Pricing With Limited Knowledge of Demand

Maxime Cohen, Georgia Perakis, Robert Pindyck

Pricing With Limited Knowledge of Demand

Introduction

In some settings, companies need to set prices for new products for which there is 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 in its new iTunes store in 2001, and more recently, pricing its new iWatch), 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 can be determined 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 price elasticities. How should firms set prices in such settings? We show that in many situations, the firm need only estimate the *maximum price* it can charge and still expect to sell at least some units. It then sets price as though the actual demand curve were linear.

Assuming constant marginal cost, the firm calculates its optimal price as half of the sum between the maximum price and the marginal cost. This price is independent of the slope of the linear demand curve, although the resulting quantity 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 drugs, music downloads, or software), knowledge of the slope, and thus the ability to predict its sales, is immaterial. As a result, the only problem at hand is to set the price.

How well can the firm expect to do if it sets the price based on a linear demand curve? Suppose that with precise knowledge of its true demand curve, the firm would set the optimal price and earn the maximum profit. The question we address is simple: How close will the profit based on the price from the linear demand curve be to the maximum profit? In other words, how well can the firm expect 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 will do very well. In what follows we will quantify this statement.

In many settings, for example, when introducing a new product in the market, there is not enough data. A way to deal with this problem in practice is through considering comparable products for which data is available and then enrich this process through learning the demand and adjusting the pricing over time. Nevertheless, a key question is how to determine the first price to set in the time horizon. In other settings, the seller is not able to dynamically adjust the price over time and hence setting the first price "accurately" is critical. In this paper, we answer this question where the seller needs to set the price of a new item with limited data on demand. In particular, we propose a way of pricing the item using only cost and maximal price information. As a result, we impose minimal assumptions on the problem and show that only with this information one can obtain a good performance relative to the true (but unknown) optimal profit.

Summary of Contributions: 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 the 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 will do very well - its profit will 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 test performance bounds computationally for some general demand models.

Literature review

The problem of pricing under limited demand information has received significant attention in both the Economics and the Operations Management communities. In the past decade, four different approaches have been considered in the literature.

The first stream posits a known parametric demand form as a function of the price. In some contexts, knowledge of the problem can allow the seller to believe that the demand admits a particular parametric form such as linear or logarithmic. The key question is how to estimate and learn the true parameters from data. That is, the seller infers the parameters from the past prices and realized demands by using methods such as least squares or maximum likelihood estimation (see for example, [4] and [7]).

A second stream of literature takes a Bayesian approach. The seller postulates a parametric demand model jointly with the prior distribution (for example, on the reservation price), using the initial knowledge of the demand. Demand observations are then used to update the prior into a posterior distribution (for more details see for example, [9] and [2]). In addition, among the first papers to adopt this approach was the work by [11]. This paper used multi-armed bandits to learn the optimal pricing strategy. A common assumption that is sometimes criticized relies on the exact knowledge of the prior distribution.

The third stream of literature considers the interplay between learning the demand curve and optimizing revenues over time without assuming any parametric form. This stream proposes an approach based on exploration and exploitation. In the first (exploration) phase, the seller starts with an initial “guesstimate” of the demand and sets a price to explore the demand. In the second (exploitation) phase, the seller optimizes using the prices and demands from the previous phase (see, e.g., [5]). A key result is to study how asymptotically the proposed approach converges to the true demand. More close to our paper, [6] investigates the revenue loss in a multi-period setting incurred if the seller uses a simple parametric demand model that differs significantly relative to the underlying demand curve (i.e., is mis-specified).

The final stream of literature considers the dynamic pricing problem using robust optimization. This approach assumes that the demand parameters lie in an uncertainty set. The seller then optimizes for the worst case parameters in this set (see, e.g., [1]). Along the same spirit, an alternative is to consider a distributionally robust approach. In this case, the seller aims to be robust with respect to a class of demand distributions with similar parameters such as mean and standard deviation (see, e.g., [3], [10] and [8]). Although the previous approaches incorporate model uncertainty may yield conservative pricing strategies.

Results

In this paper, we study the problem of pricing a product for which demand information is very costly or even not available. We impose minimal assumptions on the problem: only the constant marginal cost and the maximal price consumers are willing to pay are known. Finding a good price for such an application is crucial as it may determine the long-term success of the product and can have a significant impact on profitability. We propose a simple way of pricing by approximating the true unknown inverse demand curve by a linear function for which one can easily compute the optimal price. We show that using the price based on the linear approximation yields a good profit performance for a wide range of inverse demand curves. We derive analytical bounds on the profits and prices ratios for different demand models. In particular, we show a worst case guarantee of 8.8% for quadratic demand and 6% for semi-log (exponential) demand. We also show some tight guarantees in closed form for monomial and log-log demands. Finally, we consider several practical inverse demand curves and show computationally that the performance of our method is not far from optimal. Finally, we present a general result for any non-increasing concave demand curve that yields a 2 approximation in the worst case. We also discuss some of the limitations of this approach by describing cases where the approximation might yield a poor performance. However, in most cases, the profit guarantee is about 5 to 25%. As a result, we propose an operational and simple method to price items for which no demand information is available along with the guarantee that the profit is not far from optimal.

References

[1] Elodie Adida and Georgia Perakis. A robust optimization approach to dynamic pricing and inventory control with no backorders. *Mathematical Programming*, 107(1-2):97–129, 2006.

[2] Victor F Araman and René Caldentey. Dynamic pricing for nonperishable products with demand learning. *Operations research*, 57(5):1169–1188, 2009.

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

[4] Dimitris Bertsimas and Georgia Perakis. *Dynamic pricing: A learning approach.* Springer, 2006.

[5] Omar Besbes and Assaf Zeevi. On the minimax complexity of pricing in a changing environment. *Operations research*, 59(1):66–79, 2011.

[6] Omar Besbes and Assaf Zeevi. On the (surprising) sufficiency of linear models for dynamic pricing with demand learning. *Columbia Business School Research Paper*, (12-5), 2013.

[7] Alexandre X Carvalho and Martin L Puterman. Learning and pricing in an internet environment with binomial demands. *Journal of Revenue and Pricing Management*, 3(4):320–336, 2005.

[8] Serkan S Eren and Costis Maglaras. Monopoly pricing with limited demand information. *Journal of Revenue & Pricing Management*, 9(1):23–48, 2010.

[9] Kyle Y Lin. Dynamic pricing with real-time demand learning. *European Journal of Operational Research*, 174(1):522–538, 2006.

[10] Georgia Perakis and Guillaume Roels. Robust controls for network revenue management. *Manufacturing & Service Operations Management*, 12(1):56–76, 2010.

[11] Michael Rothschild. Models of market organization with imperfect information: A survey. *The Journal of Political Economy*, pages 1283–1308, 1973.