Marketing Colloquia

Fall 2015

Kanishka Misra
Assistant Professor of Marketing
University of Michigan, Ross School of Business

Thursday, December 17, 2015
Where: *741 JMHH*
When: *12:00 PM* to 1:20 PM

Dynamic Robust Pricing with Multi-Armed Bandits

ABSTRACT: In this paper we consider the pricing decision for a firm launching a new product with incomplete demand information. We propose a dynamic price experimentation policy where the firm considers maximize long run profits without strong parametric assumptions. Our proposed pricing algorithm sequentially sets prices to balance currently earning profits and learning about demand for future profits.

We consider a setting where prior to launch the firm only knows a non-parametric set of feasible demand curves. In each period, the firm sets prices and observes which consumers were exposed to the prices and their purchase decisions. The firm incorporates this additional information for future pricing decisions.

Our paper builds on the non-parametric multi-arm bandit (MAB) algorithms where the firm faces a tradeoff between earning profits in the current period and future learning. In non-parametric MAB algorithms (e.g. upper confidence bound), when a particular arm is played (e.g. a price is charged), the firm's learning is limited to the expected profits from that arm (price). We extend this to incorporate economic theory in the form of weakly downward sloping demand curves. When a consumer is exposed to a price, we can partially identify the consumer's underling preference. For example, if the consumer purchases at a price of $3, we infer that she will purchase for all prices below $3. We extend our approach to consider learning across consumer segments (as opposed to individual consumers). This allows a firm to consider our model in settings where consumers have limited repeat purchases. We estimate the heterogeneity of consumer preferences within a segment as well as the preferences across segments.

Our paper adds to the literature in revenue management as our method is robust to traditionally strong parametric assumptions about demand and firm learning. We also add to the partially identified econometric approaches in marketing and economics where prior approaches can solve a similar pricing problem for a small number of periods. In our paper, we solve the problem for any number of periods using scalable distribution-free algorithms.

We show theoretically that the proposed approach achieves closer to ex-post optimal profits (reduces ex-post statistical regret) than existing methods. In Monte Carlo simulations we test a range of distributions of heterogeneity and show an improved empirical performance. We show that segmentation can further expedite profit learning. The speed of this learning depends on the heterogeneity of consumer preferences within the segments (or the quality of the segmentation). Substantively we show that the optimal pricing suggests high initial prices, which provides a supply side learning alternative explanation for price patterns consistent with price skimming.

Wharton
University *of* Pennsylvania
MARKETING
DEPARTMENT