POLITECNICO MILANO 1863

ML cube

Dynamic Pricing with Volume Discounts in Online Settings

Marco Mussi*, Gianmarco Genalti*, Alessandro Nuara, Francesco Trovò, Marcello Restelli and Nicola Gatti

{marco.mussi, gianmarco.genalti, francesco1.trovo, marcello.restelli, nicola.gatti}@polimi.it, alessandro.nuara@mlcube.com

Dynamic Pricing

Dynamic pricing refers to the practice of keeping pricing schedules dynamic in time.

Why this Work?

- E-commerce websites may face different kinds of users, both customers and businesses.
- Different kinds of users present different needs in terms of item quantity.

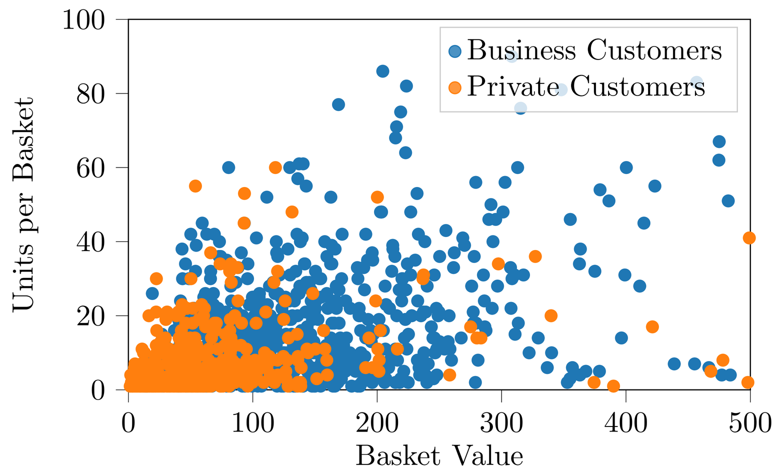

- The kind of user the website will face is not known *a priori*.
- We want to propose different prices for different volume thresholds thanks to the introduction of discounts.

Setting and Goal

Setting

- We want to price the products independently.
- We consider a setting in which only transaction data are available.
- The demand curve is assumed to be monotonic in price.

Goal

At each time t, find:

- a set of volume thresholds $\boldsymbol{\omega}_t := [\omega_{1t}, \dots, \omega_{\eta t}] \in \mathbb{N}^\eta$
- a set of prices $\boldsymbol{p}_t := [p_{1t}, \dots, p_{\eta t}] \in \mathcal{P}^\eta$

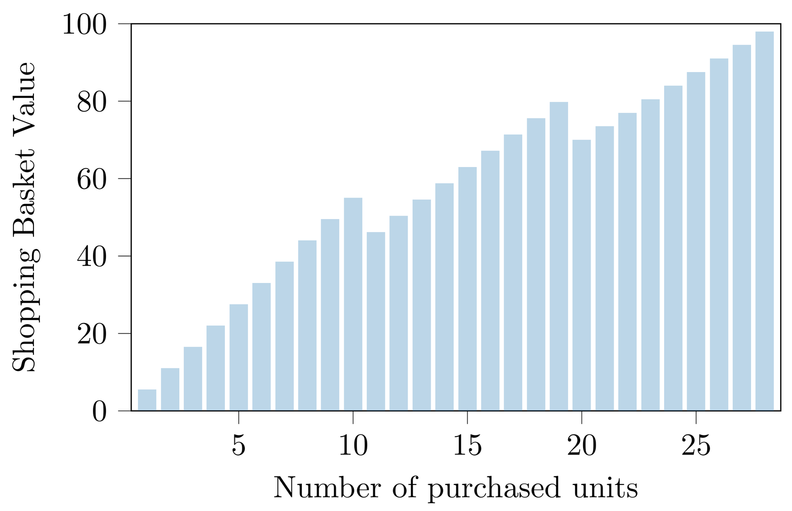

The goal is to select sets $\boldsymbol{\omega}_t$ and $\boldsymbol{p}_t$ maximizing the total profit $R(T)$ over time horizon T:

$$R(T) := \sum_{t=1}^{T} \sum_{i=1}^{\eta} (p_{it} - c) \cdot v_i(\boldsymbol{p}_t, \boldsymbol{\omega}_t, t)$$

Algorithm

Idea: If we face this problem using a single machine learning model, the learning **complexity scales exponentially in the number of thresholds**, and cannot be addressed effectively in the presence of scarce data.

We **decompose** the problem of finding both volume thresholds and corresponding optimal prices in **two sub-problems**:

1. Find the **optimal average price** for a given product **without considering volume discounts**.
2. Build an **adaptation scheme** to obtain different prices for different volumes, whose (weighted) average is equal to the optimal price previously obtained.

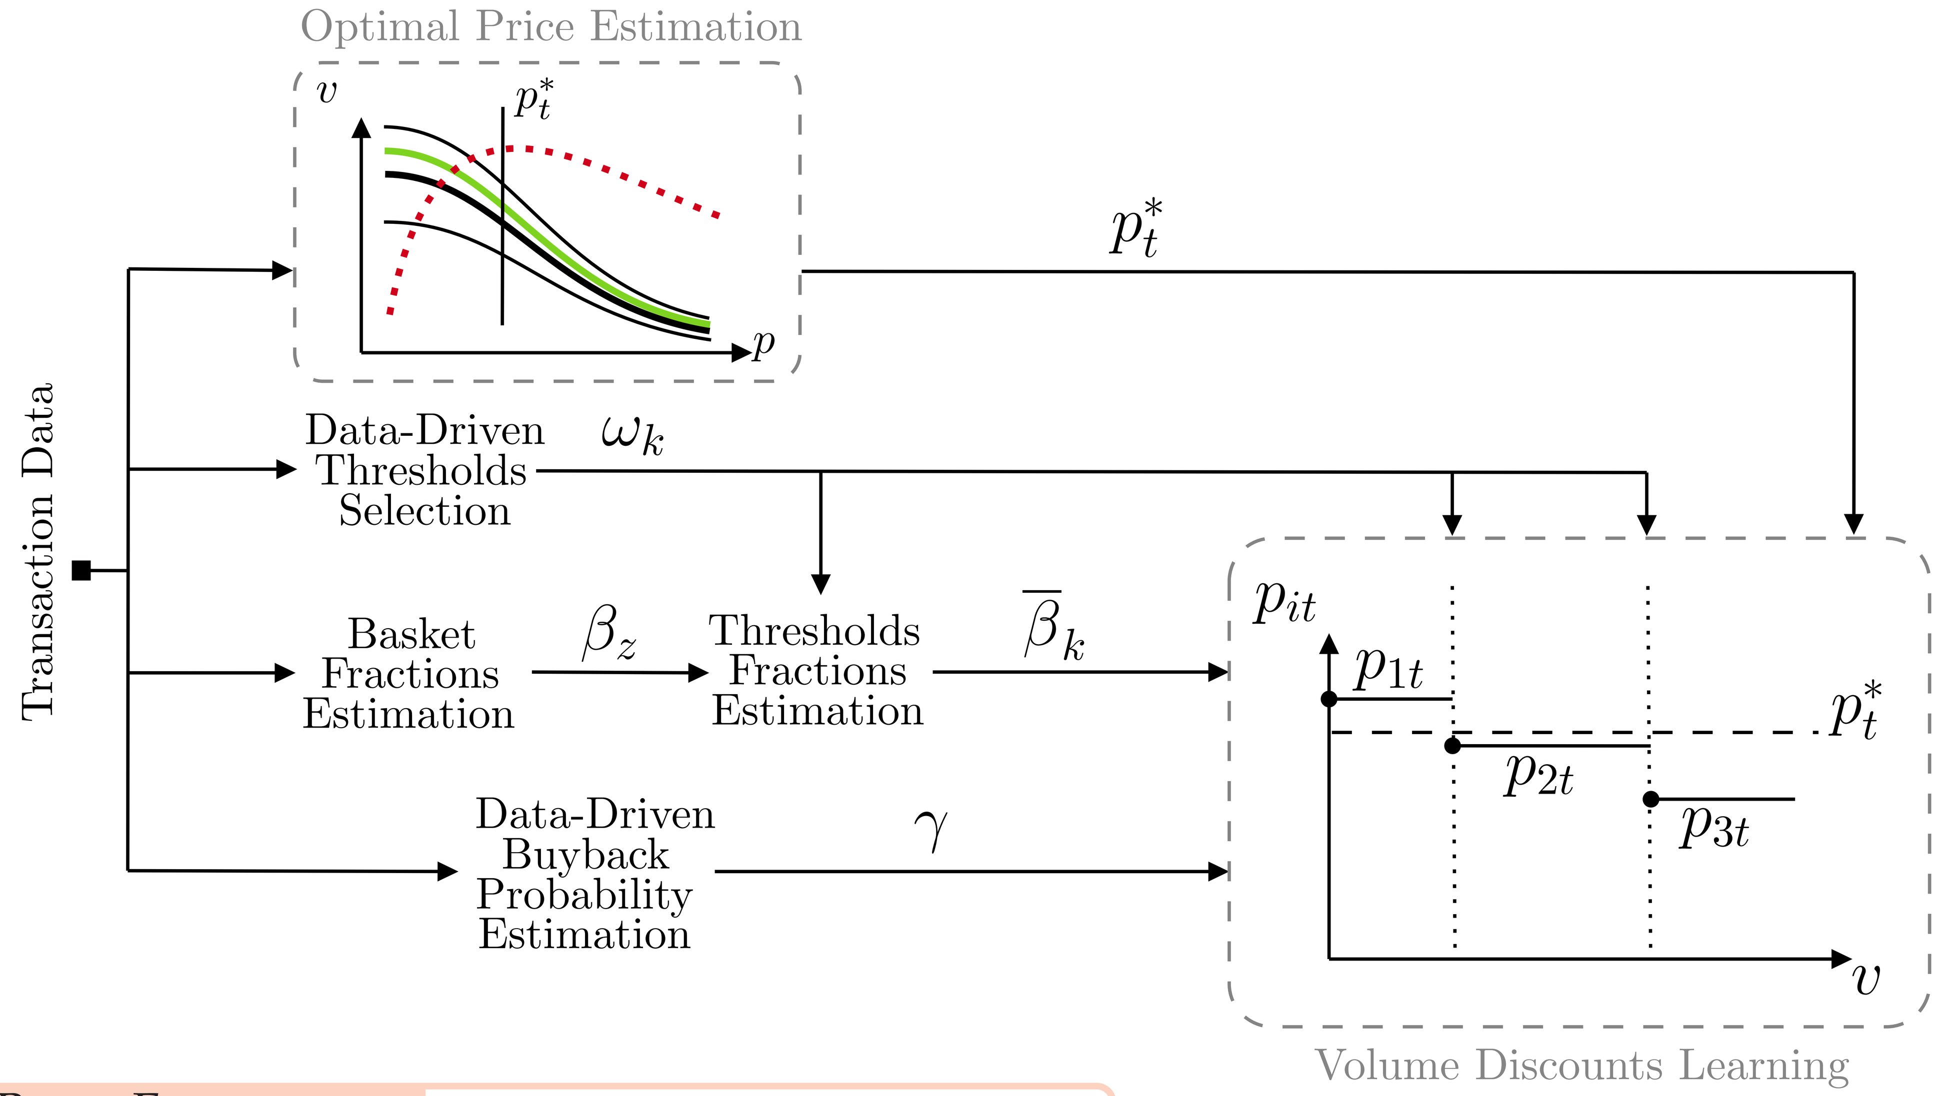

Optimal Price Estimation

- We want to find the optimal price p_t^* for every product independently
- For each t in the past, we compute the total **volumes** $\overline{v}_t$ and the average price $\overline{p}_t$
- Volume curves $\hat{v}(p,t)$ are generalized using **Bayesian Linear Regression** from price (p) and time (t) related features to generalize over demand curve and seasonality
- Price-related features are selected to be monotonic non-increasing, their weight distributions are chosen with a positive support
- At a given time τ, we can fix time-related feature, and sample using a **Thompson Sampling**-like approach a curve $\hat{v}_{TS}(p,\tau)$ binding pricing and volumes
- We estimate the best average price $p_\tau^* \in \arg\max_{p \in \mathcal{P}} (p - c) \cdot \hat{v}_{TS}(p,\tau)$

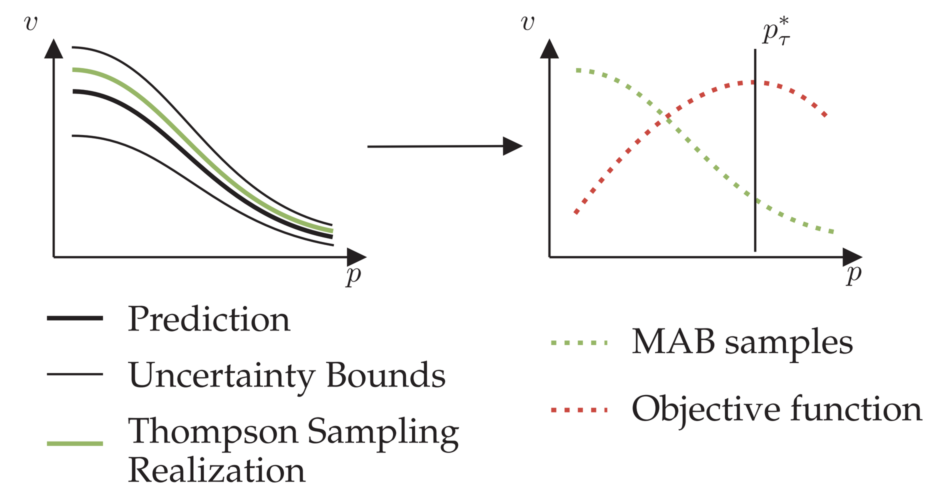

Volume Discounts Learning

- We want to find η **thresholds**, and define the related prices
- The idea of volume discount is to present a volume **discount policy** which is build on top of the optimal price p_τ^* computed so far
- The main assumption is about users' need:
  - the **need** for a given item over time is **fixed** and can be satified in one or more purchase
  - after the first purchase, a customer can **fullfill** its **need** buying from us or thanks to a competitor
- The discount policy makes use of **users' buyback probability** and other auxiliary quantities
- The proposed discount policy has as **weighted average** price p_τ^*
- The discounts are as relevant as the buyback probability decrease

Real-World Application

Real-World Setting

A/B test involving $\approx$ 300 products with $\approx$ 300 kEuros of turnover and $\approx$ 83 kEuros of margin. The test is conducted for 4 months in summer/fall 2021. The e-commerce specialists ask us to condsider $\eta = 3$ thresholds.

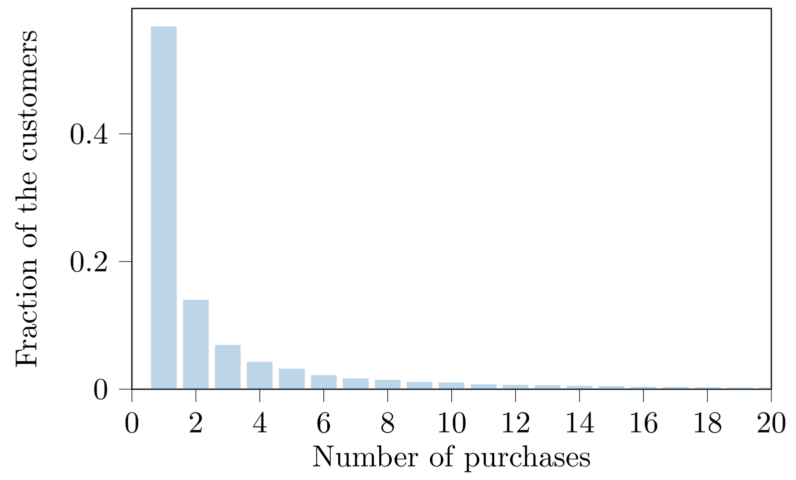

We observe that volume discounts are fundamental, given that most of the customers are not loyal

Results Overview

We improve the performance of set A of 55% w.r.t. the set B

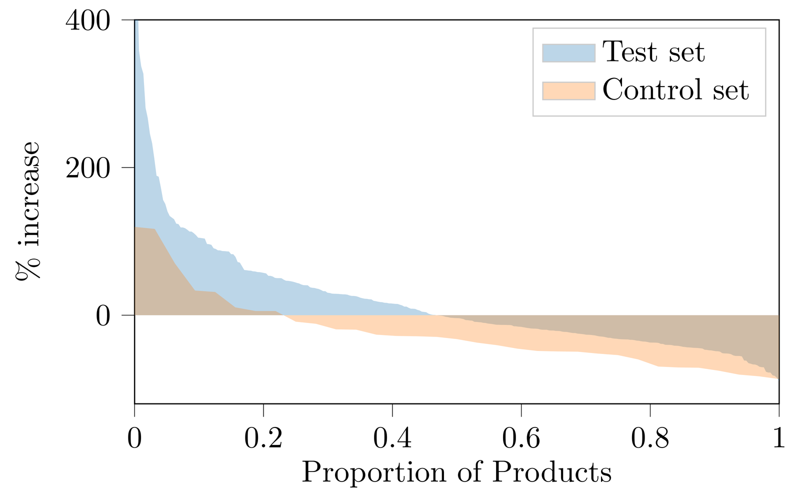

We register in set A the 47% of products which increse their turnover, against the 25% of set B

Volume Discounts Effect

The average discount applied for the algorithm are 10% for the second volume interval and 20% for the third one.

Product	$\Delta\bar{\beta}_1$	$\Delta\bar{\beta}_2$	$\Delta\bar{\beta}_3$	Δunits
1	-32%	+10%	+22%	+63%
2	-26%	+25%	+1%	+43%
3	-15%	+4%	+11%	+11%
4	-5%	+1%	+4%	+14%
Mean	-19.5%	+10%	+9.5%	+33%

Conclusions after the A/B Test

- After the A/B test, the e-commerce website decide to adopt the solution.
- The dynamic pricing algorithm is now pricing over 1200 products for a total turnover of 1.5 MEuros.

References

Arnoud V Den Boer. Dynamic pricing and learning: historical origins, current research, and new directions. *Surveys in operations research and management science*, 2015.

Robert Kleinberg and Tom Leighton. The value of knowing a demand curve: Bounds on regret for online posted-price auctions. In *FOCS*. IEEE, 2003.