Probability in the Engineering and Informational Sciences, **26**, 2012, 77–104.
doi:10.1017/S0269964811000246

OPTIMAL MARKDOWN PRICING STRATEGY WITH DEMAND LEARNING

H. Dharma Kwon
Department of Business Administration
University of Illinois at Urbana-Champaign
Champaign, IL 61820
E-mail: dhkwon@illinois.edu

Steven A. Lippman and Christopher S. Tang
UCLA Anderson School, UCLA
Los Angeles, CA 90095
E-mail: slippman@anderson.ucla.edu; ctang@anderson.ucla.edu

When launching a new product, a firm has to set an initial price with incomplete information about demand. However, after observing the demand over a period of time, the firm might decide to mark down the price, especially when the Bayesian updated belief about demand is lower than originally anticipated. We consider the case in which the manufacturer makes three decisions: initial price, when to mark down the price, and the markdown price. Modeling the cumulative demand as a Brownian motion with an unknown drift, we compute the posterior probability distribution of the unknown drift. We then show that it is optimal to mark down the price when the posterior probability is below a computable threshold. This threshold policy enables us to determine the optimal (a) regular price, (b) markdown price, and (c) markdown time. Additionally, we examine the impact of demand volatility and evaluate the value of learning.

1. INTRODUCTION

In the presence of fast-paced technological innovation, launching new products is critical for (high-tech) firms to remain competitive. In launching a new product, a firm must decide on its initial price without full knowledge of demand. To hedge against demand risk, it is common for the firm to begin with a high initial price, knowing that subsequently it can mark down the price, especially when demand falls below what was anticipated. For example, on June 29, 2007, Apple launched its iPhone

© Cambridge University Press 2011 0269-9648/12 $25.00

https://doi.org/10.1017/S0269964811000246 Published online by Cambridge University Press

by setting its initial price at $599. Sales of the new iPhone, however, did not meet Apple's expectations[1] and Apple dropped the price to $399 on September 5, 2007 (Markoff [17]).

Three factors impede the firm's ability to select an optimal new product pricing strategy. First, highly uncertain demand makes it difficult to set a wise initial price. Second, even with Bayesian learning, it is difficult to know when (if ever) enough information has been gained in order to mark down the price. Third, in addition to the significant implementation cost associated with a price change (Fisher [12]), frequent markdowns will induce strategic purchasing behavior: Some customers will postpone their purchase in anticipation of future price drops (Elmagrabhy, Lippman, Tang, and Yin [4]).[2] To capture the essence of these three challenges, we suppose that the firm can reduce the price only once. Hence, the firm makes three decisions: the initial price, the time when it will mark down the price, and the markdown price. To aid it in making these three decisions, the firm uses Bayesian learning with regard to its belief about the product's demand. Our goal is to determine the optimal initial and markdown prices as well as the (stopping) time when the markdown occurs. In addition, we examine the effect of demand uncertainty upon these three decision variables, and we evaluate the value of demand learning.

In this article, we present a continuous-time model in which the state of demand for the firm's new product is either high or low and the state never changes. Given a sales price, more is sold when the state of demand is high than when it is low. Although the firm knows the specification of the demand curve for each state, it has incomplete information: The firm knows the probability at time 0 that demand is high. The firm can, however, utilize the observed demand at each point in time t along with the initial price to update its belief and obtain the posterior probability P_t that demand is "high." There are no costs apart from the constant marginal cost of production, and all revenues come from sales. All revenues are received and all costs are incurred at the time of sales, and the firm seeks to maximize its expected discounted net profit. The firm has three decision variables at its disposable: the initial price, the markdown price, and the stopping time at which the price is marked down.

Deciding when to stop and mark down the initial price entails trade-offs. On the one hand, the advantage of an early decision to stop is that the needed price adjustment occurs early so that more profit accrues to this adjustment. On the other hand, an early decision to stop means that less information is available so that there is a larger probability that the firm will mark down the price even though demand is high, thereby leading to lower profits. Interestingly, it turns out that premature stopping does not mean stopping at a time that is small. Instead, it means stopping without enough

[1] Van Baker, an industry analyst from Gartner Group, said, "My suspicion is that they got to 750,000 really quickly, and then it started to slow down" (Markoff [17]).

[2] It is still challenging even if a firm reduces its price only once. For example, when Apple reduced the price of its iPhone by $200 three months after launch, many early buyers were angry because they had expected that iPhone would maintain its initial price $599 for a much longer period of time. In response, Apple awarded a $100 store credit to each early buyer (Hafner [13]).

https://doi.org/10.1017/S0269964811000246 Published online by Cambridge University Press

information; more specifically, it means stopping when the posterior probability of high demand has yet to fall sufficiently far.

To incorporate demand uncertainty, we model sales as a one-dimensional Brownian motion. Utilizing the encyclopedia of results for Brownian motion, we are able to determine the optimal decisions (initial price, markdown price, and markdown time), examine the impact of an increase in demand volatility on these optimal decisions, and consider the value of learning whether demand is high or low.

This article makes three contributions. First, for any pair of initial and markdown prices, a natural result of our model with demand uncertainty and Bayesian learning is that the problem of finding the optimal markdown time is a stopping time problem. We show that the optimal stopping time is a threshold rule: Mark down the price when the posterior probability that the underlying demand curve is high drops below a threshold. (We supply a closed-form solution for this threshold.) Second, we develop joint optimal decisions for the stopping time, the initial price, and the markdown price. The intuition for our third contribution rests upon the observation that the rate at which the firm learns about the underlying state of demand depends on the demand volatility σ of the Brownian motion. We show that as demand volatility increases, the firm's expected profit decrease. This is not surprising due to the fact that the observed demand process conveys less information when σ increases. Furthermore, we find that the optimal initial price decreases in σ and that the optimal markdown price and the optimal threshold increase in σ when σ is either very small or very large.

This article is organized as follows. In Section 2 we briefly review how our article fits in with the related literature. We present our model of demand and the demand learning process in Section 3, where we analyze the base model and establish that for any pair of initial and markdown prices, the optimal markdown policy is a threshold rule: Mark down the initial price when the posterior probability that demand is in the high state first hits the threshold. We use this optimal markdown policy to obtain comparative statics and evaluate the value of demand learning. Section 4 expands the base model to the case in which the regular and markdown prices are decision variables. We conclude in Section 5. To enhance the presentation, all proofs are given in the Appendix.

2. RELATIONSHIP TO THE LITERATURE

Our article is related to two streams of the operations literature: dynamic pricing with learning and optimal stopping with learning. For literature on optimal dynamic pricing, see Talluri and van Ryzin [27]. In the traditional dynamic pricing literature, the decision maker can adjust the price at any point in time. This class of dynamic pricing problems can be formulated as optimal control problems. When the firm has incomplete information about the underlying demand curve, Balvers and Consimano [2] studied a periodic-review dynamic pricing problem in which the underlying demand curve is linear in the price and demand is an AR(1) process. In their model, the noise of the realized demand at each period depends linearly on the difference between

https://doi.org/10.1017/S0269964811000246 Published online by Cambridge University Press

the selected price and the market price that is optimal for that period (ignoring the value of learning). They concluded that the incentive to learn about the future demand induces the decision maker to temper price changes he otherwise would have made. More recently, Lin [15] studied a finite-horizon continuous-time model with demand learning in which the demand arrives according to a Poisson process with uncertain arrival rate. His periodic-review dynamic pricing heuristic is based on the updated posterior of the arrival rate.

The issue of joint pricing and production planning with demand uncertainty has been examined by various researchers recently. For example, Federgruen and Heching [10] considered a joint pricing and production planning problem when the probability distribution of the underlying demand is known. Formulating this problem as a discrete-time dynamic program, they characterized the optimal price and production quantity in each time period. Yin and Rajaram [28] extended Federgruen and Heching's model to the case when the underlying demand distribution changes in accord with a Markov process. They show that the characteristics of the optimal price and production quantity are similar to those found in Federgruen and Heching [10]. Subrahmanyan and Shoemaker [26] developed a discrete-time model for the joint dynamic pricing and production planning problem with uncertain demand distribution. By using the past sales and price history and by updating the probability distribution of demand in a Bayesian manner, they used numerical examples to illustrate the value of demand learning. Despite some technical difficulties, Petruzzi and Dada [21] characterized the joint optimal price and production quantity when the demand is deterministic but unknown. Among the drawbacks associated with many models of dynamic pricing with demand learning are the facts that (a) the models are very complex and have no known closed-form solutions, (b) the implementation costs of price changes can be prohibitively large (see Fisher [12]), and (c) adjusting the price dynamically might not be practical due to strategic consumer purchasing behavior (in contrast to consumers acting as mere price takers: see Elmagrabhy et al. [9]). These observations motivate us to limit the firm in our model to at most one price markdown. By focusing on a single price change, the underlying mathematical structure of our problem is further differentiated from dynamic control problems. More importantly, the firm in our problem makes only three decisions: the initial (regular) price, the time to mark down the price, and the final (markdown) price. Given the regular and markdown prices, we obtain a closed-form solution for the optimal markdown time. The closed-form solution facilitates our examination of the impact of volatility on the optimal markdown time and our assessment of the value of demand learning. This closed-form solution for the optimal markdown time enables us to develop expressions for the optimal regular and markdown prices.

Feng and Gallego [11] formulate their dynamic pricing paper on selling a fixed stock of items over a finite horizon as a stopping time problem. In their article, the demand distribution is known but stochastic, and the seller can change the price only once. The main result consists in showing that the optimal price-change policy is a threshold rule with respect to the time remaining. Unlike our model, theirs has no learning.

https://doi.org/10.1017/S0269964811000246 Published online by Cambridge University Press

Optimal stopping with learning is the second stream of literature to which our problem is related. There is vast literature on optimal stopping under incomplete information (see Peskir and Shiryaev [20] for references). Among many applications, this literature includes financial models[3] and classical sequential analysis. Our two-state model along with many other economics models, which we now discuss, fits into this framework.

To know when to stop, it is essential to be able to compute the posterior probability P_t of being in the high state given the information available at time t. Shiryaev [25] obtained an elegant closed-form solution in a problem with continuous-time sequential testing of two simple hypotheses on the drift of a one-dimensional Brownian motion. The essential building block is the law of motion of the posterior probability P_t (see (1) in Section 3.2). Shiryaev's [25] model has been applied and extended by many researchers to problems of optimizing discounted profit streams. In particular, our problem is closely related to the optimal exit problem of Ryan and Lippman [23], which considers when to stop and abandon a project. In their problem, the firm seeks to maximize the discounted cumulative profit while operating the project; exit is costless, and cumulative profit is a Brownian motion in which the drift takes one of two known values: One of value is positive and the other is negative. Using stochastic calculus, they showed that it is optimal to abandon the project when the posterior probability that the drift is negative exceeds an explicitly computable threshold, and the threshold is time independent.[4] If the regular and the markdown prices were fixed, then our model would be nearly the same as the model of Ryan and Lippman [23]. By taking this approach in our analysis, we solve the stopping problem while examining its interrelationship with the pricing problem.

Many other economics papers apply the two-state model governed by a Brownian motion. One general context with a plethora of published work pertains to the optimal level of experimentation in several different settings. For example, in Keller and Rady [14], the monopolist selects a production quantity and observes the market price that is determined by one of two linear demand curves that contains a random component. They considered the experimentation trade-off between gaining more information about the demand state via extreme production quantities, which leads to an increase in future rewards, and maximizing the current profit rate by producing a quantity near the current optimal. Moscarini and Smith [18] studied a model in which the decision maker can control the volatility of the noise. By increasing the effort (cost) level, the decision maker achieves a higher rate of experimentation and hence a higher rate of learning of the underlying state. They show that the optimal

[3] A substantial portion of the modern finance literature can be included here. Our discussion in this second stream of literature focuses on economic models with experimentation. A closely related work in the finance literature by Bernardo and Chowdhry [4] focuses on experimentation in which the firm seeks to learn about itself. In their work, the decision maker's investment decision is based on the posterior probability distribution of the future income flow evaluated at the stopping time.

[4] See Ryan and Lippman [24] for an extension in which the drift drops to the lower known value at an unobservable exponential time. By observing the realized profit stream, the decision maker updates the posterior probability that the drift of the profit stream has dropped. They showed that it is optimal to exit when this posterior probability hits a critical value.

https://doi.org/10.1017/S0269964811000246 Published online by Cambridge University Press

level of experimentation grows with the reward from stopping and that the optimal level of experimentation peaks when the decision maker is most uncertain about the underlying state. Others works consider competitive experimentation with more than one decision maker. When more than one player engages in experimentation, information externalities rise if the observation of each player is public. Bolton and Harris [5] studied the strategic aspect of the information externalities; they found that the prospect of future experimentation by other players encourages players to increase current experimentation despite the possibility of a free-rider problem. In Bergemann and Valimaki's [3] two-firm model of competition, one firm sells a product with a known value and the other sells a new product whose value is unknown; each firm controls its own price, and all of the players in the market (firms and buyers) observe the noisy market demand for the new product and learn the probability that it has a high value. When there is one market, the information externalities result in excessive experimentation. For the most part, these works focus on decision making that leads to more rapid learning, and there is no stopping decision. Our article focuses on the optimal stopping decision—namely when to stop selling at the regular price and reduce the price.

A recent example in which the stopping problem is at the fore, like our article, is Decamps, Mariotti, and Villeneuve [7], who studied the optimal time to invest in an (noisy) asset with an unknown underlying value. In their work, the return is a function of the Brownian motion itself, and the stopping time is path dependent: It depends on the income stream rather than on the posterior probability distribution of the future income flow.

3. MODEL

This section contains the details of our two-state continuous-time model along with some insights and intuitions. We pay particular attention to the Bayesian updating and the computation of the posterior probability P_t of being in the high state at time t.

As noted in the Introduction, the firm has one opportunity to reduce the price of its new product from the regular price r to the markdown price m, and it can do so at any time. The firm seeks to maximize its expected discounted profit with judicious choices of the price schedule (r, m) and the markdown time τ. For the entirety of this section, we take the price schedule to be fixed and focus on finding the stopping time (i.e., markdown time) τ that maximizes the firm's expected discounted profit over an infinite time horizon where α is the continuous-time discount rate.

3.1. Demand

As explicated in this subsection and the next, cumulative demand has both a deterministic and a stochastic component. Whereas the realized demand, the sum of these two components, is observed without error, the stochastic component masks the state of demand.

https://doi.org/10.1017/S0269964811000246 Published online by Cambridge University Press

There are two states of demand, high and low, and the state never changes.[5] The deterministic component associated with the high and the low state are the demand curves D_h and D_ℓ. At time 0 in our continuous-time model, the firm launches its new product and sets the initial price. Let x be the initial price chosen by the firm. In accord with the names *high state* and *low state*, we have $D_h(x) > D_\ell(x)$: The deterministic component of demand per unit time in the high state exceeds the deterministic component of demand in the low state. Let D denote the actual demand curve: $D = D_h$ if the state is high and $D = D_\ell$ if the state is low. Initially, the firm does not know the state; it only knows the probability that the state is high. (Thus, D is a random variable.) The firm observes the realized demand without error and, of course, the firm knows the initial price it has chosen; nevertheless, the firm never learns the state of demand because of the stochastic component (which is independent of the state) of the realized demand.

Let $\mu_h(x) \equiv (x - c)D_h(x)$ and $\mu_\ell(x) \equiv (x - c)D_\ell(x)$ denote the expected profit per unit time in the high and low demand states, respectively, when the price is x and c is the constant marginal cost of production. For tractability, we impose mild assumptions on the demand curves:

DEMAND ASSUMPTIONS:

(i) $\mu_h(\cdot)$ *and* $\mu_\ell(\cdot)$ *are unimodal and twice continuously differentiable.*

(ii) $\mu_h(x) > \mu_\ell(x)$ *for all* x.

(iii) $\bar{x} > \underline{x}$, *where* $\bar{x}$ *and* $\underline{x}$ *are the prices that maximize* $\mu_h(\cdot)$ *and* $\mu_\ell(\cdot)$, *respectively.*

These assumptions are innocuous. In the case of linear demand, $D_h(x) = (1 - a_h x)$ and $D_\ell(x) = (1 - a_\ell x)$ with $a_\ell > a_h > 0$, we have $\bar{x} = (1 + a_h c)/2a_h > (1 + a_\ell c)/2a_\ell = \underline{x}$ so all three assumptions are satisfied. The firm must select both an initial price r, which henceforth we refer to as the *regular price*, and a markdown price m. We refer to the pair (r, m) as the *price schedule*.

3.2. Posterior Probability

Our task is to compute the posterior probability P_t of being in the high state at time t and then utilize the posterior probability process to solve the optimal pricing decisions and to determine when to switch from the initial price to the markdown price. Given the initial price r, the cumulative demand C_t at time $t \geq 0$ is given by

$$C_t = D(r)t + \sigma B_t,$$

[5] The focus of this article is when to mark down the price when the true state of demand is masked. As will be seen, this is a difficult problem. If we allow the state itself to change over time, we face an even more difficult problem that is essentially intractable. Although the state does change in Ryan and Lippman [24], they only obtain existence of an optimal threshold.

https://doi.org/10.1017/S0269964811000246 Published online by Cambridge University Press

where σ is the volatility of the realized cumulative demand, and $\mathbf{B} = \{B_t : t \geq 0\}$, the random component of demand, is a one-dimensional standard Brownian motion. Whereas the firm knows the value of σ, it does not know whether the demand state is high or low: The firm has incomplete information regarding underlying demand D or the realization of $\mathbf{B}$. Although $\mathbf{B}$ is not observable, the firm can utilize the observed value of the cumulative demand process $\mathbf{C} \equiv \{C_t : t \geq 0\}$ to update its belief regarding the probability that the demand state is high. Below we present a few technicalities necessary to define and compute the posterior probability P_t that the demand state at time t is high. In order to do this, we construct a Brownian motion $\tilde{\mathbf{B}}$ that is observable, and in (1) we produce the law of motion for the process $\mathbf{P} \equiv \{P_t : t \geq 0\}$. Because the law of motion is given in terms of observable quantities, this transforms our stopping time problem with incomplete information into one with complete information. This transformation allows us to apply the well-established machinery of stochastic calculus to solve our stopping time problems with complete information.[6] What is important is that equation (2) exhibits an easy to implement formula for updating the posterior probability P_t.

Let $\{\mathcal{F}_t : t \geq 0\}$ denote the natural filtration for the process $\mathbf{C}$ so that the σ-algebra $\mathcal{F}_t$ contains all of the information obtained through observing C_s from time $s = 0$ to time $s = t$. Neither the process B_t nor the demand state D are observable by the firm; hence, neither belongs to $\mathcal{F}_t$.

The firm knows the (prior) probability p_0 that the demand state is high. The firm needs to compute P_t, the posterior probability that the demand state is high. By definition, $P_t \equiv E^{p_0}[1_{\{D=D_h\}}|\mathcal{F}_t]$, where $E^{p_0}[\cdot]$ denotes the expectation associated with the initial condition $P_0 = p_0$ and 1_A is the indicator function of the set A. A number of authors have observed that the posterior probability process $\mathbf{P}$ is a continuous martingale with respect to $\{\mathcal{F}_t : t \geq 0\}$, $\mathbf{P}$ is a strong Markov process that is homogeneous in time, and $\mathbf{P}$ is the unique solution to the stochastic differential equation[7]

$$dP_t = \frac{D_h(r) - D_\ell(r)}{\sigma} P_t(1 - P_t)\, d\tilde{B}_t, \tag{1}$$

where r is the regular price and $\tilde{\mathbf{B}} = \{\tilde{B}_t, \mathcal{F}_t : t \geq 0\}$ is a standard Brownian motion defined by $\tilde{B}_t \equiv (1/\sigma)\left(C_t - \int_0^t E[D(r)|\mathcal{F}_s]\, ds\right)$. Unlike B_t, which is unobservable, the new Brownian motion $\tilde{B}_t$ can be deduced from the observed value of C_t over time.

[6] We hasten to point out that knowledge of stochastic calculus is not needed to understand and follow this article.

[7] See Liptser and Shiryaev [16, p. 371] and Peskir and Shiryaev [20, pp. 288–289]. For a more complete and easier to follow discussion, see Ryan and Lippman [23, pp. 452–454].

https://doi.org/10.1017/S0269964811000246 Published online by Cambridge University Press

Using the strong Markov property of the process **P** and Bayes rule (Peskir and Shiryaev, [20, pp. 288–289) we have

$$\begin{aligned} P_t &\equiv E^{p_0}[1_{\{D=D_h\}}|\mathcal{F}_t] \\ &= E^{p_0}[1_{\{D=D_h\}}|C_t] \\ &= p_0 \exp\left\{-\frac{(C_t - D_h(r)t)^2}{2\sigma^2 t}\right\} \Bigg/ \left(p_0 \exp\left\{-\frac{(C_t - D_h(r)t)^2}{2\sigma^2 t}\right\} + (1-p_0)\right. \\ &\qquad \left. \exp\left\{-\frac{(C_t - D_\ell(r)t)^2}{2\sigma^2 t}\right\}\right)^{-1} \\ &= \left[1 + \frac{1-p_0}{p_0} \exp\left\{-\frac{D_h(r) - D_\ell(r)}{\sigma^2}\left(C_t - \frac{D_h(r) + D_\ell(r)}{2}t\right)\right\}\right]^{-1}. \end{aligned} \tag{2}$$

Thus, P_t depends on p_0 and on the difference $\Delta_t \equiv C_t - t[D_h(r) + D_\ell(r)]/2$; in particular, $P_t > p_0$ if and only if $C_t > t[D_h(r) + D_\ell(r)]/2$. Observe that at time t, the entirety of the information garnered from $\{C_s : 0 \le s \le t\}$ is simply Δ_t; that is, if $x = \Delta_t$, then P_t depends on p_0 and x but not on the time t. It is both intuitive and readily evident from (2) that an increase in the volatility σ reduces the response of P_t to additional information: Higher volatility means slower learning.

3.3. Discounted Profit

The firm has one opportunity to reduce the price of its new product from the regular price r to the markdown price m, and it can do so at any time. The firm seeks to maximize its expected discounted profit with judicious choices of the price schedule and the markdown time. Let α denote the strictly positive discount rate.

Given a price schedule (r, m) and a stopping time (also called a markdown time) τ when the price is marked down to m, the expected discounted cumulative profit $\pi_\tau(r, m)$ satisfies

$$\begin{aligned} \pi_\tau(r,m) &= E^{p_0}\left[(r-c)\int_0^\tau e^{-\alpha t} D(r)\,dt + (m-c)\int_\tau^\infty e^{-\alpha t} D(m)\,dt\right] \\ &\quad + E^{p_0}\left[(r-c)\int_0^\tau e^{-\alpha t}\sigma\,dB_t + (m-c)\int_\tau^\infty e^{-\alpha t}\sigma\,dB_t\right] \\ &= \alpha^{-1}E^{p_0}[(r-c)D(r)] + \alpha^{-1}E^{p_0}[((m-c)D(m) - (r-c)D(r))e^{-\alpha\tau}], \end{aligned} \tag{3}$$

where we used the fact that $E^{p_0}[\int_0^\tau e^{-\alpha t}\sigma\,dB_t] = E^{p_0}[\int_\tau^\infty e^{-\alpha t}\sigma\,dB_t] = 0$. Observe from (3) that the first term does not depend on τ; hence, it can be simplified as $\alpha^{-1}(r-c)E^{p_0}[D(r)] = \alpha^{-1}(r-c)[p_0 D_h(r) + (1-p_0)D_\ell(r)]$. Notice that the second

https://doi.org/10.1017/S0269964811000246 Published online by Cambridge University Press

term of (3) can be rewritten as $E^{p_0}[h(P_\tau; r, m)e^{-\alpha\tau}]$, where

$$h(p; r, m) \equiv \alpha^{-1}p[\mu_h(m) - \mu_h(r)] + \alpha^{-1}(1-p)[\mu_\ell(m) - \mu_\ell(r)]. \quad \textbf{(4)}$$

The function $h(p; r, m)$ can be interpreted as the expected discounted profit generated from reducing the price from r to m when the posterior probability P_t equals p.

It turns out that the price schedule (r, m) that maximizes the expected discounted profit $\pi_\tau(r, m)$ satisfies the rather obvious property stated in Lemma 1.

LEMMA 1: *The optimal price schedule* (r^*, m^*) *satisfies* $\underline{x} < m^* < r^* < \bar{x}$.

In view of Lemma 1, we restrict our attention to price schedules (r, m) that satisfy $\underline{x} < m < r < \bar{x}$, and we use this restriction throughout the article without further mention.

3.4. Optimality of a Threshold Rule

In this subsection, we sketch the proof of the optimality of a threshold rule; the details of our proof are provided in the Appendix. A heuristic proof for a related problem is given by Ryan and Lippman [23]. For the purpose of establishing the optimality of a threshold rule, we simplify the notation and keep the price schedule (r, m) fixed. Let $\Pi(\tau)$ denote the firm's discounted profit when stopping at time τ.[8] Let $\pi^{p_0}(\tau) = E^{p_0}[\Pi(\tau)]$ denote its expected value. Application of the double expectation theorem produces

$$\begin{aligned}\pi^{p_0}(\tau) = E^{p_0}[\Pi(\tau)] = E^{p_0}\{E[\Pi(\tau)|D]\} &= p_0 E[\Pi(\tau)|D = D_h] \\ &+ (1-p_0)E[\Pi(\tau)|D = D_\ell].\end{aligned}$$

If the state were high, then the expected earnings per unit time would be $\mu_h(r)$ until time τ and $\mu_h(m)$ thereafter. Consequently,

$$\begin{aligned}E[\Pi(\tau)|D = D_h] &= \mu_h(r)E[(1 - e^{-\alpha\tau})|D = D_h]/\alpha + \mu_h(m)E[e^{-\alpha\tau}|D = D_h]/\alpha \\ &= [\mu_h(r) - \mu_h(m)]E[(1 - e^{\alpha\tau})|D = D_h]/\alpha + \mu_h(m)/\alpha.\end{aligned}$$

Combining these last two results and using the symmetric calculations for $D = D_\ell$, we obtain

$$\begin{aligned}\pi^{p_0}(\tau) = E^{p_0}[\Pi(\tau)] &= p_0[\mu_h(r) - \mu_h(m)]E[(1 - e^{-\alpha\tau})|D = D_h]/\alpha + p_0\mu_h(m)/\alpha \\ &+ (1-p_0)[\mu_\ell(r) - \mu_\ell(m)]E[(1 - e^{-\alpha\tau})|D = D_\ell]/\alpha + (1-p_0)\mu_\ell(m)/\alpha.\end{aligned}$$

[8] The firm's discounted profit has two sources of random variation: the demand state and the noise **B** in sales. The latter leads to uncertainty in the markdown time. Less important is the contribution $I = \sigma \int_0^\infty \exp[-\alpha s]\, dB_s$ to discounted profit; I has mean 0. We take the expectation over this mean 0 term to (*i*) eliminate this source of variation from $\Pi(\tau)$ and (*ii*) match our verbal description of $\Pi(\tau)$.

https://doi.org/10.1017/S0269964811000246 Published online by Cambridge University Press

Define $R(\tau,p)$ by $R(\tau,p) = \pi^p(\tau) - p\mu_h(m)/\alpha - (1-p)\mu_\ell(m)/\alpha$ so that

$$R(\tau,p) = p[\mu_h(r) - \mu_h(m)]E[(1-e^{-\alpha\tau})|D = D_h]/\alpha + (1-p)[\mu_\ell(r) - \mu_\ell(m)]E[(1-e^{-\alpha\tau})|D = D_\ell]/\alpha. \quad \textbf{(5)}$$

Let $\hat{\tau}$ be the stopping time that always stops at time 0, then $\hat{\tau} \equiv 0$. Applying (5) produces $R(\hat{\tau},p) = 0$ so that $\pi^p(\hat{\tau}) = p\mu_h(m)/\alpha + (1-p)\mu_\ell(m)/\alpha$. Of course, $\pi^p(\hat{\tau})$ is the expected discounted return when $P_0 = p$ and the firm is constrained to drop its price from r to m at time 0. In effect, $\pi^p(\hat{\tau})$ is the expected discounted return when the firm does not have the option to drop its price but must instead always sell at price m. Next, we define the important function $V(\cdot)$ by $V(p) \equiv \sup\{R(\tau,p) : \tau \in \mathcal{T}\}$, where $\mathcal{T}$ is the set of all stopping times.

Remark: $V(\cdot)$ is the option value. We briefly digress from our goal of establishing the optimality of a threshold rule to demonstrate that $V(\cdot)$ can be interpreted as the value of the option to mark down the price at a time τ chosen by the manufacturer. Let τ^* denote the optimal stopping time; the existence of an optimal stopping time is established in Proposition 1. Observe that

$$\begin{aligned}\pi^p(\tau^*) &= \sup\{\pi^p(\tau) : \tau \in \mathcal{T}\}\\ &= \sup\{R(\tau,p) + p\mu_h(m)/\alpha + (1-p)\mu_\ell(m)/\alpha : \tau \in \mathcal{T}\}\\ &= V(p) + [p\mu_h(m)/\alpha + (1-p)\mu_\ell(m)/\alpha]\\ &= V(p) + \pi^p(\hat{\tau}),\end{aligned}$$

where the penultimate equality follows because $p\mu_h(m)/\alpha + (1-p)\mu_\ell(m)/\alpha$ does not depend on the stopping time selected. Thus, $V(p) + \pi^p(\hat{\tau})$ is the optimal expected discounted return, and $V(p)$ optimizes the expected discounted return. Furthermore, because

$$V(p) = \pi^p(\tau^*) - \pi^p(\hat{\tau}),$$

it is clear that $V(p)$ can be interpreted as the value of the option to reduce the sales price when the price schedule (r,m) is given and the initial probability P_0 is p.

From the demand assumptions, $\mu_h(\cdot)$ is increasing and $\mu_\ell(\cdot)$ is decreasing on the interval $[\underline{x},\bar{x}]$. Consequently, the fact that $1 - \exp[-\alpha\tau]$ is nonnegative and the representation for $R(\tau,p_0)$ given in (5) establishes the following lemma.

LEMMA 2: *Given any stopping time τ, $R(\tau,p_0)$ is nondecreasing in p_0.*

For the moment, assume that there exists an optimal stopping time τ that maximizes $R(\tau,p_0)$. As shown by Oksendal [19, p. 220], when maximizing the expected value of a function of the form $g(Y_\tau e^{-\alpha\tau})$, where Y is a diffusion, there is an (open) set D of states such that it is optimal to stop the first time Y_t leaves D. In our case, $Y_t = P_t$ and $g(p) = \mu_L + p(\mu_H - \mu_L)$, where $\mu_L = [\mu_\ell(r) - \mu_\ell(m)]/\alpha$ and

https://doi.org/10.1017/S0269964811000246 Published online by Cambridge University Press

$\mu_H = [\mu_h(r) - \mu_h(m)]/\alpha$; that is, the optimal stopping time is stationary. This is in accord with intuition: Because **P** is time homogeneous by (1), the decision to stop and mark down the price should be a function of the posterior probability but not a function of time. We claim that D is an interval of the form $(q, 1]$, from which it is optimal to stop the first time P_t falls to q.

By continuity of $V(\cdot)$, there is a greatest number $q \leq 1$ such that $V(q) = 0$. Observe that $V(p) = \sup\{R(\tau, p)\} \geq R(\hat{\tau}, p) = 0$: $V(\cdot)$ is nonnegative. Additionally, Lemma 2 asserts that $V(\cdot)$ is nondecreasing. In light of these three facts, it follows that $V(p) \leq 0$ for $p \leq q$ and $V(p) > 0$ for $p > q$. This verifies our claim that $D = (q, 1]$ for some q; that is, a threshold rule is optimal so that

$$\tau^* = \inf\{t > 0 : P_t \leq p^*\}. \tag{6}$$

Define the stopping time τ_p by $\tau_p = \inf\{t > 0 : P_t \leq p\}$, $0 \leq p \leq 1$, and for convenience identify τ_p with p and let $\pi(p; r, m) \equiv \pi_{\tau_p}(r, m)$. From (6) we know that $\sup\{\pi(p; r, m) : 0 \leq p \leq 1\} = \sup\{\pi_\tau(r, m) : \tau \in \mathcal{T}\}$.

Applying the approach presented in Ryan and Lippman [23] to (3), we obtain an elegant formula for the firm's profit when it uses the stopping time τ_p:

$$\pi(p; r, m) = \alpha^{-1}[p_0\mu_h(r) + (1 - p_0)\mu_\ell(r)] + E^{p_0}[e^{-\alpha\tau_p}]h(p; r, m), \tag{7}$$

where the expected discount factor satisfies

$$E^{p_0}[e^{-\alpha\tau_p}] = \left(\frac{p}{p_0}\right)^{(\gamma-1)/2} \left(\frac{1 - p_0}{1 - p}\right)^{(\gamma+1)/2} \tag{8}$$

and (well-known results from Brownian motion assist in obtaining the formulas)

$$\gamma = \sqrt{1 + 8\alpha/\beta^2(r)}, \tag{9}$$

$$\beta(r) = [D_h(r) - D_\ell(r)]/\sigma. \tag{10}$$

A few simple observations will aid our understanding of the expected profit function $\pi(p; r, m)$ given in (7). To begin, notice that the first term in the right-hand side of (7) is simply the firm's discounted expected profit were it to set $p = 0$ and never mark down the price. To better understand the second term, which has two factors, we note that many textbooks (see, e.g., Borodin and Salminen [II.10) on Brownian motion provide the Laplace transform (8) of the first time a Brownian motion hits a line; in our case, the first time P_{τ_p} equals p is the first time **C** hits a line with slope $[D_h(r) + D_\ell(r)]/2$. Next, $h(p; r, m)$, defined in (4), is the expected discounted value of the change in profits when the price is marked down at time τ_p, discounted back to time τ_p. The prefactor discounts h back to time 0.

Because $\mu_h(x)$ is increasing and $\mu_\ell(x)$ is decreasing on the interval $[m, r]$, we see from (4) that $h(p; r, m) < 0$ for p large. Hence, $h(p; r, m)$, the net expected discounted (back to time τ_p) increase in profit associated with reducing the price from r to m, can be negative if the threshold p is too large: If the price reduction is made too early, the

https://doi.org/10.1017/S0269964811000246 Published online by Cambridge University Press

total expected discounted profit can decrease. The lesson here is that if the threshold is not chosen carefully, then demand learning can yield a lower expected profit; finding an optimal threshold policy is a managerial imperative.

From (7) and the first-order condition $\partial\pi(r,m;p)/\partial p = 0$, we obtain

$$p^* = \frac{[\mu_\ell(m) - \mu_\ell(r)](\gamma - 1)}{[\mu_h(r) - \mu_h(m)](\gamma + 1) + [\mu_\ell(m) - \mu_\ell(r)](\gamma - 1)} \equiv \phi(r,m;\gamma). \quad \textbf{(11)}$$

Finally, we need to justify our initial assumption that there exists an optimal policy. This final step can be implemented by verifying that $\pi(p;\ r,m)$ given in (3) with $p = p^*$ given in (11) satisfies a number of conditions stipulated by Theorem 10.4.1 of Oksendal [25]. The details of the proof are presented in the Appendix.

PROPOSITION 1: *Given the price schedule (r,m), τ_{p^*} is the optimal stopping time, and p^* is given in (*11*), mark down the price from r to m when the posterior probability hits p^*.*

As per Proposition 1, if the prior probability p_0 is less than p^*, then $\tau_{p^*} = 0$: The firm never sells at price r, and this price schedule is suboptimal. To rid ourselves of this trivial case, henceforth assume $p_0 > p^*$.

PROPOSITION 2: *The probability that P_t eventually hits p^* is $(1-p_0)/(1-p^*)$: the firm eventually marks down its price to m with probability $(1-p_0)/(1-p^*)$. More generally, if the firm were to mark down its price at time τ_p for some $p < p_0$, then $P(\tau_p < \infty) = (1-p_0)/(1-p)$.*

Whereas the proof of Proposition 2 requires the use of stochastic calculus, its meaning is straightforward. It tells us that if the firm were to use threshold $p < p_0$, then there is a positive probability (equal to $(p_0 - p)/(1-p)$) that the firm never marks down its price. In particular, if σ decreases to zero, then γ decreases to 1 by (9) and (10), from which (11) makes clear that p^* decreases to zero. Thus, Proposition 2 reveals that

$$\lim_{\sigma \to 0} P(\tau_{p^*} = \infty) = 1 - p_0.$$

This is good news: As the information becomes better and better, the probability that the firm keeps its price at r increases to $1 - p_0$, the probability that demand is in the high state (in which case, the firm "should" keep its price at r). Presumably, as σ decreases to zero, r^* increases to $\bar{x}$ and m^* decreases to $\underline{x}$. In other words, the decision maker can immediately tell whether the demand state is high or low, so it is optimal to start with $r = \bar{x}$ and then immediately change the price to $m = \underline{x}$ only if the state is low. This intuition holds, and these last results are verified in the comments immediately following Proposition 8.

3.5. Comparative Statics

Of course, the time τ_p until markdown decreases (with probability 1) as p increases. Whether this is true for p^* itself depends on what change in parameter induces p^* to

https://doi.org/10.1017/S0269964811000246 Published online by Cambridge University Press

increase. The closed-form expression (11) for p^* greatly aids our comparative statics analysis.

PROPOSITION 3: *The optimal threshold p^* increases in σ and α and it is independent of p_0.*

It is clear from (11) that p^* is independent of p_0; furthermore, γ is increasing in σ, from which (11) reveals that p^* increases in σ. This comports with intuition: As the demand volatility σ increases, the demand signal C_t becomes noisier so learning slows down and the return generated from the learning process decreases. Consequently, the value of the return generated from an earlier markdown becomes relatively more attractive; the earlier markdown is achieved by an increase in the threshold. Finally, as the discount rate α increases, the value of the future quickly decreases so it behooves the firm to make a quick decision in order to gain from its option to mark down the price.

PROPOSITION 4: *The optimal profit function $\pi(p^*;\ r, m)$ is nonincreasing in σ and in α.*

Understanding Proposition 4 is straightforward. Learning slows down as σ increases, so more time is required to attain the same level p^* of accuracy as earlier. This delay in the time of markdown reduces profits. Because γ increases in α, (7) shows that for each stopping time τ_p, the expected discounted return decreases.

At first glance, the comparative statics results of Propositions 3 and 4 appear to be in stark contrast to standard results for conventional real-options models wherein the exit threshold decreases and the firm's expected return increases in the volatility of the observable profit rate process (see Dixit [8] and Alvarez [1]). In our model, however, we see from (1) that σ is inversely proportional to the volatility of the observable process P_t. Note that the volatility of P_t is $(D_h(r) - D_\ell(r))/\sigma$, the coefficient of $d\tilde{B}_t$ in (1). Focusing on the volatility of P_t (not on σ), we see that Propositions 3 and 4 comport with the standard results of real-options theory.[9]

3.6. Value of Demand Learning

We briefly examine the value of demand learning by considering the case when the demand curves are linear—that is, $D_h(x) = (1 - a_h x)$ and $D_\ell(x) = (1 - a_\ell x)$ with $a_\ell > a_h > 0$. First, suppose that the firm does not engage in learning. Then, without updated information about demand, there is no reason for the firm to change its initial price r. The firm's expected profit is $\pi_0(r) \equiv \alpha^{-1}(r - c)[p_0(1 - a_h r) + (1 - p_0)(1 - a_\ell r)]$. The regular price r_0 that maximizes $\pi_0(r)$ solves the first-order

[9] For completeness, we note that whereas σ has no *direct* impact on the expected profit stream itself, σ has an indirect impact on profits, in that it impacts the time when the price changes from r to m.

https://doi.org/10.1017/S0269964811000246 Published online by Cambridge University Press

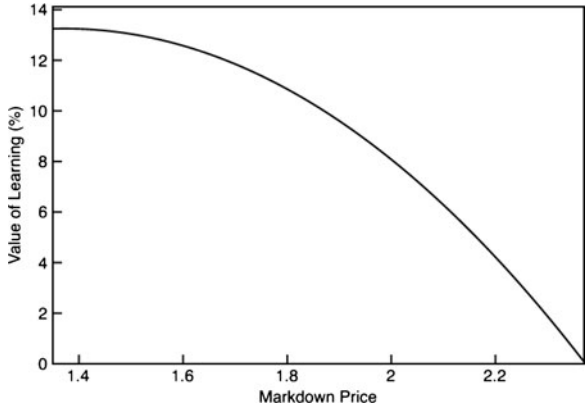


FIGURE 1. Relative profit gain obtained from demand learning.

condition $\partial \pi_0(r)/\partial r = 0$:

$$r_0 = \frac{1 + c(p_0 a_h + (1 - p_0)a_\ell)}{2(p_0 a_h + (1 - p_0)a_\ell)}. \tag{12}$$

Therefore, in the absence of learning, the optimal expected discounted profit is given by $\pi_0(r_0)$. We use this expected discounted profit as our benchmark.

One way to investigate the value of learning is to fix p_0 and use the regular price r_0 given in (12). However, for any given markdown price m with $m < r_0$, the firm engages in demand learning and follows the corresponding optimal threshold policy $p^* = \phi(r_0, m; \gamma)$ given in (11). The corresponding expected discounted profit with demand learning and price schedule (r_0, m) is given by $\pi(p^*;\, r_0, m)$.

We use $[\pi(p^*;\, r_0, m) - \pi_0(r_0)]/\pi_0(r_0)$, the percentage gain in profit, to measure the value of demand learning. We now construct a numerical example to illustrate the value of demand learning. Set $a_h = 0.1$, $a_\ell = 0.4$, $c = 0.2$, $\sigma = 0.1$, and $p_0 = 0.6$ so that $\underline{x} = 1.35$, $\bar{x} = 5.1$, and $r_0 = 2.37$. Applying Lemma 1, it suffices to evaluate the term $[\pi(p^*; r_0, m) - \pi_0(r_0)]/\pi_0(r_0)$ by focusing on a markdown price m that satisfies $\underline{x} = 1.35 < m < 2.37 = r_0$. Of course, demand learning is always beneficial; see Figure 1 for a demonstration. In our numerical example, demand learning produces a 13.3% in profit gain with $m \approx 1.37$. This demonstrates that the value of demand learning can be substantial and that choosing the optimal markdown price is critical.

4. GENERAL MODEL: PRICE SCHEDULE AND MARKDOWN TIME

Building upon the closed-form solutions presented in Section 3, we now determine the optimal price schedule (r^*, m^*) and the optimal markdown time τ^*. First, given a regular price r, we find the optimal markdown price $m^*(r)$ and the corresponding optimal

https://doi.org/10.1017/S0269964811000246 Published online by Cambridge University Press

threshold $p^*(r)$. Next, we determine the optimal regular price r^*, the corresponding optimal markdown price $m^*(r^*)$, and the corresponding optimal threshold $p^*(r^*)$.

4.1. Optimal Markdown Price

Given r, we use (11) to determine the optimal threshold $p^* = \phi(r, m; \gamma)$ as a function of m. Differentiating the return function $\pi(\phi(r, m; \gamma);\ r, m)$ given in (7) yields the first-order condition for $m^*(r)$.

LEMMA 3: *Given regular price r, the markdown price $m^*(r)$ that maximizes the return function $\pi(\phi(r, m; \gamma);\ r, m)$ is the solution to*

$$\phi(r, m; \gamma)\mu_h'(m) + [1 - \phi(r, m; \gamma)]\mu_\ell'(m) = 0. \tag{13}$$

With the price schedule $(r, m^*(r))$ given, we use (11) to determine the optimal threshold $p^*(r) \equiv \phi(r, m^*(r); \gamma)$. We now implement this approach for the case of linear demand.

LEMMA 4: *Suppose $D_h(r) = (1 - a_h r)$, $D_\ell(r) = (1 - a_\ell r)$, and $a_\ell > a_h > 0$. Then the optimal threshold $p^* \equiv \phi(r, m; \gamma)$ is increasing in m. Moreover, p^* is increasing in r if σ is sufficiently large.*

The first part of Lemma 4 is clear: When m increases, the cost of a premature markdown is reduced. The second part is far from obvious. First, by inspection, the function $\phi(\cdot, \cdot; \gamma)$ is symmetric with respect to m and r so that $\phi(\cdot, m; \gamma)$ is increasing just as $\phi(r, \cdot; \gamma)$ is increasing. This is so because the loss from overpricing is higher when r is larger, so it is better to mark down the price earlier. However, γ has further r dependence. Hence,

$$\frac{dp^*}{dr} = \frac{\partial}{\partial r}\phi(r, m; \gamma) + \frac{d\gamma}{dr}\frac{\partial}{\partial \gamma}\phi(r, m; \gamma).$$

Although the first term is positive ($\partial\phi(r, m; \gamma)/\partial r > 0$), the second term is negative because $d\gamma/dr < 0$ from the definition of γ and $\partial\phi(r, m; \gamma)/\partial\gamma > 0$ by Proposition 3. Hence, the comparative statics of p^* is determined by the trade-off between the regular price effect (just as in the first statement of Lemma 4) and the learning effect through the dependence on γ; the rate of learning increases in r, which tends to lower the markdown threshold. It turns out that the price effect is dominant when σ is sufficiently large.

PROPOSITION 5: *Given r, both $m^*(r)$ and $p^*(r)$ are increasing in σ wherever they are continuous in σ, and they are independent of p_0. Moreover, if the demand curves are linear, then $m^*(r)$ and $p^*(r)$ are increasing in r for sufficiently large σ.*

The first statement of Proposition 5 has an intuitive explanation. As σ increases, learning slows down thereby reducing the value of learning. This provides an

https://doi.org/10.1017/S0269964811000246 Published online by Cambridge University Press

inducement to hasten the markdown decision (i.e., increase the threshold $p^*(r)$). The second statement of Proposition 5 is due to Lemma 4: For sufficiently large σ, the effect of an increase in r through the change in γ (speed of learning) is negligible, so we only need to consider the price effect. If the regular (initial) price is higher, then the firm is compelled to hasten the markdown decision because the loss incurred from being in the lower state is higher. Hence, both the markdown price and the threshold are higher.

4.2. Optimal Regular Price

Using the optimal markdown price $m^*(r)$ given in (13) and the optimal threshold $p^*(r) \equiv \phi(r, m^*(r); \gamma)$, we use (7) to determine the expected return function $\pi(r) \equiv \pi(p^*(r);\ r, m^*(r))$. The proof of Proposition 6 reveals that r^* resides in the interior of the interval $(\underline{x}, \bar{x})$ from which r^* satisfies the first-order condition $d\pi(r)/dr = \partial\pi(p^*(r);\ r, m^*(r))/\partial r = 0$. Recall from Section 3.6 that when the demand curves are linear and when the firm does not engage in learning, the regular price r_0 given in (12) is the optimal regular price. When the demand curves are linear, a simple comparison of the first-order condition for r^* with the first-order condition for r_0 reveals that $r^* > r_0$.

PROPOSITION 6: *When the demand curves are linear,* $r^* > r_0$.

The veracity of Proposition 6 seems clear. If a markdown is to transpire, then surely the regular price r^* will be greater than r_0 and the markdown price $m^*(r^*)$ will be less than r_0.

4.3. Comparative Statics

In this subsection, we examine the comparative statics of the optimal policy $(p^*;\ r^*, m^*)$ and the optimal return $\pi(p^*;\ r^*, m^*)$ with respect to σ. Note that r^* that satisfies the first-order condition $\partial\pi(p^*(r); r, m^*(r))/\partial r = 0$, which is given in (A.3) in the Appendix; $m^*(r^*)$ is retrieved from (13) by setting $r = r^*$, and $p^* = \phi(r^*, m^*; \gamma)$ is retrieved from (11) by setting $r = r^*$ and $m = m^*$. The optimal return function $\pi(p^*;\ r^*, m^*)$ is retrieved from (7).

PROPOSITION 7: *The optimal expected return* $\pi(p^*;\ r^*, m^*)$ *is nonincreasing in* σ.

Proposition 7 is obvious: An increase in σ retards learning so the value of the option to mark down the price is reduced. This is consistent with Proposition 4 as presented in Section 3.5.

Because closed-form solutions do not exist for the optimal return or optimal threshold, we are unable to effect a general comparative statics analysis. Propositions 8 and 9 contain results for σ small and large, respectively. In particular, when σ increases without bound, r^* decreases to r_0 and $m^*(r^*)$ increases to r_0, where r_0 is given in (12).

https://doi.org/10.1017/S0269964811000246 Published online by Cambridge University Press

PROPOSITION 8: *For sufficiently small σ, the optimal regular price r^* is decreasing in σ whereas the optimal threshold $p^*(r^*)$ and the optimal markdown price $m^*(r^*)$ are increasing in σ.*

Clearly, if $\sigma = 0$, then the firm learns the true state instantly so that $r^* = \bar{x}$, $m^* = \underline{x}$, and the expected return is $[p_0\mu_h(\bar{x}) + (1 - p_0)\mu_\ell(\underline{x})]/\alpha$. By continuity, this suggests that as σ falls to zero, r^* must increase to $\bar{x}$ and m^* must fall to $\underline{x}$. Our intuition falters, however, regarding changes in the optimal threshold. As σ moves away from zero, perhaps the firm becomes more hesitant to mark down its price, for which $p^*(r^*)$ falls. Of course, the beauty of mathematics is that it provides the correct answers, regardless of our initial intuition. When σ is small, the firm learns about the state of demand very quickly. Hence, it behooves the firm to wait until receiving a strong signal that the state of demand is low before marking down the price. According to this intuition, if σ is made even smaller, then the firm will wait for an even stronger signal that the state of demand is low before making the markdown decision. Thus, this intuition asserts that the threshold becomes smaller as σ decreases; this agrees with Proposition 8.

PROPOSITION 9: *For sufficiently large σ, the optimal regular price r^* is decreasing in σ, whereas the optimal threshold $p^*(r^*)$ and the optimal markdown price $m^*(r^*)$ are increasing in σ. Moreover, $r^* \searrow r_0$, $m^* \nearrow r_0$, and $p^* \nearrow p_0$ as $\sigma \to \infty$, where r_0 satisfies (*12*).*

When σ is very large, there is little learning: P_t remains near p_0 unless t is very large; that is, when σ is large, the firm's option to mark down from its regular price is practically worthless, and, in any event, it will not happen for a long time. Thus, the initial price drops to r_0 as $\sigma \to \infty$. Similarly, we anticipate that $m^*(r^*)$ will increase to r_0 as σ gets large because when the price is marked down, P_t still will be close to p_0 so the markdown price should be close to the regular price.

In both limiting cases reported in Propositions 8 and 9, the markdown price $m^*(r^*)$ and the threshold $p^*(r^*)$ are increasing in σ while the regular price r^* is decreasing in σ. This result is far from obvious because there are countervailing forces. Proposition 5 tells us that given a regular price r, $m^*(r)$ and $p^*(r)$ are increasing in σ; in addition, they are increasing in r if the demand curves are linear and if σ is sufficiently large. However, as σ increases, the optimal regular price r^* decreases (per Propositions 8 and 9). The decrease in r^* generates the opposite effect on $m^*(r)$ and $p^*(r)$. This countervailing force is part of the explanation of why we are able to find analytical results only in the two extreme cases.

Finally, we have a rather intuitive result regarding the comparative statics of the optimal price schedule with respect to the initial prior p_0: A more optimistic decision maker will set a higher price.

COROLLARY 1: *When the demand curves are linear, both r^* and $m^*(r^*)$ increase in p_0 for sufficiently small and large values of σ.*

https://doi.org/10.1017/S0269964811000246 Published online by Cambridge University Press

5. SUMMARY

In this article we studied new product pricing with Bayesian learning about the state of demand. The firm can markdown the price only once, so that our problem has three decisions: the initial price, the markdown price, and the markdown time. Given a regular price and a markdown price, this problem reduces to a stopping time problem: stop and mark down the price the first time the posterior probability that demand is in the high state falls to a threshold. Using the plethora of results for Brownian motions, we exhibit an explicit closed-form expression for the optimal threshold; this leads to an explicit expression for the optimal expected discounted return. Using the closed form-expression for the optimal threshold, we carry out a comparative statics analysis with respect to σ, the demand volatility. In addition, we evaluated the value of demand learning.

The analysis becomes considerably more complicated when all three decision variables, not just the stopping time, are subject to change. In particular, in seeking the optimal values for the three decision variables, no closed-form solutions can be found. Nevertheless, we effected a nontrivial comparative statics analysis for the case when demand volatility σ is very small or very large. Under mild conditions, we showed that the optimal markdown price increases in the initial price and in σ. However, the optimal initial price decreases in σ. Hence, there are opposing forces weighing on changes in the optimal markdown price with respect to σ: A priori the direction of change is unclear. Our analysis, however, is definitive: For small and for large values of σ, the optimal regular price is decreasing in σ and the optimal markdown price is increasing in σ.

Apart from permitting multiple price reductions, there are a great many ways to generalize our model. For instance, we could incorporate a lump sum cost $k > 0$ each time the firm elects to mark down its price from r to m. Because the cost k will be incurred at time τ, this markdown cost adds $-ke^{-\alpha\tau} = -k\int_{\tau}^{\infty} \alpha e^{-\alpha s}\, ds$ to $\pi(\tau; r, m)$. Consequently, we merely need to add $-\alpha k$ to both $\mu_h(m)$ and $\mu_\ell(m)$. Inspection of (11) makes clear that an increase (from zero) in k serves to lower p^* and reduces the expected profit.[10] This is intuitive: As the cost of implementation increases, the firm seeks to delay this cost.

References

1. Alvarez, L.H.R. (2003). On the properties of r-excessive mappings for a class of diffusions. *Annals of Applied Probability* 13: 1517–1533.
2. Balvers, R.J. & Cosimano, T.F. (1990). Actively learning about demand and the dynamics of price adjustment. *The Economic Journal* 100: 882–898.

[10] The optimal threshold p_k^* satisfies

$$p_k^* = \frac{[\mu_\ell(m) - \mu_\ell(r) - k\alpha](\gamma - 1)}{[\mu_h(r) - \mu_h(m) + k\alpha](\gamma + 1) + [\mu_\ell(m) - \mu_\ell(r) - k\alpha](\gamma - 1)}.$$

https://doi.org/10.1017/S0269964811000246 Published online by Cambridge University Press

3. Bergemann, D. & Valimaki, J. (2000). Experimentation in markets. *Review of Economic Studies* 67: 213–234.
4. Bernardo, A.E. & Chowdhry, B. (2002). Resources, real options, and corporate strategy. *Journal of Financial Economics* 63: 211–234.
5. Bolton, P. & Harris, C. (1999). Strategic experimentation. *Econometrica* 67: 349–374.
6. Borodin, A.N. & Salminen, P. (2002). *Handbook of brownian motion: Facts and formulae*, 2nd ed. Basel: Birkhauser.
7. Decamps, J.-P., Mariotti, T., & Villeneuve, S. (2005). Investment timing under incomplete information. *Mathematics of Operations Research* 30: 472–500.
8. Dixit, A. (1992). Investment and hysteresis. *The Journal of Economic Perspectives* 6: 107–132.
9. Elmaghraby, W., Lippman, S.A., Tang, C.S., & Yin, R. (2008). Will more purchasing options benefit customers? *Production and Operations Management* 18: 381–401.
10. Federgruen, A. & Heching, A. (1999). Combined pricing and inventory control under uncertainty. *Operations Research* 47: 454–475.
11. Feng, Y. & Gallego, G. (1995). Optimal starting times for end-of-season sales and optimal stopping times for promotional fares. *Management Science* 41: 1371–1391.
12. Fisher, M. (2007). Personal communication.
13. Hafner, K. & Stone, B. (2007). iPhone owners crying foul over price cut. *New York Times*, September 7.
14. Keller, G. & Rady, S. (1999). Optimal experimentation in a changing environment. *Review of Economic Studies* 66: 475–507.
15. Lin, K.Y. (2006). Dynamic pricing with real-time demand learning. *European Journal of Operational Research* 174: 522–538.
16. Liptser, R.S. & Shiryaev, A.N. (1974). *Statistics of random processes I. General theory*, 2nd ed. New York: Springer.
17. Markoff, J. (2007). Apple cuts iPhone price ahead of holidays. *New York Times*, September 6.
18. Moscarini, G. & Smith, L. (2001). The optimal level of experimentation. *Econometrica* 69: 1629–1644.
19. Oksendal, B. (2003). *Stochastic differential equations: An introduction with applications*, 6th ed. New York: Springer.
20. Peskir, G. & Shiryaev, A. (2006). *Optimal stopping and free-boundary problems*. Basel: Birkhauser.
21. Petruzzi, N.C. & Dada, M. (2002). Dynamic pricing and inventory control with learning. *Naval Research Logistics* 49: 303–325.
22. Ross, S. (1983). *Introduction to stochastic dynamic programming*. New York: Academic Press.
23. Ryan, R. & Lippman, S.A. (2003). Optimal exit from a project with noisy returns. *Probability in the Engineering and Informational Sciences* 17: 435–458.
24. Ryan, R. & Lippman, S.A. (2005). Optimal exit from a deteriorating project with noisy returns. *Probability in the Engineering and Informational Sciences* 19: 327–343.
25. Shiryaev, A.N. (1978). *Optimal stopping rules*. New York: Springer-Verlag.
26. Subrahmanyan, S. & Shoemaker, R. (1996). Developing optimal pricing and inventory policies for retailers who face uncertain demand. *Journal of Retailing* 72: 7–30.
27. Talluri, K.T. & van Ryzin, G.J. (2005). *The theory and practice of revenue management*. New York: Springer.
28. Yin, R. & Rajaram, K. (2007). Joint pricing and inventory control with a Makovian demand model. *European Journal of Operational Research* 182: 113–126.

APPENDIX

Proofs

PROOF OF LEMMA 1: For any $x \in (0, 1)$, consider the function $f(r) = (r - c)[xD_h(r) + (1 - x)D_\ell(r)]$. By Assumptions (i) and (iii), $df(r)/dr > 0$ for all $r < \underline{x}$ and $df(r)/dr < 0$ for all $r > \bar{x}$.

https://doi.org/10.1017/S0269964811000246 Published online by Cambridge University Press

Hence, $df(r)/dr = 0$ occurs only strictly within the interval $(\underline{x}, \bar{x})$. This result is applied to the first-order necessary conditions for r and m to maximize $\pi_\tau(r, m)$ in (3) for any given τ. ■

PROOF OF PROPOSITION 1: If there is an optimal policy for the stopping problem of (3), then the theory of optimal stopping stipulates that there is a set of necessary and sufficient conditions that the optimal return function $\sup_\tau \pi_\tau(r, m)$ satisfies. Hence, we will need to verify that there is a function that satisfies this set of conditions. A common technique in problems of this kind is to begin with a candidate function and verify that this candidate satisfies the set of conditions. Our candidate is the function $\pi(p^*; r, m)$ defined in the next few lines. To begin, define $h(p; r, m) = \alpha^{-1}p[\mu_h(m) - \mu_h(r)] + \alpha^{-1}(1-p)[\mu_\ell(m) - \mu_\ell(r)]$. The candidate function is

$$\pi(p^*; r, m) = \begin{cases} \alpha^{-1}[p_0\mu_h(r) + (1-p_0)\mu_\ell(r)] + h(p_0; r, m) & \text{if } p_0 \le p^* \\ \alpha^{-1}[p_0\mu_h(r) + (1-p_0)\mu_\ell(r)] & \\ \quad + \left(\dfrac{p^*}{p_0}\right)^{(\gamma-1)/2} \left(\dfrac{1-p_0}{1-p^*}\right)^{(\gamma+1)/2} h(p^*; r, m) & \text{if } p_0 > p^*, \end{cases} \tag{A.1}$$

where p^* is defined in (11). Equation (A.1) is the return function of a stopping time τ_{p^*} (Ryan and Lippman, [23]), so τ_{p^*} is the optimal stopping time if (A.1) turns out to be the optimal return function. In the right-hand side of (A.1), only the second term depends on the stopping policy, so we can safely study only the second term:

$$f(x) \equiv \begin{cases} \left(\dfrac{p^*}{x}\right)^{(\gamma-1)/2} \left(\dfrac{1-x}{1-p^*}\right)^{(\gamma+1)/2} h(p^*; r, m) & \text{if } x > p^* \\ h(x; r, m) & \text{otherwise .} \end{cases}$$

Define the α-excessive infinitesimal generator for the process P_t:

$$\mathcal{A} \equiv -\alpha + \frac{1}{2}\beta^2 x^2 (1-x)^2 \frac{\partial^2}{\partial x^2} ,$$

where $\beta = [D_h(r) - D_\ell(r)]/\sigma$. By Theorem 10.4.1 of Oksendal [19], $f(x)$ is the optimal solution if the following hold:

(i) $\mathcal{A}f(x) = 0$ for $x > p^*$ and $\mathcal{A}f(x) \le 0$ for $x \le p^*$;

(ii) $f(\cdot)$ is continuously differentiable in the interval $(0, 1)$;

(iii) $f(x) \ge h(x; r, m)$ for all $x \in (0, 1)$.

(i) $\mathcal{A}f(x) = 0$ for $x > p^*$ can be algebraically verified once we identify $\gamma = \sqrt{1 + 8\alpha/\beta^2}$. Moreover, $\mathcal{A}f(x) \le 0$ for $x \le p^*$ because $\mathcal{A}h(x; r, m) = -\alpha h(x; r, m) \le 0$ for all $x \le p^*$; this is because $h(x; r, m)$ decreases in x and $h(p^*; r, m) \ge 0$ (by the fact that p^* is an optimally chosen threshold).

(ii) It is clear that $f(\cdot)$ is continuously differentiable everywhere except perhaps at p^*. To see that $f(\cdot)$ is continuously differentiable at p^*, notice that

$$\lim_{x \searrow p^*} f(x) = \lim_{x \nearrow p^*} f(x),$$

$$\lim_{x \searrow p^*} \frac{df(x)}{dx} = \lim_{x \nearrow p^*} \frac{df(x)}{dx}$$

from the definition [(11)] of p^*.

https://doi.org/10.1017/S0269964811000246 Published online by Cambridge University Press

(iii) First, we establish that $h(p^*; r, m) > 0$. Here, p^* is chosen to maximize the function $\pi(p; r, m)$. Hence, p^* maximizes

$$\left(\frac{p}{p_0}\right)^{(\gamma-1)/2} \left(\frac{1-p_0}{1-p}\right)^{(\gamma+1)/2} h(p; r, m).$$

From the functional form of $h(p; r, m)$, we see that $h(p; r, m)$ can take strictly positive values if p is sufficiently low. Hence, the optimal value of

$$\left(\frac{p}{p_0}\right)^{(\gamma-1)/2} \left(\frac{1-p_0}{1-p}\right)^{(\gamma+1)/2} h(p; r, m)$$

is strictly positive. Therefore, at the optimal value p^*, $h(p^*; r, m)$ is also positive. Next, we show that $f(x)$ is strictly convex for $x > p^*$. From (i),

$$\begin{aligned}\frac{d^2 f(x)}{dx^2} &= \frac{\alpha}{\frac{1}{2}\beta^2 x^2 (1-x)^2} f(x) \\ &= \frac{\alpha}{\frac{1}{2}\beta^2 x^2 (1-x)^2} \left(\frac{p^*}{x}\right)^{(\gamma-1)/2} \left(\frac{1-x}{1-p^*}\right)^{(\gamma+1)/2} h(p^*; r, m)\end{aligned}$$

for $x > p^*$, so the inequality $d^2 f(x)/dx^2 > 0$ follows from the fact that $h(p^*; r, m) > 0$.

For $x \le p^*$, $f(x) = h(x; r, m)$. At $x = p^*$, $df(x)/dx = \partial h(x; r, m)/\partial x$, but $df(x)/dx$ is nondecreasing due to convexity, whereas $\partial h(x; r, m)/\partial x$ is constant for all x. It follows that $f(x) \ge h(x; r, m)$ for $x > p^*$. ■

Proof of Proposition 2: Proving this proposition requires the use of stochastic calculus. In particular, we exploit the martingale properties of selected functions of P_t and obtain the properties of P_t as $t \to \infty$.

First, consider $f_1(p) = p$. Obviously, $f_1(P_t)$ is a martingale because P_t is a martingale. Let τ_p be the first time (stopping time) that P_t hits the threshold p. Then $f_1(P_{\tau_p \wedge t}) \equiv M_t$ is also a martingale by Doob's stopping time theorem (Peskir and Shiryaev [20]). Because M_t is bounded, by the martingale convergence theorem, there exists a random variable M_∞ such that $\lim_{t\to\infty} M_t = M_\infty$ and $E^{p_0}[M_\infty] = M_0 \equiv p_0$. (Note that M_t never strictly hits M_∞.)

Next, note that the infinitesimal generator (Oksendal [19, Sect. 7.3] $\mathcal{L}$ of the process $\mathbf{P}$ is given by

$$\mathcal{L} = \frac{1}{2}\beta^2(r) p^2 (1-p)^2 \partial_p^2.$$

Consider $f_2(p) = p(1-p)$ so that $\mathcal{L} f_2(p) = -\beta^2(r) p^2 (1-p)^2$. Then by Theorem 8.3.1 of Oksendal [19], the following is a martingale:

$$N_t \equiv f_2(P_{\tau_p \wedge t}) + \int_0^{\tau_p \wedge t} \beta^2(r) P_s^2 (1-P_s)^2 \, ds.$$

Because this is a strictly positive martingale, the pointwise limit exists from the martingale convergence theorem, and the limit is finite almost surely. Hence, as $t \to \infty$, either $P_t \to 1$

https://doi.org/10.1017/S0269964811000246 Published online by Cambridge University Press

(in case $\tau_p = \infty$) or $\tau_p < \infty$. Therefore, $P^{p_0}[M_\infty = 1 \text{ or } M_\infty = p] = 1$. From $E^{p_0}[M_\infty] = p_0$, we have

$$E^{p_0}[M_\infty] = pP^{p_0}[M_\infty = p] + 1(1 - P^{p_0}[M_\infty = p]) = p_0,$$

which yields

$$P^{p_0}[M_\infty = p] = \frac{1 - p_0}{1 - p}.$$

We conclude that P_t hits the threshold p with the probability $(1 - p_0)/(1 - p)$, and P_t asymptotically approaches 1 with the probability $(p_0 - p)/(1 - p)$, which is strictly less than p_0. ■

PROOF OF PROPOSITION 3: It is clear that $\phi(\cdot, \cdot; \cdot)$ given in (11) is independent of p_0. By (9), γ increases in both σ and α. Moreover,

$$\frac{\partial}{\partial \gamma}\phi(r, m; \gamma) = \frac{2[\mu_\ell(m) - \mu_\ell(r)][\mu_h(r) - \mu_h(m)]}{\{[\mu_h(r) - \mu_h(m)](\gamma + 1) + [\mu_\ell(m) - \mu_\ell(r)](\gamma - 1)\}^2} > 0,$$

so the statements of the Proposition follow. ■

PROOF OF PROPOSITION 4: From (7) and (11) and the fact that $\partial\pi(p; r, m)/\partial p = 0$ at $p = p^*$,

$$\frac{d}{d\sigma}\pi(p^*; r, m) = \frac{1}{2}\frac{d\gamma}{d\sigma}\log\left(\frac{1 - p_0}{1 - p^*}\frac{p^*}{p_0}\right)\alpha^{-1}\left(\frac{1 - p_0}{1 - p^*}\frac{p^*}{p_0}\right)^{(\gamma-1)/2}\left(\frac{1 - p_0}{1 - p^*}\right)h(p^*; r, m).$$

We note that $h(p^*; r, m) > 0$ because $h(p; r, m)$ takes positive values for some $p \le p_0$; moreover, $p^* < p_0$, so

$$\frac{1 - p_0}{1 - p^*}\frac{p^*}{p_0} < 1.$$

Therefore, $d\pi(p^*; r, m)/d\sigma \le 0$. ■

PROOF OF LEMMA 3: From the fact that $\partial_p\pi(p^*; r, m) = 0$, where $p^* = \phi(r, m; \gamma)$, the total derivative of $\pi(\phi(r, m; \gamma); r, m)$ with respect to m (at fixed r) is

$$\begin{aligned}\frac{d\pi(\phi(r, m; \gamma); r, m)}{dm} &= \partial_m\pi(\phi(r, m; \gamma); r, m) \\ &= \alpha^{-1}\left(\frac{1 - p_0}{1 - p^*}\frac{p^*}{p_0}\right)^{(\gamma-1)/2}\left(\frac{1 - p_0}{1 - p^*}\right)[p^*\mu_h'(m) + (1 - p^*)\mu_\ell'(m)]\end{aligned}$$

and, hence, the first-order condition for the optimal m is given by (13). We note that $d\pi(\phi(r, m; \gamma); r, m)/dm$ is positive at $m = \underline{x}$ because $\mu_\ell'(\underline{x}) = 0$ and $\mu_h'(\underline{x}) > 0$. Moreover,

https://doi.org/10.1017/S0269964811000246 Published online by Cambridge University Press

from the value of p^* defined in the limit $m \to r$,

$$p^* = \lim_{m \to r} \phi(r, m; \gamma) = \frac{-\mu'_\ell(r)(\gamma - 1)}{\mu'_h(r)(\gamma + 1) - \mu'_\ell(r)(\gamma - 1)},$$

we have

$$\lim_{m \to r} \frac{d\pi(\phi(r, m; \gamma); r, m)}{dm}$$
$$= \alpha^{-1} \left(\frac{1 - p_0}{1 - p^*} \frac{p^*}{p_0} \right)^{(\gamma-1)/2} \left(\frac{1 - p_0}{1 - p^*} \right) \frac{2\mu'_\ell(r)\mu'_h(r)}{\mu'_h(r)(\gamma + 1) - \mu'_\ell(r)(\gamma - 1)} < 0.$$

Thus, $d\pi(\phi(r, m; \gamma); r, m)/dm$ is negative in the limit $m \to r$. We conclude that the optimal m resides in the interval $(\underline{x}, r)$ and satisfies the first-order condition (13). ■

Proof of Lemma 4: After straightforward algebra, the following is obtained:

$$\frac{\partial}{\partial m} \phi(r, m; \gamma) = \frac{2(\gamma^2 - 1)(a_\ell - a_h)(r - m)^2}{\{[\mu_h(r) - \mu_h(m)](\gamma + 1) + [\mu_\ell(m) - \mu_\ell(r)](\gamma - 1)\}^2} > 0$$

and

$$\frac{d}{dr} \phi(r, m; \gamma) = \frac{\partial}{\partial r} \phi(r, m; \gamma) + \frac{d\gamma}{dr} \frac{\partial}{\partial \gamma} \phi(r, m; \gamma).$$

After some algebra, we can show that $d\phi(r, m; \gamma)/dr > 0$ if and only if

$$\gamma r(a_\ell - a_h)(r - m)^2 - [\mu_\ell(m) - \mu_\ell(r)][\mu_h(r) - \mu_h(m)] > 0. \qquad \textbf{(A.2)}$$

Notice that γ increases in σ, so $d\phi(r, m; \gamma)/dr > 0$ if σ is sufficiently large. ■

Proof of Proposition 5: We study the dependence of $\pi(\phi(r, m; \gamma); r, m)$ on r and σ. For convenience, define $f(\sigma, m) \equiv \pi(\phi(r, m; \gamma); r, m)$. Noting that $\partial f(\sigma, m)/\partial m = 0$ at $m = m^*(r)$, we have

$$\frac{\partial^2}{\partial \sigma \partial m} f(\sigma, m)|_{m = m^*(r)}$$
$$= \alpha^{-1} \left(\frac{1 - p_0}{1 - p^*} \frac{p^*}{p_0} \right)^{(\gamma-1)/2} \left(\frac{1 - p_0}{1 - p^*} \right) \frac{\partial p^*}{\partial \sigma} [\mu'_h(m^*(r)) - \mu'_\ell(m^*(r))] > 0,$$

because $\partial p^*/\partial \sigma > 0$ by Proposition 3 and $\mu'_h(m) - \mu'_\ell(m) > 0$. So $f(\sigma, m)$ is locally supermodular in m and σ at $m = m^*$. Thus, from the property of supermodular functions, $m^*(r)$ increases in σ (Ross [22]) wherever $m^*(r)$ is continuous in σ.

In case of linear demand functions, we define $g(r, m) = \pi(\phi(r, m; \gamma); r, m)$ and

$$\frac{\partial^2}{\partial r \partial m} g(r, m)|_{m = m^*(r)} = \alpha^{-1} \left(\frac{1 - p_0}{1 - p^*} \frac{p^*}{p_0} \right)^{(\gamma-1)/2} \left(\frac{1 - p_0}{1 - p^*} \right) \frac{\partial p^*}{\partial r} [\mu'_h(m^*(r)) - \mu'_\ell(m^*(r))].$$

Here $\partial p^*/\partial r = d\phi(r, m; \gamma)/dr > 0$ for large enough σ (Lemma 4), so we have $\partial^2 g(r, m)|_{m = m^*(r)}/\partial r \partial m > 0$, so it is locally supermodular. Thus, $m^*(r)$ increases in r wherever $m^*(r)$ is continuous in r for large enough σ. ■

https://doi.org/10.1017/S0269964811000246 Published online by Cambridge University Press

PROOF OF PROPOSITION 6: First, we determine the equation for r^*. We note that r^* satisfies the first-order condition $d\pi(r^*)/dr = 0$ because r^* is in the interior of the interval $(\underline{x}, \bar{x})$ by Lemma 1. Next, the total derivative of $\pi(\cdot)$ with respect to r is

$$\frac{d\pi(r)}{dr} = \left(\frac{dp^*}{dr}\partial_p + \partial_r + \frac{d}{dr}[m^*(r)]\partial_{m^*} + \frac{d\gamma}{dr}\partial_\gamma \right) \pi(r, m^*; p^*).$$

By the definition of p^* and m^*, we have $\partial_p \pi(p^*; r, m^*) = 0 = \partial_{m^*}\pi(p^*; r, m^*)$. Hence, $d\pi(r)/dr = \partial_r \pi(p^*; r, m^*) + d\gamma\, \partial_\gamma \pi(p^*; r, m^*)/dr$, which leads to

$$\begin{aligned}
\frac{d\pi(r)}{dr} &= \alpha^{-1}[p_0\mu_h'(r) + (1-p_0)\mu_\ell'(r)] \\
&\quad - \alpha^{-1}\left(\frac{1-p_0}{1-p^*}\frac{p^*}{p_0}\right)^{(\gamma-1)/2}\left(\frac{1-p_0}{1-p^*}\right)\Bigg[\{p^*[\mu_h(r) - \mu_h(m^*)] \\
&\quad + (1-p^*)[\mu_\ell(r) - \mu_\ell(m^*)]\} \\
&\quad \times \frac{1}{2}\log\left(\frac{1-p_0}{1-p^*}\frac{p^*}{p_0}\right)\frac{d\gamma}{dr} + p^*\mu_h'(r) + (1-p^*)\mu_\ell'(r)\Bigg] \\
&= 0. \qquad \textbf{(A.3)}
\end{aligned}$$

Then we consider two cases: $p_0 \le p^*$ and $p_0 > p^*$. If $p_0 \le p^*$, then it is optimal to change the price at $t = 0$, so the initial price does not affect the return function. Hence, we can always take r larger than r_0. If $p_0 > p^*$, then the optimal final price m^* is less than r, so that $p^*\mu_h'(r) + (1-p^*)\mu_\ell'(r) < 0$. From the condition $d\pi(r^*)/dr = 0$ for r^* in (A.3), we find that $p_0\mu_h'(r^*) + (1-p_0)\mu_\ell'(r^*) < 0$, which implies that $r^* > r_0$ from the concavity of the function $p_0\mu_h(\cdot) + (1-p_0)\mu_\ell(\cdot)$ in the case of linear demand curves. ■

PROOF OF PROPOSITION 7: For the purpose of this proof, we write the expected return function as $\pi(p^*; r^*, m^*; \sigma)$ and the threshold probability as $\phi(r, m; \sigma)$ with the explicit σ dependence. By Proposition 4, $\pi(\phi(r, m; \sigma); r, m; \sigma)$ is nonincreasing in σ. For any pair $\sigma^1 < \sigma^2$, let r^j and m^j be the optimal initial and final prices when $\sigma = \sigma^j$ for $j \in \{1, 2\}$. Then we obtain the following inequalities:

$$\pi(\phi(r^2, m^2; \sigma); r^2, m^2; \sigma^2) \le \pi(\phi(r^2, m^2; \sigma^1); r^2, m^2; \sigma^1) \le \pi(\phi(r^1, m^1; \sigma^1); r^1, m^1; \sigma^1).$$

■

In the following proofs, we study the small-σ and large-σ limits. In order to keep track of leading-order terms of power expansions of σ or σ^{-1}, we introduce a notation to denote the subleading-order terms: We say that $f(x) = O(j(x))$ if $f(x)/j(x) < \infty$ as $x \to 0$, where $f(x)$ and $j(x)$ are functions of x.

PROOF OF PROPOSITION 8: The small-σ limit coincides with the large-β or small-$(\gamma - 1)$ limit:

$$\gamma = \sqrt{1 + 8\alpha/\beta^2} = 1 + 4\frac{\alpha}{\beta^2} + O\left(\frac{\alpha^2}{\beta^4}\right).$$

https://doi.org/10.1017/S0269964811000246 Published online by Cambridge University Press

For convenience, we introduce an alternative parameter

$$g(r) \equiv \frac{4\alpha\sigma^2}{[D_h(r) - D_\ell(r)]^2},$$

so that $g(\cdot) \downarrow 0$ as $\sigma \to 0$. From (11), we obtain

$$\phi(r, m; \gamma) = g(r)\frac{\mu_\ell(m) - \mu_\ell(r)}{2[\mu_h(r) - \mu_h(m)]} - g^2(r)\frac{[\mu_\ell(m) - \mu_\ell(r)][\mu_\ell(m) - \mu_\ell(r) + \mu_h(r) - \mu_h(m)]}{4[\mu_h(r) - \mu_h(m)]^2} + O(g^3). \quad \textbf{(A.4)}$$

First, we note that if $g(\cdot) = 0$, then $m^*(r) = \underline{x}$, because the underlying demand curve will be revealed instantaneously to the decision maker and so the price will be changed only if the demand is in the lower state. Thus, in the small-σ limit, for $s \in \{h, \ell\}$,

$$\mu_s(m^*) = \mu_s(\underline{x}) + (m^* - \underline{x})\mu_s'(\underline{x}) + O((m^* - \underline{x})^2).$$

Defining $\Delta_h \equiv \mu_h(r) - \mu_h(\underline{x}) > 0$ and $\Delta_\ell \equiv \mu_\ell(\underline{x}) - \mu_\ell(r) > 0$, we obtain from (13) and (A.4) the following:

$$m^*(r) = \underline{x} + g(r)\frac{\Delta_\ell}{2\Delta_h}\frac{\mu_h'(\underline{x})}{[-\mu_\ell''(\underline{x})]} + O(\sigma^4), \quad \textbf{(A.5)}$$

$$p^*(r) = g(r)\frac{\Delta_\ell}{2\Delta_h} + O(\sigma^4).$$

By (A.3), we can obtain the optimal choice of the initial price given the prior probability p_0:

$$r^* = \bar{x} - g(\bar{x})\log\left(\frac{1}{g(\bar{x})}\right)\frac{(1 - p_0)\mu_\ell'(\bar{x})}{2p_0\mu_h''(\bar{x})} + O(g(\bar{x})). \quad \textbf{(A.6)}$$

Hence, for sufficiently small values of σ, r^* decreases in σ. Substituting (A.6) into (A.5) and keeping the terms of order $O(\sigma^2)$, we find that $m^*(r^*)$ increases in σ. Similarly, $p^*(r^*)$ increases in σ.

Finally, we study the expected return

$$\pi(r^*, m^*; p^*) = \alpha^{-1}[p_0\mu_h(\bar{x}) + (1 - p_0)\mu_\ell(\underline{x})] + \alpha^{-1}[\mu_\ell(\bar{x}) - \mu_\ell(\underline{x})]\frac{1 - p_0}{2\bar{x}^2}g(\bar{x})\log\left(\frac{1}{g(\bar{x})}\right) + O(\sigma^2).$$

Hence, the expected return decreases in σ for sufficiently small σ. ■

PROOF OF PROPOSITION 9: Let $m_0 \equiv \lim_{\sigma\to\infty} m^*$ and $r_0 \equiv \lim_{\sigma\to\infty} r^*$. The large-$\sigma$ limit coincides with the large-γ limit because $\gamma = \sigma\sqrt{8\alpha}/(D_h(r) - D_\ell(r)) + O(1)$.

https://doi.org/10.1017/S0269964811000246 Published online by Cambridge University Press

Our goal is to find a solution (p^*, r^*, m^*) to the first-order conditions (13) and (A.3) along with (11) that maximizes $\pi(p^*; r^*, m^*)$. First, from (11), p^* satisfies

$$p^*[\mu_h(r^*) - \mu_h(m^*)] + (1 - p^*)[\mu_\ell(r^*) - \mu_\ell(m^*)] = O(\sigma^{-1}). \tag{A.7}$$

From (A.3), we find that the optimal initial price r^* satisfies

$$\frac{d}{dr}\pi(r^*) = \alpha^{-1}[p_0\mu_h'(r^*) + (1 - p_0)\mu_\ell'(r^*)] + k(\sigma, r^*) = 0,$$

where

$$k(\sigma, r^*) \equiv \frac{d}{dr}\pi(r^*) - \alpha^{-1}[p_0\mu_h'(r^*) + (1 - p_0)\mu_\ell'(r^*)] = O(\sigma^{-1})$$

because of (A.7). Hence, $r^* - r_0 = O(\sigma^{-1})$, where r_0 satisfies

$$p_0\mu_h'(r_0) + (1 - p_0)\mu_\ell'(r_0) = 0. \tag{A.8}$$

Thus, $r^* \searrow r_0$ in the limit $\sigma \to \infty$.

By (13), we have

$$p^* = \frac{\mu_\ell'(m^*)}{\mu_\ell'(m^*) - \mu_h'(m^*)}.$$

Equating this expression with (11) and from the fact that $r^* - r_0 = O(\sigma^{-1})$, we have

$$\frac{\mu_\ell'(m^*)}{\mu_\ell'(m^*) - \mu_h'(m^*)} = \frac{\mu_\ell(m^*) - \mu_\ell(r_0)}{\mu_\ell(m^*) - \mu_\ell(r_0) + \mu_h(r_0) - \mu_h(m^*)} + O(\sigma^{-1}).$$

Hence, in the limit $\sigma \to \infty$, m^* converges to m_0, which satisfies

$$\frac{\mu_\ell'(m_0)}{\mu_\ell'(m_0) - \mu_h'(m_0)} = \lim_{x \to m_0} \frac{\mu_\ell(x) - \mu_\ell(r_0)}{\mu_\ell(x) - \mu_\ell(r_0) + \mu_h(r_0) - \mu_h(x)}. \tag{A.9}$$

One obvious solution to (A.9) is $m_0 = r_0$, in which case $p^* \to p_0$ by (A.8). Moreover,

$$m^* - r_0 = -\frac{D_h(r_0) - D_\ell(r_0)}{\sigma\sqrt{8\alpha}} \frac{\mu_\ell'(m_0)\mu_h'(r_0) + \mu_h'(m_0)\mu_\ell'(r_0)}{-\mu_\ell'(r_0)\mu_h''(m_0) + \mu_h'(r)\mu_\ell''(m_0)} + O(\sigma^{-2}) < 0, \tag{A.10}$$

so $m^* \nearrow r_0$.

Suppose that there is another solution $m_0 \neq r_0$ to (A.9) that maximizes $\pi(p^*; r^*, m^*)$ for large values of σ. Then there is $p_1 < p_0$ such that $p^* \to p_1$. In that case,

$$\left(\frac{1 - p_0}{1 - p^*}\frac{p^*}{p_0}\right)^{(\gamma-1)/2}$$

converges faster to zero in the limit $\gamma \to \infty$ than if $p^* \to p_0$, so $\pi(p^*; r^*, m^*)$ decreases (converges) to $p_0\mu_h(r_0) + (1 - p_0)\mu_\ell(r_0)$ faster than if $p^* \to p_0$. Therefore, the return $\pi(p^*; r^*, m^*)$ from the hypothetical solution $m_0 \neq r_0$ is actually less than the return $\pi(p^*; r^*, m^*)$ obtained from the solution $m_0 = r_0$. Hence, a solution $m_0 \neq r_0$ to (A.9) cannot

https://doi.org/10.1017/S0269964811000246 Published online by Cambridge University Press

maximize $\pi(p^*; r^*, m^*)$, and the limiting behaviors $\lim_{\sigma\to\infty} p^* = p_0$ and $\lim_{\sigma\to\infty} m^* = r_0$ are the only possibility. ■

PROOF OF COROLLARY 1:

(i) For sufficiently large σ, Proposition 9 assures that both r^* and $m^*(r^*)$ increase in p_0. (Note that this result does not depend on the assumption of linear demand curves.)

(ii) For sufficiently small σ, (A.6) shows that r^* increases in p_0. We also note that $m^*(r)$ does not have explicit dependence on p_0; however, $m^*(r^*)$ has dependence on p_0 through r^*. Hence, it is sufficient to study $m^{*\prime}(\bar{x})$ in the small-σ limit. After some algebra, we find

$$m^{*\prime}(\bar{x}) = \frac{4\alpha\sigma^2}{(a_\ell - a_h)^2\bar{x}^3(\mu_h(\bar{x}) - \mu_h(\underline{x}))}(1 + a_\ell c)(\bar{x} - \underline{x}) + O(\sigma^4).$$

The leading-order term is positive, so $m^*(r^*)$ also increases in p_0. ■

https://doi.org/10.1017/S0269964811000246 Published online by Cambridge University Press