Executive Summary of the Thesis

An Online Dynamic Pricing Algorithm for Complementary Products

Laurea Magistrale in Computer Science and Engineering

Author: Andrea Cerasani

Advisor: Prof. Marcello Restelli

Co-advisors: Marco Mussi   Alessandro Lavelli

Academic year: 2022-2023

1. Introduction

Determining the optimal price of a product is a crucial task for retailers, impacting the overall success of a retail business. *Dynamic Pricing* directly addresses this task, determining optimal selling prices of products or services in settings where prices can easily and frequently be adjusted. Product interactions are among the key factors that determine the demand for a product and its optimal selling price. Two main types of relationships can be identified between products: substitutability and complementarity. Substitutable products refer to goods that can be used in place of each other to satisfy a particular need or want while complementary products enhance the value and utility of each other when consumed or utilized simultaneously and are frequently bought together. This type of relation is exploited in pricing to take into account product relations when selecting the optimal selling price.

1.1. Goal and Challenges

Dynamic Pricing algorithms usually face the problem of finding the optimal prices of a product independently from the others and this can lead to suboptimal solutions as we miss the chance to exploit product interactions. The goal of our thesis is to fill this gap by proposing an online learning algorithm for optimizing the pricing strategy of a set of products, considering both positive and negative interactions between them. The main challenges reside in the characteristics of the complementary relationship and in the complexity of the problem. Complementarity exhibits asymmetry and non-transitivity, unlike an equality relation like substitutability. Moreover, when dynamically pricing a product, we need to tackle the exploration-exploitation tradeoff during the learning process. The task is also computationally complex, dealing with a combinatorial explosion when considering the relation a product has with all the others in a catalogue.

1.2. Original Contribution

Most works in the literature on complementary pricing focus on Discrete Choice Models [1], approaches that, compared to Multi-Armed Bandit solutions, do not effectively address the trade-off between exploration and exploitation [2]. Focusing on pricing non-luxury products in retail e-commerce with unlimited inventory, in this thesis we present `CPP` (`Complementary Product Pricing`), an online learning algorithm for optimizing the pricing strategies of a set of products, considering the substitutable and complementary relationships between them. The algorithm makes use of transaction data to learn the interaction between the different items and then optimize the pricing strategies through efficient multi-armed bandit solutions.

2. Problem Formulation

We study the scenario in which we want to find the optimal pricing strategy for a set of products $\mathcal{J}$ ($|\mathcal{J}| = N$).
Our goal is, given a time horizon T, to set for every time $t \in \{1, \dots, T\}$ a vector of percentage margins (from now on, margins) $\mathbf{m}_t = (m_{1t}, \dots, m_{Nt})$ where $m_{jt} \in \mathcal{M}$ is the price we choose for product j at time t, and $\mathcal{M}$ is the (even infinite) set of possible margins. We define the (percentage) margin m_{jt} as:

$$m_{jt} := \frac{p_{jt} - c_j}{c_j},$$

where p_{jt} is the selling price and the acquisition cost for product j at time t, and c_j is its acquisition cost. For a generic product $j \in \mathcal{J}$ we denote as $v_j(\mathbf{m}_t)$ the demand of product j which we assume to depend on the margin vector $\mathbf{m}_t$ of all products. We identify two types of relationships between products: substitutability and complementarity. Given two products $a, b \in \mathcal{J}$, we call a and b substitutable products if the increases in the sales of one product imply a decrease in the sales of the other while we call a and b complementary products if the increases in the sales of one product imply an increment also in the sales of the other. We consider non-perishable products with unlimited availability, with monotonically non-increasing demand function w.r.t. its price. The assumption of unlimited availability virtually holds for e-commerce websites adopting the dropshipping paradigm, while the assumption of the monotonic demand holds for sellers vending products that are different from *Veblen*, *Giffen* or *Luxury* ones. We assume to be in a stationary environment. We consider a scenario in which we assume to have access to transaction data reporting all the sales for every product $j \in \mathcal{J}$, divided by baskets. The only information we have about products is the groups of substitutable products. The detection of substitutable products has been thoroughly discussed in literature and is out of the scope of this work.

The goal of our learning problem is to find the vector of the optimal margin $\mathbf{m}^*$, i.e., the vector maximizing our objective function $f(\mathbf{m})$. Formally:

$$\mathbf{m}^* \in \arg\max_{\mathbf{m} \in \mathcal{M}^N} f(\mathbf{m}), \tag{1}$$

where the objective function $f(\mathbf{m})$ is profit:

$$f(\mathbf{m}) := \sum_{j \in \mathcal{J}} m_j \; c_j \; v_j(\mathbf{m}), \tag{2}$$

over all the products. We call $\boldsymbol{\pi}$ a policy returning at each time t a vector of margins $\mathbf{m}_t$. The goal of our algorithm is to find a policy minimizing the expected cumulative regret:

$$\mathbb{E}\left[R(\boldsymbol{\pi}, T)\right] := T f(\mathbf{m}^*) - \mathbb{E}\left[\sum_{t=1}^{T} f(\mathbf{m}_t)\right]. \tag{3}$$

3. Proposed Solution

Our proposed algorithm to identify and price complementary products in an online manner consists of two parts: we first deal with the aggregation of substitutable products and the discovery of complementary products, we then proceed to estimate the demand by employing an exploration strategy to tackle the exploration-exploitation tradeoff. In Figure 1 we show an outline of the proposed algorithm.

3.1. Complementary Products Discovery

Given the computational complexity of the task, a way to simplify the problem while focusing on the most significant relationships between products is necessary, in order to have an algorithm that scales properly with catalogues containing hundreds of thousands of products. We represent the relationships between products as a directed graph, given the asymmetry of complementarity, where each node is a product and each edge represents a complementarity relationship between 2 products. In Section 2 we modelled the demand of every product to depend on the margins of the entire catalogue of products, so our initial representation of the graph of complementarity relationship (from now on the complementarity graph) is a connected graph.
We propose an approach to prune the edges of this graph while maintaining the most meaningful ones. The process makes use of information

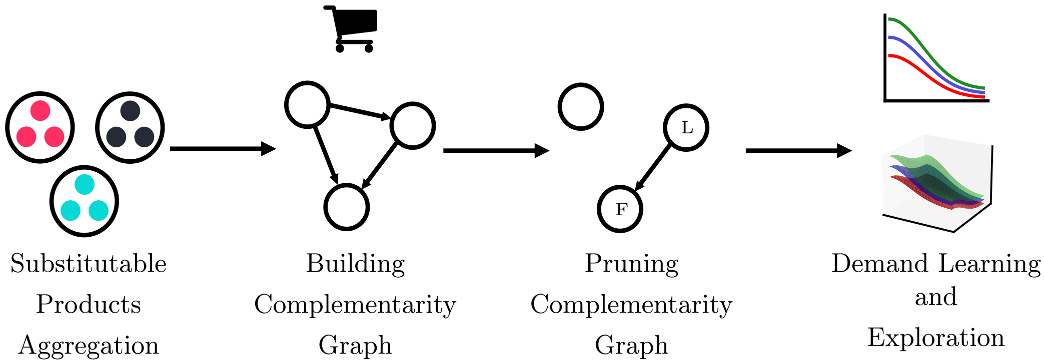

Figure 1: Algorithm outline

about substitutable products and transactional data and is divided into two steps:

1. Clustering substitutable products
2. Mining complementarity relationships.

Clustering substitutable products Two substitutable products compete with each other and are often subject to the phenomenon of cannibalization since they satisfy the same need. Applying dynamic pricing solutions to these products separately would exacerbate cannibalization and could lead to letting them compete against each other, with the effect of reducing profit. Clustering these products together allows them to be priced with the same pricing policy, minimizing cannibalization. Following [3], given the information about substitutable products, which we assume to have, and a set $\mathcal{T}$ of time instants for which historical transaction data is available, we cluster a set of substitutable products $\mathcal{K} \subseteq \mathcal{J}$ and their historical margin $m_{k\tau}$ and volume $v_{k\tau}$ for all products $k \in \mathcal{K}$ and time $\tau \in \mathcal{T}$ in the following way:

$$m_{\mathcal{K}\tau} := \sum_{k \in \mathcal{K}} m_{k\tau} \cdot \frac{v_{k\tau}}{\sum_{h \in \mathcal{K}} v_{h\tau}},$$

$$v_{\mathcal{K}\tau} := \sum_{k \in \mathcal{K}} v_{k\tau}.$$

Given a time horizon T, the margin $m_{\mathcal{K}t}$ chosen at each time $t \in \{1, \dots, T\}$ will be applied to every product $k \in \mathcal{K}$.

We refer to clustered substitutable products simply as products and each clustered product is a node in the complementarity graph. From now on we assume that there are no substitutable products in different nodes of the complementarity graph.

Mining complementarity relationships We proceed to identify meaningful complementary relationships between products. Since complementary products are frequently bought together, they can be identified using co-purchases in transactional data. We propose a way to measure complementarity between products making use of the binomial test. We want to test the independence of every pair of products in order to verify if the co-occurrence of two products in the same basket is higher than *random chance* with a given significance level. Formally, given two products $a, b \in \mathcal{J}$ we denote the total number of baskets as n, the number of baskets containing both a and b as n_{ab} and with $P(a)$ and $P(b)$ the probability of having respectively product a and product b in a basket.

Under the assumption of independence, the probability of having both products a and b in a basket is $P(a)P(b)$ and the hypothesized probability of success π is $H_0 : \pi = P(a)P(b)$.

Since we want to test if the probability of the co-occurrence of a and b is higher than if they were independent, we perform a right-tailed test and our alternative hypothesis is $H_1 : \pi > P(a)P(b)$.

The number of trials in our hypothesis test is the total number of baskets n while the number of successes k is the number of co-occurrences of products a and b: $k = n_{ab}$.

Finally, we want to perform the test with a significance level of 1%, so we reject H_0 if the p-value is smaller than 0.01. For the product pairs for which we reject the null hypothesis H_0, we have statistical evidence that the probability of products a and b appearing in the same basket is higher than if they were independent, and we can consider product a and b to be complementary. We set the direction of the complementarity relationship between two products to be from the product with the bigger number of baskets where it appears to the one with less, on the assumption that the product with more selling volumes has more influence on the product with less than vice versa. We call the product that influences the other product *leader* while the product that is influenced *follower*. By imposing this direction we obtain a Weighted Directed Acyclic Graph (DAG) of complementarity relationships between products, where the p-values of the hypothesis tests are the weights of the edges. The absence of cycles is due to the unidirectional flow of the edges given by the product volumes. In order to break down combinatorial complexity when optimizing, we want to reduce the graph structure to multiple *star structures* composed of one leader and multiple followers, where each node can only be either a leader or a follower. To do this, we first obtain tree structures from the DAG by keeping among the inbound edges of each node only the one with the highest weight, i.e. the most influential leader. Finally, to cut the trees into star structures we propose to consider each node that, starting from the leaves and going up to the roots, is both a leader and a follower to be only a leader. We have obtained a structure that allows us to perform optimization of the margins of the products breaking down the intrinsic combinatorial complexity of the problem.

3.2. Pricing Complementary Products

We aim to find the margins that maximize the total profit (2). At each time t, our algorithm outputs the vector margins $\mathbf{m}_t$ and tackles the exploration-exploitation dilemma in order to minimize the expected regret (3). We differentiate the demand estimation between the one for products whose demand does not depend on other products and the one for products whose demand depends on complementary products.

Univariate demand learning The isolated nodes in the complementarity graph resulting from the above-mentioned steps indicate products for which no complementary relationships have been identified, therefore, we model the demand for these products on the assumption that it depends only on their own margins. The same assumption is made for the demand of the leader products. Given the set of products $\mathcal{J}$, we define the set of isolated products $\mathcal{I} \subseteq \mathcal{J}$ and for the star structures we define the set of leaders $\mathcal{L} \subset \mathcal{J}$ and the set of followers for each leader $i \in \mathcal{L}$ $\mathcal{F}_i \subset \mathcal{J}$, i.e. such that there is a directed edge $(i, j) \quad \forall j \in \mathcal{F}_i, \forall i \in \mathcal{L}$.
We estimate the demand using a Bayesian Linear Regressor (BLR) because this type of conditional model allows the estimation of uncertainty and will allow us to use a MAB approach to balance exploration and exploitation.
Using the BLR we build an estimate $\hat{d}_i(\cdot)$ of the demand function for product i as a linear combination of the basis function taken as input, formally:

$$\hat{d}_i(m_i) = \sum_{h=0}^{Z} \theta_h \ \phi_h(m_i),$$

where θ_h is the h-th weight distribution and $\phi_h(m_i)$ is the h-th basis function of the margin $m_i \in \mathcal{M}$.
In order to improve the robustness of our model we employ a shape-constrained model for which we impose the shape of the function to learn to be monotonically decreasing, following the assumption of monotonicity of demand made in Section 2.
To do so we make use of the set of monotonically decreasing basis functions employed in [3] together with the choice of prior distributions with non-negative support.

Bivariate Demand Learning Unlike the demand for the leader products, we assume the demand for the follower products to depend on their own margin and on the margin of their leader.
Similarly to the univariate demand case, we employ a BLR to produce the estimate $\hat{d}_j(\cdot)$ of the demand function for product $j \in \mathcal{F}_i$, with $i \in \mathcal{L}$. For each follower product $j \in \mathcal{F}_i$ of leader product $i \in \mathcal{L}$, the demand is estimated as:

$$\hat{d}_j(m_j, m_i) = \sum_{h=0}^{Z} \theta_h \ \phi_h(m_j) + \sum_{h=0}^{Z} \theta_{h+Z+1} \ \phi_h(m_i),$$

where θ_h is the h-th weight distribution and $\phi_h(m)$ is the h-th basis function of the margin $m_j \in \mathcal{M}$.

Since complementary products exhibit negative cross-price elasticity, we can assume the additive contribution of leader products on follower volumes to be monotonically decreasing on their own margin and employ the same set of monotonically decreasing basis functions for both margins.

Exploration strategy and joint optimization We can naturally frame our problem as an online learning one, where we want to acquire new information about the function we want to learn while at the same time minimizing the cumulative regret. The use of BLR allows us to measure uncertainty and use this information in a MAB setting to balance exploration and exploitation, where the arms of the MAB are the margins our algorithm chooses at each round. To drive the exploration we employ an approach similar to Thompson Sampling (TS): in each round we draw a sample from each of the posterior distribution of the BLR weights, obtaining an estimation of the demand curve given the margins on which it depends. Given the estimation of the demand curve, we proceed to choose the margins that maximize the estimated profit. We distinguish the case where we price an isolated product and the one where we price related products.
For an isolated product $i \in \mathcal{I}$ we can compute the estimated objective function $\hat{f}_i(m), \forall m \in \mathcal{M}$ as:

$$\hat{f}_i(m) = m \ c_i \ \hat{d}_i(m) \quad \forall i \in \mathcal{I}.$$

We want to maximize $\hat{f}_i$, to do so we choose at each time $t \in \{1, \dots, T\}$ the optimal margin $\hat{m}^*_{it}$ for product i as:

$$\hat{m}^*_{it} = \arg\max_{m \in \mathcal{M}} \hat{f}_i(m).$$

With regard to the products belonging to a star structure, the estimated objective function for a leader product $i \in \mathcal{L}$ is, in the same way as the isolated products:

$$\hat{f}_i(m_i) = m_i \ c_i \ \hat{d}_i(m_i) \quad \forall i \in \mathcal{L},$$

while for the follower products $j \in \mathcal{F}_i$ of leader $i \in \mathcal{L}$ the estimated objective function is:

$$\hat{f}_j(m_j, m_i) = m_j \ c_j \ \hat{d}_j(m_j, m_i).$$

When choosing the optimal margins for the products belonging to a star structure, we jointly optimize the objective functions and obtain the vector margin $\hat{\mathbf{m}}_t$:

$$\hat{\mathbf{m}}_t = \arg\max_{\mathbf{m} \in \mathcal{M}^N} \sum_{i \in \mathcal{L}} \left[\hat{f}_i(\mathbf{m}) + \sum_{j \in \mathcal{F}_i} \hat{f}_j(\mathbf{m}) \right],$$

where N is the number of products belonging to $\mathcal{L} \cup \bigcup_{i \in \mathcal{L}} \mathcal{F}_i$.

4. Experimental Evaluation

Simulation Environment The experimental evaluation of our algorithm is carried out through a synthetic environment with which we simulate the purchasing dynamics of customers and their price sensitivity, emulating market dynamics with reasonable assumptions. The environment is characterized by a fixed number of potential customers and a catalogue of products to buy. Each product is characterized by a *conversion rate*, the probability of purchasing such product given its margin. We denote the conversion rate of product $i \in \mathcal{J}$ given the vector margin $\mathbf{m} \in \mathcal{M}^N$ as $c_i(\mathbf{m})$. Each product is related to others through fixed complementarity relationships and we assign to each follower product one leader. In order to simulate the effect that leaders have on the purchase of followers in the same basket we model the conversion rate of followers to be conditioned on the purchase of their leader and to increase in case their leader is added to the basket. In particular, for a follower $j \in \mathcal{F}_i$ with $\mathcal{F}_i$ the set of followers of leader $i \in \mathcal{L}$, we scale the conversion rate by employing a coefficient l in the following way:

$$c_j(\mathbf{m}) \quad \text{when } i \text{ is not in the same basket}$$

$$l \ c_j(\mathbf{m}) \quad \text{when } i \text{ is in the same basket.}$$

The transactions of the users are generated by sampling a *Bernoulli* distribution for each product in the catalogue and using as parameter p the conversion rate of each product, such that:

$$P(B_i = 1) = c_i(\mathbf{m})$$

$$P(B_i = 0) = 1 - c_i(\mathbf{m}),$$

where B_i is the event of a purchase of product $i \in \mathcal{J}$ by a potential customer. The conversion

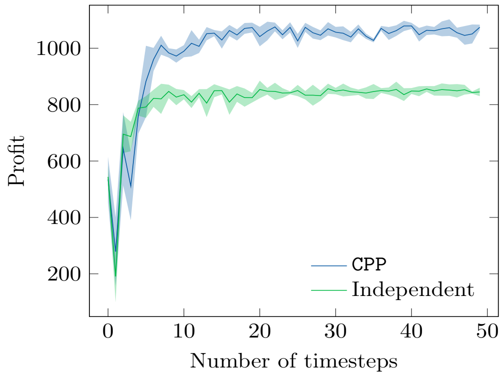

Figure 2: Instantaneous profits.

rates used to generate transactions are univariate in the case of the leaders and bivariate on the margins of the followers of the leaders in the case of the followers. Specifically, the function used for the conversion rates is $f_1(x) = e^{-(2x)^2}/3$ for the leaders and $f_2(x, y) = f_1(x) + 0.7f_1(y)$ for the followers. We limit the margins domain in $[0, 1]$. The conversion rate of the leader i is $c_i(m_i) = f_1(m_i)$ while for the follower j of leader i is $c_j(m_j, m_i)/l$ when there is not leader i in the same basket and $c_j(m_j, m_i)$ when leader i is present. Coefficient l is set to 1.5. We employ as basis functions in the demand estimators the monotonically decreasing transformed Bernstein polynomial mentioned in Section 3 with degree 10. We employ as priors of the BLRs the Lognormal distribution. At each timestep, the agent, using the transactional data observed up to that point, infers the complementarity relationships and fits the BLRs. It then chooses the optimal margins by managing exploration through Thompson Sampling.

Comparison with independently priced products We compare the performance of `CPP` with an algorithm that prices the products independently with the same BLR employed for isolated and leader products. We do so to evaluate the profit increase obtained thanks to the joint optimization of complementary product prices. In Figure 2 we show and compare the instantaneous profits obtained by jointly pricing complementary products with `CPP` and those obtained by independent pricing. We can observe that after 4 timesteps where the performance of the two approaches are comparable, `CPP` reaches a better optimum and unlocks profits up to 30% more w.r.t. independent pricing.

5. Conclusions

In this thesis, we faced the problem of finding the optimal pricing strategy for products presenting substitutable and complementary relationships. We presented the problem under analysis, the related assumptions, and the learning problem, which consists of minimizing the expected regret. Then, we proposed `Complementary Product Pricing` (`CPP`), a novel strategy for learning online in this setting. The algorithm is composed of two main phases. In the former, we provided a strategy for the online identification of complementary relationships. In the latter, we discussed a model for efficiently jointly optimizing the margin of the products. We conducted an extensive experimental campaign to assert the solution's soundness and goodness. The results showed that `CPP` effectively outperforms an independent pricing strategy, obtaining an increase of up to 30% in profits compared to independently priced products in a synthetic environment.
Future developments may consider removing the assumption of the knowledge of substitutable products. Another possible extension is dropping the non-stationary assumption on the environment, exploring the evolution of complementary relationships and demand of products over time. Finally, we considered relations of products purchased in the same basket. An extension to this is to investigate the complementarity in purchases made over time by the same users.

References

[1] Guillermo Gallego and Huseyin Topaloglu. *Revenue Management and Pricing Analytics*. International Series in Operations Research & Management Science. Springer New York, 2019.

[2] Tor Lattimore and Csaba Szepesvári. *Bandit Algorithms*. Cambridge University Press, 1 edition, July 2020.

[3] Marco Mussi, Gianmarco Genalti, Francesco Trovò, Alessandro Nuara, Nicola Gatti, and Marcello Restelli. Pricing the Long Tail by Explainable Product Aggregation and Monotonic Bandits. In *SIGKDD*, pages 3623–3633, 2022.