both cases, θ_c could be equal to $\theta_{\min}$ if $\theta_{\min}$ is sufficiently high.) To emphasize the critical threshold θ_c in the firm's product offering, we can express the firm's problem with the free outside option of quality q_0 as follows:

$$V^f(b, q_0) = \max_{p(\cdot), q(\cdot)} \left\{ \int_{\theta_c}^{\theta_{\max}} (p(\theta) - C(b, q(\theta))) f(\theta) d\theta \right\} \quad (4.1a)$$

$$\text{s.t. } \theta q(\theta) - p(\theta) \geq \theta q(\tilde{\theta}) - p(\tilde{\theta}) \quad \text{for all } \theta, \tilde{\theta} \in [\theta_c, \theta_{\max}], \quad (4.1b)$$

$$\theta q(\theta) - p(\theta) \geq \theta q_0 \quad \text{for all } \theta \in [\theta_c, \theta_{\max}], \quad (4.1c)$$

$$q(\theta) \geq 0 \quad \text{for all } \theta \in [\theta_c, \theta_{\max}]. \quad (4.1d)$$

We note that our original problem in (2.3) is a special case of (4.1) with $q_0 = 0$, which implies that $V^f(b, 0) = V(b)$. The distinguishing feature of the preceding problem is the net utility of the free outside option, θq_0 , which appears on the right hand side of the IR constraints (4.1c). To normalize the customers' utility curves, let us define the firm's excess quality allocation function:

$$x(\theta) := q(\theta) - q_0 \quad \text{for all } \theta \in [\theta_c, \theta_{\max}]. \quad (4.2)$$

Using the definition of $x(\cdot)$, we obtain the following counterpart of Proposition 1.

PROPOSITION 7. (incentive compatibility and individual rationality)

- (i) *The IC constraints (4.1b) hold if and only if $x(\cdot)$ is increasing and differentiable almost everywhere, and $p'(\theta) = \theta x'(\theta)$ for $\theta \in \Theta$, except on a set of measure zero.*
- (ii) *Assume that the IC constraints (4.1b) hold, and that $\theta_c x(\theta_c) - p(\theta_c) = 0$. Then the IR constraints (4.1c) also hold.*

As argued in (2.4), Proposition 7 implies that

$$p(\theta) = \theta x(\theta) - \int_{\theta_c}^{\theta} x(\xi) d\xi \quad \text{for all } \theta \in [\theta_c, \theta_{\max}]. \quad (4.3)$$

Therefore, by the arguments used to derive (2.6), the firm's problem becomes

$$V^f(b, q_0) = \max_{x(\cdot) \in \mathcal{A}^*} \left\{ \int_{\theta_c}^{\theta_{\max}} \left(\theta x(\theta) - \frac{\bar{F}(\theta)}{f(\theta)} x(\theta) - C(b, q_0 + x(\theta)) \right) f(\theta) d\theta \right\} \quad (4.4)$$

for all $b \in [0, 1]$ and $q_0 \geq 0$, where $\mathcal{A}^*$ is the set of functions $x : [\theta_c, \theta_{\max}] \rightarrow \mathbb{R}_+$ that are increasing and differentiable almost everywhere. Let $q^f(b, \theta, q_0)$ be the firm's optimal quality allocation for the problem in (4.4). That is, $q^f(b, \theta, q_0) = q_0 + x^f(b, \theta, q_0)$, where

$$x^f(b, \cdot, q_0) = \arg \max_{x(\cdot) \in \mathcal{A}^*} \left\{ \int_{\theta_c}^{\theta_{\max}} \left(\theta x(\theta) - \frac{\bar{F}(\theta)}{f(\theta)} x(\theta) - C(b, q_0 + x(\theta)) \right) f(\theta) d\theta \right\}. \quad (4.5)$$

Figure 6 shows a plot of $q^f(b, \theta, q_0)$ in a numerical example. Note that the critical threshold θ_c , over which the firm's quality allocation is preferred to the outside option, varies over different values of the firm's belief b . In particular, as the firm's belief on the higher cost curve $c_1(\cdot)$ increases, the firm chooses to serve a smaller set of potential customers.




Figure 6 Optimal quality allocation with free low-quality outside option. The shaded surface shows the firm's optimal quality allocation in presence of a free outside option of quality $q_0 = 0.15$. (The rest of the problem parameters are the same as in the example in Figure 2.)

Our next goal is to verify that a free outside option can be implemented as a welfare-increasing regulation. To that end, let $\bar{q}(b, \theta, q_0)$ be the quality level that a customer with sensitivity parameter θ would receive when the firm's belief is b and the quality of the free outside option is q_0 ; i.e.,

$$\bar{q}(b, \theta, q_0) = \begin{cases} q_0 & \text{if } \theta \leq \theta_c^f(b, q_0) \\ q^f(b, \theta, q_0) & \text{otherwise,} \end{cases} \quad (4.6)$$

where $\theta_c^f(b, q_0) = \inf \{ \theta \in [\theta_{\min}, \theta_{\max}] : x^f(b, \theta, q_0) \geq 0 \}$ is the critical threshold over which the firm offers its customers a quality that is preferred to the outside option. Then, the social welfare with the free outside option of quality q_0 is

$$W^f(b, q_0) = \int_{\theta_{\min}}^{\theta_{\max}} \left(\theta \bar{q}(b, \theta, q_0) - C(b, \bar{q}(b, \theta, q_0)) \right) f(\theta) d\theta. \quad (4.7)$$

In the original problem in (2.3), we have $q_0 = 0$, and hence the social welfare in that case is

$$W(b) = W^f(b, 0) = \int_{\theta_{\min}}^{\theta_{\max}} \left(\theta q^*(b, \theta) - C(b, q^*(b, \theta)) \right) f(\theta) d\theta. \quad (4.8)$$

Our next result shows that the presence of a free outside option can increase social welfare.

PROPOSITION 8. (welfare impact of free low-quality outside option) *There exists a finite and positive constant κ such that, if $0 < q_0 \leq \kappa$, then $W^f(b, q_0) > W(b)$ for all $b \in [0, 1]$.*

The preceding proposition establishes that a free outside option can increase social welfare as long as the quality of the outside option is positive and sufficiently small. As seen on Figure 7, the social welfare function $W^f(b, q_0)$ is increasing in q_0 for smaller values of q_0 and becomes decreasing in q_0 when the value of q_0 becomes excessively large. The main reason for this behavior is that, for large values of q_0 , the outside option simply provides a very high quality whose marginal cost exceeds its marginal benefit to society, resulting in a net decrease of social welfare.




Figure 7 Social welfare with free low-quality outside option. The social welfare with a free outside option of quality q_0 , namely $W^f(b, q_0)$, is increasing in q_0 for sufficiently small values of q_0 . The firm's prior belief is $b = 0.5$, and the problem parameters are the same as in the example in Figure 2.

Having shown that a free low-quality outside option can be implemented as a welfare-increasing regulation, our next goal is to study its impact on the quality distortion due to uncertainty. Given that the customers now receive qualities according to (4.6), we modify the definition of our quality degradation function in (2.9) as follows:

$$\delta^f(b, \theta, q_0) := b\bar{q}(1, \theta, q_0) + (1 - b)\bar{q}(0, \theta, q_0) - \bar{q}(b, \theta, q_0) \quad \text{for all } b \in [0, 1], \theta \in \Theta, \text{ and } q_0 > 0. \quad (4.9)$$

Note that $\delta^f(b, \theta, q_0)$ accounts for the fact that customers that are not served by the firm receive the quality of q_0 . Our next result characterizes how a free outside option with positive quality moderates the quality degradation due to uncertainty.

THEOREM 4. (quality degradation due to uncertainty) *Let $q_0 > 0$.*

- (i) *There exists an increasing function $g : \Theta \rightarrow \mathbb{R}_+$ such that $g(\theta_{\max}) > 0$ and $\delta^f(b, \theta, q_0) \geq g(\theta)b(1 - b) \geq 0$ for $b \in [0, 1]$ and $\theta \in [\theta_c^f(1, q_0), \theta_{\max}]$.*
- (ii) *There exists an open set $E \subseteq [0, 1] \times [\theta_c^f(0, q_0), \theta_c^f(1, q_0)]$ such that $\delta^f(b, \theta, q_0) < 0$ for $(b, \theta) \in E$.*

Remark The function $g(\cdot)$ that appears in part (i) of the above theorem is the same function $g(\cdot)$ in Theorem 1.

Theorem 4 describes in two parts the combined effect of the cost uncertainty and the free outside option with positive quality. Part (i) establishes that, in the presence of a free outside option with positive quality, the quality distortion result in Theorem 1 is preserved for customers that are highly sensitive to quality. On the other hand, part (ii) shows that we can have mixed results for lower levels of quality sensitivity; in particular, the firm's optimal quality allocation can exceed the expected clairvoyant quality allocation.




Figure 8 Quality degradation with free low-quality option. The bold curves in panels (a) and (b) show cross sections of the quality allocation function $q^f(b, \theta, q_0)$ for $\theta_\ell = 0.745$ and $\theta_h = 0.755$ respectively, where: $q_0 = 0.25$, the consumers' quality sensitivity parameters are uniformly distributed between 0 and 1, and the cost parameters are $a_0 = 0.1$ and $a_1 = 0.5$. The dashed curves show the expected clairvoyant quality allocation in each case. Panel (a) shows that $q^f(b, \theta, q_0)$ can exceed the expected clairvoyant quality allocation for small values of θ , whereas panel (b) shows that the quality degradation due to uncertainty will almost surely affect top consumers.

Figure 8 displays a numerical example regarding the two parts of Theorem 4. As shown in Figure 8(b), the customers with sufficiently high θ never prefer the outside option, and the convexity of the firm's optimal quality allocation leads to a distortion similar to the one we established in Theorem 1. However, the customers with relatively lower values of θ might have to purchase the low-quality outside option if the underlying cost curve happens to be the larger one, which means that the clairvoyant quality allocation for these customers would be q_0 if $c(\cdot) = c_1(\cdot)$. Consequently, as shown in Figure 8(a), the firm's uncertainty about the cost curve might result in a better quality allocation than the expected clairvoyant quality allocation for customers with relatively lower quality sensitivity.

4.2. Price Subsidy

Suppose that a regulatory authority subsidizes a fraction of the price that a customer pays for the product. For a customer with quality sensitivity parameter θ , let $\tau(\theta) \in [0, 1]$ denote the unsubsidized fraction of the price. That is, a customer with quality sensitivity θ pays only $\tau(\theta)p$ to purchase the product at price p , and the remainder of the price, namely $(1 - \tau(\theta))p$, is subsidized. In this case, the net utility of purchasing a product with quality q at price p becomes

$$U^s(\theta, q, p) = \theta q - \tau(\theta)p, \quad (4.10)$$

where $\theta \in \Theta$ is the quality sensitivity of the purchasing customer. Consequently, the firm's problem with subsidized prices can be expressed as follows:

$$V^s(b, \tau) = \max_{p(\cdot), q(\cdot)} \left\{ \int_{\theta_{\min}}^{\theta_{\max}} (p(\theta) - C(b, q(\theta))) f(\theta) d\theta \right\} \quad (4.11a)$$

$$\text{s.t. } \theta q(\theta) - \tau(\theta)p(\theta) \geq \theta q(\tilde{\theta}) - \tau(\theta)p(\tilde{\theta}) \quad \text{for all } \theta, \tilde{\theta} \in \Theta, \quad (4.11b)$$

$$\theta q(\theta) - \tau(\theta)p(\theta) \geq 0 \quad \text{for all } \theta \in \Theta, \quad (4.11c)$$

$$q(\theta) \geq 0 \quad \text{for all } \theta \in \Theta. \quad (4.11d)$$

Letting $\mathbf{1}(\cdot)$ denote the constant function that maps Θ to $\{1\}$, we note that the original problem formulation in (2.3) is a special case of (4.11) with $\tau = \mathbf{1}$. Thus, $V^s(b, \mathbf{1}) = V(b)$.

Our first goal is to find necessary and sufficient conditions for IC and IR when prices are subsidized. For that purpose, let us define a feasible set for the function τ . Let $\mathcal{T}$ be the set of functions $\tau : \Theta \rightarrow (0, 1]$ that are increasing, concave, and differentiable almost everywhere. In Figure 9, we display a piecewise linear subsidization scheme $\tau \in \mathcal{T}$.




Figure 9 Price subsidization example. Customers with $\theta \leq 0.4$ pay only a fraction of the price, whereas the remaining customers pay the full price.

In the next result, we extend Proposition 1 to the case of subsidized prices.

PROPOSITION 9. (incentive compatibility and individual rationality) *Let $\tau \in \mathcal{T}$. Then the following statements hold:*

- (i) *The IC constraints (4.11b) hold if and only if $q(\cdot)$ is increasing and differentiable almost everywhere, and $\tau(\theta)p'(\theta) = \theta q'(\theta)$ for $\theta \in \Theta$, except on a set of measure zero.*

- (ii) Assume that the IC constraints (4.11b) hold, and that $\theta_{\min} q(\theta_{\min}) - \tau(\theta_{\min})p(\theta_{\min}) = 0$. Then the IR constraints (4.11c) also hold.

Using Proposition 9 and integrating by parts, we deduce that the pricing function $p(\cdot)$ must satisfy the following to induce IC and IR:

$$p(\theta) = \nu(\theta)q(\theta) - \int_{\theta_{\min}}^{\theta} \nu'(\xi)q(\xi)d\xi \quad \text{for all } \theta \in \Theta, \quad (4.12)$$

where $\nu(\theta) := \theta/\tau(\theta)$. Here, $\nu(\theta)$ can be interpreted as the effective quality sensitivity of a customer under the price subsidization scheme τ . Recalling the arguments we used for deriving (2.6), we can re-express the firm's problem in (4.11) as follows:

$$V^s(b, \tau) = \max_{q(\cdot) \in \mathcal{A}^*} \left\{ \int_{\theta_{\min}}^{\theta_{\max}} \left(\nu(\theta)q(\theta) - \frac{\bar{F}(\theta)}{f(\theta)} \nu'(\theta)q(\theta) - C(b, q(\theta)) \right) f(\theta) d\theta \right\} \quad (4.13)$$

for all $b \in [0, 1]$ and $\tau \in \mathcal{T}$, where $\mathcal{A}^*$ is the set of functions $q : \Theta \rightarrow \mathbb{R}_+$ that are increasing and differentiable almost everywhere. Denote by $q^s(b, \theta, \tau)$ the firm's optimal quality allocation for the problem in (4.13); i.e.,

$$q^s(b, \cdot, \tau) = \arg \max_{q(\cdot) \in \mathcal{A}^*} \left\{ \int_{\theta_{\min}}^{\theta_{\max}} \left(\nu(\theta)q(\theta) - \frac{\bar{F}(\theta)}{f(\theta)} \nu'(\theta)q(\theta) - C(b, q(\theta)) \right) f(\theta) d\theta \right\}. \quad (4.14)$$

As in the preceding subsection, we will first verify that a price subsidy can be implemented as a welfare-increasing regulation. Under the quality allocation $q^s(b, \theta, \tau)$, the social welfare is

$$W^s(b, \tau) = \int_{\theta_{\min}}^{\theta_{\max}} \left(\theta q^s(b, \theta, \tau) - C(b, q^s(b, \theta, \tau)) \right) f(\theta) d\theta. \quad (4.15)$$

The following proposition shows that there exists a function τ under which $W^s(b, \tau)$ achieves its highest possible value over all feasible quality allocations.

PROPOSITION 10. (welfare impact of price subsidy) Let $\tau^*(\theta) = \theta \bar{F}(\theta) / \left(\int_{\theta}^{\theta_{\max}} \xi f(\xi) d\xi \right)$ for all $\theta \in \Theta$. Then,

$$W^s(b, \tau^*) = \max_{q(\cdot) \in \mathcal{A}^*} \left\{ \int_{\theta_{\min}}^{\theta_{\max}} \left(\theta q(\theta) - C(b, q(\theta)) \right) f(\theta) d\theta \right\} \geq W(b) \quad \text{for all } b \in [0, 1]. \quad (4.16)$$

Proposition 10 explicitly characterizes a price subsidization scheme τ^* that induces the firm to choose a welfare-maximizing quality allocation, which is usually referred to as the *first best solution*. (Figure 10 displays τ^* in a numerical example.) Naturally, the existence of τ^* implies that a price subsidy can be implemented as a welfare-increasing regulation. To derive the subsidization formula in Proposition 10, we solve an ordinary differential equation that ensures that the firm's marginal revenue for serving a customer with quality sensitivity θ is equal to the marginal social benefit of serving this customer. Once this relationship is established, the firm simply chooses its product offering to maximize social welfare.




Figure 10 Welfare-maximizing price subsidization. The function τ^* increases concavely on $\Theta = [0, 1]$. The consumers' quality sensitivity parameters are uniformly distributed over Θ .

Finally, we focus on whether the firm exercises any quality degradation in its product offering in the presence of both price subsidies and cost uncertainty. To that end, define

$$\delta^s(b, \theta, \tau) := bq^s(1, \theta, \tau) + (1-b)q^s(0, \theta, \tau) - q^s(b, \theta, \tau) \quad \text{for all } b \in [0, 1], \theta \in \Theta, \text{ and } \tau \in \mathcal{T}. \quad (4.17)$$

The last result of this subsection shows that the quality distortion due to uncertainty persists under price subsidies.

THEOREM 5. (quality degradation due to uncertainty) *Let $\tau \in \mathcal{T}$. Then there exists an increasing function $G: \Theta \rightarrow \mathbb{R}_+$ such that $G(\theta_{\max}) > 0$ and $\delta^s(b, \theta, \tau) \geq G(\theta)b(1-b) \geq 0$ for all $b \in [0, 1]$ and $\theta \in \Theta$.*

The preceding theorem states that, even if welfare-maximizing subsidization schemes such as τ^* are implemented, a firm facing uncertainty about its cost curve would still degrade the quality of its products, relative to the expected clairvoyant quality allocation. Figure 11 demonstrates this behavior of the firm in a numerical example.




Figure 11 Optimal quality allocation with price subsidy. The shaded surface shows the firm's optimal quality allocation $q^s(b, \theta, \tau^*)$ under the welfare-maximizing price subsidization scheme τ^* . While the monopoly distortion effect disappears under τ^* , the uncertainty distortion effect is still present in $q^s(b, \theta, \tau^*)$. (The problem parameters are the same as in the example in Figure 2.)

4.3. Dynamic Learning in Presence of Quality Subsidization

This subsection extends our findings on dynamic learning in Section 3 to the case of quality-subsidizing regulations. To generalize the basic problem formulation, we first consider the case of a free outside option with positive quality, and let $r^f(b, \mathcal{S}, q_0)$ denote the firm's single-period profit, expressed as a function of the firm's belief b , product offering $\mathcal{S} = \{(p(\theta), q(\theta)) : \theta \in \Theta\}$, and the quality of the free outside option q_0 . Similarly, in the case of price subsidies, let $r^s(b, \mathcal{S}, \tau)$ be the firm's single-period profit as a function of the firm's belief b , product offering $\mathcal{S}$, and the subsidization scheme τ . Under both subsidization schemes, the firm can make observations on its cost curve as explained in (3.1), and we construct admissible policies in a non-anticipating fashion such that the product offering in a given period $t = 1, \dots, T$ depends only on past observations. As a result, the firm solves the following multi-period problems under quality-subsidizing regulations:

$$V_n^f(b, q_0) := \max_{\pi \in \Pi} \mathbb{E}^\pi \left\{ \sum_{t=1}^n r^f(b_t, \mathcal{S}_t, q_0) \mid b_1 = b \right\} \quad \text{for } b \in [0, 1], q_0 > 0, \text{ and } n = 1, \dots, T, \quad (4.18)$$

$$V_n^s(b, \tau) := \max_{\pi \in \Pi} \mathbb{E}^\pi \left\{ \sum_{t=1}^n r^s(b_t, \mathcal{S}_t, \tau) \mid b_1 = b \right\} \quad \text{for } b \in [0, 1], \tau \in \mathcal{T}, \text{ and } n = 1, \dots, T, \quad (4.19)$$

where Π is the set of all admissible policies. The problem in (4.18) corresponds to the case of a free outside option with positive quality q_0 , while the problem in (4.19) corresponds to the case of price subsidization scheme τ . Let $q_n^f(b, \theta, q_0)$ and $q_n^s(b, \theta, \tau)$ be the firm's optimal quality allocation functions in the first sales period of the problems in (4.18) and (4.19), respectively. Then, we can extend the definition of the quality degradation functions in (3.19) as follows:

$$\delta_n^f(b, \theta, q_0) := bq_n^f(1, \theta, q_0) + (1-b)q_n^f(0, \theta, q_0) - q_n^f(b, \theta, q_0) \quad (4.20)$$

$$\delta_n^s(b, \theta, \tau) := bq_n^s(1, \theta, \tau) + (1-b)q_n^s(0, \theta, \tau) - q_n^s(b, \theta, \tau) \quad (4.21)$$

for all $b \in [0, 1]$, $q_0 > 0$, $\tau \in \mathcal{T}$, and $n = 1, \dots, T$. The following theorem extends Theorem 2 by showing that, under quality subsidization, the firm uniformly improves the quality in its product offering to accelerate its learning.

THEOREM 6. (uniform quality improvement for learning) *Let $b \in [0, 1]$, $n \geq 2$, $q_0 > 0$, and $\tau \in \mathcal{T}$. Then, $\delta_n^f(b, \theta, q_0) \leq \delta_1^f(b, \theta, q_0)$ and $\delta_n^s(b, \theta, \tau) \leq \delta_1^s(b, \theta, \tau)$ for all $\theta \in \Theta$.*

5. Extensions and Concluding Remarks

In this last section, we discuss possible extensions of our analysis.

Batch arrivals of customers. In many practical applications of dynamic pricing, firms might not necessarily be able to adjust their product offerings after every sales opportunity. To accommodate such cases in our analysis, let us consider an extension of our model where the firm makes multiple observations on the cost curve before the product offering can be changed. Suppose that, in every

period t , N distinct customers arrive. The firm observes their quality choices $\{Q_{jt}, j = 1, \dots, N\}$, and incurs the production costs $\{C_{jt}, j = 1, \dots, N\}$ to serve these customers, where $C_{jt} = c(Q_{jt}) + \epsilon_{jt}$, and $\epsilon_{jt} \stackrel{\text{iid}}{\sim} \mathcal{N}(0, \sigma^2)$. To extend our analysis to this case, we replace the belief updating equation in (3.3) with the following: The firm's posterior belief in period $t + 1$ is given by

$$b_{t+1} = \frac{b_t \prod_{j=1}^N L_1(C_{jt}, Q_{jt})}{b_t \prod_{j=1}^N L_1(C_{jt}, Q_{jt}) + (1 - b_t) \prod_{j=1}^N L_0(C_{jt}, Q_{jt})}, \quad (5.1)$$

where: $L_i(c, q) = \frac{1}{\sigma} \phi\left(\frac{c - c_i(q)}{\sigma}\right)$ for $i = 0, 1$, and $\phi(\cdot)$ is the standard Gaussian density. Redefining the transition density for beliefs in (3.5) based on (5.1), the firm's optimal policy in this modified dynamic learning problem can be constructed as in Section 3. We note that this modification of the dynamic learning problem can be viewed as a more constrained version of our original problem in Sections 2 and 3, where the added constraint is that the firm is allowed to use its most recent beliefs only in periods $t \in \{1, N + 1, 2N + 1, \dots\}$. Consequently, the optimal policy in this extension is a feasible (but not necessarily optimal) policy in our original problem. Therefore, the firm's loss due to uncertainty would increase in this extension, relative to our setting. More importantly, to reflect the value of learning in its quality allocation, the firm would still need to improve quality, but because every sales period provides a larger amount of information, the firm would have to make more drastic innovations in its product offering after every batch of observations.

Generalizing the belief distribution on cost uncertainty. In this paper, we have modeled the firm's uncertainty about costs using two cost hypotheses. That is, the firm has a Bernoulli belief distribution on two possible cost curves. To extend this model to more general prior beliefs, suppose that the firm has a belief distribution on $M + 1$ possible cost curves. That is, the firm initially believes that $c(\cdot) = c_i(\cdot)$ with probability $b_{i,1}$ for $i = 0, 1, \dots, M$, where $\sum_{i=0}^M b_{i,1} = 1$, $c_i(q) = a_i q^2$ for all $q \in \mathcal{Q}$, and $0 < a_0 < \dots < a_M < \infty$. Let $b_t = (b_{i,t}, i = 0, 1, \dots, M)$ be the firm's belief vector in period t . Then, the firm's belief updating equation to compute b_{t+1} becomes

$$b_{i,t+1} = \frac{b_{i,t} L_i(C_t, Q_t)}{\sum_{i=0}^M b_{i,t} L_i(C_t, Q_t)} \quad \text{for } i = 0, 1, \dots, M, \quad (5.2)$$

where: $L_i(c, q) = \frac{1}{\sigma} \phi\left(\frac{c - c_i(q)}{\sigma}\right)$ for $i = 0, 1, \dots, M$, and $\phi(\cdot)$ is the standard Gaussian density. Using (5.2), we can extend the definition of the transition density in (3.5) to the case of belief vectors, and then construct the corresponding optimal policy as in Section 3. A key issue in this construction is that the generalized dynamic programming problem has a higher dimensional state space, which might involve computational challenges for large values of M . To establish the value of learning in this case, one needs to prove the convexity of the value function associated with generalized dynamic learning problem (i.e., the firm should find the expected clairvoyant profit more valuable than its optimal profit

as function of the belief vector). For that purpose, we note that our proof of the convexity of the value function uses the linearity of the expected profit with respect to b_t , which remains valid in this generalization. Because learning is still valuable in the generalized version of our dynamic learning problem, one would expect that the firm needs to improve quality for learning purposes.

When the firm has a belief distribution on a multitude of cost curves, it is possible that the extreme cost scenarios in the firm's belief distribution could be drastically different from each other, thereby exacerbating the uncertainty in the firm's business environment. A key issue in this context is to provide the firm the right incentives to mitigate some of the quality distortion due to uncertainty. In health insurance practice, this can be achieved by covering the firm's additional cost under a high-cost scenario while charging a fee under a low-cost scenario (see, e.g., the risk corridors of the ACA, Goodell 2014). In our setting, this effectively eliminates some of the possible cost scenarios, resulting in a belief distribution concentrated on a smaller number of cost hypotheses. If a regulatory authority implements a stabilization scheme to eliminate the extreme cost curves, then the differences between the clairvoyant quality allocations for the worst and best cost scenarios would be reduced. As a result, under such an stabilization scheme, one would expect to see a smaller quality distortion due to uncertainty.

Time-varying customer population. Our model assumes that the distribution of customers' quality sensitivity parameters θ has a fairly general density $f(\theta)$, but this density remains the same over time. In the health insurance context, which serves as our motivation, a key issue is the aging population of potential customers. In particular, as the customer population gets older, one might expect that their medical needs would increase, thereby raising the quality-sensitivity of customers. To gain insights into this case, we can consider a time-varying modification of our model: Let the distribution of θ in period t be a uniform distribution between 0 and $\theta_{\max}(t)$, and assume that $\{\theta_{\max}(t), t = 1, 2, \dots\}$ evolves according to the following equation:

$$\theta_{\max}(t+1) = A(\theta_{\max}(t)) \quad \text{for } t = 1, 2, \dots, \quad (5.3)$$

where $A: \mathbb{R} \rightarrow \mathbb{R}$ is an increasing function, and $\theta_{\max}(1)$ is a positive constant. In this case, the firm's dynamic learning problem in (3.8) would be replaced by

$$\max_{\pi \in \Pi} \mathbb{E}^{\pi} \left\{ \sum_{t=1}^T r_t(b_t, \mathcal{S}_t) \mid b_1 = b \right\}, \quad (5.4)$$

where $r_t(b, \mathcal{S})$ is the firm's single-period profit function for period t . The distinguishing feature of the problem in (5.4) is that the single-period profit function is now dependent on period t . Because the firm's customer population is spread to larger intervals in later periods, the firm would be able to screen its infinitesimal customers in a finer way as t increases, meaning that $r_t^*(b) \leq r_{t+1}^*(b)$ where

$r_t^*(b) = \max_{S \in \mathcal{A}} \{r_t(b, S)\}$. As a result, later periods would provide more lucrative sales opportunities to the firm; hence, one would expect to see a more pronounced quality improvement for learning.

Cost-heterogeneity of customers. In our model, we considered a cost curve that is characterized by the product quality; hence the firm's production costs depend indirectly on the quality-sensitivities of customers. If the customers' quality-sensitivities are known to have a direct and heterogeneous effect on production costs, we can extend our model to accommodate this case as follows. Let us replace $c_i(q)$, $i = 0, 1$, with the two-dimensional cost curves $c_i(\theta, q) = a_i(\theta)q^2$, $i = 0, 1$, where $a_0(\cdot)$ and $a_1(\cdot)$ are functions from Θ to $\mathbb{R}_{++}$ such that $a_0(\theta) < a_1(\theta)$ for all $\theta \in \Theta$. In this case, the firm's problem with static cost uncertainty, which was originally expressed in (2.7), becomes

$$\max_{q(\cdot) \in \mathcal{A}^*} \left\{ \int_{\theta_{\min}}^{\theta_{\max}} \left(\theta q(\theta) - \frac{\bar{F}(\theta)}{f(\theta)} q(\theta) - C(b, \theta, q(\theta)) \right) f(\theta) d\theta \right\}, \quad (5.5)$$

where $C(b, \theta, q) = bc_1(\theta, q) + (1-b)c_0(\theta, q)$, and $\mathcal{A}^*$ is the set of functions $q: \Theta \rightarrow \mathbb{R}_+$ that are increasing and differentiable almost everywhere. Letting $q^*(b, \theta)$ be the optimal quality allocation for the problem in (5.5), we note that the proof of Theorem 1 is valid in this extension, with the exception that the function $g(\cdot)$ in the quality distortion lower bound in Theorem 1 is replaced by

$$g(\theta) = \frac{(a_1(\theta) - a_0(\theta))^2 (\theta - \bar{F}(\theta)/f(\theta))_+}{2a_0(\theta)a_1^2(\theta)}. \quad (5.6)$$

Because $g(\theta)$ is positive for all $\theta \in \Theta$, one would still observe a quality degradation due to uncertainty in this extension. Interestingly, this suggests that the firm would degrade the quality further for the customers whose quality sensitivity θ increases the difference $a_1(\theta) - a_0(\theta)$; i.e., the customers who are expected to incur substantially different costs under the hypotheses $\{c = c_0\}$ and $\{c = c_1\}$ would face more quality distortion. From a practical standpoint, this means that if a regulation or technological innovation significantly affects a particular region in the domain of the firm's cost curve $c(\theta, q)$, then the heterogenous cost parameters under the two hypotheses, namely $a_0(\theta)$ and $a_1(\theta)$, would differ substantially and result in stronger quality distortion over the affected region.

Health-contingent wellness programs and moral hazard. In the broader context of health insurance, a related issue is the implementation of health-contingent wellness programs, which can help an insurer partially control its customers' cost shocks. The Affordable Care Act allows certain insurers to implement health-contingent wellness programs to induce healthier customer behavior (see, e.g., Health Care Reform Bulletin 2013). One of the goals of such wellness programs is to mitigate the moral hazard that emerges when a customer decides to exploit a high-quality product *after* the time of sales, resulting in increased subsequent costs for the firm's higher quality products. To explicitly model the role of a customer's unobservable effort that can avoid cost inflation after the time of sales, one can consider the following extension of our setting. After a sale in period t , the firm makes multiple cost

observations that involve a sequence of cost shocks $\{\epsilon_{s,t}, s = 1, 2, \dots, S\}$ in (3.1), such that each shock $\epsilon_{s,t}$ depends on a subsequent effort process $\{e_s \in [e_{\min}, e_{\max}], s = 1, \dots, S\}$, which is to be determined by the customer. To be precise, we let $\epsilon_{s,t} = \mu(e_s) + \xi_{s,t}$ for all $s = 1, \dots, S$, where $\mu : [e_{\min}, e_{\max}] \rightarrow \mathbb{R}$ is a decreasing function with $\mu(e_{\min}) > 0$ and $\mu(e_{\max}) = 0$, and $\xi_{s,t}$ are independent and identically distributed random variables with mean zero and variance σ^2 . If it is costly for customers to exert an effort greater than $e_{\min}$, then the firm would need to incentivize its customers to reduce after-sales costs using rewards and penalties adapted to the idiosyncratic cost observations from a customer. Our basic model can be viewed as the case where the firm induces the customers to exert a fixed effort level $e_s = e^*$ for all $s = 1, 2, \dots, S$ such that the expected cost due to moral hazard is minimized.

More on regulations: Funding and labor supply. In Section 4, we have studied how commonly used quality-subsidizing regulations affect the firm's price and quality decisions. It is worth noting that two additional questions in the implementation of these regulations are: how they should be funded, and how they will influence the labor supply. For the former question, we note that both of the regulations we have studied can be funded by allocating taxes on the firm and the customers. A free low-quality outside option decreases the firm's profits, and hence needs to be funded by customers who receive larger amounts of net utility. On the other hand, price subsidies increase the firm's profits, and hence can be funded by the firm in a Pareto-efficient way. Regarding the latter question, if we view the firm as an employer that provides health insurance to its employees, then the customers who obtain a positive utility from the firm's products constitute the firm's labor supply. With this interpretation in mind, the two aforementioned regulations have contrasting effects on labor supply: While the free low-quality outside option decreases labor supply, the price subsidies increase it (see Garthwaite, Gross and Notowidigdo 2014, for further discussion of the relationship between health insurance regulations and labor supply).

A. Proof of Propositions 2-3 and Theorem 1

Proof of Proposition 2. Let us first characterize the optimal solution of (2.7). Note that we can maximize the integral in (2.7) by choosing the quality $q(\theta)$ to maximize the integrand for each value of θ . Therefore, the first order necessary conditions for (unconstrained) optimality are as follows:

$$\theta - \frac{\bar{F}(\theta)}{f(\theta)} = \frac{\partial C(b, q(\theta))}{\partial q} \quad \text{for all } \theta \in \Theta. \quad (\text{A.1})$$

Because $C(b, q) = bc_1(q) + (1 - b)c_0(q) = ba_1q^2 + (1 - b)a_0q^2$, the preceding conditions are equivalent to $\theta - \bar{F}(\theta)/f(\theta) = 2ba_1q(\theta) + 2(1 - b)a_0q(\theta)$ for all $\theta \in \Theta$. Since $C(b, q)$ is convex in q , we know that the above necessary conditions are also sufficient for optimality. Recalling that $f(\theta)$ has increasing failure rate, we deduce that $\bar{F}(\theta)/f(\theta)$ is decreasing in θ , and hence the solution to (A.1) is increasing in θ , satisfying the monotonicity constraint for $q(\cdot)$. Therefore, the optimal quality allocation $q^*(b, \theta)$ for the problem in (2.7) is given by

$$q^*(b, \theta) = \left(\frac{\theta - \bar{F}(\theta)/f(\theta)}{2a_0 + 2(a_1 - a_0)b} \right)_+ \quad \text{for all } b \in [0, 1] \text{ and } \theta \in \Theta, \quad (\text{A.2})$$

where $x_+ = \max\{x, 0\}$ for all $x \in \mathbb{R}$. Note that, if $q^*(b, \theta)$ is an interior solution, then $\partial q^*(b, \theta)/\partial b = -(\theta - \bar{F}(\theta)/f(\theta))(a_1 - a_0)/(2a_0 + 2(a_1 - a_0)b)^2 < 0$ and $\partial^2 q^*(b, \theta)/\partial b^2 = 2(\theta - \bar{F}(\theta)/f(\theta))(a_1 - a_0)^2/(2a_0 + 2(a_1 - a_0)b)^3 > 0$. As a result, in any open subset of $[0, 1] \times \Theta$ in which $q^*(b, \theta)$ is an interior solution, $q^*(b, \theta)$ is a strictly decreasing and strictly convex function of b . ■

Proof of Proposition 3. Let $\theta_c = \inf \{\theta \in [\theta_{\min}, \theta_{\max}] : \theta - \bar{F}(\theta)/f(\theta) \geq 0\}$. Let $b \in [0, 1]$. Because $q^*(b, \theta) = 0$ for all $\theta \leq \theta_c$, we have

$$V(b) = \int_{\theta_c}^{\theta_{\max}} \left(\theta q^*(b, \theta) - \frac{\bar{F}(\theta)}{f(\theta)} q^*(b, \theta) - C(b, q^*(b, \theta)) \right) f(\theta) d\theta \stackrel{(b)}{=} \frac{\kappa}{a_0 + (a_1 - a_0)b}, \quad (\text{A.3})$$

where $\kappa = \frac{1}{4} \int_{\theta_c}^{\theta_{\max}} (\theta - \bar{F}(\theta)/f(\theta))^2 f(\theta) d\theta$, and (b) follows by (A.2). Using the preceding identity and elementary algebra, we consequently get

$$\Delta(b) = \frac{b\kappa}{a_1} + \frac{(1-b)\kappa}{a_0} - \frac{\kappa}{a_0 + (a_1 - a_0)b} \geq \frac{(a_1 - a_0)^2 \kappa b(1-b)}{a_0 a_1^2} = Kb(1-b), \quad (\text{A.4})$$

where $K = (a_1 - a_0)^2 \kappa / (a_0 a_1^2)$. ■

Proof of Theorem 1. Note that

$$\delta(b, \theta) = bq^*(1, \theta) + (1-b)q^*(0, \theta) - q^*(b, \theta) \stackrel{(a)}{=} g_0(\theta) \left(\frac{b}{a_1} + \frac{1-b}{a_0} - \frac{1}{a_0 + (a_1 - a_0)b} \right) \stackrel{(b)}{\geq} \frac{g_0(\theta)(a_1 - a_0)^2 b(1-b)}{a_0 a_1^2}$$

for all $b \in [0, 1]$ and $\theta \in \Theta$, where: $g_0(\theta) = \frac{1}{2}(\theta - \bar{F}(\theta)/f(\theta))_+$, (a) follows by (A.2), and (b) follows by the argument used to derive (A.4). Therefore, letting $g(\theta) = g_0(\theta)(a_1 - a_0)^2 / (a_0 a_1^2)$, we deduce that $\delta(b, \theta) \geq g(\theta)b(1-b)$ for all $b \in [0, 1]$ and $\theta \in \Theta$. Since $f(\theta)$ has increasing failure rate, $g(\theta)$ is increasing in θ . Moreover, because $\bar{F}(\theta_{\max}) = 0$ and $\theta_{\max} > 0$, we have $g(\theta_{\max}) > 0$. ■

B. Proof of Theorem 2

Step 1: Prove that $\partial \mathcal{V}_n(b, q)/\partial q \geq 0$ for all $b \in [0, 1]$, and $q \in \mathcal{Q}$. Fix $b \in [0, 1]$, $q \in \mathcal{Q}$, and $n \geq 2$. Let us first express the marginal conditional value function in the following lemma.

LEMMA 1. (derivative of conditional value function)

$$\frac{\partial \mathcal{V}_n(b, q)}{\partial q} = b \int_{-\infty}^{\infty} V_n(\beta(b, \gamma, q)) s(\gamma, q) L_1(\gamma, q) d\gamma + (1-b) \int_{-\infty}^{\infty} V_n(\beta(b, \gamma, q)) s(\gamma, q) L_0(\gamma, q) d\gamma \quad (\text{B.1})$$

for all $b \in [0, 1]$, and $q \in \mathcal{Q}$, where $s(\gamma, q) := 2q^{-1}(\sigma^{-2}(\gamma - a_0 q^2)(\gamma - a_1 q^2) - 1)$.

Note that, in the integrands on the right hand side of (B.1), almost all terms are non-negative; the only term with a mixed sign is $s(\gamma, q)$. To be more precise, the function $s(\cdot, q)$ is a parabola with the following roots:

$$\gamma_0 = \frac{(a_0 + a_1)q^2 - \sqrt{(a_1 - a_0)^2 q^4 + 4\sigma^2}}{2} \quad \text{and} \quad \gamma_1 = \frac{(a_0 + a_1)q^2 + \sqrt{(a_1 - a_0)^2 q^4 + 4\sigma^2}}{2}. \quad (\text{B.2})$$

Thus,

$$s(\gamma, q) \begin{cases} < 0 & \text{if } \gamma \in (\gamma_0, \gamma_1) \\ \geq 0 & \text{otherwise.} \end{cases} \quad (\text{B.3})$$

Let $x_i = \beta(b, \gamma_i, q)$ and $y_i = V_n(x_i)$ for $i = 0, 1$. Because $\beta(b, \gamma, q)$ is increasing in γ , we know that

$$\begin{aligned} \beta(b, \gamma, q) &\leq x_0 && \text{if and only if } \gamma \leq \gamma_0; \\ \beta(b, \gamma, q) &\in (x_0, x_1) && \text{if and only if } \gamma \in (\gamma_0, \gamma_1); \\ \beta(b, \gamma, q) &\geq x_1 && \text{if and only if } \gamma \geq \gamma_1. \end{aligned} \quad (\text{B.4})$$

Now, define a function $\bar{h} : [0, 1] \rightarrow \mathbb{R}$ such that $\bar{h}(b) = m_0 + m_1 b$, where $m_0 = y_0 x_1 - y_1 x_0$ and $m_1 = (y_0 - y_1)/(x_0 - x_1)$. That is, $\bar{h}(\cdot)$ is the affine function whose trajectory coincides with the line passing through the

points (x_0, y_0) and (x_1, y_1) . Because $V_n(\cdot)$ is a convex function, $V_n(b) \leq \bar{h}(b)$ if $b \in (x_0, x_1)$, and $V_n(b) \geq \bar{h}(b)$ if $b \in (-\infty, x_0] \cup [x_1, \infty)$. Consequently, we deduce from (B.4) that

$$\begin{aligned} V_n(\beta(b, \gamma, q)) - \bar{h}(\beta(b, \gamma, q)) &\geq 0 \quad \text{if and only if} \quad \gamma \leq \gamma_0; \\ V_n(\beta(b, \gamma, q)) - \bar{h}(\beta(b, \gamma, q)) &\leq 0 \quad \text{if and only if} \quad \gamma \in [\gamma_0, \gamma_1]; \\ V_n(\beta(b, \gamma, q)) - \bar{h}(\beta(b, \gamma, q)) &\geq 0 \quad \text{if and only if} \quad \gamma \geq \gamma_1. \end{aligned} \quad (\text{B.5})$$

Recalling (B.3), we deduce that $(V_n(\beta(b, \gamma, q)) - \bar{h}(\beta(b, \gamma, q)))s(\gamma, q) \geq 0$ for all $\gamma \in \mathbb{R}$. This implies that

$$\begin{aligned} 0 &\leq b \int_{-\infty}^{\infty} (V_n(\beta(b, \gamma, q)) - \bar{h}(\beta(b, \gamma, q))) s(\gamma, q) L_1(\gamma, q) d\gamma + (1-b) \int_{-\infty}^{\infty} (V_n(\beta(b, \gamma, q)) - \bar{h}(\beta(b, \gamma, q))) s(\gamma, q) L_0(\gamma, q) d\gamma \\ &\stackrel{(a)}{=} \frac{\partial \mathcal{V}_n(b, q)}{\partial q} - b \int_{-\infty}^{\infty} \bar{h}(\beta(b, \gamma, q)) s(\gamma, q) L_1(\gamma, q) d\gamma - (1-b) \int_{-\infty}^{\infty} \bar{h}(\beta(b, \gamma, q)) s(\gamma, q) L_0(\gamma, q) d\gamma, \end{aligned} \quad (\text{B.6})$$

where (a) follows by Lemma 1. To eliminate the remaining integrals on the right hand side of (B.6), we state another lemma.

LEMMA 2. *Let $h : [0, 1] \rightarrow \mathbb{R}$ be an affine function. Then,*

$$0 = b \int_{-\infty}^{\infty} h(\beta(b, \gamma, q)) s(\gamma, q) L_1(\gamma, q) d\gamma + (1-b) \int_{-\infty}^{\infty} h(\beta(b, \gamma, q)) s(\gamma, q) L_0(\gamma, q) d\gamma \quad (\text{B.7})$$

for all $b \in [0, 1]$, and $q \in \mathcal{Q}$.

By Lemma 2, the right hand side of (B.6) is equal to $\partial \mathcal{V}_n(b, q)/\partial q$, from which we conclude that $\partial \mathcal{V}_n(b, q)/\partial q \geq 0$.

Step 2: Prove that $q_n^*(b, \theta) \geq q_1^*(b, \theta)$ for $n \geq 2$. Fix $b \in [0, 1]$. By (3.17), we know that

$$q_n^*(b, \cdot) = \arg \max_{q(\cdot) \in \mathcal{A}^*} \left\{ \int_{\theta_{\min}}^{\theta_{\max}} \left(R(\theta)q(\theta) - a_b(q(\theta))^2 + \mathcal{V}_{n-1}(b, q(\theta)) \right) f(\theta) d\theta \right\}, \quad (\text{B.8})$$

where $R(\theta) = \theta - \bar{F}(\theta)/f(\theta)$ and $a_b = ba_1 + (1-b)a_0$. As argued in the proof of Proposition 2, we can maximize the integral in (B.8) by maximizing its integrand for every $\theta \in \Theta$, and the (unconstrained) maximizer satisfies the following first order necessary condition:

$$\Psi_n(b, \theta, q) := R(\theta) - 2a_b q + \frac{\partial \mathcal{V}_{n-1}(b, q)}{\partial q} = 0. \quad (\text{B.9})$$

Consider the first order condition for $n = 1$; in this case, we have $\mathcal{V}_{n-1}(b, q) = \mathcal{V}_0(b, q) = 0$ for all $q \in \mathcal{Q}$, and hence $\partial \mathcal{V}_0(b, q)/\partial q = 0$ for all $q \in \mathcal{Q}$. Let $\theta_c = \inf \{\theta \in \Theta : R(\theta) \geq 0\}$. For all $\theta \leq \theta_c$, we have $q_1^*(b, \theta) = 0$, which implies that $q_n^*(b, \theta) \geq q_1^*(b, \theta)$. For all $\theta > \theta_c$, we have $q_1^*(b, \theta) = R(\theta)/(2a_b)$, and hence $R(\theta) - 2a_b q \geq 0$ for all $q \in [0, q_1^*(b, \theta)]$. Now, let $n \geq 2$. Because $\partial \mathcal{V}_{n-1}(b, q)/\partial q \geq 0$ as shown in Step 1, $\Psi_n(b, \theta, q) = R(\theta) - 2a_b q + \partial \mathcal{V}_{n-1}(b, q)/\partial q \geq 0$ for all $q \in [0, q_1^*(b, \theta)]$. This means that the maximizer that satisfies the first order condition for $n \geq 2$, namely $q_n^*(b, \theta)$, is greater than or equal to $q_1^*(b, \theta)$. As a result, $q_n^*(b, \theta) \geq q_1^*(b, \theta)$ for $\theta \in \Theta$ and $n \geq 2$.

Step 3: Prove that $\delta_n(b, \theta) \leq \delta_1(b, \theta)$ for $n \geq 2$. Fix $b \in [0, 1]$, $\theta \in \Theta$, and $n \geq 2$. Note that

$$\begin{aligned} \frac{\partial \mathcal{V}_n(i, q)}{\partial q} &\stackrel{(b)}{=} i \int_{-\infty}^{\infty} V_n(\beta(i, \gamma, q)) s(\gamma, q) L_1(\gamma, q) d\gamma + (1-i) \int_{-\infty}^{\infty} V_n(\beta(i, \gamma, q)) s(\gamma, q) L_0(\gamma, q) d\gamma \\ &\stackrel{(c)}{=} i V_n(i) \int_{-\infty}^{\infty} s(\gamma, q) L_1(\gamma, q) d\gamma + (1-i) V_n(i) \int_{-\infty}^{\infty} s(\gamma, q) L_0(\gamma, q) d\gamma \stackrel{(d)}{=} 0 \end{aligned} \quad (\text{B.10})$$

for $i = 0, 1$, and $q \in \mathcal{Q}$, where: (b) follows by Lemma 1; (c) follows because $\beta(i, \gamma, q) = i$ for $i \in \{0, 1\}$, $\gamma \in \mathbb{R}$, and $q \in \mathcal{Q}$; and (d) follows by Lemma 2 and the fact that a constant function is affine. Therefore, $q_n^*(i, \theta) = q_1^*(i, \theta)$ for $i = 0, 1$, implying that $\delta_n(b, \theta) \leq \delta_1(b, \theta)$. ■

References

Araman, V. and Caldentey, R. (2009), ‘Dynamic Pricing for Nonperishable Products with Demand Learning’, *Operations Research* **57**(5), 1169–1188.

- Aviv, Y. and Pazgal, A. (2005), 'A Partially Observed Markov Decision Process for Dynamic Pricing', *Management Science* **51**(9), 1400–1416.
- Bergemann, D., Shen, J., Xu, Y. and Yeh, E. (2012), 'Mechanism Design with Limited Information: The Case of Nonlinear Pricing', *Game Theory for Networks* pp. 1–10.
- Bergemann, D. and Välimäki, J. (2000), 'Experimentation in markets', *The Review of Economic Studies* **67**(2), 213–234.
- Besanko, D., Donnenfeld, S. and White, L. (1987), 'Monopoly and Quality Distortion: Effects and Remedies', *The Quarterly Journal of Economics* **102**(4), 743–768.
- Besanko, D., Donnenfeld, S. and White, L. (1988), 'The Multiproduct Firm, Quality Choice, and Regulation', *The Journal of Industrial Economics* **36**(4), 411–429.
- den Boer, A. (2014), 'Dynamic Pricing with Multiple Products and Partially Specified Demand Distribution', *Mathematics of Operations Research* **39**(3), 863–888.
- Farias, V. and van Roy, B. (2010), 'Dynamic Pricing with a Prior on Market Response', *Operations Research* **58**(1), 16–29.
- Garthwaite, C., Gross, T. and Notowidigdo, M. (2014), 'Public Health Insurance, Labor Supply, and Employment Lock', *Quarterly Journal of Economics* **129**(2), 653–696.
- Goodell, S. (2014), 'Health Policy Brief: Risk Corridors'. *Health Affairs*, June 26, 2014.
- Goodnough, A. (2014), 'Newly Insured, Many Now Face Learning Curve'. *The New York Times*, Published on August 3, 2014, p. A1. URL: <http://nyti.ms/1m9Xr71>.
- Harrison, J., Keskin, N. and Zeevi, A. (2012), 'Bayesian Dynamic Pricing Policies: Learning and Earning Under a Binary Prior Distribution', *Management Science* **58**(3), 570–586.
- Health Care Reform Bulletin (2013), 'Wellness Program Final Rules Increase Maximum Rewards, but Include Significant New Requirements'. *Towers Watson*, June 11, 2013.
- Jullien, B. (2000), 'Participation Constraints in Adverse Selection Models', *Journal of Economic Theory* **93**, 1–47.
- Keller, G. and Rady, S. (1999), 'Optimal Experimentation in a Changing Environment', *The Review of Economic Studies* **66**(3), 475–507.
- Keller, G. and Rady, S. (2003), 'Price Dispersion and Learning in a Dynamic Differentiated-Goods Duopoly', *RAND Journal of Economics* **34**(1), 138–165.
- Keskin, N. and Zeevi, A. (2014), 'Dynamic Pricing with an Unknown Demand Model: Asymptotically Optimal Semi-myopic Policies', *Operations Research* **62**(5), 1142–1167.
- Kowalski, A. (2014), 'The Early Impact of the Affordable Care Act'. Working Paper, Yale University, New Haven, CT.
- Lippman, S. and McCardle, K. (1987), 'Does Cheaper, Faster, or Better Imply Sooner in the Timing of Innovation Decisions?', *Management Science* **33**(8), 1058–1064.
- Maskin, E. and Riley, J. (1984), 'Monopoly with Incomplete Information', *The RAND Journal of Economics* **15**(2), 171–196.
- Mazzola, J. and McCardle, K. (1996), 'A Bayesian Approach to Managing Learning-curve Uncertainty', *Management Science* **42**(5), 680–692.
- McCardle, K. (1985), 'Information Acquisition and the Adoption of New Technology', *Management Science* **31**(11), 1372–1389.
- Mussa, M. and Rosen, S. (1978), 'Monopoly and Product Quality', *Journal of Economic Theory* **18**(2), 301–317.
- Rochet, J. and Stole, L. (2002), 'Nonlinear Pricing with Random Participation', *The Review of Economic Studies* **69**, 277–311.
- Smith, J. and Ulu, C. (2012), 'Technology Adoption with Uncertain Future Costs and Quality', *Operations Research* **60**(2), 262–274.
- Ulu, C. and Smith, J. (2009), 'Uncertainty, Information Acquisition, and Technology Adoption', *Operations Research* **57**(3), 740–752.