Price Experimentation with Strategic Buyers*

Oksana Loginova† Curtis Taylor‡

January 2003

JEL Classifications: C73, D81, and D82.
Keywords: price experimentation, learning, strategic rejections.

Abstract

There are many situations in which buyers have a significant stake in what a firm learns about their demands. Specifically, any time that price discrimination is possible on an individual bases and repeat purchases are likely, buyers possess incentives for strategic manipulation of demand information.

A simple two-period model in which a monopolist endeavors to learn about the demand parameter of a repeat buyer is presented here. It is shown that high first-period prices may lead to strategic rejections by high-valuation buyers who wish to conceal information (i.e., to pool), while low first-period prices may lead to strategic rejections by low-valuation buyers who wish to reveal information (i.e., to signal). The seller never experiments against patient buyers in any equilibrium. Indeed, the seller often charges first-period prices that reveal no information at all, and she may even set an equilibrium first-period price strictly below the buyer's lowest possible valuation.

*We thank Paul Klemperer, Tracy Lewis, Preston McAfee, Margaret Meyer, Wolfgang Koehler, and Huseyin Yildirim for helpful comments. Financial support from the National Science Foundation (grant SES-0136817) is gratefully acknowledged.

†Department of Economics, Duke University, Box 90097, Durham, NC 27708-0097, email: loginova@econ.duke.edu

‡Department of Economics, Duke University, Box 90097, Durham, NC 27708-0097, email: crtaylor@econ.duke.edu (corresponding author)

1 Introduction

In a classic paper, Rothschild (1974) showed that the pricing problem facing a monopolist with unknown demand is often analogous to a two-armed bandit problem.[1] Hence, the optimal policy for such a firm is to experiment with prices in order to learn about its unknown demand parameters. It is, however, well-known that the optimal policy may not result in complete learning because of the opportunity cost of experimentation. In addition, the learning process may be severely hampered unless the firm possesses significant prior knowledge about the type of uncertainty confronting it. For instance, even when demand is deterministic, Aghion, Bolton, Harris, and Jullien (1991) show that strong conditions such as continuity and quasi-concavity of the profit function are required to guarantee that a monopoly will eventually learn all the relevant information.

In this paper, a very different – but important – caveat is added to the list of reasons that a monopolist may have difficulty learning its demand. The firm may serve customers who do not want their demand characteristics to be known!

In the prior literature on price experimentation, the possibility of strategic buyers has been largely ignored.[2] Specifically, it has typically been assumed either that the monopolist faces a sequence of identical customers who exist in the market for only one period or that market demand is composed of a *large* number of *small* customers.[3] There are, however, many real-world situations in which buyers have a significant stake in what a firm learns about their demands. Specifically, any time that price discrimination is possible on an individual bases and repeat purchases are likely, buyers possess incentives for strategic manipulation of demand information. In any long-term supply relationship, the buyer wants the supplier to think that he has very elastic demand for the product, and the buyer may even strategically reject some price offers in order to manipulate the suppliers beliefs to this end.

In this paper, a simple two-period experimental pricing and learning environment is analyzed. Specifically, there is assumed to be a single buyer whose underlying demand parameter, λ, is private information. The seller makes a take-it-or-leave-it offer in the first period and updates her belief about the value of λ based on the buyer's acceptance decision. In particular, acceptance (rejection) of a high first-period price implies that the buyer's first-period valuation was high (low). This leads the seller to update her beliefs about λ and, therefore, to infer that the buyer's second-period valuation for her product is also likely to be high (low). Indeed, it is shown that if the buyer is myopic, then the informational value of a high first-period price can lead the seller to charge one when it would otherwise not be optimal to do so.

Things are very different, however, if the buyer is as patient as the seller. In this case, it is shown that if the buyer's first-period valuation is high, then he will often attempt to conceal information by strategically rejecting high first-period offers. In fact, a buyer with a high value of λ stands to gain the most from concealing his high first-period valuation. This gives rise to 'reverse screening' at high prices in the sense that only buyers with low values of λ (and high valuations) will accept high first-period prices. When the first-period price is low, however, two types of continuation equilibria emerge, a Good equilibrium (for the seller) in which all buyer types purchase the product and a Bad equilibrium in which a buyer with a relatively high value of λ but low valuation for the

[1] Many authors have subsequently refined and extended this observation. See, for example, Aghion, Bolton, Harris, and Jullien (1991); Mirman, Samuelson, and Urbano (1993); Rustichini and Wolinsky (1995); Keller and Rady (1999).

[2] An exception is Kennan (2000) who shows that persistent private information may lead to stochastic cycles in repeated labor negotiations.

[3] In a related paper, Segal (2002) considers a setting in which there is a finite number of buyers in the market from the outset. He shows that if the common distribution of buyers' valuations is unknown, then learning through price experimentation is dominated by a multi-unit auction which sets a price to each buyer on the basis of the demand distribution inferred statistically from other buyers' bids.

product may strategically reject an offer to signal his low valuation to the seller. This signaling behavior at low prices is the mirror image of the screening that occurs at high prices. Hence, strategic rejections at high prices conceal information while strategic rejections at low prices reveal information.

Since all of the continuation equilibria involve a lower effective first-period demand relative to the myopic buyer case, the seller generally finds it optimal to set a lower price when confronted with a strategic buyer. In fact, she never 'experiments' by charging a high price in order to obtain information. Indeed, the seller often charges first-period prices that reveal no information at all, and she may even set an equilibrium first-period price strictly below the buyer's lowest possible valuation.

The pricing and acceptance behavior exhibited in the model presented here can be viewed as a manifestation of the rachet effect familiar from the regulation and agency literature.[4] Specifically, the inability of the firm to commit not to use the information it learns in the first period harms it if the buyer is relatively patient. The strategic rejections associated with patient buyers often result in both a lower first-period equilibrium price and a lower probability of sale than would prevail if the firm could commit to price non-contingently.

This paper also contributes to the burgeoning literature on behavior-based price discrimination.[5] In Internet retailing as well as many other market settings, firms now have the ability to track the purchasing behavior of individual customers and to tailor price offers to them.[6] To the extent that consumers are aware of this, the findings presented here indicate that they possess significant incentives to manipulate the information collected. Hence, it will typically be necessary for firms to offer their customers valuable benefits in order to induce them to reveal their private information.

The basic model is presented in the next section. In Section 3, the bench-mark setting in which the buyer is not strategic is characterized. The analysis at the core of the paper is presented in Section 4, where the first-period expected demand of a strategic buyer is derived. Sections 5 and 6 deal respectively with the best and worst equilibria for the seller and contain most of the economic results. Some brief concluding remarks appear in Section 7. Proofs not appearing in the text have been relegated to the Appendix.

2 The Model

There are two risk-neutral players, a seller (S, she) and a buyer (B, he), who possess respective discount factors $\delta \in [0, 1]$ and $\beta \in [0, 1]$. In each period, $t = 1, 2$, B demands one unit of a good which S may produce and sell to him. S's production cost is normalized to zero. B's valuation for the good in period t, v_t, is high, v_H, with probability λ and low, v_L, with probability $1 - \lambda$, $v_H > v_L \geq 0$. In other words, B's valuations are independent draws from a two-point distribution with parameter λ.

The demand parameter λ, is itself the realization of a random variable which is continuously distributed on $[0, 1]$ with probability density function $f(\lambda)$. (For instance, λ might represent B's income and $f(\lambda)$ the distribution of income in the population of potential buyers.) Let $E[\lambda]$ denote its expected value under the prior. Also, it is notationally convenient to define the constant

$$\lambda^* \equiv v_L/v_H.$$

[4]See, for example, Laffont and Tirole (1988), and especially Hart and Tirole (1988).

[5]See, for example, Acquisti and Varian (2002), Taylor (2002), Shaffer and Zhang (2000), Fudenberg and Tirole (2000), and Villas Boas (1999).

[6]See Krugman (2000) and Streitfield (2000).

At the beginning of the game, B privately observes λ and v_1, and he privately observes v_2 at the beginning of the second period. Hence, in any given period, B's 'type' has two components, his permanent 'long-run' type $\lambda \in [0,1]$ and his transitory 'short-run' type $v_t \in \{v_L, v_H\}$.[7]

In each period $t = 1, 2$, B and S play an extensive-form game with the following stages.

1. B observes his valuation $v_t \in \{v_L, v_H\}$.
2. S announces price $p_t \in \Re_+$ at which she is willing to sell the good to B.
3. B either accepts ($q_t = 1$) or rejects ($q_t = 0$) S's offer.
4. B's (contemporaneous) payoff is $q_t(v_t - p_t)$, and S's payoff is $q_t p_t$.

Note that while this game has a recursive structure, it is not a repeated game due to the presence of asymmetric information. Specifically, S updates her prior beliefs about λ from the first period to the second.

In particular, let $h_S \equiv (p_1, q_1)$ be the history of first-period events observed by S, and let $h_B \equiv (\lambda, v_1, p_1, q_1)$ be the history of first-period events observed by B at the beginning of period 2. A behavior strategy for S is a pair of probability distributions, $(\Phi_1(p_1), \Phi_2(p_2; h_S))$, over all possible price offers. Similarly, a behavior strategy for B is a pair of functions, $(\gamma_1(\lambda, v_1, p_1), \gamma_2(v_2, p_2; h_B))$, where γ_t is the probability that B accepts S's offer in period t.

Let $f(\lambda|h_S)$ denote S's posterior beliefs about λ at the beginning of period 2. Likewise, let $E[\lambda|h_S]$ denote her updated expectation. The solution concept employed is efficient perfect Bayesian equilibrium (PBE); i.e., a PBE in which indifference about pricing or purchasing is resolved in favor of efficiency. (Since inefficient PBEs occur only for a non-generic set of parameter values, the 'efficient' qualifier is suppressed below.)

Observe that in the second period, B optimally accepts any price that does not exceed his valuation regardless of the history. Given this, S believes the good will be sold for price p_2 with probability

$$D_2(p_2; h_S) = \begin{cases} 0, & \text{if } p_2 > v_H \\ E[\lambda|h_S], & \text{if } p_2 \in (v_L, v_H] \\ 1, & \text{if } p_2 \leq v_L. \end{cases}$$

This is S's expectation of B's second-period demand for her product. Hence, S may either price at v_H and sell with probability $E[\lambda|h_S]$, or she can price at v_L and sell with probability one. Thus, S optimally sets $p_2 = v_L$ if $v_L > v_H E[\lambda|h_S]$ and $p_2 = v_H$ if $v_L < v_H E[\lambda|h_S]$.

This serves as proof of the following proposition that gives necessary conditions on second-period equilibrium strategies.

Proposition 1 (Second-Period Equilibrium Behavior) *In any PBE, B's strategy in period 2 is*

$$\gamma_2(v_2, p_2; h_B) = \begin{cases} 1, & \text{if } v_2 \geq p_2 \\ 0, & \text{if } v_2 < p_2. \end{cases}$$

S offers $p_2 = v_L$ with probability $\theta(h_S)$ and $p_2 = v_H$ with probability $1 - \theta(h_S)$ according to

$$\theta(h_S) = \begin{cases} 0, & \text{if } \lambda^* < E[\lambda|h_S] \\ \text{any } \theta \in [0,1], & \text{if } \lambda^* = E[\lambda|h_S] \\ 1, & \text{if } \lambda^* > [\lambda|h_S]. \end{cases}$$

The central concern of the next two sections is the first-period equilibrium behavior.

[7]While B's type is multi-dimentional, the focus here is on pricing without commitment rather than on the optimal monopolistic screening mechanism (e.g., Armstron (1996) and Rochet and Chone (1998)).

3 Myopic Buyer

In this section the setting in which B completely discounts the future (i.e., $\beta = 0$) is analyzed. There are two possible interpretations for this bench mark. First, it might occur because B is unaware that S is tracking his purchasing history in order to learn about his preferences. Second, B may actually consist of two stochastically equivalent (i.e., λ is the same), but distinct buyers, B_1 and B_2, who arrive sequentially at S's store. In this case, S learns about the preferences of B_2 by observing the purchasing behavior of B_1. Under either interpretation, B acts myopically in the first period.

Denote by E_L the expected value of λ given B's first-period valuation was v_L, and let E_H be the expected value of λ given $v_1 = v_H$. Straightforward calculations yield

$$E_L \equiv E[\lambda | v_1 = v_L] = \frac{E[\lambda] - E[\lambda^2]}{1 - E[\lambda]}$$

and

$$E_H \equiv E[\lambda | v_1 = v_H] = \frac{E[\lambda^2]}{E[\lambda]}.$$

Observe that $E_L < E[\lambda] < E_H$.

Proposition 2 (Equilibrium Behavior when the Buyer is Myopic) *Suppose $\beta = 0$. In any PBE, B accepts in period $t = 1, 2$ if and only if $v_t \geq p_t$. Equilibrium prices are given by:*

(i) *$p_1 = p_2 = v_H$, if $\lambda^* < E_L$,*

(ii) *$p_1 = v_H$ and $p_2 = q_1 v_H + (1 - q_1) v_L$, if $\lambda^* \in [E_L, \overline{\lambda})$,*

(iii) *$p_1 = p_2 = v_L$, if $\lambda^* \geq \overline{\lambda}$,*

where

$$\overline{\lambda} \equiv \frac{E[\lambda] + \delta E[\lambda^2]}{1 + \delta E[\lambda]}.$$

The intuition behind this result is easily grasped. First, $\beta = 0$ means that B treats each period as a one-shot game in which he accepts any offer yielding him a non-negative payoff. From S's prospective, the probability B accepts an offer of p_1 in the first period is

$$D_1(p_1) = \begin{cases} 0, & \text{if } p_1 > v_H \\ E[\lambda], & \text{if } p_1 \in (v_L, v_H] \\ 1, & \text{if } p_1 \leq v_L. \end{cases}$$

Second, note that charging $p_1 = v_L$ provides no information about B's second-period demand because he always accepts this offer. Pricing at $p_1 = v_H$, however, does reveal information because B accepts the high price if and only if $v_1 = v_H$. Hence, the probability that B accepts p_2 conditional on acceptance of $p_1 = v_H$ is

$$D_2(p_2; v_H, 1) = \begin{cases} 0, & \text{if } p_2 > v_H \\ E_H, & \text{if } p_2 \in (v_L, v_H] \\ 1, & \text{if } p_2 \leq v_L, \end{cases}$$

and the probability that B accepts p_2 conditional on rejection of $p_1 = v_H$ is

$$D_2(p_2; v_H, 0) = \begin{cases} 0, & \text{if } p_2 > v_H \\ E_L, & \text{if } p_2 \in (v_L, v_H] \\ 1, & \text{if } p_2 \leq v_L. \end{cases}$$

S finds it optimal to charge v_H in both periods if λ^* is low enough, and she charges v_L in both periods if it is high. In the intermediate range for λ^*, S charges $p_1 = v_H$ in the first period, $p_2 = v_H$ following acceptance and $p_2 = v_L$ following rejection in the second period.

The region where S prices contingently is $\lambda^* \in [E_L, \overline{\lambda})$. Observe that $\overline{\lambda} \in [E[\lambda], E_H)$. When $\lambda^* \in [E_L, E[\lambda]]$, the expected first-period payoff to S is maximized at $p_1 = v_H$. In contrast, when $\lambda^* \in (E[\lambda], \overline{\lambda})$, S *experiments*: she is willing to run a significant risk of losing a first-period sale (her expected first-period payoff is higher under $p_1 = v_L$) in order to obtain valuable information about B's demand parameter, λ.

Corresponding to the notion that the value of information gained through pricing at $p_1 = v_H$ increases as S becomes more patient, $\overline{\lambda}$ is an increasing function of δ:

$$\frac{\partial \overline{\lambda}}{\partial \delta} = \frac{E[\lambda^2] - (E[\lambda])^2}{(1 + \delta E[\lambda])^2} > 0.$$

When $\delta = 0$ (S completely discounts the future), the experimentation region disappears altogether.

In order to focus on settings where information is potentially valuable to S, the following necessary condition is assumed to hold throughout the remainder of the paper:

$$E_L < \lambda^* < E_H.$$

If this fails to hold, then S is either so pessimistic ($\lambda^* \geq E_H$) that she would set $p_2 = v_L$ even if she knew $v_1 = v_H$, or she is so optimistic ($\lambda^* \leq E_L$) that she would set $p_2 = v_H$ even if she knew $v_1 = v_L$. In either case, learning v_1 has no value to her.

Corollary 1 (below) compares the setting in which B is myopic with a setting where S can publicly commit to not use any information she learns in period 1. When S can commit to price non-contingently, B optimally accepts $p_t \leq v_t$ in period $t = 1, 2$, and equilibrium prices are given by:

$$p_t = \begin{cases} v_H, & \text{if } \lambda^* < E[\lambda] \\ v_L, & \text{if } \lambda^* \geq E[\lambda]. \end{cases}$$

For ease of exposition, this bench-mark case is called *the Fixed-Price setting* and the case without commitment is called *the Contingent-Price setting.*

Corollary 1 (Welfare when the Buyer is Myopic) *Suppose $\beta = 0$. Then, the following comparisons between the Contingent-Price setting and the Fixed-Price setting hold:*

(i) *S prefers the Contingent-Price setting if $\lambda^* < \overline{\lambda}$, and she is indifferent between the two settings otherwise.*

(ii) *B prefers the Fixed-Price setting if $\lambda^* \in [E[\lambda], \overline{\lambda})$, and he is indifferent between the two settings otherwise.*

(iii) *Welfare is higher under the Contingent-Price setting if $\lambda^* < E[\lambda]$; it is higher under the Fixed-Price one if $\lambda^* \in [E[\lambda], \overline{\lambda})$; and it is the same under the two settings if $\lambda^* \geq \overline{\lambda}$.*

Corollary 1 indicates that, compared with the Contingent-Price setting, S would always be (weakly) worse off and B would always be (weakly) better off if S committed herself not to learn anything about B's demand. When $\lambda^* < E[\lambda]$, S's expected payoff is higher under the Contingent-Price setting and B is indifferent between the two settings. Hence, welfare is higher under the Contingent-Price setting.

When $\lambda^* \in [E[\lambda], \overline{\lambda})$, S's expected payoff is higher under the Contingent-Price setting, because the value of information obtained from charging $p_1 = v_H$, $\delta E(\lambda)(E_H v_H - v_L)$, outweighs the *private* cost of experimentation, $v_L - E[\lambda] v_H$. B is worse off in this case under the Contingent-Price setting because he would have received a low-price offer in the first period under the Fixed-Price setting. Welfare is lower, because the *social* cost of experimentation, $(1 - E[\lambda])v_L$, outweighs the value of information.[8]

Attention now turns to situations in which $\beta > 0$ and B, therefore, has strategic considerations regarding the revelation of his private information.

4 Strategic Buyer

Now, suppose $\beta > 0$ and that S offers p_1 in the first period. By Proposition 1, the expected payoff of B with first-period valuation v_1 and long-run type λ is

$$v_1 - p_1 + \beta\lambda\theta(p_1, 1)(v_H - v_L),$$

if he accepts p_1, and

$$\beta\lambda\theta(p_1, 0)(v_H - v_L),$$

if he rejects p_1.

This simple observation serves as proof of the following claim.

Lemma 1 (Dynamic Incentives) *Suppose $\beta > 0$. In any PBE, B accepts p_1 if and only if*

$$v_1 - p_1 \geq \beta\lambda(\theta(p_1, 0) - \theta(p_1, 1))(v_H - v_L).$$

In any PBE, it must be the case that $E[\lambda|p_1, q_1]$ is derived from B's first-period behavior given $\theta(p_1, q_1)$, and $\theta(p_1, q_1)$ is optimal for S given $E[\lambda|p_1, q_1]$. This interdependence between optimal actions and beliefs is the key to the next important result.

Lemma 2 (Beliefs and Actions) *Suppose $\beta > 0$. In any PBE, if $D_1(p_1) \in (0, 1)$, then*

$$E[\lambda|p_1, 1] \geq E[\lambda|p_1, 0],$$

$$\theta(p_1, 1) \leq \theta(p_1, 0).$$

Intuitively, since B is more likely to accept when $v_1 = v_H$, S's beliefs about λ should be higher when she observes $q_1 = 1$ than when she observes $q_1 = 0$. Along with Lemma 1, Lemma 2 implies that B never accepts an offer yielding negative first-period surplus.

Corollary 2 (Honest Rejections) *Suppose $\beta > 0$. In any PBE, B always rejects $p_1 > v_1$.*

[8] To see this, note that $\lambda^* \in [E[\lambda], \overline{\lambda})$ implies $\lambda^* > E[\lambda^2]$. Then,

$$(1 - E[\lambda])v_L > E[\lambda^2]v_H - E[\lambda]v_L = E[\lambda](E_H v_H - v_L).$$

Corollary 2 indicates that B never accepts $p_1 > v_H$ and that he does not accept $p_1 \in (v_L, v_H]$ when $v_1 = v_L$.

The next step in deriving the equilibrium of the game is to consider B's acceptance decision if $p_1 \in (v_L, v_H]$ and $v_1 = v_H$. To this end, define the constant

$$\underline{p} \equiv v_H - \beta(v_H - v_L).$$

Let $p_1 \in (v_L, \underline{p}]$. Suppose, B with $v_1 = v_H$ always accepts p_1. S updates her beliefs: $E[\lambda|p_1, 0] = E_L < \lambda^*$ and $E[\lambda|p_1, 1] = E_H > \lambda^*$. Hence, S sets $p_2 = v_H$ following acceptance and $p_2 = v_L$ following rejection. Given this strategy for S, B faces a trade-off: grab the current surplus of $v_H - p_1$ and face $p_2 = v_H$, or forego $v_H - p_1$ and face $p_2 = v_L$. The discounted expected value of facing $p_2 = v_L$ is $\beta\lambda(v_H - v_L) \leq v_H - p_1$. Thus, it is optimal for B to accept p_1 when $v_1 = v_H$. This serves as a proof of part (i) of Lemma 4 (below).

Now let $p_1 \in (\underline{p}, v_H)$. First, suppose B with $v_1 = v_H$ always accepts p_1. S updates her beliefs and sets $p_2 = v_H$ following acceptance and $p_2 = v_L$ following rejection. Given this strategy for S, long-run type $\lambda > (v_1 - p_1)/(\beta(v_H - v_L))$ of B does better rejecting p_1, which contradicts the supposition.

Second, suppose B with $v_1 = v_H$ always rejects $p_1 \in (\underline{p}, v_H)$. No updating occurs, S sets $p_2 = v_L$ if $\lambda^* > E[\lambda]$ and $p_2 = v_H$ if $\lambda^* < E[\lambda]$. In either case, B does better accepting p_1, which contradicts the supposition.

Summarizing the above, when $p_1 \in (\underline{p}, v_H)$, B's purchasing decision must be based not only on his first-period valuation, but also on his long-run type λ. Determining when B with $v_1 = v_H$ accepts $p_1 \in (\underline{p}, v_H)$ requires some additional notation and machinery. To start with, for any $\mu \in [0, 1]$, define the functions

$$\alpha(\mu) \equiv E[\lambda|v_1 = v_H \cap \lambda \leq \mu] = \frac{\int_0^\mu \lambda^2 f(\lambda)\, d\lambda}{\int_0^\mu \lambda f(\lambda)\, d\lambda}$$

and

$$\rho(\mu) \equiv E[\lambda|\{v_1 = v_L\} \cup \{v_1 = v_H \cap \lambda > \mu\}] = \frac{E[\lambda] - \int_0^\mu \lambda^2 f(\lambda)\, d\lambda}{1 - \int_0^\mu \lambda f(\lambda)\, d\lambda}.$$

These functions have some important properties which are summarized in the following technical lemma.

Lemma 3 (Geometric Properties of α and ρ) *Functions α and ρ possess the following properties:*

(i) *α starts at $\alpha(0) = 0$ and increases monotonicly until it ends at $\alpha(1) = E_H$.*

(ii) *ρ starts at $\rho(0) = E[\lambda]$, increases until it crosses the 45-degree line, and then decreases until it ends at $\rho(1) = E_L$.*

(iii) *α and ρ cross once, and at their intersection, m_{min}, $\alpha(m_{min}) = \rho(m_{min}) = E[\lambda]$.*

These functions are interpreted as follows. Suppose that B accepts some price p_1 if and only if $v_1 = v_H$ and $\lambda \leq \mu$. Then, $\alpha(\mu)$ is the expected value of λ conditional on acceptance, and $\rho(\mu)$ is the expected value of λ conditional on rejection.

Lemma 3 implies that there exists a unique number $m \in [m_{min}, 1)$ defined as follows:

$$m \equiv \begin{cases} \rho^{-1}(\lambda^*), & \text{if } \lambda^* \in (E_L, E[\lambda]] \\ \alpha^{-1}(\lambda^*), & \text{if } \lambda^* \in (E[\lambda], E_H). \end{cases}$$

Figure 1: Geometric Properties of α and ρ

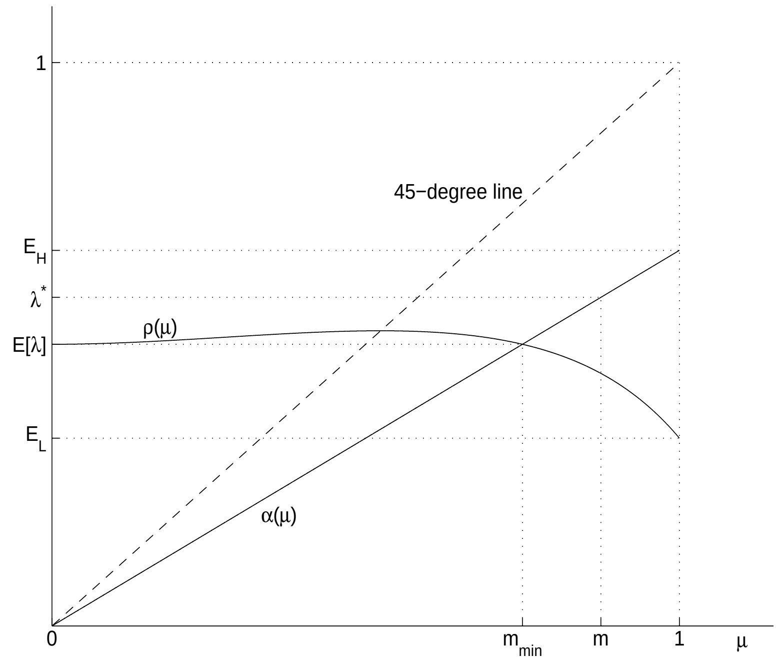

In Figure 1, m is shown for the case $\lambda^* \in (E[\lambda], E_H)$.

Given m, define the constant

$$\overline{p} \equiv v_H - \beta m(v_H - v_L)$$

and the function

$$\mu(p_1) \equiv \frac{v_H - p_1}{\beta(v_H - v_L)}, \quad p_1 \in [\underline{p}, \overline{p}].$$

Observe that $\mu(p_1)$ is monotone decreasing with $\mu(\underline{p}) = 1$ and $\mu(\overline{p}) = m$.

Lemma 4 (Strategic Rejections) *Suppose $\beta > 0$ and $v_1 = v_H$. In any PBE, the following must hold:*

(i) *If $p_1 \in (v_L, \underline{p}]$, then B always accepts the price. S sets $p_2 = v_H$ following acceptance and $p_2 = v_L$ following rejection.*

(ii) *If $p_1 \in (\underline{p}, \overline{p}]$, then B accepts the price if and only if $\lambda \leq \mu(p_1)$. S sets $p_2 = v_H$ following acceptance and $p_2 = v_L$ following rejection.*

(iii) *If $p_1 \in (\overline{p}, v_H]$, then B accepts the price if and only if $\lambda \leq m$. When $\lambda^* < E[\lambda]$, S sets $p_2 = v_H$ following acceptance and randomizes between $p_2 = v_L$ and $p_2 = v_H$ following rejection,*

$$\theta(p_1, 0) = \frac{v_H - p_1}{\beta m(v_H - v_L)}.$$

When $\lambda^* > E[\lambda]$, S sets $p_2 = v_L$ following rejection and randomizes between $p_2 = v_H$ and $p_2 = v_L$ following acceptance,

$$\theta(p_1, 1) = 1 - \frac{v_H - p_1}{\beta m(v_H - v_L)}.$$

To understand this result, first consider $p_1 \in (\underline{p}, \overline{p}]$. Given that B accepts p_1 if and only if $v_1 = v_H$ and $\lambda \leq \mu(p_1)$, S updates her beliefs: $E[\lambda|p_1, 1] = \alpha(\mu(p_1)) \geq \lambda^*$ and $E[\lambda|p_1, 0] = \rho(\mu(p_1)) \leq \lambda^*$. Hence, S optimally sets $p_2 = v_H$ following acceptance and $p_2 = v_L$ following rejection. Given this strategy for S, B with $v_1 = v_H$ faces a trade-off: grab the current surplus of $v_H - p_1$ and face $p_2 = v_H$, or forego $v_H - p_1$ and face $p_2 = v_L$. The discounted expected value of facing $p_2 = v_L$ is $\beta\lambda(v_H - v_L)$. Thus, if $\lambda > \mu(p_1)$, then B strategically rejects p_1.

Notice that 'reverse screening' occurs over this range in the sense that high prices induce rejection by high types. Indeed, as p_1 increases, the 'marginal' long-run type, $\mu(p_1)$, falls and the set of long-run types willing to strategically reject increases. Acceptance and rejection, therefore, become less informative. Once $p_1 = \overline{p}$, further increases in p_1 cannot induce more strategic rejection.

In particular, for $p_1 \in (\overline{p}, v_H]$, the marginal type must remain at m. This requires S to randomize between $p_2 = v_L$ and $p_2 = v_H$. Specifically, if $\lambda^* > E[\lambda]$, then $\alpha(m) = \lambda^* > \rho(m)$. S optimally sets $p_2 = v_L$ following rejection and mixes following acceptance, $\phi(p_1, 1)$ is calibrated to make B with $\lambda = m$ and $v_1 = v_H$ indifferent about accepting p_1. Similarly, if $\lambda^* < E[\lambda]$, then $\alpha(m) > \lambda^* = \rho(m)$. S optimally sets $p_2 = v_H$ following acceptance and mixes following rejection, $\phi(p_1, 0)$ is calibrated to make B with $\lambda = m$ and $v_1 = v_H$ indifferent about accepting p_1.

To complete the characterization of equilibrium play, it remains to consider $p_1 \leq v_L$. The analysis over this range, however, is complicated by the fact that equilibrium behavior is not unique. In particular, there exists a lower bound $\tilde{p} < v_L$ such that B accepts in any PBE if $p_1 \leq \tilde{p}$. For any $p_1 \in (\tilde{p}, v_L]$, however, either pooling or signaling may occur in equilibrium. To illustrate this, the two extreme equilibria that involve minimal signaling (i.e., all types accept $p_1 \leq v_L$) and maximal signaling (i.e., strategic rejections by some types for all $p_1 \in (\tilde{p}, v_L]$) are derived. For ease of exposition, these equilibria are called (using S's perspective) respectively *the Good PBE* and *the Bad PBE*, with the recognition that they actually bracket a continuum of intermediate cases.

Interestingly, B's behavior in the Bad PBE for prices $p_1 \in (\tilde{p}, v_L]$ is the mirror image of his behavior for prices in $(\underline{p}, \overline{p}]$. Specifically, for $p_1 \in (\underline{p}, \overline{p}]$, B always rejects if $v_1 = v_L$ and rejects strategically if $v_1 = v_H$ and $\lambda > \mu(p_1)$. In other words, strategic rejections *conceal* information. By contrast, for $p_1 \in (\tilde{p}, v_L]$, it is shown below that B always accepts if $v_1 = v_H$ and rejects strategically if $v_1 = v_L$ and $\lambda > \hat{\mu}(p_1)$ (defined below). In other words, strategic rejections *reveal* information.

The following result is proved by applying Lemma 1 and Lemma 2.

Corollary 3 (Honest Acceptances) *Suppose $\beta > 0$. In any PBE, B with $v_1 = v_H$ always accepts $p_1 \leq v_L$.*

Next, for any $\hat{\mu} \in [0, 1]$, define

$$\hat{\alpha}(\hat{\mu}) \equiv E[\lambda|\{v_1 = v_H\} \cup \{v_1 = v_L \cap \lambda \leq \hat{\mu}\}] = \frac{E[\lambda] - \int_{\hat{\mu}}^1 \lambda(1-\lambda)f(\lambda)\, d\lambda}{1 - \int_{\hat{\mu}}^1 (1-\lambda)f(\lambda)\, d\lambda}$$

and

$$\hat{\rho}(\hat{\mu}) \equiv E[\lambda|v_1 = v_L \cap \lambda > \hat{\mu}] = \frac{\int_{\hat{\mu}}^1 \lambda(1-\lambda)f(\lambda)\, d\lambda}{\int_{\hat{\mu}}^1 (1-\lambda)f(\lambda)\, d\lambda}.$$

Lemma 5 establishes some important geometric properties of $\hat{\alpha}$ and $\hat{\rho}$.

Lemma 5 (Geometric Properties of $\hat{\alpha}$ and $\hat{\rho}$) *Functions $\hat{\alpha}$ and $\hat{\rho}$ possess the following properties:*

(i) *$\hat{\alpha}$ starts at $\hat{\alpha}(0) = E_H$, decreases until it crosses the 45-degree line, and then increases until it ends at $\hat{\alpha}(1) = E[\lambda]$.*

Figure 2: Geometric Properties of $\hat{\alpha}$ and $\hat{\rho}$

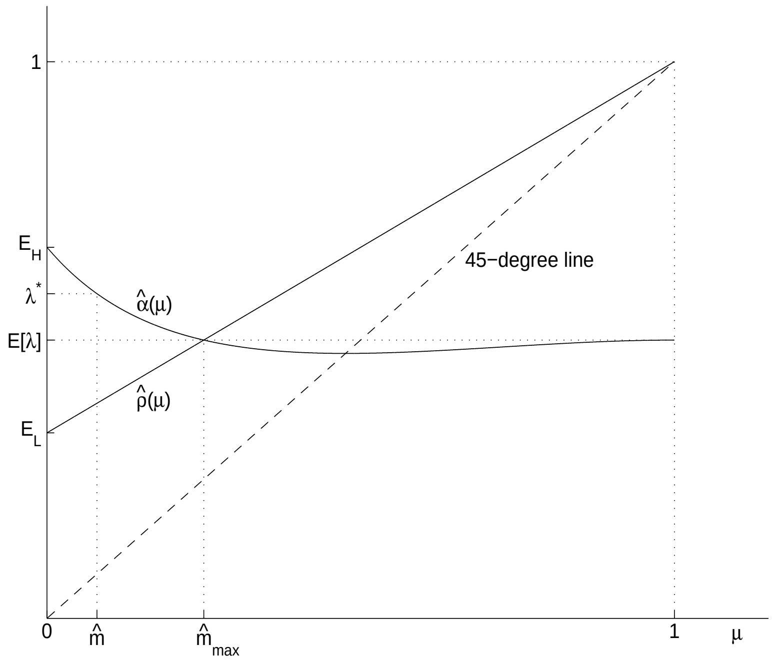

(ii) *$\hat{\rho}$ starts at $\hat{\rho}(0) = E_L$ and increases monotonicly until it ends at $\rho(1) = 1$.*

(iii) *$\hat{\alpha}$ and $\hat{\rho}$ cross once, and at their intersection, $\hat{m}_{max}$, $\hat{\alpha}(\hat{m}_{max}) = \hat{\rho}(\hat{m}_{max}) = E[\lambda]$.*

These functions are interpreted as follows. Suppose that B rejects some price p_1 if and only if $v_1 = v_L$ and $\lambda > \hat{\mu}$. Then, $\hat{\alpha}(\hat{\mu})$ is the expected value of λ conditional on acceptance, and $\hat{\rho}(\hat{\mu})$ is the expected value of λ conditional on rejection.

Lemma 5 implies that there exists a unique number $\hat{m} \in (0, \hat{m}_{max}]$ defined as follows:

$$\hat{m} \equiv \begin{cases} \hat{\rho}^{-1}(\lambda^*), & \text{if } \lambda^* \in (E_L, E[\lambda]] \\ \hat{\alpha}^{-1}(\lambda^*), & \text{if } \lambda^* \in (E[\lambda], E_H). \end{cases}$$

In Figure 2, $\hat{m}$ is shown for the case $\lambda^* \in (E[\lambda], E_H)$.

Given $\hat{m}$, define the constant

$$\tilde{p} \equiv v_L - \beta \hat{m}(v_H - v_L)$$

and the function

$$\hat{\mu}(p_1) \equiv \frac{v_L - p_1}{\beta(v_H - v_L)}, \quad p_1 \in [\tilde{p}, v_L].$$

Observe that $\hat{\mu}(p_1)$ is monotone decreasing with $\hat{\mu}(\tilde{p}) = \hat{m}$ and $\hat{\mu}(v_L) = 0$.

Corollary 3 indicates that B with $v_1 = v_H$ always accepts $p_1 \leq v_L$. However, it is silent about B's purchasing decision when $v_1 = v_L$.

Lemma 6 (The Good, the Bad, and the Ugly) *Suppose $\beta > 0$ and $v_1 = v_L$. There is a Good PBE in which B always accepts $p_1 \leq v_L$, S sets $p_2 = v_H$ if $\lambda^* < E[\lambda]$ and $p_2 = v_L$ if $\lambda^* > E[\lambda]$. There is also a Bad PBE in which the following holds:*

(i) *If $p_1 \le \tilde{p}$, then B always accepts. S sets $p_2 = v_H$ if $\lambda^* < E[\lambda]$ and $p_2 = v_L$ if $\lambda^* > E[\lambda]$.*

(ii) *If $p_1 \in (\tilde{p}, v_L]$, then B accepts if and only if $\lambda \le \hat{\mu}(p_1)$. S sets $p_2 = v_H$ following acceptance and $p_2 = v_L$ following rejection.*

The strategic rejection exhibited in the Bad PBE for prices $p_1 \in (\tilde{p}, v_L]$ warrants some discussion. Given that B rejects if and only if $v_1 = v_L$ and $\lambda > \hat{\mu}(p_1)$, S updates her beliefs: $E[\lambda|p_1, 0] = \hat{\rho}(\hat{\mu}(p_1)) \le \lambda^*$ and $E[\lambda|p_1, 1] = \hat{\alpha}(\hat{\mu}(p_1)) \ge \lambda^*$. Hence, S optimally sets $p_2 = v_L$ following rejection and $p_2 = v_H$ following acceptance. Given this strategy for S, B faces a trade-off: grab the current surplus of $v_L - p_1$ and face $p_2 = v_H$, or forego $v_L - p_1$ and face $p_2 = v_L$. The discounted expected value of facing $p_2 = v_L$ is $\beta\lambda(v_H - v_L)$. Thus, if $\lambda > \hat{\mu}(p_1)$, then B strategically rejects p_1.

Notice that as p_1 falls, $\hat{\mu}(p_1)$ increases and the set of long-run types willing to strategically reject shrinks. Acceptance and rejection, thus, provide weaker information about λ. For prices $p_1 \le \tilde{p}$, signaling is not possible and Lemma 6 shows that any such price must induce acceptance by all types in any PBE.

With Lemmas 1 through 6 in hand, it is now possible to calculate the expected first-period demand (which is depicted in Figure 3, and the expected values of λ conditional on acceptance and rejection of the first-period price.

Proposition 3 (Expected First-Period Demand and Posterior Beliefs) *Suppose $\beta > 0$. Then, the expected probability that B accepts p_1 and the expected values of λ conditional on acceptance and rejection of p_1 are as follows:*

(i) *If $p_1 > v_H$, then $D_1(p_1) = 0$ and $E[\lambda|p_1, 0] = E[\lambda]$.*

(ii) *If $p_1 \in (\overline{p}, v_H]$, then $D_1(p_1) = \int_0^m \lambda f(\lambda)\, d\lambda$ and*

$$E[\lambda|p_1, q_1] = \begin{cases} \alpha(m), & \text{if } q_1 = 1 \\ \rho(m), & \text{if } q_1 = 0. \end{cases}$$

(iii) *If $p_1 \in (\underline{p}, \overline{p}]$, then $D_1(p_1) = \int_0^{\mu(p_1)} \lambda f(\lambda)\, d\lambda$ and*

$$E[\lambda|p_1, q_1] = \begin{cases} \alpha(\mu(p_1)), & \text{if } q_1 = 1 \\ \rho(\mu(p_1)), & \text{if } q_1 = 0. \end{cases}$$

(iv) *If $p_1 \in (v_L, \underline{p}]$, then $D_1(p_1) = E[\lambda]$ and*

$$E[\lambda|p_1, q_1] = \begin{cases} E_H, & \text{if } q_1 = 1 \\ E_L, & \text{if } q_1 = 0. \end{cases}$$

(v) *If $p_1 \le v_L$ and the Good PBE obtains, then $D_1(p_1) = 1$ and $E[\lambda|p_1, 1] = E[\lambda]$. If the Bad PBE obtains, then*

$$D_1(p_1) = \begin{cases} E[\lambda] + \int_0^{\hat{\mu}(p_1)} (1 - \lambda) f(\lambda)\, d\lambda, & \text{if } p_1 \in (\tilde{p}, v_L] \\ 1, & \text{if } p_1 \le \tilde{p}, \end{cases}$$

$$E[\lambda|p_1, q_1] = \begin{cases} \hat{\alpha}(\hat{\mu}(p_1)), & \text{if } q_1 = 1 \text{ and } p_1 \in (\tilde{p}, v_L] \\ \hat{\rho}(\hat{\mu}(p_1)), & \text{if } q_1 = 0 \text{ and } p_1 \in (\tilde{p}, v_L] \\ E[\lambda], & \text{if } q_1 = 1 \text{ and } p_1 \le \tilde{p}. \end{cases}$$

Figure 3: Expected First-Period Demand

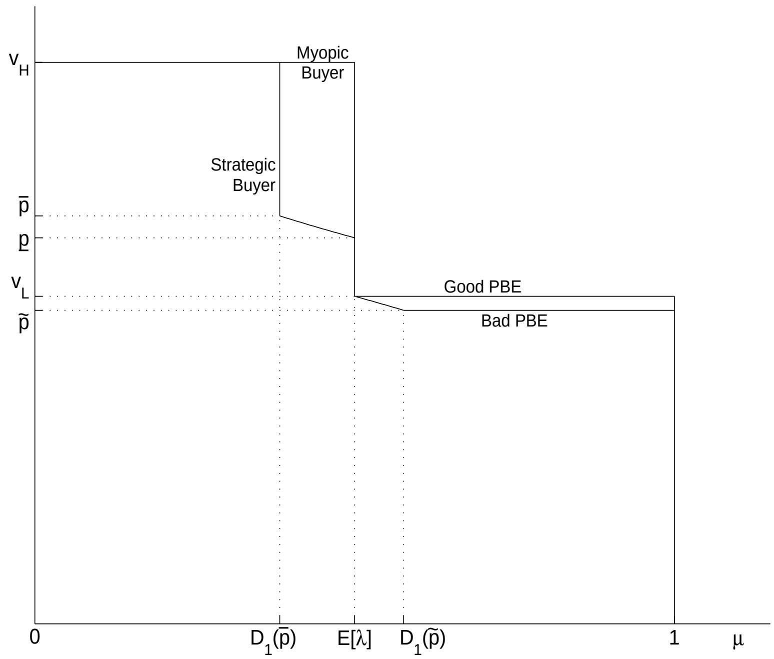

Given the above, S's expected payoff is

$$\Pi_S(p_1) \equiv D_1(p_1)p_1 + \delta[D_1(p_1)\max\{E[\lambda|p_1,1],\lambda^*\} + (1 - D_1(p_1))\max\{E[\lambda|p_1,0],\lambda^*\}]v_H.$$

The final step in deriving the players' equilibrium behavior is to determine the value of p_1 that maximizes $\Pi_S(p_1)$. Of course, the solution to this problem depends on whether the Good or the Bad PBE obtains. These cases are investigated respectively in the next two sections.

5 The Good Equilibrium

In this section, the Good PBE in which B always accepts $p_1 \leq v_L$ is explored. For any $\mu \in [0,1]$, it is notationally convenient to define

$$I(\mu) \equiv \int_0^\mu \lambda f(\lambda)\, d\lambda.$$

First, suppose S offers $p_1 \in [\overline{p}, v_H]$, then her expected payoff can be written as

$$\Pi_S(p_1) = I(m)p_1 + \delta \max\{E[\lambda],\lambda^*\}v_H.$$

The value of the information to S conveyed by B's acceptance/rejection decision is zero in this case. Indeed, if $\lambda^* > E[\lambda]$, then $E[\lambda|p_1,1] = \lambda^* > E[\lambda|p_1,0]$. Thus, charging $p_2 = v_L$ (which is optimal under the prior beliefs) maximizes S's expected payoff following acceptance as well as rejection. If $\lambda^* < E[\lambda]$, then $E[\lambda|p_1,1] > \lambda^* = E[\lambda|p_1,0]$. Thus, charging $p_2 = v_H$ (which is optimal under the prior beliefs) maximizes S's expected payoff following acceptance as well as rejection. Observe that S's expected payoff over $p_1 \in [\overline{p}, v_H]$ is maximized at $p_1 = v_H$.

Second, suppose S offers $p_1 \in (\underline{p}, \overline{p})$, then

$$\Pi_S(p_1) = I(\mu(p_1))p_1 + \delta[I(\mu(p_1))E[\lambda|p_1, 1]v_H + (1 - I(\mu(p_1)))v_L].$$

The information conveyed by B's purchasing decision is valuable to S. In the second period, she optimally sets $p_2 = v_H$ following acceptance and $p_2 = v_L$ following rejection, as $E[\lambda|p_1, 1] > \lambda^* > E[\lambda|p_1, 0]$ in this case.

Next, suppose S offers $p_1 \in (v_L, \underline{p}]$, then

$$\Pi_S(p_1) = E[\lambda]p_1 + \delta[E[\lambda]E_H v_H + (1 - E[\lambda])v_L].$$

The value of information is maximal, as no strategic rejections occur over this range of prices. Obviously, S's expected payoff over $p_1 \in (v_L, \underline{p}]$ is maximized at $p_1 = \underline{p}$.

Finally, suppose S offers $p_1 \leq v_L$, then

$$\Pi_S(p_1) = p_1 + \delta \max\{E[\lambda], \lambda^*\}v_H.$$

Observe that S's expected payoff is maximized at $p_1 = v_L$ in this case.

Summarizing the above, in equilibrium S charges some $p_1 \in \{v_L, v_H\} \cup [\underline{p}, \overline{p})$, depending on the parameters of the model, λ^*, β, δ and $f(\lambda)$. The following result establishes that if S's prior beliefs are 'pessimistic', $\lambda^* \geq E[\lambda]$, and B is at least as patient as S, then S charges v_L in both periods.

Proposition 4 (Optimal Prices with a Patient Buyer and Pessimistic Beliefs) *Suppose $\lambda^* \geq E[\lambda]$ and $\beta \geq \delta > 0$. Then, S offers $p_1 = v_L$ and $p_2 = v_L$ in the Good PBE.*

Recall from Proposition 2 that when $\beta = 0$, it is optimal for S to experiment by charging $p_1 = v_H$ if $\lambda^* \in [E[\lambda], \overline{\lambda})$. When $\beta > 0$, S must either pay B high information rent (set $p_1 = \underline{p}$), or run the risk that he will strategically reject her offer (marginal type, $\mu(p_1)$, falls as p_1 goes from $\underline{p}$ to $\overline{p}$).

For any $p_1 \in [\underline{p}, \overline{p})$, Proposition 4 shows that as β approaches δ, the information rent dominates the value of the information obtained through experimentation, and S, therefore, opts not to experiment. In addition, since S learns nothing by setting $p_1 = v_L$ in the Good PBE, she also sets $p_2 = v_L$. This, of course, maximizes welfare because B buys the good in both periods with probability one.

When $\lambda^* < E[\lambda]$ ('optimistic' beliefs), the story is somewhat more complicated. The following result shows that in this case, S also has incentives to select a price that generates no valuable information.

Proposition 5 (Optimal Prices with a Patient Buyer and Optimistic Beliefs) *Suppose $\lambda^* < E[\lambda]$ and $\beta \geq \delta > 0$.*

(i) *There exists $\epsilon > 0$ such that if $\lambda^* \in [E[\lambda] - \epsilon, E[\lambda])$, then S sets $p_1 = v_L$ and $p_2 = v_H$ in the Good PBE.*

(ii) *There exists $\xi > 0$ such that if $\lambda^* \in (E_L, E_L + \xi]$, then S sets $p_1 = v_H$ and $p_2 = v_H$ in the Good PBE.*

The story behind this result is similar to the previous one. When B is as patient as S, the information rent S must pay outweighs the value of information she obtains. Hence, S prefers prices that obtain information that has zero value. In the Good PBE, there are two potentially optimal prices for which this is true, $p_1 = v_L$ and $p_1 = v_H$. When λ^* is close to $E[\lambda]$, then m

is relatively small and $v_L > I(m)v_H$ (i.e., the first-period return to $p_1 = v_L$ is greater than the expected first-period return to $p_1 = v_H$). On the other hand, when λ^* is close to E_L, then m is relatively large and $I(m)v_H > v_L$.

The next result follows naturally in light of Propositions 4 and 5.

Corollary 4 (Welfare in the Good PBE) *Suppose $\beta \geq \delta > 0$ and the Good PBE obtains. Then, the following comparisons between the Contingent-Price setting and the Fixed-Price setting hold:*

(i) *S strictly prefers the Fixed-Price setting to the Contingent-Price one if $\lambda^* < E[\lambda]$, and she is indifferent between the two settings otherwise.*

(ii) *B is indifferent between the two settings if $\lambda^* \geq E[\lambda]$ or if $\lambda^* \leq E_L + \xi$ for some $\xi > 0$, and he prefers the Contingent-Price setting to the Fixed-Price one otherwise.*

(iii) *Welfare is higher under the Contingent-Price setting if $\lambda^* \in [E[\lambda] - \epsilon, E[\lambda])$ for some $\epsilon > 0$; it is higher under the Fixed-Price one if $\lambda^* \leq E_L + \xi$; and it is the same under the two settings if $\lambda^* \geq E[\lambda]$.*

Comparing Corollaries 1 and 4 reveals some striking welfare reversals. Specifically, when B is myopic, B is always (weakly) worse off and S is always (weakly) better off under the Contingent-Price setting. Moreover, experimental pricing reduces welfare because the value of the information obtained by S is outweighed by the social cost of experimentation. By contrast, when B is relatively patient, B is always (weakly) worse off and S is always (weakly) better off under the Fixed-Price setting.

The reason for this reversal is clear. When $\beta \geq \delta$, B aggressively protects his private information by rejecting offers that do not provide him with sufficient information rent. In this case, the direct cost of acquiring information typically outweighs its value to S. Hence, when $\beta \geq \delta$ and $\lambda^* \geq E[\lambda]$, S sets $p_1 = v_L$ and learns nothing about B's preferences. This is, of course, the same outcome as in the Fixed-Price setting.

Now, suppose $\lambda^* < E[\lambda]$. In the Fixed-Price setting, S sets $p_1 = v_H$. Since S cannot use information obtained, B demands no information rent, and he, therefore, accepts this offer with expected probability $E[\lambda]$. In the Contingent-Price setting, however, S cannot commit not to use information, and she must either lower the price price or face a substantially lower expected probability of sale. In either case, S is worse off and B is weakly better off as compared with the Fixed-Price setting. It is not surprising that the Contingent-Price setting generates lower welfare than the Fixed-Price one when λ^* is close to E_L. Indeed, the prices are the same under the two settings, $p_1 = p_2 = v_H$, but B with $v_1 = v_H$ strategically rejects p_1 when $\lambda > m$.

Corollary 4 raises the question how S might commit not to learn about B's preferences or not to use any information she learns. First, it seems unlikely that S could commit not to learn because such a commitment is difficult (if not impossible) to verify. Second, while it seems plausible that S could commit not to raise prices, a commitment not to lower them is not renegotiation proof. Note, however, that it is the commitment not to lower prices that has strategic value to S.

If S is actually composed of two sellers, S_1 and S_2, selling distinct but related products, then it might be possible for S_1 to commit not to share information with S_2. This commitment may, however, still be difficult for B to verify. In other words, strategic rejections by B and the concomitant inefficiency may be difficult for S to eliminate through commitment.

6 The Bad Equilibrium

In this section, the Bad PBE described in Lemma 6 is investigated. For any $\hat{\mu} \in [0, 1]$ it is convenient to define

$$\hat{I}(\hat{\mu}) \equiv E[\lambda] + \int_0^{\hat{\mu}} (1 - \lambda) f(\lambda)\, d\lambda.$$

First, suppose S offers $p_1 \in (\tilde{p}, v_L]$, then S's expected payoff can be written as

$$\Pi_S(p_1) = \hat{I}(\hat{\mu}(p_1))p_1 + \delta[\hat{I}(\hat{\mu}(p_1))\hat{\alpha}(\hat{\mu}(p_1))v_H + (1 - \hat{I}(\hat{\mu}(p_1)))v_L].$$

The information conveyed by B's purchasing decision is valuable to S. In the second period, she optimally sets $p_2 = v_H$ following acceptance and $p_2 = v_L$ following rejection, as $E[\lambda|p_1, 1] < \lambda^* < E[\lambda|p_1, 0]$ in this case.

Second, suppose S offers $p_1 \leq \tilde{p}$, then S's expected payoff can be written as

$$\Pi_S(p_1) = p_1 + \delta \max\{E[\lambda, \lambda^*\}.$$

Obviously, $p_1 = \tilde{p}$ dominates any $p_1 < \tilde{p}$.

Summarizing the above, in equilibrium S charges some $p_1 \in \{\tilde{p}, v_H\} \cup (\tilde{p}, v_L] \cup [\underline{p}, \overline{p})$, depending on the parameters of the model. The following result establishes that the Bad PBE is really bad. That is, worse for S than the Good PBE.

Proposition 6 (The Bad PBE v.s. the Good PBE) *The expected payoff to S is (weakly) higher in the Good PBE than in the Bad one.*

This result is easily understood. If the Bad PBE involves a first-period price of $p_1 > v_L$, then S can get the same payoff in the Good PBE by adopting the same strategy. If the Bad PBE involves $p_1 \in (\tilde{p}, v_L]$, then the probability of a first-period sale is so low relative to the Good PBE that it outweighes the value of the information obtained in the Bad PBE.

Combining Corollary 4 and Proposition 6 yields the following result.

Corollary 5 (The Seller's Welfare in the Bad PBE) *Suppose $\beta \geq \delta > 0$ and the Bad PBE obtains. Then, S always strictly prefers the Fixed-Price setting to the Contingent-Price setting.*

Corollary 4 indicates that if the Good PBE obtains, the Contingent-Price setting is as good for S as the Fixed-Price one if $\lambda^* \geq E[\lambda]$. In this range for parameter λ^*, S's expected payoff from charging any $p_1 \in (\tilde{p}, v_L]$ in the Bad PBE is strictly lower than her expected equilibrium payoff in the Good PBE. Thus, S strictly prefers the Fixed-Price setting.

If $\lambda^* < E[\lambda]$, S prefers the Fixed-Price setting to the Contingent-Price setting when the Good PBE obtains, and, thus, when the Bad PBE obtains. Hence, she is strictly worse off under the Contingent-Price setting when the Bad PBE obtains relative to the Fixed-Price setting for all values of parameter λ^*.

The final result in this section demonstrates that the Bad PBE is not necessarily bad for B. In particular, in order to preempt signaling by B, S may find it optimal to induce complete pooling by setting a price lower than v_L.

Proposition 7 (Optimal Prices when Beliefs are Pessimistic) *Suppose $\beta > 0$. There exists $\psi > 0$ such that if $\lambda^* \in [E_H - \psi, E_H)$, then in the Bad PBE, S offers $p_1 = \tilde{p}$ and $p_2 = v_L$, and B accepts both offers with probability one.*

The intuition here is easily understood. When λ^* is close to E_H, the value of any information S can obtain is small. Also, $\tilde{p}$ is close to v_L in this case (because $\hat{m}$ is small). S, therefore, preferrs to sell at $\tilde{p}$ with certainty rather than at a slightly higher price with much lower probability.

7 Conclusion

This paper was concerned with learning from a strategic agent in the context of monopoly pricing. Two settings were analyzed: the Contingent-Price setting, where the monopolist learns information about a demand parameter of a specific customer and uses this information to tailor future offers to him, and the Fixed-Price setting, where the monopolist publicly commits to price non-contingently.

It was shown that the buyer fared poorly and the firm well under the Contingent-Price setting when the buyer was myopic. Indeed, the opportunity to price contingently can give the firm an incentive to charge a high 'experimental' price in the first period. Such experimentation unambiguously lowers welfare because the value of information obtained by the firm is outweighed by the loss in expected consumer surplus.

When the buyer is non-myopic, he may strategically reject the firm's first-period offers for one of two reasons. First, in order to conceal information (i.e., to pool), a high-valuation buyer may reject high prices that would never be accepted by a low-valuation buyer. Second, in order to reveal information (i.e., to signal), a low-valuation buyer may reject low prices that would always be accepted by a high-value buyer. Given these strategic reactions, the firm often finds it optimal to post prices that generate no useful information. It was shown that the firm did better committing to price non-contingently when the buyer was farsighted. Lacking this commitment, the buyer possessed strong incentives to manipulate the information acquired by the firm, and this manipulation typically results in either low prices or low sales or both. In short, the ability of a monopolist to learn about the demand characteristics of a strategic consumer through experimental pricing appears to be very limitted.

Appendix

Proof of Proposition 2: When $\beta = 0$, it is evidently optimal for B to accept in period $t = 1, 2$ if and only if $v_t \geq p_t$. Given this behavior of B, S would never charge $p_t \neq v_L$ or v_H.

Case 1: $\lambda^* < E_L$**.** If S sets $p_1 = v_H$, then the expected values of λ conditional on acceptance and rejection are E_H and E_L, respectively. In either case, S sets $p_2 = v_H$ by Proposition 1. On the other hand, if S sets $p_1 = v_L$, then the expected value of λ conditional on acceptance is $E[\lambda]$, and the expected value conditional on rejection is immaterial since rejection does not occur in equilibrium. Hence, S sets $p_2 = v_H$. It is optimal for S to set $p_1 = p_2 = v_H$ rather than $p_1 = v_L$ and $p_2 = v_H$ since

$$v_L + \delta E[\lambda] v_H > E[\lambda] v_H + \delta E[\lambda] v_H$$

holds for $\lambda^* < E_L$.

Case 2: $\lambda^* \in [E_L, \overline{\lambda})$**.** If S sets $p_1 = v_H$, then she optimally sets $p_2 = v_L$ following rejection and $p_2 = v_H$ following acceptance. On the other hand, if S sets $p_1 = v_L$, then she optimally sets $p_2 = v_H$ if $\lambda^* \in [E_L, E[\lambda])$ and $p_2 = v_L$ if $\lambda^* \in (E[\lambda], \overline{\lambda})$.

First, suppose $\lambda^* \in [E_L, E[\lambda])$. It is optimal for S to charge $p_1 = v_H$ and $p_2 = q_1 v_H + (1-q_1) v_L$ rather than $p_1 = v_L$ and $p_2 = v_H$, since

$$E[\lambda] v_H + \delta[E[\lambda] E_H v_H + (1 - E[\lambda]) v_L] > v_L + \delta E[\lambda] v_H$$

holds for $\lambda^* \in [E_L, E[\lambda])$.

Second, suppose $\lambda^* \in (E[\lambda], \overline{\lambda})$. It is optimal for S to charge $p_1 = v_H$ and $p_2 = q_1 v_H + (1-q_1) v_L$ rather than $p_1 = p_2 = v_L$, since

$$E[\lambda] v_H + \delta[E[\lambda] E_H v_H + (1 - E[\lambda]) v_L] > v_L + \delta v_L$$

holds for $\lambda^* \in [E[\lambda], \overline{\lambda})$.

Case 3: $\lambda^* \geq \overline{\lambda}$**.** If S sets $p_1 = v_L$, then she optimally sets $p_2 = v_L$. First, suppose $\lambda^* \in [\overline{\lambda}, E_H)$. If S sets $p_1 = v_H$, then she optimally sets $p_2 = v_H$ following acceptance and $p_2 = v_L$ following rejection. It is optimal for S to charge $p_1 = p_2 = v_L$ rather than $p_1 = v_H$ and $p_2 = q_1 v_H + (1 - q_1) v_L$, since

$$v_L + \delta v_L \geq E[\lambda] v_H + \delta[E[\lambda] E_H v_H + (1 - E[\lambda]) v_L]$$

holds for $\lambda^* \in [\overline{\lambda}, E_H)$.

Second, suppose $\lambda^* \geq E_H$. If S sets $p_1 = v_H$, then she optimally sets $p_2 = v_L$ whether p_1 was accepted or rejected. It is optimal for S to charge $p_1 = p_2 = v_L$ rather than $p_1 = v_H$ and $p_2 = v_L$, since

$$v_L + \delta v_L > E[\lambda] v_H + \delta v_L$$

holds $\lambda^* \geq E_H$. □

Proof of Lemma 2: The proof consists of 3 steps.

Step 1. By way of contradiction, suppose $\theta(p_1, 1) = \theta(p_1, 0)$ and $E[\lambda | p_1, 1] < E[\lambda | p_1, 0]$. Then, Lemma 1 implies the following:

1. If $p_1 > v_H$, then B always rejects the first-period offer. $D_1(p_1) = 0$ in this case.
2. If $p_1 \in (v_L, v_H]$, then B accepts *iff* $v_1 = v_H$. Hence, $E[\lambda|p_1, 1] = E_H > E[\lambda|p_1, 0] = E_L$, which contradicts the supposition.
3. If $p_1 \le v_L$, then always B accepts the first-period offer. $D_1(p_1) = 1$ in this case.

Step 2. By way of contradiction, suppose $\theta(p_1, 1) > \theta(p_1, 0)$. Then Lemma 1 implies the following:

1. If $p_1 > v_H$, then B accepts *iff* $v_1 = v_H$ and $\lambda \ge \lambda'$, where

$$\lambda' = \frac{p_1 - v_H}{\beta(\theta(p_1, 1) - \theta(p_1, 0))(v_H - v_L)}.$$

(If calculated $\lambda' \ge 1$, then B always rejects the first-period offer, $D_1(p_1) = 0$.) Hence, $E[\lambda|p_1, 1] \ge E_H > E[\lambda] \ge E[\lambda|p_1, 0]$. This along with Proposition 1 implies that $\theta(p_1, 1) = 0 \le \theta(p_1, 0)$ must hold, which contradicts the supposition.

2. If $p_1 \in (v_L, v_H]$, then B rejects *iff* $v_1 = v_L$ and $\lambda < \lambda'$, where

$$\lambda' = \frac{p_1 - v_L}{\beta(\phi(p_1, 1) - \phi(p_1, 0))(v_H - v_L)}.$$

(If calculated $\lambda' \ge 1$, then B accepts p_1 *iff* $v_1 = v_H$.) Hence, $E[\lambda|p_1, 1] \ge E[\lambda] > E_L \ge E[\lambda|p_1, 0]$. This along with Proposition 1 implies that $\theta(p_1, 1) \le \theta(p_1, 0) = 1$ must hold, which contradicts the supposition.

3. If $p_1 \le v_L$, then B always accepts the first-period offer. $D_1(p_1) = 1$ in this case.

Step 3. By way of contradiction, suppose $\theta(p_1, 1) \ne \theta(p_1, 0)$ and $E[\lambda|p_1, 1] < E[\lambda|p_1, 0]$. This implies either $E[\lambda|p_1, 1] < \lambda^* \le E[\lambda|p_1, 0]$ or $E[\lambda|p_1, 1] \le \lambda^* < E[\lambda|p_1, 0]$. Thus, it must be $\theta(p_1, 1) > \theta(p_1, 0)$, which cannot happen in equilibrium by Step 2. □

Proof of Lemma 3: Each part is proven in turn.

(i) Differentiating α gives

$$\alpha'(\mu) = \frac{\mu f(\mu) \int_0^\mu (\mu - \lambda)\lambda f(\lambda)\, d\lambda}{(\int_0^\mu \lambda f(\lambda)\, d\lambda)^2}.$$

This is positive for $\mu > 0$.

(ii) Differentiating ρ gives

$$\rho'(\mu) = \frac{\mu f(\mu)[(E[\lambda] - \int_0^\mu \lambda^2 f(\lambda)\, d\lambda) - \mu(1 - \int_0^\mu \lambda f(\lambda)\, d\lambda)]}{(1 - \int_0^\mu \lambda f(\lambda)\, d\lambda)^2}.$$

This is positive for sufficiently small $\mu > 0$. Hence, ρ is initially increasing. Moreover, setting the above expression equal to zero establishes that ρ has a unique critical point where it crosses the 45-degree line. Hence, ρ is increasing up to this point and decreasing thereafter.

(iii) Equating $\alpha(m_{min})$ and $\rho(m_{min})$ and performing simple algebra reveals $\alpha(m_{min}) = E[\lambda]$. □

Proof of Lemma 4: Each part is proven in turn.

(i) Consider $p_1 \in (v_L, \underline{p}]$. Given B's strategy to accept p_1 *iff* $v_1 = v_H$, S's posterior beliefs are $E[\lambda|p_1, 0] = E_L < \lambda^*$ and $E[\lambda|p_1, 1] = E_H > \lambda^*$. It then follows from Proposition 1 that setting $\theta(p_1, 0) = 1$ and $\theta(p_1, 1) = 0$ is optimal. By Lemma 1, B accepts p_1 *iff*

$$v_1 - p_1 \geq \beta\lambda(v_H - v_L),$$

or $v_1 = v_H$.

(ii) Consider $p_1 \in (\underline{p}, \overline{p}]$. Given B's strategy to accept p_1 *iff* $v_1 = v_H$ and $\lambda \leq \mu(p_1)$, S's posterior beliefs are $E[\lambda|p_1, 0] = \rho(\mu(p_1)) \leq \lambda^*$ and $E[\lambda|p_1, 1] = \alpha(\mu(p_1)) \geq \lambda^*$. It then follows from Proposition 1 that setting $\theta(p_1, 0) = 1$ and $\theta(p_1, 1) = 0$ is optimal. By Lemma 1, B accepts p_1 *iff*

$$v_1 - p_1 \geq \beta\lambda(v_H - v_L),$$

or $v_1 = v_H$ and $\lambda \leq \mu(p_1)$.

(iii) Consider $p_1 \in (\overline{p}, v_H]$ and suppose $\lambda^* > E[\lambda]$. Given B's strategy to accept p_1 *iff* $v_1 = v_H$ and $\lambda \leq m$, S's posterior beliefs are $E[\lambda|p_1, 0] = \rho(m) < \lambda^*$ and $E[\lambda|p_1, 1] = \alpha(m) = \lambda^*$. It then follows from Proposition 1 that setting $\theta(p_1, 0) = 1$ and any $\theta(p_1, 1) \in [0, 1]$ is optimal. Mixing probability $\theta(p_1, 1)$ is calibrated to make B with $v_1 = v_H$ and $\lambda = m$ indifferent between accepting and rejecting,

$$v_1 - p_1 = \beta m(1 - \theta(p_1, 1))(v_H - v_L).$$

Now suppose $\lambda^* < E[\lambda]$. Given B's strategy to accept p_1 *iff* $v_1 = v_H$ and $\lambda \leq m$, S's posterior beliefs are $E[\lambda|p_1, 0] = \rho(m) = \lambda^*$ and $E[\lambda|p_1, 1] = \alpha(m) > \lambda^*$. It then follows from Proposition 1 that setting $\theta(p_1, 1) = 0$ and any $\theta(p_1, 0) \in [0, 1]$ is optimal. Mixing probability $\theta(p_1, 0)$ is calibrated to make B with $v_1 = v_H$ and $\lambda = m$ indifferent between accepting and rejecting,

$$v_1 - p_1 = \beta m\theta(p_1, 0)(v_H - v_L).$$

□

Proof of Lemma 5: Each part is proven in turn.

(i) Differentiating $\hat{\alpha}$ gives

$$\hat{\alpha}'(\hat{\mu}) = \frac{-(1-\hat{\mu})f(\hat{\mu})\left[\left(E[\lambda] - \int_{\hat{\mu}}^{1} \lambda(1-\lambda)f(\lambda)\,d\lambda\right) - \hat{\mu}\left(1 - \int_{\hat{\mu}}^{1}(1-\lambda)f(\lambda)\,d\lambda\right)\right]}{\left(1 - \int_{\hat{\mu}}^{1}(1-\lambda)f(\lambda)\,d\lambda\right)^2}.$$

This is negative for small $\hat{\mu} > 0$. Hence, $\hat{\alpha}$ is initially decreasing. Moreover, setting the above expression equal to zero establishes that $\hat{\alpha}$ has a unique critical point where it crosses the 45-degree line. Hence, $\hat{\alpha}$ is decreasing up to this point and increasing thereafter.

(ii) Differentiating $\hat{\rho}$ gives

$$\hat{\rho}'(\hat{\mu}) = \frac{(1-\hat{\mu})f(\hat{\mu})\int_{\hat{\mu}}^{1}(\lambda - \hat{\mu})(1-\lambda)f(\lambda)\,d\lambda}{\left(\int_{\hat{\mu}}^{1}(1-\lambda)f(\lambda)\,d\lambda\right)^2}.$$

This is strictly positive for $\hat{\mu} < 1$.

(ii) Equating $\hat{\alpha}$ and $\hat{\rho}$ and performing simple algebra reveals $\hat{\alpha}(\hat{m}_{max}) = E[\lambda]$. □

Proof of Lemma 6: First, consider the Good PBE. Given B's strategy to accept $p_1 \leq v_L$, S's beliefs remains unchanged, $E[\lambda|p_1, 1] = E[\lambda]$ and $E[\lambda|p_1, 0] = E[\lambda]$ (immaterial). It then follows from Proposition 1 that setting $\theta(p_1, 1) = \theta(p_1, 0) = 0$ if $E[\lambda] > \lambda^*$ and $\theta(p_1, 1) = \theta(p_1, 0) = 1$ if $E[\lambda] < \lambda^*$ is optimal. By Lemma 1, B has no incentives to reject p_1 as

$$v_L - p_1 \geq \beta\lambda(\theta(p_1, 0) - \theta(p_1, 1))(v_L - v_H)$$

holds for all $\lambda \in [0, 1]$.

Second, consider the Bad PBE.

(i) Suppose $p_1 \leq \tilde{p}$. Given B's strategy to accept the price, S's beliefs are $E[\lambda|p_1, 1] = E[\lambda]$ and $E[\lambda|p_1, 0] = E[\lambda]$ (immaterial). It then follows from Proposition 1 that setting $\theta(p_1, 1) = \theta(p_1, 0) = 0$ if $E[\lambda] > \lambda^*$ and $\theta(p_1, 1) = \theta(p_1, 0) = 1$ if $E[\lambda] < \lambda^*$ is optimal. By Lemma 1, B has no incentives to reject p_1 as

$$v_L - p_1 \geq \beta\lambda(\theta(p_1, 0) - \theta(p_1, 1))(v_L - v_H)$$

holds for all $\lambda \in [0, 1]$.

(ii) Suppose $p_1 \in (\tilde{p}, v_L]$. Given B's strategy to reject p_1 *iff* $v_1 = v_L$ and $\lambda > \hat{\mu}(p_1)$, S's beliefs are $E[\lambda|p_1, 1] = \hat{\alpha}(\hat{\mu}(p_1)) \geq \lambda^*$ and $E[\lambda|p_1, 0] = \hat{\rho}(\hat{\mu}(p_1)) \leq \lambda^*$. It then follows from Proposition 1 that setting $\theta(p_1, 1) = 0$ and $\theta(p_1, 0) = 1$ is optimal. By Lemma 1, B accepts p_1 *iff*

$$v_L - p_1 \geq \beta\lambda(v_L - v_H),$$

or $v_1 = v_H$ and $\lambda \leq \hat{\mu}(p_1)$. □

Proof of Proposition 4: Offering $p_1 = v_L$ dominates $p_1 = v_H$ if

$$v_L + \delta v_L > I(m)v_H + \delta v_L.$$

But, this follows from $\lambda^* \geq E[\lambda] > I(m)$.

Thus, it is left to show that $p_1 = v_L$ dominates all $p_1 \in [\underline{p}, \overline{p})$, or

$$v_L + \delta v_L > I(\mu)(v_H - \beta\mu(v_H - v_L)) + \delta[I(\mu)\alpha(\mu)v_H + (1 - I(\mu))v_L]$$

for all $\mu \in (m, 1]$. Note that the right side of this inequality is increasing in β. Hence, if it holds for $\beta = \delta$, then it holds for all $\beta \geq \delta$. The condition may, therefore, be recast as

$$(v_L - I(\mu)v_H) + \delta I(\mu)(1 - \mu)v_L + \delta[I(\mu)\mu - I(\mu)\alpha(\mu)]v_H > 0.$$

The first two terms of this expression are non-negative. The third term is positive, as

$$I(\mu)\mu - I(\mu)\alpha(\mu) = \int_0^\mu (\mu - \lambda)\lambda f(\lambda)\, d\lambda > 0$$

for all $\mu \in (m, 1]$. □

Proof of Proposition 5: Each part is proven in turn.

(i) Let $\lambda^* \equiv E[\lambda] - \epsilon$ for some $\epsilon > 0$. First, offering $p_1 = v_L$ dominates $p_1 = v_H$ if

$$v_L + \delta E[\lambda] v_H > I(m) v_H + \delta E[\lambda] v_H,$$

or

$$v_L > I(m) v_H.$$

Plugging in for $\lambda^* \equiv E[\lambda] - \epsilon$ gives

$$E[\lambda] - \epsilon \geq I(m(\epsilon)),$$

where

$$m(\epsilon) = \rho^{-1}(E[\lambda] - \epsilon).$$

Observe that $\lim_{\epsilon \to 0} m(\epsilon) = m_{min}$. Taking the limit of the above inequality as ϵ goes to zero gives

$$E[\lambda] > I(m_{min}).$$

Thus, the condition is satisfied for $\epsilon > 0$ sufficiently small.

Second, offering $p_1 = v_L$ dominates all $p_1 \in [\underline{p}, \overline{p})$ if

$$v_L + \delta E[\lambda] v_H > I(\mu)(v_H - \beta\mu(v_H - v_L)) + \delta[I(\mu)\alpha(\mu) v_H + (1 - I(\mu)) v_L]$$

for all $\mu \in (m, 1]$. Note that the right side of this inequality is increasing in β. Hence, if it holds for $\beta = \delta$, then it holds for all $\beta \geq \delta$. The condition may, therefore, be recast as

$$\delta(E[\lambda] v_H - v_L) + \delta I(\mu)(1 - \mu) v_L + \delta[I(\mu)\mu - I(\mu)\alpha(\mu)] v_H + (v_L - I(\mu) v_H) > 0.$$

The first two terms of this expression are non-negative. The third term is positive (see the proof of Proposition 4). The forth term is positive for $\epsilon > 0$ sufficiently small.

(ii) Let $\lambda^* \equiv E_L + \xi$ for some $\xi > 0$. First, offering $p_1 = v_H$ dominates offering $p_1 = v_L$ if

$$I(m) v_H + \delta E[\lambda] v_H > v_L + \delta E[\lambda] v_H,$$

or

$$I(m) v_H > v_L.$$

Plugging in for $\lambda^* \equiv E_L + \xi$ gives

$$I(m(\xi)) \geq E_L + \xi,$$

where

$$m(\xi) \equiv \rho^{-1}(E_L + \xi).$$

Observe that $\lim_{\xi \to 0} m(\xi) = 1$. Taking the limit of the above inequality as ξ goes to zero gives

$$E[\lambda] > E_L.$$

Thus, the condition is satisfied for $\xi > 0$ sufficiently small.

Second, offering $p_1 = v_H$ dominates all $p_1 \in [\underline{p}, \overline{p})$ if

$$I(m) v_H + \delta E[\lambda] v_H > I(\mu)(v_H - \beta\mu(v_H - v_L)) + \delta[I(\mu)\alpha(\mu) v_H + (1 - I(\mu)) v_L]$$

for all $\mu \in (m, 1]$. Note that the right side of this inequality is increasing in β. Hence, if it holds for $\beta = \delta$, then it holds for all $\beta \geq \delta$. The condition may, therefore, be recast as

$$\delta[E[\lambda]v_H + I(\mu)\mu(v_H - v_L) - I(\mu)\alpha(\mu)v_H - (1 - I(\mu))v_L] > (I(\mu) - I(m))v_H.$$

Observe that as ξ goes to zero, interval $(m, 1]$ shrinks to $\{1\}$. Thus, if the above condition holds strictly at the limit as μ goes to 1, then it holds for all $\mu \in (m, 1]$ when ξ is sufficiently small. Taking the limit as μ goes to 1 gives

$$\delta(2E[\lambda] - E[\lambda^2] - E_L) > 0,$$

which evidently holds since $E[\lambda] > E[\lambda^2]$ and $E[\lambda] > E_L$. □

Proof of Proposition 6: Suppose that in the Bad PBE, S offers $p_1 > v_L$, then she can get the same payoff in the Good PBE by adopting the same strategy. Hence, suppose $p_1 \in \{\tilde{p}\} \cup (\tilde{p}, v_L]$. There are two cases to consider.

Case 1: $\lambda^* \geq E[\lambda]$**.** S's payoff in the Good PBE is at least $v_L + \delta v_L$. First, suppose S offers $p_1 = \tilde{p}$ in the Bad PBE, then her expected payoff is

$$\tilde{p} + \delta v_L < v_L + \delta v_L.$$

Second, suppose S offers $p_1 \in (\tilde{p}, v_L]$. It must be shown that

$$\hat{I}(\hat{\mu})(v_L - \beta\hat{\mu}(v_H - v_L)) + \delta[\hat{I}(\hat{\mu})\hat{\alpha}(\hat{\mu})v_H + (1 - \hat{I}(\hat{\mu}))v_L] \leq v_L + \delta v_L$$

for all $\hat{\mu} \in (0, \hat{m}]$. If this holds for all $\delta \leq 1$, then it holds for $\delta = 1$. The condition may, therefore, be recast as

$$\beta\hat{I}(\hat{\mu})\hat{\mu}(v_H - v_L) + (v_L - \hat{I}(\hat{\mu})\hat{\alpha}(\hat{\mu})v_H) \geq 0.$$

The first term is non-negative, the second term is also non-negative, as

$$v_L - \hat{I}(\hat{\mu})\hat{\alpha}(\hat{\mu})v_H = v_L - \left(E[\lambda] - \int_{\hat{\mu}}^{1} \lambda(1 - \lambda)f(\lambda)\, d\lambda\right) v_H \geq (\lambda^* - E[\lambda])v_H \geq 0.$$

Case 2: $\lambda^* < E[\lambda]$**.** S's payoff in the Good PBE is at least $v_L + \delta v_H$. First, suppose S offers $p_1 = \tilde{p}$ in the Bad PBE, then her expected payoff is

$$\tilde{p} + \delta v_H < v_L + \delta v_H.$$

Second, suppose S offers $p_1 \in (\tilde{p}, v_L]$. It must be shown that

$$\hat{I}(\hat{\mu})(v_L - \beta\hat{\mu}(v_H - v_L)) + \delta[\hat{I}(\hat{\mu})\hat{\alpha}(\hat{\mu})v_H + (1 - \hat{I}(\hat{\mu}))v_L] \leq v_L + \delta v_H,$$

for all $\hat{\mu} \in (0, \hat{m}]$. If this holds for all $\delta \leq 1$, then it holds for $\delta = 1$. The condition may, therefore, be recast as

$$\beta\hat{I}(\hat{\mu})\hat{\mu}(v_H - v_L) + v_H(1 - \hat{I}(\hat{\mu})\hat{\alpha}(\hat{\mu})) \geq 0.$$

Obviously, both terms of the left side of the condition are non-negative. □

Proof of Proposition 7: Let $\lambda^* \equiv E_H - \psi$ for some $\psi > o$. Also, let

$$m(\psi) \equiv \alpha^{-1}(E_H - \psi)$$

and

$$\hat{m}(\psi) \equiv \hat{\alpha}^{-1}(E_H - \psi).$$

Observe that $\lim_{\psi \to 0} m(\psi) = 1$ and $\lim_{\psi \to 0} \hat{m}(\psi) = 0$.

First, offering $p_1 = \tilde{p}$ dominates all $p_1 \in (\tilde{p}, v_L]$ if

$$\tilde{p} + \delta v_L > \hat{I}(\hat{\mu})(v_L - \beta\hat{\mu}(v_H - v_L)) + \delta[\hat{I}(\hat{\mu})\hat{\alpha}(\hat{\mu})v_H + (1 - \hat{I}(\hat{\mu}))v_L]$$

for all $\hat{\mu} \in (0, \hat{m}]$. Note that if the condition holds for all $\delta \leq 1$, then it holds for $\delta = 1$. The condition may, therefore, be recast as

$$-\beta\hat{I}(\hat{\mu})\hat{\mu}(v_H - v_L)) + \hat{I}(\hat{\mu})\hat{\alpha}(\hat{\mu})v_H < \tilde{p}.$$

Taking the limit as ψ goes to zero gives

$$E[\lambda]E_H < E_H.$$

Thus, the condition is satisfied for $\psi > 0$ sufficiently small.

Second, offering $p_1 = \tilde{p}$ dominates all $p_1 \in [\underline{p}, \overline{p})$ if

$$\tilde{p} + \delta v_L > I(\mu)(v_H - \beta\mu(v_H - v_L)) + \delta(I(\mu)\alpha(\mu)v_H + (1 - I(\mu))v_L)$$

for all $\mu \in (m, 1]$. Note that if the condition holds for all $\delta \leq 1$, then it holds for $\delta = 1$. The condition may, therefore, be recast as

$$-\beta I(\mu)\mu(v_H - v_L)) + I(\mu)\alpha(\mu)v_H < \tilde{p}.$$

Taking the limit as ψ goes to zero gives

$$-\beta E[\lambda](1 - E_H) + E[\lambda]E_H < E_H.$$

Thus, the condition is satisfied for $\psi > 0$ sufficiently small.

Finally, offering $p_1 = \tilde{p}$ dominates $p_1 = v_H$ if

$$\tilde{p} + \delta v_L > I(m)v_H + \delta v_L,$$

or

$$\tilde{p} > I(m)v_H.$$

Taking the limit as ψ goes to zero gives

$$E_H > E[\lambda].$$

Thus, the condition is satisfied for $\psi > 0$ sufficiently small. □

References

Acquisti, Alessandro, and Hal R. Varian: 2002, "Conditioning Prices on Purchase History," Mimeo, SIMS UC Berkeley.

Aghion, Philippe, Patrick Bolton, Christopher Harris, and Bruno Jullien: 1991, "Optimal Learning by Experimentation," *Review of Economic Studies,* 58, 621–54.

Armstrong, Mark: 1996, "Multiproduct Nonlinear Pricing," Econometrica, 64, 51–75.

Fudenberg, Drew, and Jean Tirole: 2000, "Customer Poaching and Brand Switching," *RAND,* 31, 634–57.

Hart, Oliver D., and Jean Tirole: 1988, "Contract Renegotiation and Coasian Dynamics," *Review of Economic Studies,* 55, 509–40.

Keller, Godfrey, and Sven Rady: 1999, "Optimal Experimentation in a Changing Environment," *Review of Economic Studies,* 66, 475–507.

Kennan, John: 2001, "Repeated Bargaining with Persistent Private Information," *Review of Economic Studies,* 68, 719–55.

Krugman, Paul: 2000, "Reckonings; What Price Fairness?" *The New York Times,* Oct, 4, A35.

Laffont, Jean-Jacques, and Jean Tirole: 1988, "The Dynamics of Incentive Contracts," *Econometrica* 56, 1153–75.

Lazear, Edward P.: 1986, "Retail Pricing and Clearance Sales," *American Economic Review,* 76, 14–32.

Mirman, Leonard J., Samuelson, Larry, and Amparo Urbano: 1993, "Monopoly Experimentation," *International Economic Review,* 34, 549–63.

Rochet, Jean-Charles, and Philippe Chone: 1998, "Ironing, Sweeping, and Multidimensional Screening," *Econometrica,* 66, 783–826.

Rothschild, Michael: 1974, "A Two-Armed Bandit Theory of Market Pricing," *Journal of Economic Theory,* 9, 185–202.

Rustichini, Aldo, and Asher Wolinsky: 1995, "Learning about Variable Demand in the Long Run," *Journal of Economic Dynamics and Control,* 19, 1283–92.

Segal, Ilya R.: 2002, "Optimal Pricing Mechanisms with Unknown Demand," *American Economic Review,* forthcoming.

Shaffer, Greg, and Z. John Zhang: 2000, "Pay to Switch or Pay to Stay: Preference-Based Price Discrimination in Markets with Switching Costs," *Journal of Economics and Management Strategy,* 9, 397–424.

Streitfield, David: 2000 "On the Web Price Tags Blur: What You Pay Could Depend on Who You Are," *The Washington Post,* sep. 27, A1.

Taylor, Curtis R.: 2000 "Supplier Surfing: Competition and Consumer Behavior in Subscription Markets," *RAND,* forthcoming.

Villas-Boas, J. Miguel: 1999, "Dynamic Competition with Customer Recognition," *RAND,* 30, 604–31.