

NBER WORKING PAPER SERIES

PRICING WITH LIMITED KNOWLEDGE OF DEMAND

Maxime C. Cohen
Georgia Perakis
Robert S. Pindyck

Working Paper 21679
<http://www.nber.org/papers/w21679>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue
Cambridge, MA 02138
October 2015

We thank Daron Acemoglu, Gerard Cachon, Dennis Carlton, Adam Elmachtoub, John Hauser, Teck Ho, Mahesh Nagarajan, Olivier Rubel, Richard Schmalensee, David Schmittlein, Gal Shulkind, Duncan Simester, Chris Tang, Catherine Tucker, Hal Varian, Dimitri Vayanos, Rakesh Vohra and Lawrence White for helpful comments and suggestions. The authors declare that they have no material financial interests that relate to the research described in this paper. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2015 by Maxime C. Cohen, Georgia Perakis, and Robert S. Pindyck. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Pricing with Limited Knowledge of Demand
Maxime C. Cohen, Georgia Perakis, and Robert S. Pindyck
NBER Working Paper No. 21679
October 2015
JEL No. D40,D80,D81

ABSTRACT

How should a firm price a new product for which little is known about demand? We propose a simple pricing rule: the firm only estimates the maximum price it can charge and still expect to sell at least some units, and then sets price as though the actual demand curve were linear. We show that if the true demand curve is one of many commonly used demand functions, or even if it is a more complex function, and if marginal cost is known and constant, the firm can expect its profit to be close to what it would earn if it knew the true demand curve. We derive analytical performance bounds for a variety of demand functions, and calculate expected profit performance for randomly generated demand curves.

Maxime C. Cohen
Google Research
New York, NY 10011
maxccohen@gmail.com

Georgia Perakis
Sloan School of Management
M.I.T.
Cambridge, MA 02142
georgiap@mit.edu

Robert S. Pindyck
MIT Sloan School of Management
100 Main Street, E62-522
Cambridge, MA 02142
and NBER
RPINDYCK@MIT.EDU

1 Introduction

Firms that introduce new products often set price with little or no knowledge of demand, and no data on which to base elasticity estimates. Examples include pharmaceutical companies introducing new types of drugs (Lilly’s Prozac in 1987), technology companies introducing new products or services (Apple setting the price of music downloads when launching its iTunes store in 2002, and more recently, Apple’s iWatch and Google’s Chromecast), or a company introducing an existing product in an emerging market (P&G launching Pampers in China in 1998). Although marginal cost may be easy to estimate (it is close to zero for most drugs and music or software downloads, and is known from experience for diapers), the firms are likely to know little or nothing about the demand curves they face, and may not even be able to estimate arc elasticities. How should firms set prices in such settings?

As discussed below, this problem has been the subject of a variety of studies, most of which focus on experimentation and learning, e.g., setting different prices and observing the outcomes. Experimenting with price, however, is usually not feasible or desirable; it is more common for firms to choose and maintain a particular price for at least a year.¹ We examine a much simpler approach to this pricing problem that involves no experimentation.

We show that in many situations, the firm can use a remarkably simple pricing rule. The use of the rule requires that (i) the firm’s marginal cost is known and constant, (ii) the firm need not know in advance the quantity it will sell, and (iii) the firm can estimate the *maximum price* it can charge and still expect to sell at least some units. (We relax this last assumption in Section 3.) Denoting that maximum price by P_m , the firm sets price as though the actual demand curve were linear, i.e.,

$$P = P_m - bQ . \tag{1}$$

Assuming marginal cost, c , is constant, the firm’s profit-maximizing price is then $P^* = (P_m + c)/2$, which we refer to as the “linear price.” This price is independent of the slope b of the linear demand curve, although the resulting quantity, $Q_L^* = (P_m - c)/2b$, is not. But as long as the firm does not need to invest in production capacity or otherwise plan on a particular sales level (as would be the case for most new drugs, music downloads, or

¹This was the case when pharmaceutical companies introduced new drugs, when Apple launched the iTunes store, when Google introduced Chromecast, and when P&G launched Pampers in China.

software), knowledge of b , and thus the ability to predict its sales, is immaterial. More sales are better than less, but the only problem at hand is to set the price. Denote the resulting price and profit from using eqn. (1) by P^* and Π^* respectively.

How well can the firm expect to do if it sets the price P^* ? Suppose that with precise knowledge of its true demand curve, the firm would have set a different price P^{**} and earn a (maximum) profit Π^{**} . The question we address is simple: How close can we expect Π^* to be relative to Π^{**} , i.e., how well is the firm likely to do using this simple pricing rule? As we will show, if the true demand curve is one of many commonly used demand functions, or even if it is a more complex function, the firm can expect to do very well.

Determining the maximum price P_m might not be easy, but it is a much less difficult task than estimating an entire demand curve. A pharmaceutical company might estimate P_m by comparing a new drug to existing therapies (including non-drug therapies).² And when it planned to sell music through iTunes, Apple might have estimated P_m to be around \$2 or \$3 per song, as a multiple of the per-song price of compact discs.³

The basic idea behind this paper is quite simple, and is illustrated in Figure 1. The demand curve labeled “Actual Demand” was drawn so it might apply to a new drug, or to music downloads in the early years of the iTunes store. A pharmaceutical company might estimate a price P_m at which some doctors will prescribe and some consumers will buy its new drug, even if insurance companies refuse to pay for it. As the price is lowered and the drug receives insurance coverage, the quantity demanded expands considerably. At some point the market saturates so that even if the price is reduced to zero there will be no further increase in sales. For music downloads, at prices above P_m it is more economical to buy the CD and “rip” the desired songs to one’s computer. At lower prices demand expands rapidly, and at some point the market saturates.

If the firm knew this curve, it would set the profit-maximizing price P^{**} and expect to sell the quantity Q^{**} . (In the figure, P^{**} and Q^{**} are computed numerically by maximizing profit.) But the firm does not know the actual demand curve. A linear demand curve that starts at P_m has also been drawn, and labeled D_L . This linear demand curve implies a

²For example, when pricing Prilosec, the first proton-pump inhibitor anti-ulcer drug, Astra-Merck could expect P_m to be two or three times higher than the price of Zantac, an older generation anti-ulcer drug.

³A CD with 12 songs might cost \$12 to \$15, but most consumers would want only a few of those songs.




Figure 1: Illustration for a representative demand curve

profit-maximizing price P^* and quantity Q_L^* , where the subscript L refers to the quantity sold if D_L were the true demand curve. (Note that the slope of D_L is immaterial for the pricing decision.) The actual quantity that would be sold given the price P^* is Q^* (under the actual demand curve). How badly would the firm do by pricing at P^* instead of P^{**} ? For the demand curve and marginal cost shown in Figure 1, the profit and price ratios (determined numerically) are $\Pi^{**}/\Pi^* = 1.023$ and $P^{**}/P^* = 1.069$, i.e., the resulting profit is within a few percent of what the firm could earn if knew the actual demand curve and used it to set price. (The firm would do worse if $c = 0$, in which case $\Pi^{**}/\Pi^* = 1.084$.)

There are certainly demand curves for which this pricing rule will perform very poorly. For example, suppose the true demand curve is a rectangle, i.e., $P = P_m$ for $0 \leq Q \leq Q_m$ and $P = 0$ for $Q > Q_m$. Then the profit-maximizing price is clearly P_m , and the resulting profit is $\Pi^{**} = (P_m - c)Q_m$. Setting a price $P^* = (P_m + c)/2$ will yield a much lower profit; in fact $\Pi^{**}/\Pi^* = 2.0$. We want to know how well our pricing rule will perform — i.e., what is Π^{**}/Π^* — for alternative “true” demand curves.

There is a large literature on optimal pricing with limited knowledge of demand, much of which deals with experimentation and learning. An early example is Rothschild (1974),

who assumes that a firm chooses from a finite set of prices, observes outcomes, and because each trial is costly, eventually settles on the price that it thinks (perhaps incorrectly) is optimal. The firm’s choice is then the solution of a multi-armed bandit problem. (In the simplest version of the model, the firm prices “high” or “low.”) The solution does not involve estimating a demand curve.

Other studies focus on learning in a parametric or non-parametric context. Several address the use of learning to update estimates of parameters of a known demand function; see, e.g., Aviv and Pazgal (2005), Bertsimas and Perakis (2006), Lin (2006) and Farias and Van Roy (2010). A second stream examines the interplay between learning and optimizing revenues over time without assuming a parametric form. Following Rothschild (1974), several authors assume the seller first sets a price to learn about demand, and then adjusts the price to optimize revenues (see, e.g., Besbes and Zeevi (2009) and Araman and Caldentey (2011)).

The operations research literature examines dynamic pricing using robust optimization, where the functional form of the demand curve is known but one or more parameters are only known to lie in an “uncertainty set.” For example, demand might depend on two unknown parameters α_1 and α_2 , so the profit function is $\Pi(\alpha_1, \alpha_2, p)$. The price p is chosen to maximize the worst possible outcome over the uncertainty set, i.e., $\max_p \min_{\alpha_1, \alpha_2} \Pi(\alpha_1, \alpha_2, p)$.⁴ In related work, Bergemann and Schlag (2011) consider a single consumer’s valuation, with a distribution that is unknown but assumed to be in a neighborhood of a given model distribution. The authors characterize robust pricing policies that maximize the seller’s minimum profit (maximin), or that minimize worst-case regret (difference between the true valuation and the realized profit). Although robust optimization incorporates uncertainty, its focus on worst-case scenarios yields conservative pricing strategies.

Our paper is also related to studies of model misspecification. In particular, we show that a simple linear demand model can perform well even if the true demand curve is far from linear. Others have likewise shown that linear models can perform well (e.g., Dawes (1979) in clinical prediction and Carroll (2014) in contract theory). Besbes and Zeevi (2013) study the “price of misspecification” for dynamic pricing with demand learning. They investigate the revenue loss in a multi-period setting incurred if the seller uses a simple parametric demand

⁴See, e.g., Adida and Perakis (2006) and Thiele (2009). An alternative is the “distributionally robust” approach, where price is robust with respect to a class of demand distributions with similar parameters such as mean and variance. See, e.g., Lim and Shanthikumar (2007), and Ball and Queyranne (2009).

model that differs significantly from the true demand curve.

Our approach to pricing is quite different from the studies cited above, and is related to the prescriptive rules of thumb found, for example, in Shy (2006). Managers often seek simple and robust rules for pricing (and other decisions such as levels of advertising or R&D), and other studies have shown that simple rules can be very effective.⁵ The pricing rule we suggest is certainly simple; the extent to which it is effective is the focus of this paper.

In the following sections, we characterize the performance of our pricing rule by deriving analytical bounds for the profit ratio Π^{**}/Π^* for several classes of demand curves: quadratic, monomial, semi-log and log-log. We also find bounds for the ratio Π^{**}/Π^* in the case of a general concave demand, and treat the case of a maximum price that is not known exactly. We then examine randomly generated “true” demand curves, and determine computationally the expected profit ratio Π^{**}/Π^* for our pricing rule, and confidence bounds for the ratio. Finally, we examine the welfare implications of our pricing rule.

2 Common Demand Functions

Here we examine several demand models — quadratic, monomial, polynomial, semi-log and log-log. For each we derive performance bounds comparing the profits from our simple pricing rule to the profits that would result if the actual demand function were known. We will see that in most cases the profit ratio is close to one.

Before proceeding, we note that the relationship between the linear price P^* and the optimal price P^{**} depends on the convexity properties of the actual demand function. In the Appendix we show that if the actual inverse demand curve is convex (concave) with respect to Q , the linear price is greater (smaller) than the optimal price:

Result 1. *If the actual inverse demand curve $P_A(Q)$ is convex with respect to Q , then $P^{**} \leq P^*$, and if $P_A(Q)$ is concave, $P^{**} \geq P^*$.*

⁵In related work, Chu, Leslie and Sorensen (2011) show how “Bundle Size Pricing” (BSP) provides a close approximation to optimal mixed bundling. In BSP, a price is set for each good, for any bundle of two, for any bundle of three, etc., up to a bundle of all the goods produced. Profits are close to what would be obtained from mixed bundling. Also Carroll (2014) examines a principal who has only limited knowledge of what an agent can do, and wants to write a contract robust to this uncertainty. He shows that the most robust contract is a linear one — e.g., the agent is paid a fixed fraction of output. Hansen and Sargent (2008) provide a general treatment of robust control, i.e., optimal control with model uncertainty.

Note that we only need $P_A(Q)$ to be convex (or concave) in the range $[0, Q^{**}]$ and not everywhere. The value of Q^{**} might not be known, but this result can still be useful in that it tells us whether our simple rule will over- or under-price relative to the optimal price, and it might be possible to correct for this error by adjusting the price up or down.

2.1 Quadratic Demand

Suppose the actual inverse demand function is a quadratic:

$$P_A(Q) = P_m - b_1Q + b_2Q^2, \quad (2)$$

where, as before, P_m is the maximum price. We want analytical bounds for the profit ratio Π^{**}/Π^* and price ratio P^{**}/P^* . The bounds depend on the convexity properties of the function in (2) and are summarized in the following result. (Proofs are in the Appendix.)

Result 2. *For the quadratic demand curve of eqn. (2), the profit and price ratios satisfy:*

- **Convex case:** $b_1, b_2 \geq 0$ and $b_2 \leq b_1^2/4P_m$

$$1 \leq \frac{\Pi^{**}}{\Pi^*} \leq \frac{8\sqrt{2}}{27(\sqrt{2}-1)} = 1.0116$$

$$\frac{8}{9} \leq \frac{P^{**}}{P^*} \leq 1$$

- **Concave case:** $b_1 \geq 0$ and $b_2 \leq 0$

$$1 \leq \frac{\Pi^{**}}{\Pi^*} \leq \frac{4\sqrt{2}}{3\sqrt{3}} = 1.0887$$

$$1 \leq \frac{P^{**}}{P^*} \leq \frac{2}{3} \left(\frac{2P_m + c}{P_m + c} \right) \leq \frac{4}{3} = 1.33$$

Note that the restrictions on the values of b_1 and b_2 are necessary and sufficient conditions to guarantee that the inverse demand curve is non-negative and non-increasing everywhere.

If demand is convex, the simple pricing rule yields a profit that is only about 1% less than what the firm could achieve if it knew the true demand curve. Also, this is a “worst case” result that applies when $c = 0$; if $c > 0$, the ratio Π^{**}/Π^* is even closer to 1. The price P^* can be as much as 12% lower than the optimal price P^{**} , but the concern of the firm is (or should be) its profit performance. (Also, P^{**}/P^* deviates the most from 1 when $c = 0$.)

If demand is concave, the resulting profit Π^* is within 9% of the optimal profit, irrespective of the parameters b_1 and b_2 . In the proof of Result 2 in the Appendix, we show that the largest value of Π^{**}/Π^* (1.0887) occurs when $b_1 = 0$; for positive values of b_1 , the profit ratio is closer to 1. The reason is that when b_1 increases, the curve becomes closer to a linear function. In addition, one can show that the profit ratio becomes closer to 1 for the concave case when either c or b_2 increase (recall that $b_2 \leq 0$).

2.2 Monomial Demand

Now suppose the actual inverse demand curve is a monomial of order n :

$$P_A(Q) = P_m - \gamma Q^n, \quad \gamma > 0. \quad (3)$$

Note that all functions of the form (3) are concave and decreasing, given that $\gamma > 0$. The Appendix shows that the profit and price ratios are now:

Result 3. *For the inverse demand curve of eqn. (3), the profit and price ratios satisfy:*

$$1 \leq \frac{\Pi^{**}}{\Pi^*} = \frac{2^{\frac{1}{n}+1}n}{(n+1)^{\frac{1}{n}+1}} \leq 2$$

$$1 \leq \frac{P^{**}}{P^*} = \frac{2(nP_m + c)}{(n+1)(P_m + c)} \leq 2$$

Note that the above results are not bounds but equalities. Also, for any monomial demand curve, the profit ratio only depends on the order of the monomial n ; it does not depend on the values of P_m , c or γ . (The price ratio does depend on P_m , c and n , but not on γ .) Both ratios are monotonically increasing with the degree of the monomial n and converge to 2 and $2P_m/(P_m + c) \leq 2$ respectively, as $n \rightarrow \infty$. For monomials of order 3 and 4, the profit ratios are 1.19 and 1.27 respectively.

These results can be extended to the case of polynomials with non-positive coefficients. In particular, one can show that for a polynomial demand function of order n with non-positive coefficients, the profit and price ratios are the farthest from 1 for the corresponding monomials. By adding non-positive terms to the monomial, the demand curve becomes more linear, improving the performance of our pricing rule. In short, if the true demand curve is a low-order monomial or polynomial, our pricing rule yields profits that are reasonably close to what the firm could achieve if it knew the true demand curve.

2.3 Semi-Log Demand

Now consider the semi-log inverse demand curve:

$$P_A(Q) = P_m e^{-\alpha Q}, \quad \alpha > 0. \quad (4)$$

The following result (proof in Appendix) shows the profit and price ratios for the case where marginal cost $c = 0$ and then for $c > 0$.

Result 4. *For the semi-log inverse demand curve of eqn. (4),*

- *When $c = 0$, the ratios of profits and prices are:*

$$\Pi^{**}/\Pi^* = 2e^{-1}/\log(2) = 1.0615$$

$$P^{**}/P^* = 2e^{-1} = 0.7357$$

- *When $c > 0$, the ratios are closer to 1:*

$$1 \leq \Pi^{**}/\Pi^* < 1.0615$$

$$0.7357 < P^{**}/P^* \leq 1$$

If $c = 0$ both ratios can be computed exactly and do not depend on α or P_m ; in this worst case, the simple pricing rule yields a profit that differs from the optimal by only 6.15%, even though the prices differ by 26.5%. When $c > 0$, one cannot compute the ratios in closed form. Instead, we solve numerically for Π^{**} and P^{**} and present the results in Figure 2, where we plot the ratios as a function of c/P_m . (Numerical tests show the ratios to be independent of α .) Note that as c increases both ratios approach 1.

2.4 Log-Log Demand

We turn now to the commonly used log-log demand model:

$$P_A(Q) = A_0 Q^{-1/\beta}; \quad \beta > 1, \quad (5)$$

where $-\beta$ is the (constant) elasticity of demand. Because this demand curve has no maximum price, we truncate it so that $P(0) = P_m$. Setting $P_A(Q_0) = P_m$, the corresponding quantity is $Q_0 = (P_m/A_0)^{-\beta}$. We therefore work with the following modified version of eqn (5):

$$P_A(Q) = \begin{cases} P_m; & \text{if } Q < Q_0 \\ P_m(Q/Q_0)^{-1/\beta}; & \text{if } Q \geq Q_0 \end{cases} \quad (6)$$




Figure 2: Profit and price ratios for the semi-log inverse demand curve as a function of c/P_m

We require that $\beta > \beta_{min} = P_m/(P_m - c)$ in order for the optimal price P^{**} to be smaller than the maximum price P_m . The performance of our pricing rule in this case is:

Result 5. *For the demand curve of eqn. (6), the profit and price ratios are:*

$$\frac{\Pi^{**}}{\Pi^*} = \frac{2}{(P_m/c - 1)(\beta - 1)} \left[\frac{2\beta}{(P_m/c + 1)(\beta - 1)} \right]^{-\beta}$$

$$\frac{P^{**}}{P^*} = \frac{2\beta}{(P_m/c + 1)(\beta - 1)}$$

Note that these ratios are exact, and depend only on the elasticity β and P_m/c . Also, there is a unique value of $\beta^* = (P_m + c)/(P_m - c)$ for which both ratios equal 1.⁶

There are two limiting cases to note: c large and c very small. If c is large, i.e., $c \rightarrow P_m$, $\beta_{min} \rightarrow \infty$. If β_{min} is very large, $\beta (> \beta_{min})$ is very large (i.e., demand is very elastic), so both the profit and price ratios will be close to 1. At the other extreme, as $c \rightarrow 0$, $P^{**} \rightarrow 0$, whereas $P^* \rightarrow .5P_m$, and Π^{**}/Π^* is unbounded. But an isoelastic demand curve would then make little sense, because $Q^{**} \rightarrow \infty$.

⁶If $\beta = \beta^*$, the elasticity of the isoelastic demand equals the elasticity of the linear demand at the optimal price. The latter elasticity is $E_d = bP^*/Q_L^* = (P_m + c)/(P_m - c)$, so if $\beta = \beta^*$, both the linear and log-log demand curves have the same profit-maximizing price and output.




Figure 3: Profit and price ratios as a function of P_m/c for the log-log demand curve

The general case is illustrated in Figure 3, which shows the profit and price ratios as a function of P_m/c for $\beta = 1.5, 2.0,$ and 2.5 . If $\beta = 1.5$, Π^{**}/Π^* is always close to 1. But if $\beta = 2.5$, Π^{**}/Π^* can exceed 2 for large enough values of P_m/c . (Note that P^{**} can be larger or smaller than P^* .⁷) Thus if demand is very elastic (i.e., β is large) or marginal cost c is very small, our pricing rule may not perform well.

Table 1 summarizes these results. It shows that our pricing rule works well for a variety of underlying demand functions — but not all. If the true demand is a truncated log-log function, Π^{**}/Π^* can deviate substantially from 1 if demand is very elastic and/or marginal cost is small. This follows from the convexity of this function, and the fact that (unrealistically) the quantity demanded expands without limit as the price is reduced towards zero.

3 Uncertain Maximum Price

We have assumed that while the firm does not know its true demand curve, it *does* know the maximum price P_m it can charge and still expect to sell at least some units. Suppose instead that the firm only has an estimate of the maximum price:

$$\hat{P}_m = P_m(1 + \varepsilon), \quad (7)$$

⁷The log-log demand curve is convex, but truncating modifies its convexity properties, which affects the relationship between P^{**} and P^* (see Result 1). If either β or P_m/c is small, the optimal quantity Q^{**} is small and can lie on the truncated — and non-convex — part of the curve.

Inverse demand function	P^{**}/P^*	Π^{**}/Π^*
Quadratic convex: $P_A(Q) = P_m - b_1Q + b_2Q^2$ $b_1, b_2 \geq 0$ and $b_2 < b_1^2/4P_m$	$\frac{8}{9} \leq P^{**}/P^* \leq 1$	≤ 1.0116
Quadratic concave: $P_A(Q) = P_m - b_1Q + b_2Q^2$ $b_1 \geq 0$ and $b_2 \leq 0$	$1 \leq P^{**}/P^* \leq \frac{4P_m+2c}{3P_m+3c} \leq 1.33$	≤ 1.0887
Monomial: $P_A(Q) = P_m - \gamma Q^n$ $n = 3$ $n = 4$	$2(nP_m + c)/(n + 1)(P_m + c)$ ≤ 1.5 ≤ 1.6	$2^{(n+1)/n}n/(n + 1)^{(n+1)/n}$ 1.19 1.27
Semi-log: $P_A(Q) = P_m e^{-\alpha Q}$ $c = 0$ $c > 0$	0.7357 < 0.7357	1.0615 < 1.0615
Log-log (truncated): $P_A(Q) = \begin{cases} P_m; & \text{if } Q < Q_0 \\ P_m(Q/Q_0)^{-1/\beta}; & \text{if } Q \geq Q_0 \end{cases}$ $\beta \geq \beta_{min} = P_m/(P_m - c)$	$2\beta/(P_m/c + 1)(\beta - 1)$	$\frac{2}{(P_m/c-1)(\beta-1)} \left(\frac{2\beta}{(P_m/c+1)(\beta-1)} \right)^{-\beta}$

Table 1: Price and profit ratios for several “true” demand curves

where ε lies in some interval $[-B, B]$, with $0 \leq B \leq 1$. Our pricing rule is now $P^* = (\hat{P}_m + c)/2$, and suffers from two miss-specifications: the form of the demand curve and the value of the intercept. To see how this second source of uncertainty affects Π^{**}/Π^* , we assume that ε is uniformly distributed with $B = 0.2$ (i.e., $\varepsilon \sim U[-0.2, 0.2]$).

To simplify matters, we assume that $c = 0$. (Recall from the previous section that Π^{**}/Π^* deviates from 1 the most when $c = 0$ for the demand curves we considered.) The Appendix provides closed form expressions for Π^{**}/Π^* as a function of ε for the demand curves examined in Section 2. We also compute the expected value of the profit ratio when $\varepsilon \sim U[-0.2, 0.2]$, as well as the profit ratio that results when the true P_m is known exactly (i.e., $\varepsilon = 0$). The results are shown in Table 2.

Note that Table 2 includes the linear demand $P_A(Q) = P_m - bQ$, but with the pricing rule based on \hat{P}_m . This misspecification yields an expected profit loss of about 1.4%. For the quadratic, monomial and semi-log inverse demand functions, $\mathbb{E}[\Pi^{**}/\Pi^*]$ is very close to

“True” inverse demand function	$\mathbb{E}[\Pi^{**}/\Pi^*]$	$\Pi^{**}/\Pi^*(\varepsilon = 0)$
Linear: $P_A(Q) = P_m - bQ$	1.0137	1
Quadratic convex: $P_A(Q) = P_m - b_1Q + b_2Q^2$ $b_1, b_2 \geq 0$ and $b_2 < b_1^2/4P_m$	1.0254	1.0116
Quadratic concave: $P_A(Q) = P_m - b_1Q + b_2Q^2$ $b_1 \geq 0$ and $b_2 \leq 0$	1.1023	1.0887
Monomial: $P_A(Q) = P_m - \gamma Q^n$ $n = 3$ $n = 4$	1.205 1.2882	1.19 1.27
Semi-log: $P_A(Q) = P_m e^{-\alpha Q}$	1.0748	1.0615

Table 2: Expected profit ratios when $\hat{P}_m = P_m(1 + \varepsilon)$ and $\varepsilon \sim U[-0.2, 0.2]$

the ratio when the true value of P_m is known. (We show Π^{**}/Π^* for the truncated log-log curve in the Appendix, but omit its expectation as it depends on P_m/c and β .)

Of course, depending on the “draw” for ε , the actual profit ratio could be farther from 1. To see how much farther, we used the closed-form expressions in the Appendix to plot the profit ratios as a function of ε for $-0.2 \leq \varepsilon \leq 2$. As Figure 4 shows, the monomial demand (with $n = 3$) is most sensitive to the value of ε , with Π^{**}/Π^* reaching 1.5 when $\varepsilon = -0.2$. For the other demand curves, $\Pi^{**}/\Pi^* < 1.25$ over the range of ε we consider. Thus a misspecification of the maximum price increases Π^{**}/Π^* , but only moderately.

4 General Demand Functions

We turn now to more general demand functions. We begin with a result for concave demand curves that will be useful elsewhere:

Result 6. *For any concave inverse demand curve, we have:*

$$1 \leq \Pi^{**}/\Pi^* \leq 2; \quad 1 \leq P^{**}/P^* \leq 2$$




Figure 4: Profit ratios as a function of ε for different demand curves

In the worst case, the profit and price ratios will equal 2 if the true demand curve is a rectangle. For other concave functions, $\Pi^{**}/\Pi^* < 2$, but except for specific functional forms, we cannot say how much less.

We might expect that in some cases the inverse demand curve will not be concave and may even have a flat area (plateau), as in Figure 1. In this case, Π^{**}/Π^* will be sensitive to whether the plateau is below or above P^* . If the plateau is below P^* and very long, Π^{**}/Π^* can be arbitrarily large; by pricing at P^* , the firm is missing a large mass of consumers. But if the plateau is above P^* , Π^{**}/Π^* will usually be close to 1. Thus if the firm believes there is such a plateau, it might set price below P^* in order to capture it.

In practice, a firm introducing a new product may know little or nothing about the shape of the demand curve. Indeed, that is the motivation for this paper. The firm might have no reason to expect that demand is characterized by one of the commonly used functions we examined earlier, or any other particular function. If the firm uses our pricing rule — with no knowledge at all of the true demand curve, other than the maximum price P_m — how well can it expect to do?

4.1 General Random Demand Curves

We address this question by randomly generating a set of “true” demand curves. For each randomly generated curve we compute (numerically) the profit-maximizing price and profit, P^{**} and Π^{**} , and we compare Π^{**} to the profit Π^* the firm would earn by using our pricing rule, i.e., by setting the price $P^* = (P_m + c)/2$. We generate 100,000 such demand curves, and then examine the resulting distribution of Π^{**}/Π^* . The only restriction we impose on these demand curves is that they are non-increasing everywhere.

We generate each demand curve as follows. We assume the maximum price P_m is known, as is the maximum quantity Q_{max} that can be sold at a price of zero (i.e., the maximum potential size of the market). We divide the segment $[0, Q_{max}]$ into S equally spaced intervals, and generate a piece-wise non-increasing demand curve by drawing random values for the different pieces. Since $P(0) = P_m$ and $P(Q_{max}) = 0$, there are $S - 1$ breaking points between 0 and P_m . (One might interpret this partition of the market as representing customer segments, or simply an approximation to a continuous curve.) With this partition, we draw a value for the end of the first segment from a uniform distribution between 0 and P_m . Call this random value P_1 (see one realization for P_1 in Figure 5). Next, we draw a value for the end of the second segment, again from a uniform distribution but now between 0 and P_1 . Call this P_2 . We repeat this process until we have $S - 1$ non-increasing prices, i.e., a random demand curve that has S segments. Figure 5 shows an example of such a randomly generated demand curve that has 5 segments (for $P_m = 500$ and $Q_{max} = 5$). Given this demand curve, we calculate P^{**} , Π^{**} , and the profit ratio Π^{**}/Π^* for $c = 0$ and $c = 0.5P_m$.

We generate 100,000 demand curves and compute 100,000 corresponding values for Π^{**}/Π^* . We calculate the mean value of Π^{**}/Π^* , as well as the 80% and 90% points (i.e., the value of Π^{**}/Π^* , such that 80% or 90% of the randomly generated ratios are below this number). The number of segments S can affect the resulting Π^{**} , so in Table 3 we show results for different values of S and for c/P_m equal to 0 and 0.5.

One can see that whatever the number of segments, S , the average profit ratio is less than 1.14 if $c = 0$ and less than 1.08 if $c = 0.5P_m$. Also, 80% (90%) of the demand curves yield profit ratios less than 1.22 (1.41) if $c = 0$ and less than 1.13 (1.37) if $c = 0.5P_m$. In Figure 6 we plot histograms of the 100,000 profit ratios for $S = 5$ and both $c = 0$ and $c = 0.5P_m$. When $c = 0$ ($c = 0.5P_m$), more than 40% (75%) of the ratios are less than 1.01, and 54%




Figure 5: Randomly generated inverse demand curve with $S = 5$ pieces

(79%) are less than 1.05. Thus it is likely that our simple pricing rule will yield a profit close to what would result if the firm knew its demand curve.

4.2 Uncertain Maximum Price

What if the maximum price P_m is not known exactly? As in Section 3, we assume that the firm only has an estimate \hat{P}_m of P_m , given by eqn. (7). We generate 10,000 random demand curves of S segments each for different values of S , but for each demand curve, we now assume that the firm must base its price on its estimate \hat{P}_m of P_m in eqn. (7), with ε uniformly distributed over $[-0.2, 0.2]$. For each random demand curve we draw independently 100 values of ε , yielding a total of 1 million random \hat{P}_m -demand curve combinations. We compute the expectation of the profit ratio Π^{**}/Π^* (over both the random variable ε and the randomly generated demand curves). The results are shown in Table 4, along with the corresponding expected profit ratio when the maximum price is known exactly (i.e., $\varepsilon = 0$).

Note from Table 4 that the expected profit ratio is always less than 1.15. We also calculated the 80% and 90% percentiles (as in Table 3) and the values are very close to those when the maximum price is known. This is consistent with the results in Section 3. Once again, our pricing rule is generally robust to miss-specification of the maximum price.

$c = 0$				$c = 0.5P_m$			
S	Mean	80%	90%	S	Mean	80%	90%
2	1.0672	1.1625	1.2442	2	1.0748	1.1255	1.3696
5	1.1332	1.2057	1.3926	5	1.0525	1.0645	1.2271
10	1.1351	1.2081	1.3979	10	1.0523	1.0647	1.2254
50	1.1379	1.2161	1.4071	50	1.0525	1.0621	1.2264
100	1.1344	1.2124	1.4045	100	1.0525	1.0628	1.2265

Table 3: Profit ratios for randomly generated demands




Figure 6: Histogram of profit ratios when $S = 5$ for $c = 0$ and $c = 0.5P_m$

5 Welfare Implications

Here we compare the welfare (consumer plus producer surplus) obtained from our pricing rule $P^* = 0.5(P_m + c)$ to the welfare that would have resulted if the firm knew the true demand curve and set the price to P^{**} . Welfare, denoted by $W(P)$, is:

$$W(P) = \Pi(P) + CS(P) = \int_0^Q P_A(y)dy - cQ. \quad (8)$$

We are interested in $W(P^{**})/W(P^*) \equiv W^{**}/W^*$. Note that this ratio can be less than one, i.e., our pricing rule can increase total welfare relative to that when the profit maximizing price P^{**} is used. In particular, assuming that $W(P)$ in (8) is concave (which is the case for

$c = 0$			$c = 0.5P_m$		
S	Mean	$\varepsilon = 0$	S	Mean	$\varepsilon = 0$
2	1.0806	1.0672	2	1.1406	1.0748
5	1.145	1.1332	5	1.115	1.0525
10	1.149	1.1351	10	1.1159	1.0523
50	1.149	1.1379	50	1.115	1.0525
100	1.149	1.1344	100	1.116	1.0525

Table 4: Expected profit ratios when the maximum price is unknown

any strictly monotone demand curve), we show in the Appendix that:

Result 7. *If the actual inverse demand curve $P_A(Q)$ is convex with respect to Q , then $W^{**} \geq W^*$, and if $P_A(Q)$ is concave, $W^{**} \leq W^*$.*

If the demand curve is concave, we know from Result 1 that $P^* \leq P^{**}$, so in this case using the “wrong” price P^* improves total welfare. Next, we calculate the welfare ratio for randomly generated demand curves following the approach of Section 4.1. As before, we fix the values of P_m , c and Q_{max} and compute the ratio W^{**}/W^* for $c/P_m = 0$ and 0.5, using 100,000 randomly generated demand curves. We did this for different values of S , and found the results to be very similar. In Figure 7, we show the histograms for $S = 5$.

The average welfare ratios for $c = 0$ and $c = 0.5P_m$ are 1.139 and 0.993 respectively. Also, for 50% of the instances when $c = 0$ and for 83% when $c = 0.5P_m$, W^{**}/W^* is between 0.9 and 1.1. But when $c = 0$, W^{**}/W^* can be large; it exceeds 1.25 in 23.5% of the cases.

6 Conclusions

Setting price is one of the most basic economic decisions firms make. Introductory economics courses make this decision seem easy; just write down the demand curve and set marginal revenue equal to marginal cost. But of course firms rarely have precise knowledge of their demand curves. When introducing new products (or existing products in new markets), firms may know little or nothing about demand, but must still set a price. Price experimentation is often not feasible, and the price a firm sets is often the one it sticks with for some time.




Figure 7: Histogram of welfare ratios when $S = 5$ for $c = 0$ and $c = 0.5P_m$

We have shown that under certain conditions the firm can use a simple pricing rule. The conditions are: (i) marginal cost c is known and constant, (ii) the firm need not predict the quantity it will sell, and (iii) the firm can estimate the maximum price P_m it can charge and still expect to sell some units. For new products or services introduced by technology companies, these conditions often hold. The firm then sets a price of $P^* = (P_m + c)/2$.

How well can the firm expect to do if it follows this pricing rule? We have shown that if the true demand curve is one of several commonly used demand functions, or even if it is a more complex function, the firm can expect to do quite well. In most cases it will earn a profit Π^* reasonably close to the profit Π^{**} it could earn if it knew the true demand curve.

Some caveats are in order. Perhaps most important, our analysis is entirely static. We assumed the true demand curve is fixed; it does not shift over time, perhaps in response to network effects (which can be important for new products). We also assumed that the firm sets and maintains a single price; it does not change price over time to inter-temporally price discriminate or to respond to changing market conditions, nor does it offer different prices to different groups of customers. We have also ruled out learning about demand, either passively or via experimentation, which has been the focus of the earlier literature on pricing with uncertain demand. To the extent that such dynamic considerations are important, our pricing rule can be viewed as a starting point. Managers often seek simple and robust rules for pricing; the rule we suggest is certainly simple, and we have seen that it is also effective.

Appendix

Proof of Result 1

The actual inverse demand curve is $P_A(Q)$, and $P_A(0) = P_m$. One can write: $P_A(Q) = P_m - bQ + f(Q)$, with $f(0) = 0$ and $P'_A(Q) = f'(Q)$. Equating marginal revenue with marginal cost:

$$Q^{**} = \frac{P_m - c + f(Q^{**}) + f'(Q^{**})Q^{**}}{2b}.$$

This yields an expression for the optimal price as a function of Q^{**} :

$$P^{**} = P_A(Q^{**}) = P_m - \frac{1}{2} \left(P_m - c + f(Q^{**}) + f'(Q^{**})Q^{**} \right) + f(Q^{**}).$$

Recall that $P^* = (P_m + c)/2$ and therefore: $P^{**} = P^* + 0.5[f(Q^{**}) - f'(Q^{**})Q^{**}]$. From the first order Taylor expansion, we have for any differentiable function $f(\cdot)$: $f(x) = f(a) + f'(a)(x - a) + R_1$, where $R_1 = 0.5f''(\zeta)(x - a)^2$, for some $\zeta \in [x, a]$. Then:

$$f(Q^{**}) - f'(Q^{**})Q^{**} = -R_1 = \frac{f''(\zeta)}{2}(Q^{**})^2 = \frac{P''_A(\zeta)}{2}(Q^{**})^2.$$

Therefore $P^{**} - P^* = -P''_A(\zeta)(Q^{**})^2/4$, for some $\zeta \in [0, Q^{**}]$. As a result, if $P_A(Q)$ is convex, $P''_A(\cdot) \geq 0$ so that $P^{**} \leq P^*$, and if $P_A(Q)$ is concave, $P''_A(\cdot) \leq 0$ so that $P^{**} \geq P^*$.

Proof of Result 2

Convex case: The profit Π^* , the optimal price P^{**} and quantity Q^{**} are:

$$\begin{aligned} \Pi^* &= \frac{P_m - c}{2} \left[\frac{1}{2b_2} \left(b_1 - \sqrt{b_1^2 - 2b_2(P_m - c)} \right) \right] \\ P^{**} &= P_m - \frac{b_1}{3b_2} \left[b_1 - \sqrt{b_1^2 - 3b_2(P_m - c)} \right] + \frac{1}{9b_2} \left[b_1 - \sqrt{b_1^2 - 3b_2(P_m - c)} \right]^2 \\ Q^{**} &= \frac{b_1 - \sqrt{b_1^2 - 3b_2(P_m - c)}}{3b_2} \end{aligned}$$

The optimal profit is $\Pi^{**} = (P^{**} - c)Q^{**}$. One can express the profit and price ratios as functions of c and b_2 and check the monotonicity to conclude that the profit and price ratios are largest when $c = 0$ and $b_2 = b_1^2/4P_m$, in which case Π^* and Π^{**} are: $P^{**} = (4/9)P_m$ and $Q^{**} = (2P_m)/(3b_1)$. Finally, we can now compute both profits:

$$\Pi^* = \frac{b_1 P_m}{4b_2} \left(1 - \frac{1}{\sqrt{2}} \right) = \frac{P_m^2}{b_1} \left(1 - \frac{1}{\sqrt{2}} \right); \quad \Pi^{**} = \frac{2b_1 P_m}{27b_2} = \frac{8P_m^2}{27b_1}$$

Then the profit and price ratios are: $\Pi^{**}/\Pi^* = 8\sqrt{2}/[27(\sqrt{2} - 1)] = 1.0116$; $P^{**}/P^* = 8/9$. These are the largest values for the ratios, so in general we have inequalities.

Concave case: The optimal quantity Q^{**} equates marginal revenue with marginal cost:

$$Q^{**} = \frac{-b_1 \pm \sqrt{b_1^2 - 3b_2(P_m - c)}}{-3b_2}.$$

Since $Q^{**} > 0$, the positive root applies. Then the optimal price is :

$$P^{**} = P_A(Q^{**}) = \frac{2P_m + c}{3} - \frac{b_1^2}{9b_2} + \frac{1}{9b_2}b_1\sqrt{b_1^2 - 3b_2(P_m - c)}.$$

Finally, the optimal profit as a function of P_m , c , b_1 and b_2 follows from $\Pi^{**} = (P^{**} - c)Q^{**}$. Our pricing rule is $P^* = (P_m + c)/2$, so

$$Q_A(P^*) = \frac{-b_1 \pm \sqrt{b_1^2 - 2b_2(P_m - c)}}{-2b_2}.$$

We select the positive root in order to satisfy $Q^* > 0$. The profit is:

$$\Pi^* = (P^* - c)Q_A(P^*) = \frac{P_m - c}{2} \left[\frac{1}{-2b_2} \left(\sqrt{b_1^2 - 2b_2(P_m - c)} - b_1 \right) \right].$$

Expressing the profit and price ratios in terms of b_1 and checking the monotonicity, one can see that the worst case for both ratios occurs when $b_1 = 0$. Intuitively, the larger is b_1 , the more linear is the function, making the ratios closer to 1. If $b_1 = 0$, $P^{**} = (2P_m + c)/3$ and $P^* = (P_m + c)/2$, so

$$\Pi^{**} = \frac{2(P_m - c)}{3} \frac{\sqrt{-3b_2(P_m - c)}}{-3b_2}; \quad \Pi^* = \frac{P_m - c}{2} \frac{\sqrt{-2b_2(P_m - c)}}{-2b_2}$$

Then, the profit and price ratios are:

$$\frac{\Pi^{**}}{\Pi^*} = \frac{4\sqrt{2}}{3\sqrt{3}} = 1.0887; \quad \frac{P^{**}}{P^*} = \frac{2}{3} \frac{2P_m + c}{P_m + c} \leq \frac{4}{3} = 1.33$$

For $b_1 > 0$, we have inequalities for both ratios.

Proof of Result 3

Equating marginal revenue and marginal cost, $MR_A(Q^{**}) = P_m - (n+1)\gamma(Q^{**})^n = c$. Thus: $Q^{**} = [(P_m - c)/(n+1)\gamma]^{1/n}$, and the optimal price is: $P^{**} = P_A(Q^{**}) = (nP_m - c)/(n+1)$. Note that P^{**} is independent of γ . Next, the optimal profit is:

$$\Pi^{**} = (P^{**} - c)Q^{**} = \frac{n}{(n+1)^{\frac{1}{n}+1}\gamma^{1/n}}(P_m - c)^{\frac{1}{n}+1}.$$

Recall that $P^* = (P_m + c)/2$, so the corresponding quantity is $Q_A(P^*) = [(P_m - c)/(2\gamma)]^{1/n}$. Therefore, $\Pi^* = (P^* - c)Q_A(P^*) = [(P_m - c)^{\frac{1}{n}+1}]/[2^{\frac{1}{n}+1}\gamma^{1/n}]$. We can now compute both ratios:

$$\frac{\Pi^{**}}{\Pi^*} = \frac{2^{\frac{1}{n}+1}n}{(n+1)^{\frac{1}{n}+1}} \leq 2; \quad 1 \leq \frac{P^{**}}{P^*} = \frac{2(nP_m + c)}{(n+1)(P_m + c)} \leq 2$$

Proof of Result 4

First, suppose $c = 0$. Equating marginal revenue and marginal cost, $MR_A(Q^{**}) = P_m e^{-\alpha Q^{**}} - \alpha P_m Q^{**} e^{-\alpha Q^{**}} = 0$, so $Q^{**} = 1/\alpha$. Then $P^{**} = P_m e^{-1}$; $\Pi^{**} = P_m e^{-1} \alpha^{-1}$. If the firm prices at P^* , profit is $\Pi^* = (P^* - c)Q_A(P^*) = 0.5P_m Q_A(P^*)$. Since $c = 0$ and $P^* = 0.5P_m$, we obtain: $Q_A(P^*) = -(1/\alpha) \log(0.5)$. Therefore: $\Pi^* = 0.5P_m \log(2)/\alpha$. Therefore:

$$\frac{\Pi^{**}}{\Pi^*} = \frac{P_m e^{-1} 2\alpha}{\alpha P_m \log(2)} = \frac{2e^{-1}}{\log(2)} = 1.0615; \quad \frac{P^{**}}{P^*} = \frac{P_m e^{-1}}{P_m/2} = 2e^{-1} = 0.7357$$

We now show that when $c > 0$, both ratios are closer to 1. Start with the price ratio; we show that $\frac{\partial}{\partial c} \left[\frac{P^{**}}{P^*} \right] \geq 0$; $\forall 0 \leq c \leq P_m$. We have:

$$\frac{\partial}{\partial c} \left[\frac{P^{**}}{P^*} \right] = \frac{\frac{\partial P^{**}}{\partial c} P^* - \frac{\partial P^*}{\partial c} P^{**}}{(P^*)^2}. \quad (9)$$

For eqn. (9) to be nonnegative, we need: $\frac{\partial P^{**}}{\partial c} \frac{1}{P^{**}} \geq \frac{\partial P^*}{\partial c} \frac{1}{P^*}$. Recall that $P^* = (P_m + c)/2$ and therefore: $\partial P^*/\partial c = 0.5$. As a result, we need to show:

$$\frac{\partial P^{**}}{\partial c} \geq \frac{P^{**}}{P_m + c}. \quad (10)$$

From the first order condition: $MR_A(Q^{**}) = P_m e^{-\alpha Q^{**}} - \alpha P_m Q^{**} e^{-\alpha Q^{**}} = P^{**}(1 - \alpha Q^{**}) = c$. By differentiating both sides with respect to c and isolating $\partial P^{**}/\partial c$:

$$\frac{\partial P^{**}}{\partial c} = \frac{1 + \alpha P^{**} \frac{\partial Q^{**}}{\partial c}}{1 - \alpha Q^{**}}. \quad (11)$$

Recall that $P^{**} = P_m e^{-\alpha Q^{**}}$ and hence by differentiating with respect to c :

$$\frac{\partial P^{**}}{\partial c} = -\alpha P^{**} \frac{\partial Q^{**}}{\partial c}. \quad (12)$$

By combining (11) and (12), we obtain: $\partial P^{**}/\partial c = 1/(2 - \alpha Q^{**})$. Since the demand curve is convex, from Result 1: $P^{**} \leq P^* = (P_m + c)/2$ and therefore: $P^{**}/(P_m + c) \leq 0.5$. From the FOC, $0 \leq 1 - \alpha Q^{**} \leq 1$ (so that $P^{**} \geq c$). Thus $1 \leq 2 - \alpha Q^{**} \leq 2$, so $1/(2 - \alpha Q^{**}) \geq 0.5$, implying that (10) is satisfied. This concludes the proof for the price ratio.

The same logic applies for the profit ratio, i.e., $\partial[\frac{\Pi^{**}}{\Pi^*}]/\partial c \leq 0$; $\forall 0 \leq c \leq P_m$.

Proof of Result 5

Equating marginal revenue to marginal cost, $MR_A(Q^{**}) = P_m \left(1 - \frac{1}{\beta}\right) (Q^{**}/Q_0)^{-1/\beta} = c$. Thus: $Q^{**} = Q_0 \left[\frac{\beta c}{(\beta-1)P_m} \right]^{-\beta}$. Note that Q^{**} is larger than the truncation value Q_0 . The optimal price and profit are: $P^{**} = \beta c/(\beta-1)$; $\Pi^{**} = Q_0 c/(\beta-1) \left[\frac{\beta c}{(\beta-1)P_m} \right]^{-\beta}$. By requiring

$\beta \geq P_m/(P_m - c)$ we ensure that $P^{**} \leq P_m$. We next compute the profit by using P^* : $\Pi^* = (P^* - c)Q_A(P^*) = 0.5(P_m - c)Q_A(P^*)$. We have: $Q_A(P^*) = Q_0 \left(\frac{P_m + c}{2P_m} \right)^{-\beta} \geq Q_0$. Then: $\Pi^* = 0.5Q_0(P_m - c) \left(\frac{P_m + c}{2P_m} \right)^{-\beta}$. We can now compute both ratios:

$$\frac{\Pi^{**}}{\Pi^*} = \frac{2}{(P_m/c - 1)(\beta - 1)} \left(\frac{2\beta}{(P_m/c + 1)(\beta - 1)} \right)^{-\beta}; \quad \frac{P^{**}}{P^*} = \frac{2\beta}{(P_m/c + 1)(\beta - 1)}$$

Proof of Result 6

Consider any non-increasing concave inverse demand curve. We know from Result 1: $P^* \leq P^{**}$. Recall that $P^* = 0.5(P_m + c)$ and therefore, $P^* \leq P^{**} \leq P_m = 2P^* - c \leq 2P^*$. We next show the inequality for the profits: $\Pi^{**} = (P^{**} - c)Q_A(P^{**}) \leq 2(P^* - c)Q_A(P^{**}) \leq 2(P^* - c)Q_A(P^*) = 2\Pi^*$, where the last inequality follows from the fact that $Q_A(\cdot)$ is non-increasing. In conclusion, we have $1 \leq \Pi^{**}/\Pi^* \leq 2$ and $1 \leq P^{**}/P^* \leq 2$.

Expressions for Section 3

Here are the closed form expressions of the profit ratio Π^{**}/Π^* as a function of ε for the demand models we considered, for $c = 0$. (Setting $\varepsilon = 0$ yields the expressions in Section 2.)

- **Linear:** $P_A(Q) = P_m - bQ$ $\Pi^{**}/\Pi^*(\varepsilon) = 1/(1 - \varepsilon^2)$
- **Quadratic convex:** $P_A(Q) = P_m - b_1Q + b_2Q^2$; $b_1, b_2 \geq 0$ and $b_2 \leq b_1^2/4P_m$

$$\frac{\Pi^{**}}{\Pi^*}(\varepsilon) \leq \frac{8\sqrt{2}}{27(1 + \varepsilon)} \frac{1}{\sqrt{2} - \sqrt{1 + \varepsilon}}$$

- **Quadratic concave:** $P_A(Q) = P_m - b_1Q + b_2Q^2$; $b_1 \geq 0$ and $b_2 \leq 0$

$$\frac{\Pi^{**}}{\Pi^*}(\varepsilon) \leq \frac{4\sqrt{2}}{3\sqrt{3}} \frac{1}{(1 + \varepsilon)\sqrt{(1 - \varepsilon)}}$$

- **Monomial:** $P_A(Q) = P_m - \gamma Q^n$ $\Pi^{**}/\Pi^*(\varepsilon) = \frac{2^{\frac{1}{n}+1}n}{(n+1)^{\frac{1}{n}+1}} \frac{1}{(1+\varepsilon)(1-\varepsilon)^{1/n}}$
- **Semi-log:** $P_A(Q) = P_m e^{-\alpha Q}$ $\Pi^{**}/\Pi^*(\varepsilon) = \frac{2e^{-1}}{(1+\varepsilon)\log(\frac{2}{1+\varepsilon})}$

- **Log-log (truncated):** $P_A(Q) = \begin{cases} P_m; & \text{if } Q < Q_0 \\ P_m(Q/Q_0)^{-1/\beta}; & \text{if } Q \geq Q_0 \end{cases}$

$$\frac{\Pi^{**}}{\Pi^*}(\varepsilon) = \frac{2}{\left[\frac{P_m}{c}(1 + \varepsilon) - 1 \right](\beta - 1)} \left[\frac{2\beta}{(\beta - 1) \left[\frac{P_m}{c}(1 + \varepsilon) + 1 \right]} \right]^{-\beta}$$

Proof of Result 7

Since $W(P)$ is concave, there exists a unique maximizer denoted by P_W^* . The profit maximizer price P^{**} is such that $\partial\Pi/\partial P = 0$. Note that the consumer surplus is always non-decreasing with respect to the price. Indeed, one can write the consumer surplus in the price space (using integration by parts): $CS(P) = \int_P^{P_m} Q(z)dz$, where $Q(z)$ is the demand price function. Since $Q(z) \geq 0$ and $P \leq P_m$, one can see that $CS(P)$ is non-increasing. As a result: $\partial W/\partial P(P^{**}) = \partial CS/\partial P(P^{**}) \leq 0$. Therefore, $P_W^* \leq P^{**}$.

Using Result 1, if the actual inverse demand curve $P_A(Q)$ is convex, then $P^{**} \leq P^*$ so that $P_W^* \leq P^{**} \leq P^*$. By using the concavity of $W(P)$, we conclude that $W^{**} \geq W^*$.

We next consider the case where $P_A(Q)$ is concave. From Result 1, we have $P^* \leq P^{**}$. Using equation (8), the first order condition for maximizing welfare is: $W'(P) = (P - c)Q'(P) = 0$. Therefore, either $P_W^* = c$ or $Q'(P_W^*) = 0$. If $P_W^* = c$, then it is clear that $W^* \geq W^{**}$. We next address the case where $Q'(P_W^*) = 0$. Since $P_A(Q)$ is non-increasing, concave and has a bounded image $[c, P_m]$, $Q(P)$ is concave as well (see, e.g., Mršević (2008)). Since $P_W^* \geq c$ and $Q'(P)$ is non-increasing (and non-positive), every point in $[c, P_W^*]$ also satisfies $Q'(P) = 0$. As a result, for any $P \geq P_W^*$, the welfare function is non-increasing. Since $P^* \leq P^{**}$, we have: $W^* \geq W^{**}$.

References

- Adida, Elodie, and Georgia Perakis.** 2006. “A robust optimization approach to dynamic pricing and inventory control with no backorders.” *Mathematical Programming*, 107(1-2): 97–129.
- Araman, Victor F, and René Caldentey.** 2011. “Revenue management with incomplete demand information.” *Wiley Encyclopedia of Operations Research and Management Science*.
- Aviv, Yossi, and Amit Pazgal.** 2005. “A partially observed Markov decision process for dynamic pricing.” *Management Science*, 51(9): 1400–1416.
- Ball, Michael O, and Maurice Queyranne.** 2009. “Toward robust revenue management: Competitive analysis of online booking.” *Operations Research*, 57(4): 950–963.
- Bergemann, Dirk, and Karl Schlag.** 2011. “Robust monopoly pricing.” *Journal of Economic Theory*, 146(6): 2527–2543.
- Bertsimas, Dimitris, and Georgia Perakis.** 2006. *Dynamic pricing: A learning approach*. Springer.
- Besbes, Omar, and Assaf Zeevi.** 2009. “Dynamic pricing without knowing the demand function: Risk bounds and near-optimal algorithms.” *Operations Research*, 57(6): 1407–1420.
- Besbes, Omar, and Assaf Zeevi.** 2013. “On the (surprising) sufficiency of linear models for dynamic pricing with demand learning.” *Columbia Business School Research Paper*, , (12-5).
- Carroll, Gabriel.** 2014. “Robustness and Linear Contracts.” *American Economic Review*. forthcoming.
- Chu, Chenghuan Sean, Phillip Leslie, and Alan Sorensen.** 2011. “Bundle-Size Pricing as an Approximation to Mixed Bundling.” *American Economic Review*, 101: 262–303.
- Dawes, Robyn M.** 1979. “The robust beauty of improper linear models in decision making.” *American psychologist*, 34(7): 571.
- Farias, Vivek F, and Benjamin Van Roy.** 2010. “Dynamic pricing with a prior on market response.” *Operations Research*, 58(1): 16–29.
- Hansen, Lars Peter, and Thomas J. Sargent.** 2008. *Robustness*. Princeton University Press.

- Lim, Andrew EB, and J George Shanthikumar.** 2007. “Relative entropy, exponential utility, and robust dynamic pricing.” *Operations Research*, 55(2): 198–214.
- Lin, Kyle Y.** 2006. “Dynamic pricing with real-time demand learning.” *European Journal of Operational Research*, 174(1): 522–538.
- Mršević, Mila.** 2008. “Convexity of the inverse function.” *The Teaching of Mathematics*, XI, 1: 21–24.
- Rothschild, Michael.** 1974. “A two-armed bandit theory of market pricing.” *Journal of Economic Theory*, 9(2): 185–202.
- Shy, Oz.** 2006. *How to Price: A Guide to Pricing Techniques and Yield Management*. Cambridge University Press.
- Thiele, Aurélie.** 2009. “Multi-product pricing via robust optimisation.” *Journal of Revenue & Pricing Management*, 8(1): 67–80.