Pricing With Limited Knowledge of Demand

Maxime C. Cohen
NYU Stern School of Business, New York, NY 10012, maxime.cohen@stern.nyu.edu

Georgia Perakis
Sloan School of Management, MIT, georgiap@mit.edu

Robert S. Pindyck
Sloan School of Management, MIT, rpindyck@mit.edu

How should a firm price a new product for which little is known about demand? We propose a simple and practical pricing rule for new products where demand information is limited. The rule is simple: Set price as though the demand curve were linear. Our pricing rule can be used if three conditions hold: the firm can estimate the maximum price it can charge and still expect to sell some units, the firm need not plan in advance the quantity it will sell, and marginal cost is known and constant. We show that if the true demand curve is one of many commonly used demand functions, or even a more complex (randomly generated) function, 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, calculate expected profit performance for randomly generated demand curves, and evaluate the welfare implications of our pricing rule. We show that with limited demand information (maximum price and marginal cost), our simple pricing rule can be used for new products while often achieving a near-optimal performance. We also discuss the limitations of our method by identifying cases where our pricing rule does not perform well.

Key words: Pricing, new products, unknown demand, pricing heuristics, linear demand approximation

1. Introduction

Firms that introduce new products must often set a price with little or no knowledge of demand, and no data from which to estimate elasticities. How should firms set prices in such settings? 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 (we discuss this literature below). Experimenting with price, however, is often not feasible or desirable; it is often common for firms to choose an introductory price and maintain that price for a year or more. We examine a much simpler approach to this pricing problem that does not involve any price experimentation.

We show that under certain conditions, the firm can use a very simple pricing rule. The conditions are that (i) the firm's marginal cost, c, is known and constant, (ii) the firm can estimate the *maximum price* P_m it can charge and still expect to sell some units, and (iii) the firm need not know or plan in advance the quantity it will sell. (We partially relax the second assumption in

Section 3.) These conditions will not always hold—but they do for many new products, particularly those that involve new technologies.

Examples of new product introductions for which these conditions hold include new types of drugs introduced by pharmaceutical companies (Lilly's Prozac in 1987, AstraZeneca's Crestor in 2003, or Merck's Vytorin in 2004), technology companies introducing new products or services (Apple setting the price of music downloads when launching its iTunes store in 2002, Intuit setting the price of TurboTax in 2001, and Adobe introducing Acrobat Distiller in 1993), or a company introducing an existing product in a new and 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 in these examples knew very little about the demand curves they faced, and had no data to estimate elasticities. They could, however, roughly estimate the maximum price they could charge, and there was no need for them to know in advance how much they could sell.

By "maximum price" we do not mean the price at which the firm can sell *any units* (that price might be extraordinarily high), but rather a price at which the firm can still expect to serve a few percent of its potential market. Determining the maximum price P_m might not be easy but it is a much less difficult task than estimating the entire demand curve.

In practice, there are several ways of estimating P_m. For some products and services, the firm can intentionally create a scarcity situation (when a product is first introduced) and then observe the highest prices paid for the product or service through secondary channels such as eBay (for products) or StubHub (for events). Alternatively, the firm could first introduce a product or service by selling it in a highest bid auction format. This is done for instance for prototypes of luxury items. Finally, it is now common for firms to hire focus groups (either reputable Marketing specialists or loyal/passionate customers). In such focus groups, one of the main topics is the maximal price that can be charged for the upcoming product. This approach is commonly used for the introduction of new software and digital services (e.g., paid subscriptions for online dating services).

A pharmaceutical company might estimate P_m by comparing a new drug to existing therapies (including non-drug therapies). For 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. 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. (A CD with 12 songs might cost \$12 to \$15 but most consumers would want only a few of those songs.) Likewise, Intuit might have used a simple survey to learn how much at least some consumers would pay for software to prepare their tax returns.

How can a firm set price without also having an estimate of the quantity it will sell? For new (biochemical) drugs, marginal cost is near zero, so the firm can produce a large amount of pills and discard whatever is not sold. Thus the firm only needs to decide the price to charge. For music or video downloads and streaming services, as well as software, no factories have to be built and marginal production cost (excluding royalties) is zero. This assumption is also satisfied in settings where capacity is cheap and production lead times are very short relative to product life.

We propose that if these three conditions hold, the firm can use the following rule: Given the maximum price P_m, set price as though the actual demand curve were linear, i.e.,

$$P(Q) = P_m - bQ, \tag{1}$$

With constant marginal cost c, the firm's profit-maximizing price is $P^* = (P_m + c)/2$, 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 plan on a particular sales level (which we have argued would be the case for most new drugs, music downloads, or 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. The price P^* can also be viewed as a convex combination of the two known parameters (P_m and c) with equal weights. We 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 P^*? Suppose that with precise knowledge of its true demand curve, the firm would 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, the answer depends on the true demand function. In this paper, we derive closed form bounds on the profit performance for several common demand functions and compute numerically the performance for randomly generated demands. We will show that in many cases this simple pricing rule performs well, i.e., Π^* is close to Π^{**}. We will also identify cases where the rule does not perform well.

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.[1] 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 reimburse 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

[1] For illustration purposes, we plot the inverse demand curve (instead of the demand) as it is common in economics.

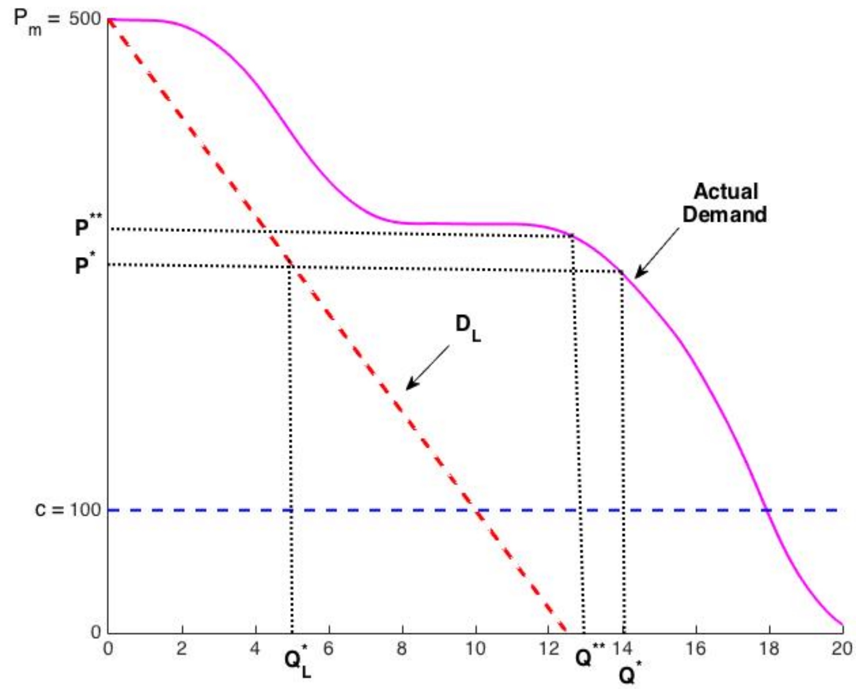

Figure 1 Illustration for a representative demand curve

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.) 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 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. (As discussed, 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 it knew the actual demand curve and used it to set price. (The firm would do a bit worse if $c = 0$, in which case $\Pi^{**}/\Pi^* = 1.084$.)

There are certainly demand curves for which this pricing rule will perform poorly. For example, suppose the true demand curve is a rectangle, i.e., $P = P_m$ for $0 \leq Q \leq Q_{max}$ and $P = 0$ for $Q > Q_{max}$. Then, the profit-maximizing price is clearly P_m and the resulting profit is $\Pi^{**} = (P_m - c)Q_{max}$. Setting a price $P^* = (P_m + c)/2$ yields 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 (exploration phase), observes outcomes, and because each trial is costly, eventually settles on the price that it thinks (perhaps incorrectly) is optimal (exploitation phase). 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.

The marketing literature has also considered the problem of pricing for new products. Examples include classical works such as Urban et al. (1996) that consider pre-market forecasting, as well as Krishnan et al. (1999) that consider a variation of the generalized Bass model that yields optimal pricing policies that are consistent with empirical data. More recently, Handel and Misra (2015) introduce a dynamic non-Bayesian framework for robust pricing of new products.

Other studies focus on learning in a parametric or non-parametric context. Several papers 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 demand and optimizing revenues over time without imposing 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, Araman and Caldentey 2011, Balvers and Cosimano 1990).

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., $\max_p \min_{\alpha_1,\alpha_2} \Pi(\alpha_1, \alpha_2, p)$.[2] 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 may yield conservative pricing strategies.

[2] See, for example, 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, Ball and Queyranne 2009).

Our paper is also related to studies of model misspecification. In particular, we study the performance of a simple linear demand model even if the true demand curve is far from linear. Others have shown that linear models can perform well (e.g., Dawes (1979) in clinical prediction and Carroll (2015) in contract theory). Besbes and Zeevi (2015) study the "price of misspecification" for dynamic pricing with demand learning. The authors propose a dynamic pricing algorithm in which the seller assumes demand is linear, and chooses a price to maximize revenue based on this linear demand function. They show that although the model is misspecified, one can achieve a good asymptotic regret performance. In our setting, however, the firm chooses a price and does not have the option to experiment over time. In addition, our paper investigates how consumer welfare is affected by demand misspecification.

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.[3] The pricing rule we suggest is certainly simple; the extent to which it is effective is the focus of this paper.

The pricing rule $P^* = (P_m + c)/2$ follows from a linear approximation to the true demand curve. Note, however, that the same pricing rule can be obtained from a different set of modeling assumptions. Suppose the firm plans to sell the product to a representative consumer with a random valuation that is uniformly distributed $U[c, P_m]$. In this case, to maximize expected profits, the optimal price is known to also equal $P^* = (P_m + c)/2$ (the proof is presented in the Appendix). The equivalence of the two models (linear demand curve and uniform consumer valuation) is well known, and provides an alternative way of justifying the pricing rule studied in this paper.

In some cases, estimating the maximum price P_m is difficult or impractical. However, one can extend the results and analysis of this paper to situations where the firm can determine a price $\bar{P} < P_m$, such that at $\bar{P}$ the firm can still sell to a small set of customers who are not very price sensitive. The pricing rule then becomes $P^* = (\bar{P} + c)/2$. The basic insights of this paper will still hold, although at the expense of a pricing rule that does not perform quite as well. Obviously, the performance deteriorates when the gap between $\bar{P}$ and P_m increases.

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. We also find bounds

[3] In related work, Chu et al. (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 (2015) 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.

for Π^{**}/Π^{*} for a general concave or convex 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^{*} 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, semi-log, and log-log. These demand models are used in many operations management and economics applications.[4] For each we compare the profits from our pricing rule to the profits that would result if the actual demand function were known. We will see that the profit ratio is often close to one.

Before proceeding, 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,[5] the linear price is greater (smaller) than the optimal price:

THEOREM 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, $P^{**} \geq P^{*}$.*

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 quadratic:

$$P_A(Q) = P_m - b_1 Q + b_2 Q^2, \tag{2}$$

where, as before, P_m is the maximum price. Equivalently, the actual demand function is given by: $Q_A(P) = 0.5[b - \sqrt{b_1^2 - 4b_2(P_m - P)}]/b_2$. 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.)

PROPOSITION 1. *For the quadratic demand curve of eqn.* (2)*, the profit and price ratios satisfy:*

[4] The log-log model is widely used in many retail applications (see, e.g., Montgomery 1997, Mulugeta et al. 2013, Cohen et al. 2017). The quadratic and semi-log functions are often used in the context of hedonic pricing (see, e.g., Wilman 1981, Milon et al. 1984).

[5] Assuming that the inverse demand is a continuous decreasing convex (resp. concave), then the demand is also convex (resp. concave).

- ***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. (Also, P^{**}/P^{*} deviates the most from 1 when $c = 0$.)

If demand is concave, the resulting profit Π^{*} is within 8.87% of the optimal profit, irrespective of the parameters b_1 and b_2. In the proof of Proposition 1 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 than $b_2 \leq 0$).

2.2. Monomial Demand

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

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

Equivalently, the actual demand is given by: $Q_A(P) = [(P_m - P)/\gamma]^{1/n}$. 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:

PROPOSITION 2. *For the inverse demand curve of eqn.* (3)*, the profit and price ratios satisfy:*

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

Thus 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 $n \to \infty$. For monomials of order 3 and 4, the profit ratios are 1.19 and 1.27 respectively.

2.3. Semi-Log Demand

Now consider the semi-log inverse demand curve:

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

Or equivalently, $Q(P) = \frac{1}{\alpha}\log(P_m/P)$. The following result (proof in Appendix) shows the profit and price ratios when the marginal cost $c = 0$ and when $c > 0$.

Proposition 3. *For the semi-log inverse demand curve of eqn.* (4)*,*

- *When $c = 0$, the profit and price ratios are:*

$$\Pi^{**}/\Pi^{*} = 2e^{-1}/\log(2) = 1.0615,$$
$$P^{**}/P^{*} = 2e^{-1} = 0.7357.$$

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

$$1 \leq \Pi^{**}/\Pi^{*} < 1.0615,$$
$$0.7357 < P^{**}/P^{*} \leq 1.$$

When $c = 0$ 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 profit 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. (The ratios are 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 (isoelastic) demand model:

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

where $-\beta$ is the (constant) elasticity of demand. We can also write this as: $\log Q_A = \beta(\log A_0 - \log P)$. Because this demand curve has no maximum price, we truncate it so that $P(0) = P_m$.

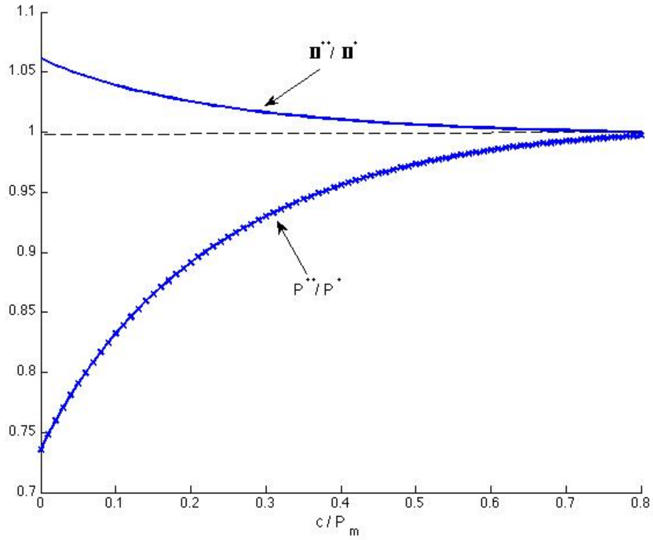

Figure 2 Profit and price ratios for the semi-log inverse demand curve as a function of c/P_m

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):

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

We require that $\beta > \beta_{min} = P_m/(P_m - c)$ for the optimal price P^{**} to be less than the maximum price P_m. In this case:

PROPOSITION 4. *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., $c \to P_m$, $\beta_{min} \to \infty$. If β_{min} is very large, β ($> \beta_{min}$) is also very large (i.e., demand is elastic), so both the profit and price ratios will be close to 1. At the other extreme, as $c \to 0$, $P^{**} \to 0$, whereas $P^* \to 0.5P_m$, and Π^{**}/Π^* is unbounded. But an isoelastic demand curve would then make little sense, because $Q^{**} \to \infty$.

[6] If $\beta = \beta^*$, 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.

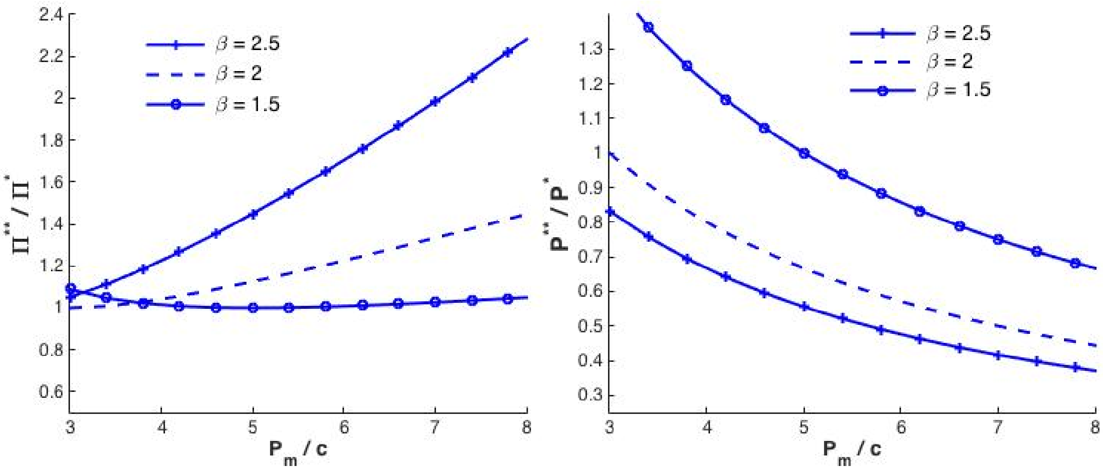

Figure 3 **Profit and price ratios as a function of P_m/c for the log-log demand curve**

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 of P_m/c. Note that P^{**} can be larger or smaller than P^*.[7] As a result, if demand is very elastic (i.e., β is large) or the marginal cost c is very small, our pricing rule will not perform well. One limitation of our pricing rule for the log-log demand model is the fact that the performance crucially depends on the price elasticity, which is not known by firm.

Table 1 summarizes these results. It shows that our pricing rule works well for a variety of underlying demand functions—but not all. For example, if the true demand is a truncated log-log function, Π^{**}/Π^* can deviate substantially from 1 if demand is very elastic and/or the 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

So far, 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 some units. Suppose instead that the firm only has an estimate of the maximum price:

$$\hat{P}_m = P_m(1+\epsilon), \tag{7}$$

where ϵ lies in some interval $[-B, B]$, with $0 \le B \le 1$. Our pricing rule is now $P^* = (\hat{P}_m + c)/2$, and suffers from two misspecifications: the form of the demand curve and the value of the intercept. To see how this second source of uncertainty affects the profit ratio Π^{**}/Π^*, we derive the profit ratios

[7] The log-log demand curve is convex but truncating it modifies its convexity properties, which affects the relationship between P^{**} and P^* (see Theorem 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: $P_A(Q) = P_m - \gamma Q^n$	$2(nP_m + c)/(n+1)(P_m + c)$	$2^{(n+1)/n}n/(n+1)^{(n+1)/n}$
$n = 3$	≤ 1.5	1.19
$n = 4$	≤ 1.6	1.27
Semi-log: $P_A(Q) = P_m e^{-\alpha Q}$		
$c = 0$	0.7357	1.0615
$c > 0$	< 0.7357	< 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

as closed-form functions of ϵ for the demand models we considered in Section 2. (Details are in the Appendix.) To simplify matters, we assume here 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.)

To further examine how uncertainty over P_m can affect Π^{**}/Π^{*}, we first compute the expected value of the profit ratio when ϵ is uniformly distributed with $B = 0.2$ (i.e., $\epsilon \sim U[-0.2, 0.2]$). We then repeat the calculation but making ϵ a truncated normal variable with $B = 0.2$ and $\sigma = 0.1$. We compare the expected value of the profit ratio relative to the profit ratio that results when the true P_m is known exactly (i.e., $\epsilon = 0$). The results are 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 the ratio when the true value of P_m is known. (We omit the truncated log-log demand as the profit ratio depends on P_m/c, β, and b.) For the truncated normal distribution, the performance depends on the value of σ. For a large σ, the results are very close to the uniform distribution since the values are spread out over the range $[-0.2, 0.2]$. On the other hand, when σ is small, the results become close to the case where we know the exact value of P_m. In some sense, the uniform distribution corresponds to the case where P_m is known to lie in some confidence interval without any additional information. Finally, we test a more volatile noise ϵ that is truncated normal with $B = 0.4$ and $\sigma = 0.1$. In this case, the expected profit ratios are (same order as in Table 2): 1.0103, 1.0215, 1.0985, 1.2015,

"True" inverse demand function	$\mathbb{E}[\Pi^{**}/\Pi^{*}]$(uniform)	$\mathbb{E}[\Pi^{**}/\Pi^{*}]$(normal)	$\Pi^{**}/\Pi^{*}(\epsilon=0)$
Linear: $P_A(Q)=P_m-bQ$	1.0137	1.008	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	1.0192	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	1.0961	1.0887
Monomial: $P_A(Q)=P_m-\gamma Q^n$			
$n=3$	1.205	1.199	1.19
$n=4$	1.2882	1.2815	1.27
Semi-log: $P_A(Q)=P_me^{-\alpha Q}$	1.0748	1.069	1.0615

Table 2 **Expected profit ratios with** $\hat{P}_m=P_m(1+\epsilon)$ **when** $\epsilon\sim U[-0.2,0.2]$ **and** $\epsilon\sim N[0,0.1^2]$

1.284, and 1.0715. **THERE IS NO σ IN TABLE 2. VERY CONFUSING, AND I DON'T THINK WE LEARN ANYTHING FROM THESE CALCULATIONS. I WOULD DROP THE TABLE, AND REWRITE TO JUST REFER TO THE FIGURE.**

Of course, depending on the "draw" for ϵ, the actual profit ratio could be farther from 1. To see how much farther, we use the closed-form expressions in the Appendix to plot the profit ratios as a function of ϵ for $-0.2\leq\epsilon\leq 2$. 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 $\epsilon=-0.2$. 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

In this section, we consider general concave and convex demand curves, and then turn to randomly generated demand functions.

4.1. Concave Demand

We begin with a result for concave demand curves:

Proposition 5. *For any concave demand curve, we have:*

$$1\leq\Pi^{**}/\Pi^{*}\leq 2\,,\qquad 1\leq P^{**}/P^{*}\leq 2.$$

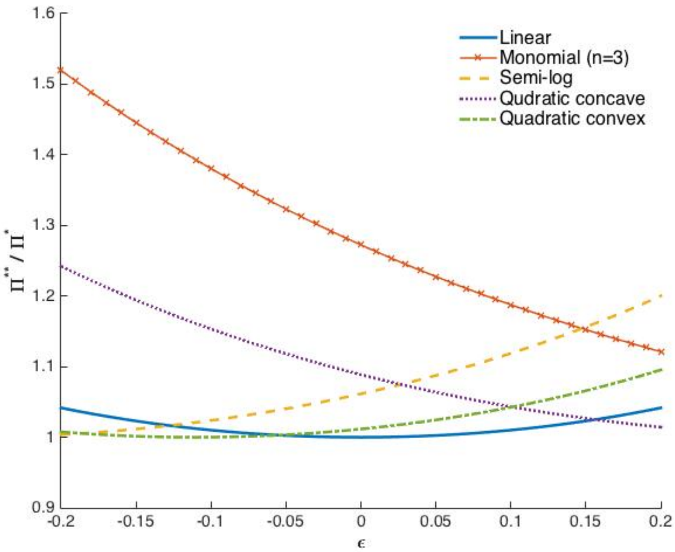

Figure 4 Profit ratios as a function of ϵ for different demand curves

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 account for it.

4.2. Convex Demand

Now consider a general demand curve $Q_A(P)$ that is non-increasing and convex. In this case, the profit and price ratios are unconstrained. Nevertheless, we identify a sufficient condition under which the ratios are bounded away from zero. The methodology for this case relates to the Lipschitz continuity of the demand function. In particular, since $Q_A(P)$ is non-increasing and convex on a bounded support, it is Lipschitz continuous, that is:

$$\frac{|Q_A(P^{**}) - Q_A(P^{*})|}{|P^{*} - P^{**}|} = \frac{Q^{**} - Q^{*}}{P^{*} - P^{**}} \leq K,$$

for some constant $K > 0$. Specifically, K can be any number larger than the maximal derivative in absolute value, i.e., $K \geq \max_x |Q'_A(x)|$. We next impose the following assumption on the function $Q_A(\cdot)$.

Assumption 1. *We assume that $Q_A(P^*) = Q^* \geq \alpha P^*$ for a given $\alpha > 0$.*

When marginal cost is small, P^* is approximately equal to $P_m/2$ so that this assumption states that when the price is about half the maximal price (or $(P_m + c)/2$ when the cost is not negligible), the demand is at least αP^*. An alternative looser condition that implies Assumption 1 is $|Q_A'(P^*)| \geq \alpha \frac{P_m+c}{P_m-c} \geq \alpha$. In other words, the derivative of the demand cannot be too small so that at P^*, demand is not inelastic. Accordingly, this allows us to avoid cases where the plateau area is located below P^*. We are now ready to state the result for convex demand curves.

Proposition 6. *For any convex demand curve that satisfies Assumption 1, we have:*

$$1 \leq \frac{\Pi^{**}}{\Pi^*} \leq \frac{K}{\alpha},$$
$$\frac{\alpha}{K} \leq \frac{P^{**}}{P^*} \leq 2.$$

Note that K is a direct property of the demand function and is guaranteed to exist. Note also that $\alpha \leq K$. The larger α is (closer to K), the better both ratios are. In particular, the ratio K/α characterizes how far the derivative at P^* is relative to the maximal value of the derivative (in absolute value). As expected, for a function close to linear, this ratio is close to 1. When the function is steep around P^* and flattens afterwards, the profit ratio might be far from optimal. More generally, when the demand curve is convex, the performance depends on how far the derivative at P^* is from the highest derivative (i.e., at $P = 0$).

THIS SECTION ON CONVEX DEMAND IS VERY CONFUSING, AND I DON'T SEE WHAT WE LEARN FROM IT. WHAT DOES IT TELL US ABOUT HOW WELL THE PRICING RULE PERFORMS? THERE ARE TWO OPTIONS: (1) DROP THIS SECTION. (2) COMPLETELY REWRITE IT SO THAT IT COHERENT, CLEAR, AND INFORMATIVE.

4.3. General Random Demand Curves

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?

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 compare Π^{**} to the profit Π^* the firm would earn by using our pricing rule, i.e., by setting $P^* = (P_m + c)/2$. We generate 100,000 such demand curves and 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 and so 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 random value for the end of the first segment from a distribution between 0 and P_m, which we will call P_1 (see one realization for P_1 in Figure 5). More precisely, we draw a random variable X_1 between 0 and 1 and $P_1 = P_m X_1$. Next, we independently draw a value for the end of the second segment, but now between 0 and P_1. Call this $P_2 = P_1 X_2$, where X_2 is drawn between 0 and 1. We repeat this process $S-1$ times, drawing a total of $S-1$ independent random variables $X_i;\ i = 1, \ldots, S-1$ between 0 and 1, generating a random demand curve with 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 numerically calculate P^{**}, Π^{**}, and the profit ratio Π^{**}/Π^{*} for $c = 0$ and $c = 0.5P_m$.

We draw the random variables $X_i;\ i = 1, \ldots, S-1$ using a power distribution of the form $X^{1/\alpha}$, where $\alpha \geq 1$ is a skewing parameter and X is uniformly distributed between 0 and 1. Note that when $\alpha = 1$, this reduces to the uniform distribution. For simplicity, we present the results for the case of a uniform distribution (i.e., $\alpha = 1$); we obtained very similar results when $\alpha = 1.5$.

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.

Observe 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, for 80% (90%) of the demand curves, the profit ratios are 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% (79%) are less than 1.05. Thus it is reasonable to expect our pricing rule to yield a profit close to what would result if the firm knew its actual demand curve.

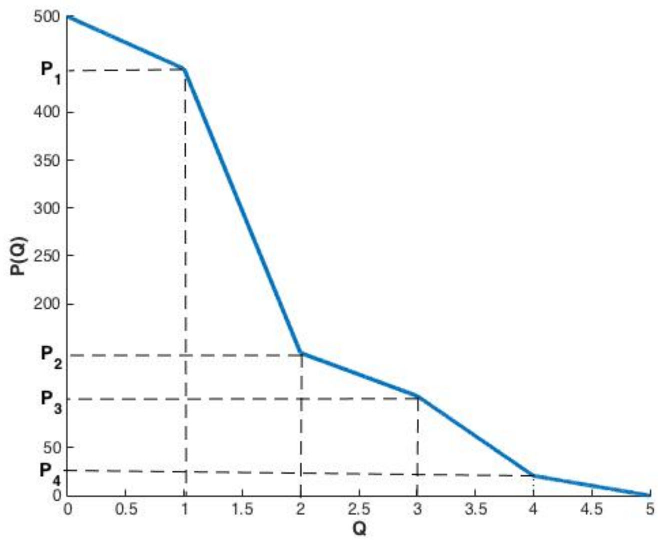

Figure 5 **Randomly generated inverse demand curve with $S = 5$ pieces**

$c = 0$				$c = 0.5P_m$			
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**

5. Welfare Implications

we now compare the total 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 at P^{**}. We also look at consumer surplus separately to see how our pricing rule affects consumers. The total welfare, denoted by $W(P)$, is:

$$W(P) = \Pi(P) + CS(P) = (P - c)Q + \Big[\int_0^Q P_A(y)dy - PQ\Big]. \tag{8}$$

We are interested in $W(P^{**})/W(P^*) \equiv W^{**}/W^*$ and $CS(P^{**})/CS(P^*) \equiv CS^{**}/CS^*$. Note that these ratios can be less than one, i.e., our pricing rule can increase the total welfare and/or the consumer surplus relative to that when the profit maximizing price P^{**} is used. In particular, we have the following result.

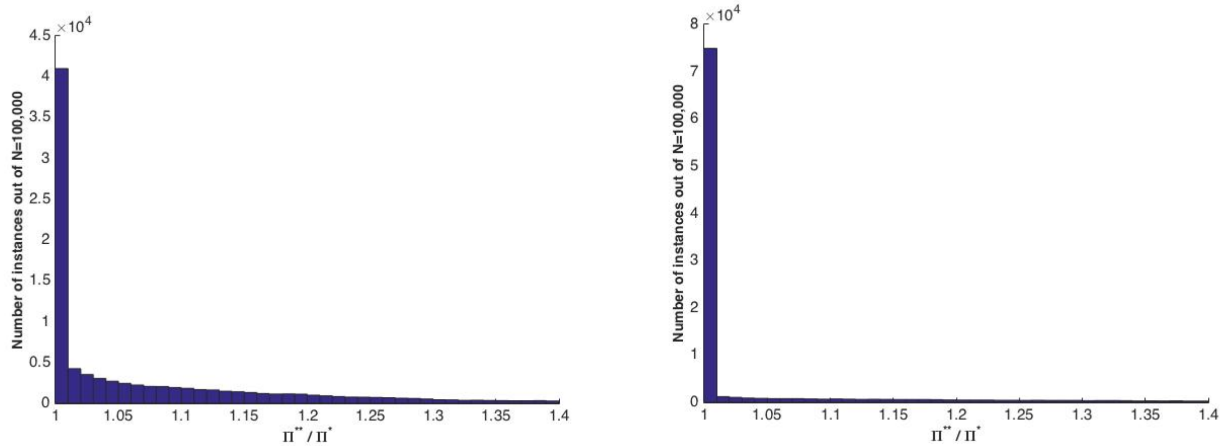

(a) $c=0$ (b) $c=0.5P_m$

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

Proposition 7. *If the actual demand is convex, then $W^{**} \geq W^{*}$ and $CS^{**} \geq CS^{*}$; whereas if it is concave, $W^{**} \leq W^{*}$ and $CS^{**} \leq CS^{*}$.*

Indeed, as long as $P \geq c$, $W(P)$ is non-increasing, so that the inequalities on W follow immediately from Theorem 1. The inequalities on CS also follow from Theorem 1 and from the fact that the consumer surplus is a non-increasing function of the price (for $P \leq P_m$). If the demand is concave, we know from Theorem 1 that $P^{*} \leq P^{**}$, so in this case using the "wrong" price P^{*} improves both the total welfare and consumer surplus (i.e., the benefit to consumers exceeds the loss to the firm). However, if the demand is convex, both the firm and consumers are worse off.

Next, we calculate the welfare and consumer surplus ratios (i) analytically for the demand models in Section 2 (see Proposition 8) and (ii) computationally for randomly generated demand curves following the approach of Section 4.3. To simplify matters, we assume that $c=0$. We do not report the details of the derivations for conciseness. The closed form expressions are as follows:

Proposition 8. *The welfare and consumer surplus ratios, W^{**}/W^{*} and CS^{**}/CS^{*}, for the different demand models are.*

- ***Quadratic convex:*** $P_A(Q) = P_m - b_1 Q + b_2 Q^2$; $b_1, b_2 \geq 0$ *and* $b_2 \leq b_1^2/4P_m$

$$\frac{W^{**}}{W^{*}} \leq 1.26045 \quad \textit{and} \quad \frac{CS^{**}}{CS^{*}} \leq 1.5756$$

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

$$\frac{W^{**}}{W^{*}} \leq \frac{16\sqrt{2}}{15\sqrt{3}} = 0.8709 \quad \textit{and} \quad \frac{CS^{**}}{CS^{*}} \leq 0.544$$

- ***Monomial:*** $P_A(Q) = P_m - \gamma Q^n$

$$W^{**}/W^{*} = \frac{2^{\frac{1}{n}+1} n}{(n+1)^{\frac{1}{n}+1}} \frac{n+1-1/(n+1)}{(2-1/(n+1))} \quad \textit{and} \quad \frac{CS^{**}}{CS^{*}} = \left(\frac{2}{n+1}\right)^{\frac{1}{n}+1}$$

- ***Semi-log:*** $P_A(Q) = P_m e^{-\alpha Q}$

$$W^{**}/W^{*} = 2(1 - e^{-1}) = 1.2642 \quad \textit{and} \quad \frac{CS^{**}}{CS^{*}} = 1.722$$

We omitted the truncated log-log demand as the welfare and consumer surplus ratios are complicated expressions that depend on P_m/c and β. For the logit demand, no closed form expressions exist. For the exponential demand, one can obtain: $CS^{**}/CS^{*} \leq 1.6487$ (the welfare ratio is unbounded). For the monomial demand, when $n = 3$ and $n = 4$, W^{**}/W^{*} is 0.85 and 0.888 respectively, and CS^{**}/CS^{*} is 0.397 and 0.318 respectively. Also, as the order of the monomial n increases, W^{**}/W^{*} approaches 1 whereas CS^{**}/CS^{*} approaches 0. Indeed, as n increases, the inverse demand curves becomes more concave so there is a greater transfer of welfare from the firm to consumers. One can see that for these demand models, the loss in total welfare from using our approximation is at most 26%, but in some cases the loss (or gain) in consumer surplus can be quite large. (For example, for the semi-log demand, when $c = 0$ the loss in profit is 6.16% but the decrease in consumer surplus is 72%.)

We next calculate W^{**}/W^{*} and CS^{**}/CS^{*} for randomly generated demand curves following the approach of Section 4.3. As before, we fix P_m, c, and Q_{max} and compute the ratios for $c/P_m = 0$ and 0.5, using 100,000 randomly generated demand curves. (The results are very similar for different values of S.) The average welfare ratio for $c = 0$ and $c = 0.5P_m$ are 1.139 and 0.993 respectively. As for CS^{**}/CS^{*}, the average ratios for $c = 0$ and $c = 0.5P_m$ are 1.1885 and 0.9148 respectively (see Figure 7 in the Appendix). As one can see from the histograms, although CS^{**}/CS^{*} is close to 1 on average, for a significant fraction of demand curves, consumers will be either better off or worse off. Thus although our pricing rule often yields profits close to optimal, the misspecification of demand can have a significant (positive or negative) impact on consumer surplus.

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.

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 can estimate the maximum price P_m it can charge and still expect to sell some units, and (iii) the firm need not predict the quantity it will sell. These conditions hold for many new products and services, especially those introduced by technology companies. 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 studied this question when the true demand curve is one of several commonly used demand functions, or even if it is a more complex function (e.g., randomly generated). Often, the firm will earn a profit reasonably close to the optimal profit it could earn if it knew the true demand curve. We also identified cases where the profit loss is significant. Our results can help us understand why linear demand functions are so popular in many applications (e.g., Koushik et al. 2012, Pekgün et al. 2013).

As mentioned in the Introduction, the results of this paper can be extended to cases where the firm does not know P_m, but it can estimate a price $\bar{P} < P_m$, such that at $\bar{P}$ the firm can sell to a small set of customers. However, as one would expect, the pricing rule $P^* = (\bar{P} + c)/2$ will not perform as well. For example, the profit ratio for the semi-log demand used in Section 2.3 will be now at most 1.181 for any value of $\bar{P}$ (the proof follows the same logic as the proof of Proposition 3 by translating the inverse 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, potentially 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 intertemporally price discriminate or to respond to changing market conditions, nor does it offer different prices to different groups of customers. An additional limitation of our analysis is the fact that we do not capture competitors' effects as well as multiple products with substitution or complementarity effects. We also imposed the assumption that capacity is not a relevant operational decision by focusing on specific type of products. Extending the approach for products that require capacity planning is an interesting avenue for future research. Finally, we have ruled out learning about demand, either passively or via experimentation, which has been the focus of the earlier literature on pricing with uncertain demand (learning and earning). 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 often effective.

Acknowledgments

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.

References

Adida, Elodie, 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, René Caldentey. 2011. Revenue management with incomplete demand information. Wiley Encyclopedia of Operations Research and Management Science.

Aviv, Yossi, Amit Pazgal. 2005. A partially observed markov decision process for dynamic pricing. *Management Science* **51**(9) 1400–1416.

Ball, Michael O, Maurice Queyranne. 2009. Toward robust revenue management: Competitive analysis of online booking. *Operations Research* **57**(4) 950–963.

Balvers, Ronald J, Thomas F Cosimano. 1990. Actively learning about demand and the dynamics of price adjustment. *The Economic Journal* 882–898.

Bergemann, Dirk, Karl Schlag. 2011. Robust monopoly pricing. *Journal of Economic Theory* **146**(6) 2527–2543.

Bertsimas, Dimitris, Georgia Perakis. 2006. *Dynamic pricing: A learning approach*. Springer.

Besbes, Omar, Assaf Zeevi. 2009. Dynamic pricing without knowing the demand function: Risk bounds and near-optimal algorithms. *Operations Research* **57**(6) 1407–1420.

Besbes, Omar, Assaf Zeevi. 2015. On the (surprising) sufficiency of linear models for dynamic pricing with demand learning. *Management Science* **61**(4) 723–739.

Carroll, Gabriel. 2015. Robustness and linear contracts. *American Economic Review* **105**(2) 536–63.

Chu, Chenghuan Sean, Phillip Leslie, Alan Sorensen. 2011. Bundle-size pricing as an approximation to mixed bundling. *American Economic Review* **101** 262–303.

Cohen, Maxime C, Ngai-Hang Zachary Leung, Kiran Panchamgam, Georgia Perakis, Anthony Smith. 2017. The impact of linear optimization on promotion planning. *Operations Research* **65**(2) 446–468.

Dawes, Robyn M. 1979. The robust beauty of improper linear models in decision making. *American psychologist* **34**(7) 571.

Farias, Vivek F, Benjamin Van Roy. 2010. Dynamic pricing with a prior on market response. *Operations Research* **58**(1) 16–29.

Handel, Benjamin R, Kanishka Misra. 2015. Robust new product pricing. *Marketing Science* **34**(6) 864–881.

Hansen, Lars Peter, Thomas J. Sargent. 2008. *Robustness*. Princeton University Press.

Koushik, Dev, Jon A Higbie, Craig Eister. 2012. Retail price optimization at intercontinental hotels group. *Interfaces* **42**(1) 45–57.

Krishnan, Trichy V, Frank M Bass, Dipak C Jain. 1999. Optimal pricing strategy for new products. *Management Science* **45**(12) 1650–1663.

Lim, Andrew EB, 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.

Milon, J Walter, Jonathan Gressel, David Mulkey. 1984. Hedonic amenity valuation and functional form specification. *Land Economics* **60**(4) 378–387.

Montgomery, Alan L. 1997. Creating micro-marketing pricing strategies using supermarket scanner data. *Marketing science* **16**(4) 315–337.

Mulugeta, Dawit, Jason Greenfield, Tison Bolen, Lisa Conley. 2013. Price- and cross-price elasticity estimation using sas. SAS Global Forum.

Pekgün, Pelin, Ronald P Menich, Suresh Acharya, Phillip G Finch, Frederic Deschamps, Kathleen Mallery, Jim Van Sistine, Kyle Christianson, James Fuller. 2013. Carlson rezidor hotel group maximizes revenue through improved demand management and price optimization. *Interfaces* **43**(1) 21–36.

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.

Urban, Glen L, Bruce D Weinberg, John R Hauser. 1996. Premarket forecasting of really-new products. *The Journal of Marketing* 47–60.

Wilman, Elizabeth A. 1981. Hedonic prices and beach recreational values. *Advances in Applied Microeconomics* 77–103.

Appendix

Equivalence with the valuation model

Consider a representative consumer with a random valuation for the product. The valuation is assumed to be between the cost c and the maximal price P_m. We assume that the seller knows the valuation distribution, represented by the cdf $\mathbb{F}(\cdot)$ and the pdf $f(\cdot)$. We also assume that the seller seeks to maximize its expected profit, given by:

$$\Pi(p) = (p-c)\mathbb{P}(p < v) = (p-c)[1-\mathbb{F}(p)], \tag{9}$$

where p denotes the price set by the seller and v denotes the valuation of the consumer (unknown to the seller). If the price is below the valuation, the consumer will purchase the item and the seller extracts a profit of $p-c$. Otherwise, there is no sale and 0 profit for the seller. One can take the first order condition and obtain:

$$P^* = c + \frac{1-\mathbb{F}(P^*)}{f(P^*)}. \tag{10}$$

Eqn. (10) is a well-known result called the virtual valuation. If we further assume that the valuation distribution is uniform in $[c, P_m]$, we have: $f(p) = 1/(P_m - c)$ and $\mathbb{F}(p) = (p-c)/(P_m-c)$. Therefore, eqn. (10) becomes: $P^* = (P_m + c)/2$.

Proof of Theorem 1

The actual inverse demand curve, $P_A(Q)$, satisfies $P_A(0) = P_m$. 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^{**}:

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

Recall that $P^* = (P_m + c)/2$ and thus: $P^{**} = P^* + 0.5\big[f(Q^{**}) - f'(Q^{**})Q^{**}\big]$. 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 $R_1 = 0.5 f''(\zeta)(x-a)^2$, for some $\zeta \in [x, a]$. Then:

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

Consequently, $P^{**} - P^* = -P_A''(\zeta)(Q^{**})^2/4$, for some $\zeta \in [0, Q^{**}]$. Therefore, 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 Proposition 1

Convex case: We have:

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

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 $P^{**} = (4/9)P_m$ and $Q^{**} = (2P_m)/(3b_1)$. Finally, we can now compute both profits:

$$\Pi^* = \frac{b_1 P_m}{4b_2}\Big(1 - \frac{1}{\sqrt{2}}\Big) = \frac{P_m^2}{b_1}\Big(1 - \frac{1}{\sqrt{2}}\Big), \quad \Pi^{**} = \frac{2b_1 P_m}{27b_2} = \frac{8P_m^2}{27b_1}.$$

Then the profit and price ratios are: $\Pi^{**}/\Pi^* = 8\sqrt{2}/\big[27(\sqrt{2}-1)\big] = 1.0116$, $\quad P^{**}/P^* = 8/9$. These are the largest values for the ratios, so that the corresponding inequalities hold.

Concave case: The optimal Q^{**} is obtained by equating marginal revenue to 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 given by:

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

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 so as to satisfy $Q^* > 0$. The profit is then:

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

Expressing the profit and price ratios as functions 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 b_1 is, the more linear the function is, 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:

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

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

Proof of Proposition 2

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 $P^{**} = P_A(Q^{**}) = (nP_m - c)/(n+1)$. 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 $Q_A(P^*) = \big[(P_m - c)/(2\gamma)\big]^{1/n}$. Therefore, $\Pi^* = (P^* - c)Q_A(P^*) = \big[(P_m - c)^{\frac{1}{n}+1}\big]/\big[2^{\frac{1}{n}+1}\gamma^{1/n}\big]$. 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 Proposition 3

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 $P^{**} = P_m e^{-1}$ and $\Pi^{**} = P_m e^{-1}\alpha^{-1}$. If the firm prices at P^*, the profit is $\Pi^* = (P^* - c)Q_A(P^*) = 0.5P_m Q_A(P^*)$. Since $c = 0$ and $P^* = 0.5P_m$, we obtain: $Q_A(P^*) = -(1/\alpha)\log(0.5)$, and hence $\Pi^* = 0.5P_m \log(2)/\alpha$. We then have:

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

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

$$\frac{\partial}{\partial c}\Big[\frac{P^{**}}{P^*}\Big] = \frac{\frac{\partial P^{**}}{\partial c}P^* - \frac{\partial P^*}{\partial c}P^{**}}{(P^*)^2}. \tag{11}$$

For eqn. (11) 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:

$$\frac{\partial P^{**}}{\partial c} \geq \frac{P^{**}}{P_m + c}. \tag{12}$$

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 $\partial P^{**}/\partial c$:

$$\frac{\partial P^{**}}{\partial c} = \frac{1 + \alpha P^{**}\frac{\partial Q^{**}}{\partial c}}{1 - \alpha Q^{**}}. \tag{13}$$

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

$$\frac{\partial P^{**}}{\partial c} = -\alpha P^{**}\frac{\partial Q^{**}}{\partial c}. \tag{14}$$

By combining (13) and (14), we obtain $\partial P^{**}/\partial c = 1/(2 - \alpha Q^{**})$. Since the demand curve is convex, from Theorem 1: $P^{**} \leq P^* = (P_m + c)/2$ and therefore: $P^{**}/(P_m + c) \leq 0.5$. From the first order condition, $0 \leq 1 - \alpha Q^{**} \leq 1$ (so that $P^{**} \geq c$). Thus $1 \leq 2 - \alpha Q^{**} \leq 2$, so $1/(2 - \alpha Q^{**}) \geq 0.5$, implying that (12) is satisfied. This concludes the proof for the price ratio.

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

Proof of Proposition 4

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: $P^{**} = \beta c/(\beta-1)$ and $\Pi^{**} = Q_0 c/(\beta-1)\left[\frac{\beta c}{(\beta-1)P_m}\right]^{-\beta}$. By requiring $\beta \geq P_m/(P_m - c)$ we ensure that $P^{**} \leq P_m$. We next compute the profit under 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:

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

Expressions for Section 3

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

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

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

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

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

- **Monomial:** $P_A(Q) = P_m - \gamma Q^n \quad \Pi^{**}/\Pi^*(\epsilon) = \frac{2^{\frac{1}{n}+1}n}{(n+1)^{\frac{1}{n}+1}}\frac{1}{(1+\epsilon)(1-\epsilon)^{1/n}}$
- **Semi-log:** $P_A(Q) = P_m e^{-\alpha Q} \quad \Pi^{**}/\Pi^*(\epsilon) = \frac{2e^{-1}}{(1+\epsilon)\log(\frac{2}{1+\epsilon})}$
- **Log-log (truncated):** $P_A(Q) = \begin{cases} P_m\,; & \text{if } Q < Q_0 \\ P_m(Q/Q_0)^{-1/\beta}\,; & \text{if } Q \geq Q_0 \end{cases}$

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

Proof of Proposition 5

Consider any non-increasing concave demand curve. We know from Theorem 1 that $P^* \leq P^{**}$. 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 profit: $\Pi^{**} = (P^{**} - c)Q_A(P^{**}) \leq 2(P^* - c)Q_A(P^{**}) \leq 2(P^* - c)Q_A(P^*) = 2\Pi^*$, where the last inequality follows form the fact that $Q_A(\cdot)$ is non-increasing. In conclusion, we have $1 \leq \Pi^{**}/\Pi^* \leq 2$ and $1 \leq P^{**}/P^* \leq 2$.

Proof of Proposition 6

Since $Q_A(P)$ is convex and non-increasing on a bounded support, it is also Lipschitz continuous, and therefore:

$$\frac{|Q_A(P^{**}) - Q_A(P^*)|}{|P^* - P^{**}|} \leq K,$$

for some constant $K > 0$. In addition, from Assumption 1, we know that $Q_A(P^*) = Q^* \geq \alpha P^*$ for some given $\alpha > 0$. We assume without loss of generality that the marginal cost is equal to zero, i.e., $c = 0$. Since the demand curve is convex and non-increasing (and so is its inverse), we have: $P^{**} \leq P^*$ and $Q^* \leq Q^{**}$. Then, the profit ratio inequality we want to show can be written as follows:

$$\frac{\Pi^{**}}{\Pi^*} = \frac{Q^{**}P^{**}}{Q^*P^*} = \frac{Q^*P^{**} + (Q^{**} - Q^*)P^{**}}{Q^*P^{**} + Q^*(P^* - P^{**})} \leq \frac{K}{\alpha}.$$

Since $K/\alpha \geq 1$, it is sufficient to show:

$$\frac{Q^{**} - Q^*}{P^* - P^{**}} \leq \frac{K}{\alpha}\frac{Q^*}{P^{**}}. \tag{15}$$

Since $P^{**} \leq P^*$, a sufficient condition that implies eqn. (15) is given by:

$$\frac{Q^{**} - Q^*}{P^* - P^{**}} \leq \frac{K}{\alpha}\frac{Q^*}{P^*}. \tag{16}$$

Since $Q^* \geq \alpha P^*$, we have: $\frac{K}{\alpha}\frac{Q^*}{P^*} \geq K$ and therefore eqn. (16) follows from the Lipschitz continuity. We then conclude that: $1 \leq \Pi^{**}/\Pi^* \leq K/\alpha$.

We next show the inequalities for the price ratio. We know that the optimal price P^{**} satisfies the first order condition:

$$P^{**} = \frac{Q_A(P^{**})}{|Q'_A(P^{**})|}.$$

Then, the prices ratio can be written as:

$$\frac{P^{**}}{P^*} = \frac{Q_A(P^{**})}{|Q'_A(P^{**})|P^*} \geq \frac{Q_A(P^*)}{|Q'_A(P^{**})|P^*} \geq \alpha\frac{1}{|Q'_A(P^{**})|} \geq \frac{\alpha}{K},$$

where the first inequality follows from the monotonicity of $Q_A(\cdot)$ and the fact that $P^* \geq P^{**}$, the second inequality from Assumption 1, and the third inequality from convexity. Finally, note that since $P^* = (P_m + c)/2$, we have $2P^* \geq P_m \geq P^{**}$, so that $P^{**}/P^* \leq 2$.

Histograms for the consumer surplus

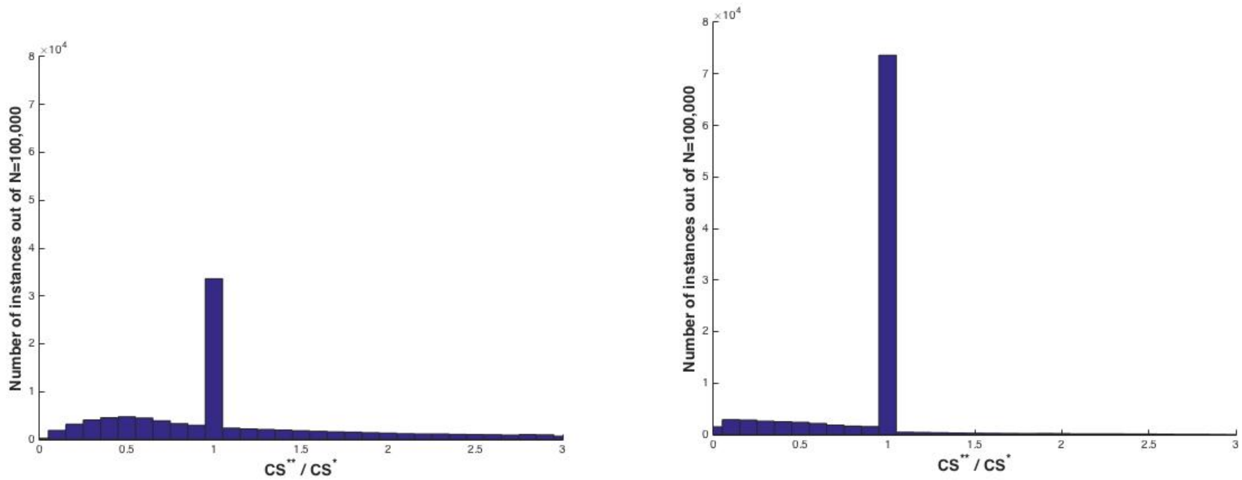

(a) $c = 0$ (b) $c = 0.5P_m$

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