Robust Dynamic Pricing With Strategic Customers

Yiwei Chen, Vivek Farias

Robust Dynamic Pricing With Strategic Customers

The discipline of Revenue or Yield Management (RM) has, over the last two decades come to occupy a place of prominence in applied Operations Research. Today, applications of revenue management run the gamut from dynamic pricing in the airline industry to hospitality to retail. The following dynamic pricing problem is one of the central problems in revenue management: A seller is endowed with an inventory of a single product that she must sell over a finite horizon. She cannot acquire additional inventory over the course of the horizon and unsold inventory has negligible salvage value. Customers arrive randomly over the course of the selling horizon with the intent of purchasing a single unit of the product. Should the posted price upon a customers arrival exceed his valuation he leaves the system for good; otherwise he purchases a single unit of the product. The seller seeks to dynamically adjust prices with a view to maximizing expected revenue. For typical assumptions on the customer arrival process – assuming, for instance, a renewal process – this problem admits a tractable dynamic programming solution. Despite its simplicity, the canonical nature of this problem serves to highlight an important view of the role of dynamic pricing in revenue management as a *tool to hedge against uncertain demand.*

In the past decade, it has become amply clear that for a number of principal RM applications, assuming myopic customer behavior, as in the problem above, is no longer a tenable assumption. In spite of this realization, optimal dynamic mechanisms proposed for the

version of this central problem that assumes *strategic* customers, face the following critique:

1. They do not admit pure pricing implementations requiring instead devices such as lotteries or end of season 'fire-sale' auctions. This typically rules out applying these mechanisms in scenarios where an anonymous posted price mechanism is the norm (unfortunately, the majority of RM applications).

2. They require the seller to calibrate a model of the customers strategicity, by learning for instance, inter-temporal preferences or search costs. This latter learning problem is non-trivial given the naturally censored data the seller has access to. The customer utility models assumed also typically place strong restrictions on inter-temporal preferences.

3. These mechanisms frequently impute sophisticated purchase timing decisions in equilibrium that are arguably as untenable as the myopic assumption given the burden they place on the customer from a computational and data standpoint.

The present paper seeks to make progress on these fronts at the expense of optimality. In particular, we propose a class of dynamic pricing policies which may be interpreted as solving the simple dynamic pricing problem for myopic customers with the additional restriction that the pricing policy satisfy what we call a 'restricted sub-martingale constraint'. We dub such policies 'robust dynamic pricing' policies and exhibit a simple to compute policy within this class with attractive properties:

1. Computing a robust policy require minimal data on customers beyond what is already required by the standard dynamic pricing problem assuming myopic customers. The specific class of robust policies we compute require *no* additional data.

2. Robust pricing policies induce customers to behave myopically under mild assumptions on customer utility.

3. We exhibit a simple to compute robust pricing policy that is guaranteed to garner revenues that are within 29% of those garnered under the optimal dynamic mechanism.

In addition to the features above, numerical results suggest that the performance of our robust pricing policy can be expected to be substantially superior to what the uniform theoretical guarantee we prove suggests. These numerical experiments also show that the loss in revenue due to an incorrectly calibrated, but otherwise optimal, dynamic mechanism can be substantial over and above the issues raised earlier.

In a nutshell, the present paper provides a tractable, provably robust approach to dynamic pricing in the face of forward looking customers. The approach is robust in that it provides revenue guarantees while making minimal assumptions of customers' inter temporal utilities and search costs, and requires minimal data about the same.