

Assortment Optimization with Visibility Constraints

Théo Barré* Omar El Housni† Marouane Ibn Brahim†
Andrea Lodi† Danny Segev‡

Abstract

Motivated by applications in e-retail and online advertising, we study the problem of assortment optimization under visibility constraints, that we refer to as [APV](#). We are given a universe of substitutable products and a stream of T customers. The objective is to determine the optimal assortment of products to offer to each customer in order to maximize the total expected revenue, subject to the constraint that each product is required to be shown to a minimum number of customers. The minimum display requirement for each product is given exogenously and we refer to these constraints as *visibility constraints*. We assume that customer choices follow a Multinomial Logit model (MNL).

We provide a characterization of the structure of the optimal assortments and present an efficient polynomial time algorithm for solving [APV](#). To accomplish this, we introduce a novel function called the “expanded revenue” of an assortment and establish its supermodularity. Our algorithm takes advantage of this structural property. Additionally, we demonstrate that [APV](#) can be formulated as a compact linear program. Next, we consider [APV](#) with cardinality constraints, which we prove to be strongly NP-hard and not admitting a Fully Polynomial Time Approximation Scheme (FPTAS), even in the special case where all the products have identical prices. Subsequently, we devise a Polynomial Time Approximation Scheme (PTAS) for [APV](#) under cardinality constraints with identical prices. Our algorithm starts by linearizing the objective function through a carefully crafted guessing procedure, then solves the linearized program, and finally randomly rounds the obtained solution to derive a near optimal solution for [APV](#) with cardinality constraints. We also examine the revenue loss resulting from the enforcement of visibility constraints, comparing it to the unconstrained version of the problem. To offset this loss, we propose a novel strategy to distribute the loss among the products subject to visibility constraints. Each vendor is charged an amount proportional to their product’s contribution to the revenue loss. Finally, we present the results of our numerical experiments providing illustration of the obtained outcomes.

Keywords: Assortment Optimization, Multinomial Logit model, Visibility Constraints, Supermodularity.

*Department of Industrial Engineering and Operations Research, University of California Berkeley, CA, USA.
Email: theo_barre@berkeley.edu.

†School of Operations Research and Information Engineering, Cornell Tech, Cornell University, NY, USA.
Email: {oe46,mi262,a1748}@cornell.edu.

‡Department of Statistics and Operations Research, School of Mathematical Sciences, Tel Aviv University, Tel Aviv 69978, Israel. Email: segevdanny@tauex.tau.ac.il.

Contents

1	Introduction	1
1.1	Main Contributions and Technical Ideas	2
1.2	Related Literature	4
2	Model Formulation	5
3	Polynomial Time Algorithm for APV	6
3.1	Expanded Revenue and Expanded Set	6
3.2	Properties of the Expanded Revenue	7
3.3	Optimal Solution for APV	10
3.4	Linear Program for APV	12
4	APV with Cardinality Constraints	14
4.1	Discretizing the universe of products	15
4.2	Linearization of the objective function and algorithm	17
4.2.1	Guessing procedure	18
4.2.2	Integer linear programming formulation	19
4.2.3	The dependent rounding scheme	20
4.2.4	Construction of the bipartite graph and feasibility of the rounded solution	21
4.3	Near-optimality of the rounded solution	22
5	Price of Visibility	25
5.1	The loss due to Visibility constraints	26
5.2	Sharing the loss	26
5.3	Numerical Study	29
6	Conclusions and future directions	32
	Bibliography	32
A	Revenue Maximization Under MNL	35
B	Additional Proofs from Section 4	36

1 Introduction

Assortment optimization is a crucial aspect of decision making in many industries such as e-retail and online advertising. In this domain, our goal is to select a subset of available products to offer to customers in order to maximize a context-appropriate objective function, such as revenue, profit, or market share. For example, e-retailers seek to strategically select which products should be displayed to customers in order to maximize their expected revenue. Online advertisers strategically select the most effective combination of advertisements to maximize user engagement and desired outcomes, such as click-through rates. Consequently, the choice of a well formed assortment is crucial due to inherent substitution effects, where a product’s attractiveness depend not only on its intrinsic value but also on the concurrent alternatives presented at that time. For instance, offering a high-quality, high-priced product alongside a comparable product at a significantly lower price may result in diminishing sales for the higher-priced product, leading in turn to an unsatisfactory platform revenue. This highlights the importance of carefully selecting assortments.

Traditionally, assortment optimization frameworks often overlook a crucial element in contemporary e-commerce: product visibility. In today’s complex business landscape, where companies adhere to Service-Level Agreements (SLAs) with suppliers and prioritize sponsored product promotion, product visibility within an assortment is pivotal. SLAs often define conditions for product representation, ensuring equitable visibility for each supplier’s products on the platform. Moreover, the concept of sponsored products has gained traction, with brands willing to pay for prominent display and increased visibility. While these strategies influence consumer behavior, solely focusing on products visibility without considering broader assortment optimization can lead to an imbalanced product mix, resulting in reduced customer satisfaction and overall revenue.

In this paper, we introduce the notion of *visibility constraints* in the context of Assortment Optimization. The purpose is to enforce a minimum display of each product, i.e., each product has to be shown at least a certain number of times in the displayed assortments. This constraint models both Service-Level Agreements and sponsored products. It can also capture the settings where the platform would like to ensure some fairness notion among vendors by ensuring that each product is given a “fair” chance, i.e., it is shown at least to a certain number of customers. Specifically, we are given a universe of substitutable products and a stream of T customers. For each customer, we have to offer an assortment from the universe of products. The customer decides to purchase one of these products, or to leave without purchasing any product (no-purchase option). We assume that the choice of the customer is governed by a Multinomial Logit (MNL) choice model. We enforce the constraint that each product in the universe has to be shown a minimum number of times among the T assortments offered. The minimum display requirement for each product is given exogenously. Our objective is to maximize the total expected revenue from the T customers. We refer to this combinatorial optimization problem as *Assortment optimization Problem with Visibility constraints*, concisely denoted as [APV](#).

A first natural question concerns the complexity of [APV](#). In fact, without visibility constraints, the problem reduces to the classic unconstrained revenue maximization problem under MNL, for which we know that the optimal assortment is revenue-ordered ([Talluri and Van Ryzin](#)

(2004)) and therefore can be solved in polynomial time. However, by enforcing the visibility constraints, we might have to include certain products in the assortment that cannibalize the sales of other more profitable products because of the substitution effect. Consequently, determining an optimal sequence of assortments for APV is not immediately clear. Subsequently, a natural follow-up problem concerns the introduction of cardinality constraints on the offered assortments of products. These form of constraints are widely used in the literature (Davis et al., 2013; Gallego and Topaloglu, 2014; Désir et al., 2020), as they model multiple applications where vendors have limited shelf space, or limited screen size for online vendors. Finally, an interesting question emerging from the visibility problem is the task of quantifying the revenue loss incurred by enforcing visibility constraints compared to the relaxed unconstrained problem. This challenge compels us to develop a pricing strategy that appropriately apportions the loss to different vendors based on the impact of their product on the overall revenue. Such a scenario frequently occurs within the framework of SLAs. Typically, a contract between a platform and a vendor includes a clause that guarantees a certain level of visibility to the vendor’s product. In return for this visibility, the vendor compensates the platform with a fee.

1.1 Main Contributions and Technical Ideas

In this paper, we introduce and study the assortment optimization problem with visibility constraints under the Multinomial Logit choice model. Our first goal is to settle the complexity question on the positive side by developing a polynomial time algorithm for APV. Next, we consider a natural extension of this problem, where a cardinality constraint on the offered assortment is enforced. We show that this additional constraints makes the problem fundamentally harder, rendering it strongly NP-hard, even in the special case where all the products have identical prices. Additionally, we leverage the NP-Hardness to prove that the existence of a Fully Polynomial Time Approximation Scheme (FPTAS) is precluded. Following this assessment on the hardness of the problem, we design a Polynomial Time Approximation Scheme (PTAS) for the special case of equal prices. A subsequent goal is to quantify the revenue loss caused by the visibility constraints compared to the unconstrained assortment optimization problem and design a strategy to share this loss among the different vendors of the products for which visibility constraints have been enforced. Our contributions are organized and summarized as follows.

1. **Polynomial time algorithm for APV.** Our main technical contribution is to design a polynomial time algorithm for APV. We introduce the notion of expanded revenue and expanded set of an assortment, which will play a pivotal role in designing our algorithm. We leverage structural properties of the expanded revenue function to characterize the structure of an optimal solution of APV, and consequently design an efficient algorithm to compute it.
 - (a) **Expanded Revenue.** In Section 3.1, we introduce the expanded revenue function. Given a universe of products \mathcal{N} and an assortment $A \subseteq \mathcal{N}$. The expanded revenue of assortment A is defined as the maximum revenue of any assortment in \mathcal{N} that

contains A . The expanded set is the assortment that achieves this maximum revenue. We show that the expanded revenue function is closely related to the objective function of **APV** in the case of a single customer. We provide a linear time algorithm to compute the expanded revenue.

- (b) **Monotonicity and supermodularity¹ of the expanded revenue.** We show that the expanded revenue, as defined above, possesses some useful properties. Namely, we prove in Lemma 3.4 a monotonicity property, i.e., we show that the expanded revenue of an assortment decreases as the assortment gets large. Then, in Lemma 3.5, we prove the main theoretical result on which our final algorithm relies: the supermodularity of the expanded revenue function.
- (c) **Our algorithm and LP formulation.** Building on the previous properties of the expanded revenue, we finally identify in Theorem 3.7 a very simple nested structure for an optimal solution of **APV**, and devise a polynomial time algorithm to efficiently solve the problem. Additionally, we demonstrate in Theorem 3.9 that **APV** can be formulated as a compact linear program.

2. **APV with cardinality constraints.** We consider in Section 4 the natural extension of **APV** where there is an upper bound on the number of products that we can display in each assortment.

- (a) **Hardness.** In Theorem 4.1, we prove that **APV** with cardinality constraints is strongly NP-hard, even in the case of equal prices, by linking its resolution to the 3-PARTITION decision problem. Moreover, we extend our proof to show that the existence of an FPTAS is precluded for this problem, unless $P = NP$.
- (b) **Polynomial-time approximation scheme (PTAS).** Our cornerstone algorithmic result in the case of equal prices is the design of a polynomial-time approximation scheme (PTAS), which for any fixed desired precision, returns a sequence of assortments within that degree of precision, in polynomial time. Since the strong NP-hardness rules out the existence of a Fully Polynomial Time Approximation Scheme (FPTAS), unless $P = NP$, a PTAS is the best approximation algorithm we can achieve. Our PTAS relies on a linearization of the objective function through the guessing of a carefully chosen set of parameters of the problem. Then, we solve the relaxed linear program, and we leverage its solution to compute a random approximate (integer) solution for the linearized problem, using a dependent rounding scheme. This rounded solution is subsequently shown to be, in expectation, near optimal for **APV** with cardinality constraints.

3. **Price of visibility.** The introduction of visibility constraints results in a reduction in the total expected revenue compared to the unconstrained version of the assortment problem. Therefore, we aim to evaluate this revenue loss and propose a fair strategy for distributing it among vendors based on their respective contributions to the loss.

¹A function $f : \Omega \rightarrow \mathbb{R}$ is supermodular if $\forall A, B \in \Omega, f(A \cup B) + f(A \cap B) \geq f(A) + f(B)$.

- (a) **Pricing the loss.** We devise in Section 5.2, a pricing strategy as follows: For each product that negatively impacts the overall revenue, we charge the vendor a fraction of the loss proportional to the ratio between the negative contribution of the product and the sum of the negative contributions of all products. We demonstrate that this strategy satisfies natural fairness properties and exhibits favorable computational tractability.
- (b) **Numerical experiments.** Additionally, we conduct in Section 5.3 some numerical experiments to illustrate our findings. We analyze the influence of visibility constraints on expected revenue and sales, examine the individual effect of a single product’s visibility constraint, and explore the trade-off between revenue and fees. Moreover, we show how our pricing strategy accurately captures the revenue loss attributed to the products made visible.

1.2 Related Literature

Assortment optimization under the Multinomial Logit (MNL) model is a well-established problem in scientific literature. Initially introduced by Luce (1959), with subsequent works by McFadden (1973) and Hausman and McFadden (1984), the MNL choice model has gained popularity for modeling customer choices due to its simplicity in computing the choice probabilities, its predictive power and its computational tractability compared to more complex choice models. It has been extensively used in various research works such as Mahajan and Van Ryzin (2001); Talluri and Van Ryzin (2004); El Housni and Topaloglu (2022); Sumida et al. (2021); Gao et al. (2021); El Housni et al. (2023), to mention a few. The MNL model proves particularly useful in assortment optimization, as demonstrated by Talluri and Van Ryzin (2004), who showed that under the MNL model, the optimal assortment in the unconstrained setting is revenue-ordered. This means that it contains all products whose revenues exceed a certain threshold, simplifying the optimization problem by avoiding the consideration of exponentially numerous potential subsets. Moreover, Gallego et al. (2015) give a linear programming formulation for the unconstrained assortment problem under MNL. Rusmevichientong et al. (2010) solved the version of the problem with a cardinality constraint, proving it is still solvable in polynomial time, and Désir et al. (2022) studied more general capacity constraints, showing it is NP-hard to solve in the general case. Sumida et al. (2021) and Davis et al. (2013) studied totally unimodular constraint structures for the assortment and showed that the resulting problem can be reformulated as a linear program. However, when considering mixtures of MNL models (MMNL), the assortment optimization problem becomes NP-hard even in the unconstrained setting with two classes of customers as shown in Rusmevichientong et al. (2014).

To the best of our knowledge, this paper is the first to study assortment optimization under MNL with visibility constraints. The topic of visibility in assortment planning has barely been covered: Chen et al. (2022) studied visibility under a fairness approach, trying to enforce similar visibility for products with similar characteristics, while Wang et al. (2021) studied a version of the assortment optimization problem in which they can increase the attractiveness of some products through an advertising budget. Recently, in a concurrent work, Lu et al. (2023) considers an assortment optimization problem under the MNL model, subject to some

fairness constraints similar to the visibility constraints presented in our paper. In contrast to our model, which is deterministic and involves a predetermined number of customers, their study investigates a probabilistic version with a single customer, where random assortments can be offered. This probabilistic version is less complex than our framework. For instance, we establish that our model with cardinality constraints is strongly NP-hard and does not admit an FPTAS, whereas their probabilistic model does admit an FPTAS under cardinality constraints. Hence, the need for different optimization techniques and approaches.

In addition, the topic of assortment optimization for a stream of customers is more often studied from an online perspective, where decisions are made sequentially such as in [Davis et al. \(2015\)](#) and [Cheung and Simchi-Levi \(2017\)](#). In contrast, we study a static version of the problem, where we plan the entirety of our assortments in advance. Other versions of the problem, such as in [Li \(2009\)](#), consider a flow of customers with randomized preferences, to which we offer a common assortment. Finally, in revenue management, pricing problems are often considered in the sense of optimizing the selling price of each product ([Wang, 2012](#); [Miao and Chao, 2021](#); [Alptekinoglu and Semple, 2016](#)), while selling prices are fixed for our problem, and we instead study the question of pricing the loss generated by enforcing visibility of each product.

Outline. The remainder of the paper is organized as follows. We first introduce the mathematical framework and define our problem [APV](#) in [Section 2](#). Then, in [Section 3](#), we devise a polynomial time algorithm for the [APV](#) problem. In [Section 4](#), we consider [APV](#) with cardinality constraints. We show that adding these constraints makes the problem strongly NP-hard even in a setting with identical prices, then we provide a PTAS for the latter setting. In [Section 5](#), we propose a pricing strategy to charge the revenue loss to the vendors proportionally to their contribution, and illustrate our results numerically. Finally, in [Section 6](#), we draw some conclusions and outline future work.

2 Model Formulation

The MNL choice model. Let $\mathcal{N} := \{1, \dots, n\}$ be a universe of substitutable products at our disposal. Each product $i \in \mathcal{N}$ has a price $p_i \geq 0$. Without loss of generality, we order the products by non-increasing prices, i.e., $p_1 \geq p_2 \geq \dots \geq p_n$. An assortment of products or an offer set, is simply a subset of products $S \subseteq \mathcal{N}$. Additionally, the option of not selecting any product is symbolically represented as product 0, and referred to it as the no-purchase option.

We assume that customers make choices according to a Multinomial Logit model. Under this model, each product $i \in \mathcal{N}$ is associated with a preference weight $v_i > 0$. Note that v_i captures the attractiveness of product i , meaning a high preference weight indicates a high popularity. Without loss of generality, we use the standard convention that the no-purchase preference weight is normalized to $v_0 = 1$. We use the notation $V(S) := \sum_{i \in S} v_i$, which is the total weight of a subset $S \subseteq \mathcal{N}$. Under the MNL model, if we offer an assortment $S \subseteq \mathcal{N}$, the customer chooses product i with probability

$$\phi(i, S) := \frac{v_i}{1 + V(S)}.$$

We refer to $\phi(i, S)$ as the choice probability of product i given assortment S . Alternatively, the customer may decide to not purchase any product, which happens with the complementary probability

$$\phi(0, S) := \frac{1}{1 + V(S)}.$$

Let $R(S)$ be the expected revenue we get from a customer if we offer assortment S . In particular, we have

$$R(S) := \sum_{i \in S} p_i \phi(i, S) = \frac{\sum_{i \in S} p_i v_i}{1 + \sum_{i \in S} v_i}.$$

Assortment Optimization with Visibility constraints. We are presented with a finite stream of T customers. Each customer t will be offered an assortment S_t . Customers make choices according to the same MNL model, i.e., a customer decides to purchase product i from assortment S_t with a probability $\phi(i, S_t)$, or they may choose the no-purchase option with probability $\phi(0, S_t)$. The expected revenue we obtain from customer t is $R(S_t)$. To ensure visibility, we impose constraints that require each product $i \in \mathcal{N}$ to be shown to at least ℓ_i customers. Note that the parameters ℓ_i are exogenous and satisfy $\ell_i \in \{0, \dots, T\}$ for all $i \in \mathcal{N}$. Our objective is to determine the assortment S_t to offer to each customer t in order to maximize the total expected revenue while satisfying the visibility constraints. We refer to this problem as the *Assortment optimization Problem with Visibility constraints* (APV). It can be formulated as follows:

$$\begin{aligned} \max_{S_1, \dots, S_T \subseteq \mathcal{N}} \quad & \sum_{t=1}^T R(S_t) \\ \text{s.t.} \quad & \sum_{t=1}^T \mathbb{1}(i \in S_t) \geq \ell_i, \quad \forall i \in \mathcal{N}. \end{aligned} \tag{APV}$$

3 Polynomial Time Algorithm for APV

The primary contribution of this paper is the development of a polynomial time algorithm for APV. To achieve this, we introduce in Section 3.1 the concepts of the ‘‘Expanded Revenue’’ and ‘‘Expanded Set’’ of an assortment. In Section 3.2, we present a polynomial time algorithm to compute the expanded set and expanded revenue and demonstrate the monotonicity and supermodularity of the expanded revenue function. Leveraging these properties, we characterize the structure of an optimal solution for APV and present an algorithm that computes it in $O(n+T)$ time (Section 3.3). Finally, in Section 3.4, we demonstrate that APV can be formulated as a compact linear program.

3.1 Expanded Revenue and Expanded Set

We begin our analysis by examining APV in the context of a single customer. In this particular scenario, the visibility constraints are given such that either $\ell_i = 0$ or $\ell_i = 1$. Let A denote the subset of all products where $\ell_i = 1$. Consequently, APV is transformed into the problem of identifying the assortment that maximizes revenue while including A . This particular problem

will serve as the building block for our analysis, as it lays the foundation for understanding the general case involving T customers. Thus, it leads us to introduce the subsequent definitions that will aid us in our analysis.

DEFINITION 3.1 (Expanded revenue). *Let $A \subseteq \mathcal{N}$. The expanded revenue of A , denoted as $\bar{R}(A)$, is defined as the maximum expected revenue achieved by any assortment in \mathcal{N} that contains A . In particular, it is given by*

$$\bar{R}(A) := \max_{S \subseteq \mathcal{N}, A \subseteq S} R(S). \quad (1)$$

The optimal solution of the maximization problem in (1) is referred to as the expanded set of A . In case multiple optimal solutions exist, we break ties by selecting the optimal assortment with the largest cardinality, which we show is unique in Lemma 3.3, hence proving that the expanded set is well defined. Formally, we have the following definition.

DEFINITION 3.2 (Expanded set). *The expanded set of A , denoted as \bar{A} , is defined as the assortment within \mathcal{N} that maximizes the expected revenue among all assortments containing A . If multiple assortments achieve the same maximum expected revenue, \bar{A} is selected as the assortment with the largest cardinality. Mathematically, \bar{A} is given by*

$$\bar{A} := \arg \max_{S \subseteq \mathcal{N}, A \subseteq S} \{|S| : R(S) = \bar{R}(A)\}.$$

Problem (1) can be viewed as equivalent to APV when considering a single customer scenario ($T = 1$) and defining A as the set of products that need to be shown once, i.e., $A = \{i \in \mathcal{N} : \ell_i = 1\}$. Thus, \bar{A} represents the optimal assortment, with the largest cardinality, for the problem.

In our analysis, we consider \bar{R} as a set function that takes an assortment $A \subseteq \mathcal{N}$ as input and returns $\bar{R}(A)$. Note that $\bar{R}(A) = R(\bar{A})$. In the subsequent section, we delve into examining various properties of this function, as well as properties associated with the expanded set.

3.2 Properties of the Expanded Revenue

In this section, we first show that we can compute the expanded revenue and the expanded set of a given assortment in polynomial time (Lemma 3.3). Then, we show that the expanded revenue is a non-increasing function and the expanded set is a non-decreasing function (Lemma 3.4). Finally, we show that the expanded revenue function is supermodular (Lemma 3.5) which is the most fundamental property for our algorithm design later in the paper.

Computing the expanded revenue and expanded set. Recall without loss of generality that $p_1 \geq \dots \geq p_n$. We define an assortment S to be price-ordered if $S = \{1, \dots, k\}$ for some $1 \leq k \leq n$. Essentially, a price-ordered assortment prioritizes products with high prices. It is worth noting there are only n possible price-ordered assortments. Consider an assortment $A \subseteq \mathcal{N}$, and its expanded set \bar{A} . In the following lemma, we demonstrate that \bar{A} is the union of A and a price-ordered assortment. Since there are only n possible price-ordered assortments, it is sufficient to compute the expected revenue of the assortments $A \cup \{1, \dots, k\}$ for each $k \in \{1, \dots, n\}$. The expanded set corresponds to the assortment with the highest expected

revenue. In the case of multiple assortments with the same maximum revenue, we break ties by selecting the one with the largest cardinality. Thus, the expanded set \bar{A} can be computed in linear time, specifically $O(n)$. The expanded revenue is simply $\bar{R}(A) = R(\bar{A})$. The proof of Lemma 3.3 leverages some structural properties of the revenue function under MNL that are presented in Appendix A.

LEMMA 3.3. *For any $A \subseteq \mathcal{N}$, the expanded set of A is given by $\bar{A} = A \cup \{i \in \mathcal{N} : p_i \geq \bar{R}(A)\}$. Furthermore, $\bar{R}(A)$ and \bar{A} can be computed in time $O(n)$.*

Proof. By definition of the expanded set, we have $A \subseteq \bar{A}$. Hence, there exists an assortment $B \subseteq \mathcal{N} \setminus A$, such that $\bar{A} = A \cup B$. Let us show that

$$B = \{i \in \mathcal{N} \setminus A : p_i \geq \bar{R}(A)\}.$$

- *Direct inclusion:* Let $i \in B$, and assume by contradiction that $p_i < \bar{R}(A) = R(A \cup B)$. It is known that, under the MNL model, when we add a product j to an assortment S , the revenue of this assortment increases if and only if $p_j \geq R(S)$. For completeness, we provide the statement and the proof of this result in Lemma A.2 in Appendix A. Using this lemma implies that removing i from B would strictly increase the expected revenue $R(A)$, which contradicts the optimality of $A \cup B$.
- *Indirect inclusion:* Let $i \in \mathcal{N} \setminus A$ such that $p_i \geq \bar{R}(A)$, and assume by contradiction that $i \notin B$. By Lemma A.2, adding i to B would increase the revenue. If this increase is strict, it contradicts the optimality of $A \cup B$. If the revenue stays the same, it contradicts the definition of $\bar{A} = A \cup B$ as the optimal solution with maximum cardinality.

Thus,

$$\bar{A} = A \cup \{i \in \mathcal{N} \setminus A, p_i \geq \bar{R}(A)\}.$$

Finally, \bar{A} can be computed in time $O(n)$. Indeed, we start from A , then we sequentially add elements by decreasing price. At each iteration, we can compute the new revenue from the previous one in constant time by storing the current numerator and denominator, since we only need to add $p_i v_i$ to the former and v_i to the latter when we reach element i . Finally, we pick the highest revenue set among the n computed sets. ■

LEMMA 3.4 (**Monotonicity**). *If $A \subseteq B \subseteq \mathcal{N}$, then $\bar{A} \subseteq \bar{B}$ and $\bar{R}(A) \geq \bar{R}(B)$.*

Proof. For $A \subseteq B \subseteq \mathcal{N}$, we have $\{S \subseteq \mathcal{N} : B \subseteq S\} \subseteq \{S \subseteq \mathcal{N} : A \subseteq S\}$. Therefore, every feasible solution for $\max_{S \subseteq \mathcal{N}, B \subseteq S} R(S)$ is a feasible solution for $\max_{S \subseteq \mathcal{N}, A \subseteq S} R(S)$. Hence, $\bar{R}(A) \geq \bar{R}(B)$. It follows that $\{i \in \mathcal{N}, p_i \geq \bar{R}(A)\} \subseteq \{i \in \mathcal{N}, p_i \geq \bar{R}(B)\}$, and therefore $\bar{A} = A \cup \{i \in \mathcal{N}, p_i \geq \bar{R}(A)\} \subseteq B \cup \{i \in \mathcal{N}, p_i \geq \bar{R}(B)\} = \bar{B}$. ■

LEMMA 3.5 (**Supermodularity**). *The expanded revenue function \bar{R} is supermodular, i.e.,*

$$\forall A, B \subseteq \mathcal{N}, \bar{R}(A \cup B) + \bar{R}(A \cap B) \geq \bar{R}(A) + \bar{R}(B).$$

Proof for Lemma 3.5. In this proof, we use the following alternative definition of supermodularity. A function $f: \Omega \rightarrow \mathbb{R}$ is supermodular if and only if for all $A, B \subseteq \Omega$ such that $A \subseteq B$, and each $i \in \Omega \setminus B$, $f(B \cup \{i\}) - f(B) \geq f(A \cup \{i\}) - f(A)$.

Following this definition, let $A, B \subseteq \mathcal{N}$ such that $A \subseteq B$, and let $i \in \mathcal{N} \setminus B$. The proof of this result is separated into two steps. In the first step, we show that

$$R(\overline{B}) - R(\overline{B \cup A \cup \{i\}}) \leq R(\overline{A}) - R(\overline{A \cup \{i\}}). \quad (2)$$

Subsequently, we show in the second step that

$$R(\overline{B \cup A \cup \{i\}}) - R(\overline{B \cup \{i\}}) \leq 0. \quad (3)$$

The result follows directly by summing the inequalities (2) and (3) term by term.

Step 1. Let us start with the following claim which follows from simple algebra.

CLAIM 3.6. *For any assortment $S_1, S_2 \subseteq \mathcal{N}$ such that $S_1 \subseteq S_2$, we have*

$$R(S_1) - R(S_2) = \frac{1}{1 + V(S_2)} \cdot \sum_{j \in S_2 \setminus S_1} (R(S_1) - p_j) \cdot v_j.$$

Using this claim, we have

$$R(\overline{B}) - R(\overline{B \cup A \cup \{i\}}) = \frac{1}{1 + V(\overline{B \cup A \cup \{i\}})} \cdot \sum_{j \in \overline{A \cup \{i\}} \setminus \overline{B}} (R(\overline{B}) - p_j) v_j.$$

Next, we know by definition of the extended set of B that $R(\overline{B}) \geq p_j$ for all $j \notin \overline{B}$. Therefore, since $V(\overline{B \cup A \cup \{i\}})$ is trivially greater than or equal to $V(\overline{A \cup \{i\}})$, we have,

$$\begin{aligned} R(\overline{B}) - R(\overline{B \cup A \cup \{i\}}) &\leq \frac{1}{1 + V(\overline{A \cup \{i\}})} \cdot \sum_{j \in \overline{A \cup \{i\}} \setminus \overline{B}} (R(\overline{B}) - p_j) v_j \\ &\leq \frac{1}{1 + V(\overline{A \cup \{i\}})} \cdot \sum_{j \in \overline{A \cup \{i\}} \setminus \overline{B}} (R(\overline{A}) - p_j) v_j, \end{aligned} \quad (4)$$

where the second inequality follows from Lemma 3.4. Finally, we have

$$\overline{A \cup \{i\}} \setminus \overline{B} \subseteq \overline{A \cup \{i\}} \setminus \overline{A}.$$

Moreover, for every $j \in \overline{A \cup \{i\}} \setminus \overline{A}$, $j \notin \overline{A}$, and in particular, $p_j \leq R(\overline{A})$. Therefore, by adding the missing terms to the sum in Equation (4), we have

$$\begin{aligned} R(\overline{B}) - R(\overline{B \cup A \cup \{i\}}) &\leq \frac{1}{1 + V(\overline{A \cup \{i\}})} \cdot \sum_{j \in \overline{A \cup \{i\}} \setminus \overline{A}} (R(\overline{A}) - p_j) v_j \\ &= R(\overline{A}) - R(\overline{A \cup \{i\}}), \end{aligned}$$

where the equality follows from Claim 3.6. This concludes the first step.

Step 2. On one hand, we have $\{i\} \subseteq \overline{A \cup \{i\}}$. Therefore, we have in particular $\{i\} \subseteq \overline{B \cup A \cup \{i\}}$. On the other hand, $B \subseteq \overline{B}$ and therefore $B \subseteq \overline{B \cup A \cup \{i\}}$. Hence $B \cup \{i\} \subseteq \overline{B \cup A \cup \{i\}}$. Recalling that $\overline{B \cup \{i\}}$ is by definition the maximum revenue assortment containing $B \cup \{i\}$, we have

$$R(\overline{B \cup \{i\}}) \geq R(\overline{B \cup A \cup \{i\}}),$$

which concludes the second step, and thereby the proof of the lemma. ■

3.3 Optimal Solution for APV

In this section, we present the main technical result in this paper. In particular, we characterize the structure of an optimal solution of APV. Our characterization relies on the supermodularity property of the expanded revenue function. Moreover, we show that we can compute such a solution in $O(n + T)$, which gives us a polynomial time algorithm to solve APV.

Optimal solution. Consider an instance of APV. Recall that for all $i \in \mathcal{N}$, ℓ_i is the lower bound on the minimum number of customers for which we must offer product i . For $t \in \{0, 1, \dots, T\}$, we define the following sets

$$L_t = \{i \in \mathcal{N}, \ell_i = t\}. \quad (5)$$

Our candidate solution for APV is given by

$$S_t^* = \overline{\bigcup_{t \leq u \leq T} L_u}, \quad \forall t \in \{1, \dots, T\}. \quad (6)$$

Note that $(L_t)_{0 \leq t \leq T}$ is a partition of \mathcal{N} . Moreover, since $\bigcup_{t+1 \leq u \leq T} L_u \subseteq \bigcup_{t \leq u \leq T} L_u$, the monotonicity property in Lemma 3.4 implies that $S_{t+1}^* \subseteq S_t^*$ for any $t = 0, \dots, T-1$. Therefore, our solution has a nested structure, i.e., $S_T^* \subseteq S_{T-1}^* \dots \subseteq S_1^*$. In the following, we prove that the assortments given by (6) are optimal for APV. Moreover, they can be computed in polynomial time. Indeed, Lemma 3.3 shows that each of them can be computed in time $O(n)$, so the entire solution can be computed in time $O(nT)$. We can further improve the running time to $O(n + T)$ as shown below.

THEOREM 3.7. *The sequence of assortments $(S_t^*)_{1 \leq t \leq T}$ defined in (6) is optimal for APV. Moreover, such a solution can be computed in $O(n + T)$ time.*

Proof. We prove the result by induction. First, for $T = 1$, for any $i \in \mathcal{N}$, we either have $\ell_i = 0$ or $\ell_i = 1$. Noting that L_1 is the set of products i such that $\ell_i = 1$, APV reduces to the problem of finding the optimal assortment that contains L_1 , i.e.,

$$\max_{S \subseteq \mathcal{N} \text{ s.t. } L_1 \subseteq S} R(S),$$

whose solution is $\overline{L_1}$ by definition of the expanded set. The result follows for $T = 1$ by noticing that $S_1^* = \overline{L_1}$.

Let us now prove the result for a general number of customers. Let $T \geq 2$, and assume by induction that the result is true for $T - 1$, in other words, given $T - 1$ customers, and for any set of visibility constraints, the optimal solution of **APV** is given by the assortments defined in Equation (6). Let us show that the result holds for T customers. We denote by A the set of all products which must be shown to at least one customer due to the visibility constraints, i.e., $A := \bigcup_{t=1}^T L_t$. We start by providing the following crucial intermediary claim.

CLAIM 3.8. *There exists an optimal solution S_1, \dots, S_T to **APV** such that $S_1 \supseteq A$. In particular, there exists an optimal solution S_1, \dots, S_T to **APV** such that $S_1 = S_1^*$.*

In other words, Claim 3.8 states that there exists an optimal solution S_1, \dots, S_T to **APV** such that S_1 contains all the products that must be shown at least once due to the visibility constraints, i.e., products i such that $\ell_i \geq 1$. Consequently, this allows us to focus only on those feasible solutions of **APV**, which offer S_1^* to customer 1. Maximizing the revenue amongst said solutions thereby guarantees attaining the optimal objective. In the remainder of this proof, we start by showing Claim 3.8, before leveraging it to conclude our induction.

Proof of Claim 3.8. The proof of this result mainly relies on a judicious exploitation of the supermodularity property. Assume by contradiction that there exists no optimal solution of **APV** such that $S_1 \supseteq A$. Let $\hat{S}_1, \dots, \hat{S}_T$ be the optimal solution of **APV** that maximizes $|\hat{S}_1 \cap A|$. In the case of ties, we pick any arbitrary solution that maximizes $|\hat{S}_1 \cap A|$. By the contradiction hypothesis, $A \not\subseteq \hat{S}_1$. In particular, there exists some product $j \in \mathcal{N}$ such that $j \in A$ and $j \notin \hat{S}_1$. Moreover, since $j \in A$, we know that $\ell_j \geq 1$, and therefore, j must be shown at least to 1 customer, which means that there exists some $u \in [T] \setminus \{1\}$ such that $j \in \hat{S}_u$. We define the following new solution to **APV**: $S_1 = \hat{S}_1 \cup \hat{S}_u$, $S_u = \hat{S}_1 \cap \hat{S}_u$, and $S_t = \hat{S}_t$ for all $t \notin \{1, u\}$. First, this newly defined solution is also feasible since any product offered once in either \hat{S}_1 or \hat{S}_u is also shown in S_1 , and each product shown in both \hat{S}_1 and \hat{S}_u is also shown in both S_1 and S_u . Second, S_1, \dots, S_T is also an optimal solution. Indeed, by the supermodularity property, we have

$$R(\hat{S}_1 \cup \hat{S}_u) + R(\hat{S}_1 \cap \hat{S}_u) \geq R(\hat{S}_1) + R(\hat{S}_u).$$

Therefore,

$$\sum_{t=1}^T R(S_t) \geq \sum_{t=1}^T R(\hat{S}_t).$$

Third, we have $|S_1 \cap A| \geq |\hat{S}_1 \cap A| + 1$ since $\hat{S}_1 \cap A \subsetneq S_1 \cap A$, as the latter set contains j but the former does not. This contradicts the definition of $\hat{S}_1, \dots, \hat{S}_T$, as the optimal solution that maximizes $|\hat{S}_1 \cap A|$, and thereby proves by contradiction that there exists a solution S_1, \dots, S_T such that $S_1 \supseteq A$.

Finally we show that we can take $S_1 = S_1^*$ in particular. On one hand, the solution A, S_2, \dots, S_T is also feasible, as it is obtained by removing all the products j such that $\ell_j = 0$ from S_1 , which cannot break any constraints. Therefore, noting that $\bar{A} = S_1^*$, the solution

S_1^*, S_2, \dots, S_T is also feasible, since $A \subseteq \bar{A}$. Finally, we have

$$R(S_1^*) + \sum_{t=2}^T R(S_t) = R(\bar{A}) + \sum_{t=2}^T R(S_t) \geq \sum_{t=1}^T R(S_t),$$

where the inequality follows from the definition of the expanded set of A , and the fact that $A \subseteq S_1$. In conclusion, there exists an optimal solution of [APV](#) such that $S_1 = S_1^*$.

Concluding the proof of the theorem. In Claim [3.8](#), we showed the existence of an optimal solution which offers S_1^* to the first customer. Therefore, restricting the search space to only such solutions still guarantees obtaining an optimal solution. Therefore, [APV](#) is equivalent to the following optimization problem:

$$\begin{aligned} \max_{S_2, \dots, S_T \subseteq \mathcal{N}} \quad & R(S_1^*) + \sum_{t=2}^T R(S_t) \\ \text{s.t.} \quad & 1 + \sum_{t=2}^T \mathbb{1}(i \in S_t) \geq \ell_i, \quad \forall i \in A, \end{aligned}$$

which itself reduces to the following different instance of [APV](#).

$$\begin{aligned} \max_{S_2, \dots, S_T \subseteq \mathcal{N}} \quad & \sum_{t=2}^T R(S_t) \\ \text{s.t.} \quad & \sum_{t=2}^T \mathbb{1}(i \in S_t) \geq \tilde{\ell}_i, \quad \forall i \in \mathcal{N}, \end{aligned} \tag{7}$$

where $\tilde{\ell}_i = \ell_i - \mathbb{1}(i \in A)$, for all $i \in \mathcal{N}$. Noting that this last problem is an instance of [APV](#) with $T - 1$ customers, we can apply the induction hypothesis, which implies that the optimal solution is given by Equations [\(6\)](#). Directly applying the formulas to this new instance implies that S_2^*, \dots, S_T^* is an optimal solution to [\(7\)](#), and hence that S_1^*, \dots, S_T^* is an optimal solution for [APV](#).

Running time. The running time of the algorithm for [APV](#) can be improved from $O(nT)$ to $O(n+T)$. In fact, we proceed by induction. We start by computing S_T^* . Then, when computing S_t^* , we do not need to directly compute $\overline{\bigcup_{t \leq u \leq T} L_u}$. Instead, since $S_t^* \supseteq S_{t+1}^*$, we just need to check the products $\{i \in \mathcal{N} \setminus (S_{t+1}^*) : p_i \geq \bar{R}(S_t^*)\}$, which is achieved in $O(1 + |S_t^*| - |S_{t+1}^*|)$. Therefore, the running time is $O(1 + |S_T^*|) + \sum_{t=1}^{T-1} O(1 + |S_t^*| - |S_{t+1}^*|) = O(|S_1^*| + T) = O(n+T)$. ■

3.4 Linear Program for [APV](#)

Consider the classic unconstrained assortment problem under MNL model for a single customer

$$\max_{S \subseteq \mathcal{N}} R(S). \tag{AP}$$

It is known that [AP](#) can be formulated as the following LP ([Gallego et al. \(2015\)](#)),

$$\max_{S \subseteq \mathcal{N}} R(S) = \max_{(\alpha_i)_{0 \leq i \leq n}} \left\{ \sum_{i=1}^n p_i \alpha_i \quad \text{s.t.} \quad \forall i \in \mathcal{N}, 0 \leq \frac{\alpha_i}{v_i} \leq \alpha_0, \quad \sum_{i=0}^n \alpha_i = 1 \right\}.$$

Motivated by the structure of the above LP and the structure of our optimal solution of APV given in Equation (6), we propose the following linear formulation for APV.

THEOREM 3.9 (LP for APV). APV is equivalent to the following linear program:

$$\begin{aligned} \max_{\alpha} \quad & \sum_{i=1}^n p_i \sum_{t=1}^T \alpha_i^t \\ \text{s.t.} \quad & \sum_{i=0}^n \alpha_i^t = 1, \quad \forall t \in [T], \\ & \alpha_i^t = v_i \alpha_0^t, \quad \forall i \in \mathcal{N}, \quad \forall t \in \{1, \dots, \ell_i\}, \\ & 0 \leq \alpha_i^t \leq v_i \alpha_0^t, \quad \forall i \in \mathcal{N}, \quad \forall t \in \{\ell_i + 1, \dots, T\}. \end{aligned} \tag{LP}$$

Proof. In this proof, we use OPT^{LP} and OPT^{APV} to denote the values of LP and APV respectively. Our objective is to show that these values are equal. We demonstrate, on one hand, that $\text{OPT}^{LP} \geq \text{OPT}^{APV}$, by constructing a feasible solution to LP whose objective is greater than or equal to OPT^{APV} . On the other hand, we demonstrate using the same technique that $\text{OPT}^{LP} \leq \text{OPT}^{APV}$. The combination of these two inequalities directly implies the desired result.

First inequality. Recall that (S_1^*, \dots, S_T^*) is the optimal solution of APV, whose expression is stated in Equations (6). We introduce a solution α for LP, defined as follows, for every $t \in [T]$,

$$\begin{aligned} \alpha_0^t &= \frac{1}{1 + V(S_t^*)}, \\ \alpha_i^t &= \frac{v_i}{1 + V(S_t^*)} \cdot \mathbb{1}(i \in S_t^*) \quad \text{for } i \in \mathcal{N}. \end{aligned}$$

Let us show that α is feasible. The first and third constraints are straightforward. Indeed, for the first constraint, we have for all $t \in [T]$,

$$\sum_{i=0}^n \alpha_i^t = \frac{1}{1 + V(S_t^*)} + \sum_{i \in S_t^*} \frac{v_i}{1 + V(S_t^*)} = 1.$$

The third constraint is also directly verified by construction, as for all $i \in \mathcal{N}$ and $t \in [T]$, we have

$$\alpha_i^t = \frac{v_i}{1 + V(S_t^*)} \cdot \mathbb{1}(i \in S_t^*) \leq \frac{v_i}{1 + V(S_t^*)} = v_i \alpha_0^t.$$

Regarding the second constraint, let $i \in \mathcal{N}$ and $t \in \{1, \dots, \ell_i\}$. We have

$$i \in L_{\ell_i} \subseteq \bigcup_{t \leq u \leq T} L_u \subseteq \overline{\bigcup_{t \leq u \leq T} L_u} = S_t^*.$$

Therefore,

$$\alpha_i^t = \frac{v_i}{1 + V(S_t^*)} = v_i \alpha_0^t,$$

which proves that the second constraint is respected, and hence shows the feasibility of the constructed solution. Finally, its objective is equal to the revenue of the sequence of assortments S_1^*, \dots, S_T^* as demonstrated by the following easy computation:

$$\sum_{i=1}^n p_i \sum_{t=1}^T \alpha_i^t = \sum_{t=1}^T \sum_{i=1}^n p_i \alpha_i^t = \sum_{t=1}^T \sum_{i \in S_t^*} p_i \phi(i, S_t^*) = \sum_{t=1}^T R(S_t^*) = \text{OPT}^{APV}.$$

Since the objective of α is trivially upper bounded by OPT^{LP} , we deduce that $\text{OPT}^{LP} \geq \text{OPT}^{APV}$.

Second inequality. Consider an optimal basic feasible solution for **LP**, denoted by $\alpha := (\alpha_i^t : i \in \mathcal{N} \cup \{0\}, t \in [T])$. **LP** is a linear program with $T \cdot (n + 1)$ variables. Any basic solution has at least $T \cdot (n + 1)$ active constraints. Since **LP** already contains $T + \sum_{i \in \mathcal{N}} \ell_i$ equality constraints, α activates at least $\sum_{i \in \mathcal{N}} (T - \ell_i)$ constraints from the remaining $2 \cdot \sum_{i \in \mathcal{N}} (T - \ell_i)$ inequality constraints. Moreover, since each pair of constraints $\alpha_i^t \geq 0$ and $\alpha_i^t \leq v_i \alpha_0^t$ consists on two incompatible constraints, at most $\sum_{i \in \mathcal{N}} (T - \ell_i)$ (i.e., half of the inequality constraints) can be active. Thus, α activates exactly $\sum_{i \in \mathcal{N}} (T - \ell_i)$ inequality constraints, and we have for all $i \in \mathcal{N}$ and $t \in \{\ell_i + 1, \dots, T\}$, $\alpha_i^t \in \{0, v_i \alpha_0^t\}$. Using this new observation, we construct a feasible solution (S_1, \dots, S_T) for **APV** as follows: for all $t \in [T]$, let $S_t := \{i \in \mathcal{N} : \alpha_i^t \neq 0\}$. In particular, if $i \in S_t$, then $\alpha_i^t = v_i \alpha_0^t$. It is then easy to see that α_i^t is exactly the choice probability of product i in assortment S_t^* . Indeed, for all $t \in [T]$, we have

$$\alpha_0^t + \sum_{i \in \mathcal{N}} \alpha_i^t = \alpha_0^t + \sum_{i \in S_t} v_i \alpha_0^t = \alpha_0^t \cdot (1 + V(S_t)),$$

which implies using the first constraint of **LP** that $\alpha_0^0 = \phi(0, S_t)$, and thereby that $\alpha_i^t = \phi(i, S_t)$ for all $i \in S_t$. Next, we can easily see that the (S_1, \dots, S_T) is feasible for **APV**, as any product $i \in \mathcal{N}$ is included in the first ℓ_i assortments, which guarantees the visibility constraints. Finally we have

$$\text{OPT}^{LP} = \sum_{i \in \mathcal{N}} p_i \sum_{t=1}^T \alpha_i^t = \sum_{t=1}^T \sum_{i \in \mathcal{N}} p_i \phi(i, S_t) = \sum_{t=1}^T R(S_t) \leq \text{OPT}^{APV}.$$

Conclusion. Combining the inequalities from the two cases shows that $\text{OPT}^{APV} = \text{OPT}^{LP}$, thereby demonstrating the equivalence between the two optimization problems **LP** and **APV**. ■

4 **APV** with Cardinality Constraints

A natural extension to **APV** consists on considering cardinality constraints on the offered assortment. Cardinality constraints arise in multiple real world applications such as shelf space in brick-and-mortar stores or screen size in online stores, in which case, offering large assortments of products is infeasible. In this section we consider the constrained version of our problem, which we refer to as *Assortment Problem with Visibility and Cardinality constraints (APVC)*, where each customer can be shown at most k products, where k is a positive integer specified

as part of the input. Formally, **APVC** is defined as follows,

$$\begin{aligned}
& \max_{S_1, \dots, S_T \subseteq \mathcal{N}} \sum_{t=1}^T \frac{\sum_{i \in S_t} p_i v_i}{1 + \sum_{i \in S_t} v_i} \\
& \text{s.t.} \quad \sum_{t=1}^T \mathbb{1}(i \in S_t) \geq \ell_i, \quad \forall i \in \mathcal{N}, \\
& \quad \quad |S_t| \leq k, \quad \forall t \in [T].
\end{aligned} \tag{APVC}$$

First, we study the complexity of **APVC**. In the next theorem, we show that **APVC** is strongly NP-Hard.

THEOREM 4.1. ***APVC** is strongly NP-hard, even when all prices p_i are equal. Moreover, there is no FPTAS for **APVC**, even with equal prices.*

The proof of this theorem relies on a reduction of the **3-PARTITION** problem, where the objective is to partition a set of cardinality $3T$ into T subsets, each of size 3, such that the sum of elements within each subset is identical. The full proof of this reduction is deferred to Appendix B.1.

Relying on the strong NP-hardness of this problem, we show that the existence of a Fully Polynomial Time Approximation Scheme (FPTAS) is precluded, unless $P = NP$. Therefore, we focus in the remainder of this section on the question of designing a Polynomial Time Approximation Scheme (PTAS) for the problem, in the case of equal prices. This setting corresponds to a sales maximization problem, which is particularly useful in the case where the platform or the company’s goal is to maximize the captured portion of customers.

THEOREM 4.2. *There exists a PTAS for **APVC** with equal prices.*

The remainder of this section focuses on proving Theorem 4.2, by devising a PTAS for **APVC** with equal weights. The proof is organized as follows. In Section 4.1, we consider an instance of **APVC** with equal prices, which we reduce to a “friendlier” instance of the problem by slightly altering the weights of a particular subset of products. Importantly, we show that we only incur an ϵ loss due to this reduction, and therefore only focus on solving the modified instance. In Section 4.2, we linearize our optimization problem through a carefully crafted guessing procedure, then we provide a method to obtain approximate solutions for the linearized program, thereby completing the formulation of our PTAS. Finally, we show in Section 4.3 that the solution yielded by our algorithm is indeed a $(1 - \epsilon)$ -approximation for our problem.

4.1 Discretizing the universe of products

The setting with identical prices is formally defined as follows

$$\begin{aligned}
& \max_{S_1, \dots, S_T \subseteq \mathcal{N}} \sum_{t=1}^T \frac{V(S_t)}{1 + V(S_t)} \\
& \text{s.t.} \quad \sum_{t=1}^T \mathbb{1}(i \in S_t) \geq \ell_i, \quad \forall i \in \mathcal{N}, \\
& \quad \quad |S_t| \leq k, \quad \forall t \in [T].
\end{aligned} \tag{SPVC}$$

In this section, we reduce the original instance of our problem into a modified instance, by rounding down the weights of a particular subset of products in our original universe. Then, we show that we incur at most an ϵ -loss by performing this reduction. For any product $i \in \mathcal{N}$,

- If $v_i \geq 1/\epsilon$, then v_i is rounded down to $v_i^\downarrow = 1/\epsilon$. We denote the product associated with the weight v_i^\downarrow simply by i^\downarrow .
- If $v_i < \epsilon^5$, then v_i remains unchanged and we define $v_i^\downarrow = v_i$. Similarly, the associated product is denoted by i^\downarrow .
- If $\epsilon^5 \leq v_i < 1/\epsilon$, then v_i is rounded down such that v_i/ϵ^5 meets the nearest power of $(1 + \epsilon)$. In other words, since $\epsilon^5 \leq v_i < 1/\epsilon$, there exists some $q \geq 1$ such that $\epsilon^5 \cdot (1 + \epsilon)^{q-1} \leq v_i < \epsilon^5 \cdot (1 + \epsilon)^q$, and v_i is rounded down to $v_i^\downarrow = \epsilon^5 \cdot (1 + \epsilon)^{q-1}$. We denote the associated product by i^\downarrow . In particular, note that by this definition, we have

$$v_i^\downarrow \leq v_i \leq (1 + \epsilon) \cdot v_i^\downarrow. \quad (8)$$

Let us denote this new universe of products by $\mathcal{N}^\downarrow = \{i^\downarrow : i \in \mathcal{N}\}$. We invite the reader to think of \cdot^\downarrow , when applied to a product, as a one-to-one map from the universe \mathcal{N} , to the modified universe \mathcal{N}^\downarrow , where a subset of weights have been rounded down. In the subsequent result, we show that solving a modified instance of **SPVC** using the products from \mathcal{N}^\downarrow yields a $(1 - \epsilon)$ -approximation of the solution on the original instance. To this purpose, we define the following new instance of **SPVC**, which only uses products from the universe \mathcal{N}^\downarrow :

$$\begin{aligned} & \max_{S_1, \dots, S_T \subseteq \mathcal{N}^\downarrow} \sum_{t=1}^T \frac{V(S_t)}{1 + V(S_t)} \\ & \text{s.t.} \quad \sum_{t=1}^T \mathbb{1}(i^\downarrow \in S_t) \geq \ell_i, \quad \forall i^\downarrow \in \mathcal{N}^\downarrow, \\ & \quad \quad |S_t| \leq k, \quad \forall t \in [T]. \end{aligned} \quad (\text{SPVC}^\downarrow)$$

Let OPT (resp. OPT^\downarrow) denote the objective of an optimal solution of **SPVC** (resp. **SPVC** $^\downarrow$). In the following, we show that OPT and OPT^\downarrow are within a $1 - \epsilon$ fraction of one another. To this purpose, we show a stronger result in Lemma 4.3. Indeed, for any sequence of assortments S_1, \dots, S_t , we define

$$\mathcal{E}(S_1, \dots, S_T) = \sum_{t=1}^T \frac{V(S_t)}{1 + V(S_t)}.$$

For any assortment $S \subseteq \mathcal{N}$, we denote by S^\downarrow its rounded counterpart, defined as $S^\downarrow = \{i^\downarrow : i \in S\}$. In the next lemma, we show that the objective achieved by any sequence of T assortments in \mathcal{N} is within $1 - \epsilon$ of their rounded counterpart. In particular, letting OPT denote the value of **SPVC** (i.e., the objective of an optimal feasible sequence of assortments), and similarly OPT^\downarrow denote the value of **SPVC** $^\downarrow$, this implies that OPT and OPT^\downarrow are also within $1 - \epsilon$.

LEMMA 4.3. For any sequence of assortments $S_1, \dots, S_T \subseteq \mathcal{N}$, we have

$$\mathcal{E}(S_1^\downarrow, \dots, S_T^\downarrow) \leq \mathcal{E}(S_1, \dots, S_T) \leq (1 + \epsilon) \cdot \mathcal{E}(S_1^\downarrow, \dots, S_T^\downarrow).$$

Consequently, we have $\text{OPT}^\downarrow \leq \text{OPT} \leq (1 + \epsilon) \cdot \text{OPT}^\downarrow$.

Proof of Lemma 4.3. First, since all the weights of the products in \mathcal{N} are either rounded down or remain the same, then for all $t \in [T]$, $V(S_t^\downarrow) \leq V(S_t)$. Consequently, by the monotonicity of the map $x \mapsto x/(1+x)$, we have $V(S_t^\downarrow)/(1+V(S_t^\downarrow)) \leq V(S_t)/(1+V(S_t))$. By summing over t , we obtain the first inequality $\mathcal{E}(S_1^\downarrow, \dots, S_T^\downarrow) \leq \mathcal{E}(S_1, \dots, S_T)$. Let us prove the second inequality. We show that $V(S_t)/(1+V(S_t)) \leq (1+\epsilon) \cdot V(S_t^\downarrow)/(1+V(S_t^\downarrow))$, for every $t \in [T]$. Let $t \in [T]$.

- If there exists a product $i \in S_t$ such that $v_i \geq 1/\epsilon$, then $v_i^\downarrow = 1/\epsilon$, and in turn, $V(S_t^\downarrow) \geq 1/\epsilon$. Therefore,

$$(1 + \epsilon) \frac{V(S_t^\downarrow)}{1 + V(S_t^\downarrow)} \geq (1 + \epsilon) \cdot \frac{1/\epsilon}{1 + 1/\epsilon} = 1 \geq \frac{V(S_t)}{1 + V(S_t)}.$$

- If for all $i \in S_t$, we have $v_i < 1/\epsilon$. Note that Equation (8) holds trivially when $v_i < \epsilon$, since $v_i^\downarrow = v_i$ by definition. Therefore, Equation (8) holds for every $i \in S_t$. Hence,

$$\frac{v_i}{1 + v_i} \leq \frac{(1 + \epsilon) \cdot v_i^\downarrow}{1 + v_i^\downarrow}.$$

The result follows by summing over $t = 1, \dots, T$.

Noticing that if S_1, \dots, S_T is feasible for **SPVC**, then its rounded counterpart is also feasible for **SPVC**[↓], we also have $\text{OPT}^\downarrow \leq \text{OPT} \leq (1 + \epsilon) \cdot \text{OPT}^\downarrow$. ■

In light of Lemma 4.3, we deduce that obtaining an approximation to the reduced problem **SPVC**[↓] automatically yields an approximation to the original problem **SPVC**, with only an additional ϵ loss. In the remainder of this section, we will only focus on solving the reduced problem **SPVC**[↓], therefore, for simplicity of notation, we denote the products $1^\downarrow, \dots, n^\downarrow$ simply by $1, \dots, n$ respectively, and we denote their respective weights simply by the original v_1, \dots, v_n instead of the heavy notation $v_1^\downarrow, \dots, v_n^\downarrow$. So, products $1, \dots, n$ and their respective weights v_1, \dots, v_n will refer to those of the modified instance **SPVC**[↓].

4.2 Linearization of the objective function and algorithm

In this section, we aim at linearizing the optimization problem **SPVC**, through a guessing procedure of a certain number of parameters pertaining to the optimal solution. We start by presenting our detailed guessing procedure in Section 4.2.1, before leveraging the guessed parameters to construct a linearized formulation in Section 4.2.2. In Section 4.2.3, we present a dependent rounding scheme introduced in Gandhi et al. (2002), which we use later in Section 4.2.4 to round a fractional solution of the linearized program, into a feasible integral solution for **SPVC**[↓].

4.2.1 Guessing procedure

Let S_1^*, \dots, S_T^* be an (unknown) optimal solution for our problem [SPVC](#)↓.

Customer types. In this definition, we aim at categorizing customers with respect to the total weight of the products offered to them in an optimal solution. First we say that a customer $t \in [T]$ is *light* if $V(S_t^*) < \epsilon$. Similarly, we say that a customer is *heavy* if $V(S_t^*) \geq 1/\epsilon$. Otherwise (i.e., if $\epsilon \leq V(S_t^*) < 1/\epsilon$), we say that the customer is *medium*. We further partition medium customers into L classes G_1, \dots, G_L , depending on the value of $V(S_t^*)$. Specifically, we define every class $\ell = 1, \dots, L$ as follows,

$$G_\ell = \left\{ t \text{ medium} : \epsilon \cdot (1 + \epsilon)^{\ell-1} \leq V(S_t^*) < \epsilon \cdot (1 + \epsilon)^\ell \right\}.$$

We fix L to be the smallest integer such that $\epsilon \cdot (1 + \epsilon)^L \geq 1/\epsilon$. In particular, $L = O(\frac{1}{\epsilon} \log(\frac{1}{\epsilon}))$. Therefore, the classes G_1, \dots, G_L form a partition of medium customers. Additionally, we denote by G_{light} and G_{heavy} the sets of light and heavy customers respectively.

Packing patterns. First let us start by defining classes of products. Recall from Section 4.1 that for every product $i \in \mathcal{N}$ such that $v_i \geq \epsilon^5$, either there exists some $1 \leq q \leq Q - 1$ such that $v_i = \epsilon^5 \cdot (1 + \epsilon)^{q-1}$ or $v_i = 1/\epsilon$, where Q is the smallest integer such that $\epsilon^5 \cdot (1 + \epsilon)^{Q-1} \geq 1/\epsilon$. Let D_q denote the set of all the products of \mathcal{N} such that $v_i = \epsilon^5 \cdot (1 + \epsilon)^{q-1}$, for every $1 \leq q \leq Q - 1$, and let D_Q denote the set of all products of \mathcal{N} such that $v_i = 1/\epsilon$. Similarly, let D_0 be the class of all products $i \in \mathcal{N}$ whose weight $v_i < \epsilon^5$. Note that by construction, the classes D_0, \dots, D_Q form a partition of the universe of products \mathcal{N} . A packing pattern is a $(Q + 1)$ -dimensional vector, taking values in $\{0, 1, 2, \dots, 1/\epsilon^6, \star\}$, where $1/\epsilon^6$ is assumed to be an integer without loss of generality. For a given customer, each entry $q = 0, \dots, Q$ of its associated packing pattern characterizes the number of products from class D_q in the assortment S_t^* , and \star should be viewed simply as a symbol which characterizes the fact that the number of products from class q is (strictly) greater than $1/\epsilon^6$. Using this definition, the number of packing patterns is

$$\left(\frac{1}{\epsilon^6} + 2 \right)^{Q+1} = 2^{O((Q+1) \log(\frac{1}{\epsilon}))} = 2^{O(\frac{1}{\epsilon} \log^2(\frac{1}{\epsilon}))}.$$

Guessing. With all the previous definitions, we are now able to present the guessing procedure. For each couple consisting of a class of customers $\ell \in \{\text{light}, \text{heavy}, 1, \dots, L\}$, and a packing pattern P , we guess the number of customers in G_ℓ that use the packing pattern P . Let us call this guessed number $K_{\ell, P}$.

It is important to note that using the above guessing procedure allows us to determine, for any given customer $t \in [T]$, both her type and her packing pattern in the optimal solution, up to a permutation of $\{1, \dots, T\}$. Indeed, since for each class $\ell \in \{\text{light}, \text{heavy}, 1, \dots, L\}$, we know the number of customers that use each packing pattern, in particular, we know the quantity $|G_\ell|$. Hence, due to the fact that all customers are initially interchangeable, we arbitrarily assign each customer to a type according to our guess, without loss of generality. Subsequently, given our guess $K_{\ell, P}$ for every customer type ℓ , and every packing pattern P , and since the customers

within the same type are all interchangeable a priori, we arbitrarily assign each customer to a packing pattern, according to our guess, without loss of generality. Therefore, for each customer $t \in [T]$, we have determined both her type and her packing pattern, up to a permutation of $\{1, \dots, T\}$.

Complexity of the guessing procedure. For any couple consisting of a specific class of customers and a specific packing pattern, we guess the number of customers from said class who use said packing pattern. There are $L+2 = O(\frac{1}{\epsilon} \log(\frac{1}{\epsilon}))$ classes of customers, and $O(2^{O(\frac{1}{\epsilon} \log^2(\frac{1}{\epsilon}))})$ packing patterns. The number of couples is therefore given by $O(2^{O(\frac{1}{\epsilon} \log^2(\frac{1}{\epsilon}))})$. Since there are T customers, there are $T+1$ possible guesses for each couple consisting of a class of customers and a packing pattern. Therefore, the total number of guesses is $O(T^{2^{O(\frac{1}{\epsilon} \log^2(\frac{1}{\epsilon}))}})$.

4.2.2 Integer linear programming formulation

Following the detailed presentation of the guessing procedure, we are now ready to leverage the guessed parameters in order to introduce our integer linear programming formulation **ILP**. This formulation is crucial to our algorithm, since its approximate solution is later (Section 4.3) shown to be near optimal for **SPVC**[↓]. In [Segev and Singla \(2021\)](#), the authors use a similar idea in order to obtain rounded solutions for the Santa Claus problem.

Let us start with some useful notation. For each customer t , we define the lower bound V_t on the $V(S_t^*)$ as follows. If t is light (resp. heavy) then $V_t = 0$ (resp. $V_t = 1/\epsilon$), and if V_t is of type ℓ for some $\ell = 1, \dots, L$, then $V_t = \epsilon \cdot (1 + \epsilon)^{\ell-1}$. Finally, we denote the packing pattern of customer t by $(k_{0,t}, k_{1,t}, \dots, k_{Q,t})$. Using these newly defined parameters, we define the following integer linear program:

$$\begin{aligned}
& \max_{\mathbf{x}} && \sum_{t \in G_{light}} \sum_{i \in \mathcal{N}} v_i x_{it} \\
& \text{s.t.} && \sum_{t=1}^T x_{it} \geq \ell_i, && \forall i \in \mathcal{N}, \\
& && \sum_{i \in \mathcal{N}} x_{it} \leq k, && \forall t \in [T], \\
& && \sum_{i \in D_q} x_{it} = k_{q,t}, && \forall t \in [T], \forall q = 0, \dots, Q \text{ such that } k_{q,t} \neq \star, \\
& && \sum_{i \in D_q} x_{it} > \frac{1}{\epsilon^6}, && \forall t \in [T], \forall q = 0, \dots, Q \text{ such that } k_{q,t} = \star, \\
& && \sum_{i \in \mathcal{N}} v_i x_{it} \geq V_t, && \forall t \in [T], \\
& && \sum_{i \in \mathcal{N}} v_i x_{it} \leq \epsilon, && \forall t \in G_{light}, \\
& && x_{it} \in \{0, 1\}, && \forall i \in \mathcal{N}, \forall t \in [T].
\end{aligned} \tag{ILP}$$

The last step is to present a way to approximate **ILP**. To this purpose, the high level idea of our method is to start by solving the linear relaxation of **ILP**, then to use a specific rounding of the obtained fractional solution, in order to derive a near-optimal integral solution.

Let \mathbf{x}^* be an optimal solution to the linear relaxation of **ILP**, that is, to the linear program

obtained by replacing the constraint $x_{it} \in \{0, 1\}$ by the relaxed constraint $x_{it} \in [0, 1]$. Similarly, let S_1^*, \dots, S_T^* be an optimal sequence of assortments for [SPVC](#)[↓]. The following lemma relates \mathbf{x}^* to the optimal solution S_1^*, \dots, S_T^* .

LEMMA 4.4.

$$\sum_{t=1}^T \frac{\sum_{i \in \mathcal{N}} v_i x_{it}^*}{1 + \sum_{i \in \mathcal{N}} v_i x_{it}^*} \geq (1 - \epsilon) \cdot \sum_{t=1}^T \frac{V(S_t^*)}{1 + V(S_t^*)}.$$

The proof of this Lemma is included in [Appendix B.2](#). The next step is to round \mathbf{x}^* in order to obtain a feasible solution for [ILP](#).

4.2.3 The dependent rounding scheme

In this section, we briefly present the dependent rounding scheme in [Gandhi et al. \(2002\)](#), which we use to round our fractional solution into an integral one. Given a bipartite graph (L, R, E) and values $x_{i,j} \in [0, 1]$ for every edge $(i, j) \in E$, the authors in [Gandhi et al. \(2002\)](#) present a polynomial time rounding scheme which returns a sequence of random variables $X_{i,j} \in \{0, 1\}$ representing a rounding of the values $x_{i,j}$, and which verify the following properties:

- (P1) Marginal distributions:** For every edge (i, j) , $\mathbb{E}(X_{i,j}) = x_{i,j}$.
- (P2) Degree-preservation:** The fractional degree of every edge is rounded to its floor or its ceiling, i.e., for any $i \in L \cup R$, if $d_i = \sum_{j \in L \cup R} x_{ij}$, then $\sum_{j \in L \cup R} X_{ij} \in \{\lfloor d_i \rfloor, \lceil d_i \rceil\}$.
- (P3) Negative correlation:** For any node $i \in L \cup R$, and any subset $S \subseteq N(i)$, where $N(i)$ is the set of neighbors of i , we have

$$\mathbb{P} \left[\bigwedge_{j \in S} (X_{i,j} = 0) \right] \leq \prod_{j \in S} \mathbb{P} [X_{i,j} = 0]$$

and

$$\mathbb{P} \left[\bigwedge_{j \in S} (X_{i,j} = 1) \right] \leq \prod_{j \in S} \mathbb{P} [X_{i,j} = 1].$$

Negative correlation **(P3)** is a strong property that allows us to derive Chernoff bounds, as stated in the following lemma.

LEMMA 4.5 ((THEOREM 1.10.24 [DOERR, 2020](#), 66)). *If X_1, \dots, X_n are negatively correlated random variables, then for any vector $a = (a_1, \dots, a_n) \in [0, 1]^n$ we have*

$$\mathbb{P}[a(X) \leq (1 - \epsilon) \cdot \mathbb{E}[a(X)]] \leq \exp \left(-\frac{\epsilon^2 \cdot \mathbb{E}[a(X)]}{2} \right),$$

where $a(X) = \sum_{i=1}^n a_i X_i$.

Next, we construct a framework that allows us to use the dependent rounding scheme in order to derive a feasible solution for [ILP](#).

4.2.4 Construction of the bipartite graph and feasibility of the rounded solution

Recall that the high level idea is to use the rounding procedure presented in 4.2.3, in order to derive an integer solution to ILP. Noticing that the rounding procedure takes as input a bipartite graph and a sequence of edge-associated values, the first step is to interpret the solution of the linear relaxation of ILP as a sequence of edge-associated values in a carefully constructed bipartite graph. Before presenting our proposed bipartite graph, let us introduce the two following definitions. We say that a customer $t \in [T]$ is *bounded* if $k_{q,t} \neq \star$, for all $q = 1, \dots, Q$ (note that q starts from 1 in this definition and not from 0). Otherwise, we say that customer t is *unbounded*. In other words, customer t is bounded if she is not offered more than $1/\epsilon^5$ products from any class D_q where $q \geq 1$, in the optimal solution.

The bipartite graph. We start by defining the node sets L and R . The set L is simply the set of all products in \mathcal{N} . Let us now describe the set of right nodes R . First, for every unbounded customer t , we introduce a node t in the set R . Alternatively, for every bounded customer t , we introduce $Q + 1$ nodes in the set R , specifically, one node (t, q) for every class of products $q = 0, \dots, Q$. Therefore,

$$R = \{t : t \text{ unbounded}\} \cup \{(t, q) : t \text{ bounded}, q = 0, \dots, Q\}.$$

Let us now describe the edge set E of the bipartite graph, as well as the associated value x_e of each edge $e \in E$. Let us fix a node $i \in L$ and describe all of its adjacent edges and associated values. First, i neighbors every node t where t is an unbounded customer, and the edge (i, t) is associated with the value x_{it}^* . Next, assume q is the class of products containing i , then the node i neighbors every node (t, q) for $t \in [T]$, and the edge $(i, (t, q))$ is associated with the value x_{it}^* . This completes our description of the bipartite graph.

PTAS. Let us summarize the steps of our PTAS. The first step is to perform the guessing procedure presented in Section 4.2.1, and then use our guess to construct ILP. The second step is to solve the linear relaxation of ILP using standard linear programming techniques, and obtain a fractional solution \mathbf{x}^* . The third step is to use \mathbf{x}^* to construct the bipartite graph as described above. We can now apply the dependent rounding presented in Section 4.2.3 to derive a sequence of random variables $(X_{it} : i \in \mathcal{N}, t \in [T])$ representing rounded values. Finally, we denote by S_1, \dots, S_T the associated sequence of assortments, that is, for each $t \in [T]$,

$$S_t = \{i \in \mathcal{N} : X_{it} = 1\},$$

and the algorithm returns this sequence (S_1, \dots, S_T) .

Complexity analysis. Since our algorithm is exhaustive, the steps presented above are applied for every possible guess. Solving the linear relaxation of ILP takes $O(n^3 T^3)$ using standard linear programming algorithms (interior point method for example). Constructing the bipartite graph takes $O(nTQ)$ running time, and running the dependent rounding scheme takes $O((nTQ)^3)$. Therefore, processing each guess takes $O(n^3 T^3 Q^3) = O(n^3 T^3 \frac{1}{\epsilon^3} \log^3 \frac{1}{\epsilon})$ running

time. Finally, since there are $O(T^{2^{O(\frac{1}{\epsilon} \log^2(\frac{1}{\epsilon}))}})$ possible guesses as argued in Section 4.2.1, the total running time of our algorithm is:

$$O\left(T^{2^{O(\frac{1}{\epsilon} \log^2(\frac{1}{\epsilon}))}} \cdot n^3 T^3 \frac{1}{\epsilon^3} \log^3 \frac{1}{\epsilon}\right) = O\left(n^3 T^{2^{O(\frac{1}{\epsilon} \log^2(\frac{1}{\epsilon}))}}\right),$$

which is polynomial in the size of the input for any fixed ϵ , confirming that our algorithm is indeed a PTAS.

4.3 Near-optimality of the rounded solution

In this section, we show that the obtained solution (S_1, \dots, S_T) is near-optimal for SPVC^\downarrow . The main result is stated in the following theorem.

THEOREM 4.6. *The sequence of assortments S_1, \dots, S_T is feasible for SPVC^\downarrow and we have*

$$\mathbb{E} \left[\sum_{t=1}^T \frac{V(S_t)}{1 + V(S_t)} \right] \geq (1 - 3\epsilon) \cdot \text{OPT}^\downarrow.$$

Proof. In this proof, we start by showing that the sequence S_1, \dots, S_T is feasible for SPVC^\downarrow . Next, we show that the objective of the sequence of assortments S_1, \dots, S_T is within $1 - 3\epsilon$ of the optimal objective. Recall that S_1^*, \dots, S_T^* is an optimal sequence of assortments for SPVC^\downarrow .

Feasibility. First, the visibility constraint is straightforward. Indeed, let $i \in \mathcal{N}$. By the first constraint of the linear relaxation of ILP , we have: $\sum_{t=1}^T x_{it}^* \geq \ell_i$. Noticing that $\sum_{t=1}^T x_{it}^*$ is the fractional degree of node i in the bipartite graph, we have by property **(P2)**, $\sum_{t=1}^T X_{it} \geq \lfloor \ell_i \rfloor = \ell_i$, which proves that the visibility constraint is respected. We now show that the cardinality constraint is respected after the rounding. Let $t \in [T]$. First, if t is unbounded, then the cardinality of S_t is simply given by the degree of node t . Since $\sum_{i=1}^n x_{it}^* \geq k$, then by property **(P2)**, we have $|S_t| = \sum_{i \in \mathcal{N}} X_{it} \geq k$. Alternatively, if t is bounded, for every class $q = 1, \dots, Q$, the degree of node (t, q) is $\sum_{i \in D_q} x_{it}^*$ which is equal to $k_{q,t}$, according to the third constraint of ILP . Since $k_{q,t}$ is an integer, the degree of node (t, q) remains unchanged after the rounding, according to property **(P2)**. Therefore, we have almost surely:

$$\begin{aligned} \sum_{i \in \mathcal{N}} v_i X_{it} &= \sum_{i \in D_0} v_i X_{it} + \sum_{q=1}^Q \sum_{i \in D_q} v_i X_{it} \\ &\leq \left\lceil \sum_{i \in D_0} v_i x_{it}^* \right\rceil + \sum_{q=1}^Q k_{q,t} \\ &= \left\lceil \sum_{i \in D_0} v_i x_{it}^* + \sum_{q=1}^Q k_{q,t} \right\rceil \\ &= \left\lceil \sum_{i \in \mathcal{N}} v_i x_{it}^* \right\rceil \\ &\leq \lceil k \rceil = k. \end{aligned}$$

The first inequality follows from Property **(P2)**. In the equality in the third line, we use the fact that the second summand is an integer. This proves the cardinality constraint and concludes the feasibility proof.

Near-optimality. We show that the contribution of every customer $t \in [T]$ to the objective function, i.e., the quantity $V(S_t)/(1 + V(S_t))$ is within $1 - 2\epsilon$ of the contribution of t in the optimal solution \mathbf{x}^* , in expectation. Let $t \in [T]$.

Case 1: *If t unbounded.* Let $q_t \in \{1, \dots, Q\}$ be the class of products such that $k_{q_t, t} = \star$. We have

$$\sum_{i \in D_{q_t}} v_i x_{it}^* \geq \epsilon^5 \cdot \sum_{i \in D_{q_t}} x_{it}^* \geq \epsilon^5 \cdot \frac{1}{\epsilon^6} = \frac{1}{\epsilon}.$$

Therefore,

$$\begin{aligned} \mathbb{P} \left[V(S_t) \leq \frac{1 - \epsilon}{\epsilon} \right] &= \mathbb{P} \left[\sum_{i \in \mathcal{N}} v_i X_{it} \leq \frac{1 - \epsilon}{\epsilon} \right] \\ &\leq \mathbb{P} \left[\sum_{i \in \mathcal{N}} v_i X_{it} \leq (1 - \epsilon) \cdot \sum_{i \in D_{q_t}} v_i x_{it}^* \right] \\ &\leq \mathbb{P} \left[\sum_{i \in D_{q_t}} v_i X_{it} \leq (1 - \epsilon) \cdot \sum_{i \in D_{q_t}} v_i x_{it}^* \right] \\ &= \mathbb{P} \left[\sum_{i \in D_{q_t}} X_{it} \leq (1 - \epsilon) \cdot \sum_{i \in D_{q_t}} x_{it}^* \right] \\ &\leq \exp \left(-\frac{\epsilon^2 \cdot \sum_{i \in D_{q_t}} x_{it}^*}{2} \right) \\ &\leq \exp \left(-\frac{1}{2\epsilon^4} \right) \\ &\leq \epsilon. \end{aligned} \tag{9}$$

In the second equality, we use the fact that products within the same class have the same preference weight to simplify all the weights. The inequality in the fifth line is a direct application of Lemma 4.5, and in the one in the sixth line, we use the fact that $\sum_{i \in D_{q_t}} x_{it}^* > 1/\epsilon^6$, since $k_{q_t, t} = \star$. Therefore, by conditioning on the event $\{V(S_t) > (1 - \epsilon)/\epsilon\}$, and neglecting one of the terms, we have

$$\begin{aligned} \mathbb{E} \left[\frac{V(S_t)}{1 + V(S_t)} \right] &\geq \mathbb{P} \left[V(S_t) > \frac{1 - \epsilon}{\epsilon} \right] \cdot \mathbb{E} \left[\frac{V(S_t)}{1 + V(S_t)} \mid V(S_t) > \frac{1 - \epsilon}{\epsilon} \right] \\ &\geq (1 - \epsilon) \cdot \frac{\frac{1 - \epsilon}{\epsilon}}{1 + \frac{1 - \epsilon}{\epsilon}} \\ &= (1 - \epsilon)^2 \\ &\geq (1 - 2\epsilon) \cdot \frac{\sum_{i \in \mathcal{N}} v_i x_{it}^*}{1 + \sum_{i \in \mathcal{N}} v_i x_{it}^*}. \end{aligned}$$

The second inequality follows from Equation (9) and the monotonicity of $x \mapsto x/(1+x)$ on $[0, +\infty)$.

Case 2: If t is bounded and heavy or if t is medium. Then we have

$$\sum_{i \in \mathcal{N}} v_i X_{it} = \sum_{i \in D_0} v_i X_{it} + \sum_{q=1}^{Q-1} \sum_{i \in D_q} v_i X_{it}.$$

Let $W_t^l = \sum_{i \in \mathcal{N} \setminus D_0} v_i X_{it}$, and $W_t^s = \sum_{i \in D_0} v_i X_{it}$. Similarly, let $w_t^l = \sum_{i \in \mathcal{N} \setminus D_0} v_i x_{it}^*$, and $w_t^s = \sum_{i \in D_0} v_i x_{it}^*$.

CLAIM 4.7. *With probability 1, we have $W_t^l = w_t^l$.*

The proof of this result is a direct application of the degree preservation property, and is deferred to Appendix B.3.

- If $w_t^s \leq \epsilon \cdot \sum_{i \in \mathcal{N}} v_i x_{it}^*$, then intuitively, the contribution of the products in D_0 represent an ϵ -fraction of the total contribution and can therefore be neglected. Formally, we have with probability 1

$$V(S_t) \geq W_t^l = w_t^l = \sum_{i \in \mathcal{N}} v_i x_{it}^* - w_t^s \geq (1 - \epsilon) \cdot \sum_{i \in \mathcal{N}} v_i x_{it}^*.$$

The first equality is a consequence of Claim 4.7, and the second inequality follows from the case hypothesis. Therefore, by the monotonicity of $x \mapsto x/(1+x)$, we have with probability 1

$$\frac{V(S_t)}{1 + V(S_t)} \geq (1 - \epsilon) \cdot \frac{\sum_{i \in \mathcal{N}} v_i x_{it}^*}{1 + \sum_{i \in \mathcal{N}} v_i x_{it}^*}. \quad (10)$$

- Otherwise, we have

$$\begin{aligned} \mathbb{P}[W_t^s \leq (1 - \epsilon) \cdot w_t^s] &= \mathbb{P}\left[\sum_{i \in D_0} \frac{v_i}{\epsilon^5} X_{it} \leq (1 - \epsilon) \cdot \sum_{i \in D_0} \frac{v_i}{\epsilon^5} x_{it}^*\right] \\ &\leq \exp\left(-\frac{\epsilon^2 \cdot \sum_{i \in D_0} v_i x_{it}^*}{2\epsilon^5}\right) \\ &\leq \exp\left(-\frac{\epsilon \cdot \sum_{i \in \mathcal{N}} v_i x_{it}^*}{2\epsilon^3}\right) \\ &\leq \exp\left(-\frac{1}{2\epsilon}\right) \\ &\leq \epsilon. \end{aligned} \quad (11)$$

The first inequality follows from Lemma 4.5. The second inequality is a consequence of the case hypothesis, and third one follows from the fact that t is not light, and hence that $\sum_{i \in \mathcal{N}} v_i x_{it}^* \geq \epsilon$. Therefore, by conditioning on the event $\{W_t^s \leq (1 - \epsilon) \cdot w_t^s\}$ and neglecting one of the terms, we have

$$\mathbb{E}\left[\frac{V(S_t)}{1 + V(S_t)}\right] \geq \mathbb{P}[W_t^s > (1 - \epsilon) \cdot w_t^s] \cdot \mathbb{E}\left[\frac{V(S_t)}{1 + V(S_t)} \mid W_t^s > (1 - \epsilon) \cdot w_t^s\right]$$

$$\begin{aligned}
&= \mathbb{P}[W_t^s > (1 - \epsilon) \cdot w_t^s] \cdot \mathbb{E} \left[\frac{w_t^l + W_t^s}{1 + w_t^l + W_t^s} \mid W_t^s > (1 - \epsilon) \cdot w_t^s \right] \\
&\geq (1 - \epsilon) \cdot \frac{w_t^l + (1 - \epsilon) \cdot w_t^s}{1 + w_t^l + (1 - \epsilon) \cdot w_t^s} \\
&\geq (1 - \epsilon)^2 \cdot \frac{\sum_{i \in \mathcal{N}} v_i x_{it}^*}{1 + \sum_{i \in \mathcal{N}} v_i x_{it}^*} \\
&\geq (1 - 2\epsilon) \cdot \frac{\sum_{i \in \mathcal{N}} v_i x_{it}^*}{1 + \sum_{i \in \mathcal{N}} v_i x_{it}^*}.
\end{aligned}$$

In the equality, we apply Lemma 4.7. In the second inequality, we use Equation (11), as well as the monotonicity of $x \mapsto x/(1+x)$.

Case 3: If t is light.

CLAIM 4.8. *We have*

$$\mathbb{E} \left[\frac{V(S_t)}{1 + V(S_t)} \right] \geq (1 - 2\epsilon) \cdot \frac{\sum_{i \in \mathcal{N}} v_i x_{it}^*}{1 + \sum_{i \in \mathcal{N}} v_i x_{it}^*}.$$

Let us present the high level idea for the proof of Claim 4.8. If $t \in [T]$ is a light customer then $\sum_{i \in \mathcal{N}} v_i x_{it}^* \leq \epsilon$. Therefore, the denominator $1 + V(S_t)$ is in expectation equal to $1 + \sum_{i \in \mathcal{N}} v_i x_{it}^* \leq 1 + \epsilon$. The high level idea is to use this remark to approximate $V(S_t)/(1 + V(S_t))$ to simply $V(S_t)$, which in expectation is equal to $\sum_{i \in \mathcal{N}} v_i x_{it}^*$. Finally, this last term is trivially lower bounded by $\sum_{i \in \mathcal{N}} v_i x_{it}^*/(1 + \sum_{i \in \mathcal{N}} v_i x_{it}^*)$. The formal proof is deferred to Appendix B.3.

Combining cases 1, 2 and 3, and by linearity of expectation, we have

$$\begin{aligned}
\mathbb{E} \left[\sum_{t=1}^T \frac{V(S_t)}{1 + V(S_t)} \right] &\geq (1 - 2\epsilon) \cdot \sum_{t=1}^T \frac{\sum_{i \in \mathcal{N}} v_i x_{it}^*}{1 + \sum_{i \in \mathcal{N}} v_i x_{it}^*} \\
&\geq (1 - 2\epsilon) \cdot (1 - \epsilon) \cdot \sum_{t=1}^T \frac{V(S_t^*)}{1 + V(S_t^*)} \\
&\geq (1 - 3\epsilon) \cdot \sum_{t=1}^T \frac{V(S_t^*)}{1 + V(S_t^*)} \\
&= (1 - 3\epsilon) \cdot \text{OPT}^\downarrow,
\end{aligned}$$

where the second inequality is an application of Lemma 4.4. ■

5 Price of Visibility

In this section, we investigate the impact of visibility constraints on the total expected revenue, comparing it to the unconstrained setting where there are no visibility constraints. In Section 5.1, we quantify the loss resulting from enforcing the visibility constraints. In Section 5.2, we introduce a novel method to distribute the loss among different products in proportion to their contribution to the overall loss. Finally, in Section 5.3, we illustrate our method through a series of numerical experiments.

5.1 The loss due to Visibility constraints

Consider the unconstrained assortment optimization [AP](#) and let S^* be its optimal solution. It is known that there exists an optimal assortment that is price-ordered. This is a standard result in assortment optimization under the MNL model (see [Talluri and Van Ryzin Talluri and Van Ryzin \(2004\)](#)).² In the absence of visibility constraints and with T customers, it is optimal to offer assortment S^* to each customer. Consequently, the total expected revenue in the unconstrained setting can be expressed as $T \cdot R(S^*)$. As the unconstrained problem serves as a relaxation of [APV](#), it possesses a higher objective function. In the following example, we show that enforcing the visibility constraints can imply a gap that is arbitrary bad as compared to the unconstrained setting.

Example. Consider an instance of [APV](#) with two products $n = 2$ and T customers. Let $p_1 = 1, v_1 = 1, \ell_1 = 0$ and $p_2 = 0, v_2 = M, \ell_2 = T$. We consider the setting where M is large. To compute the optimal assortment for [AP](#). As mentioned earlier, it sufficient to evaluate the revenue of the price-ordered assortments and choose the one with the highest revenue. We have $R(\{1\}) = \frac{1}{2}$ and $R(\{1, 2\}) = \frac{1}{2+M} < R(\{1\})$. Therefore, the optimal assortment is $S^* = \{1\}$ and $R(S^*) = \frac{1}{2}$.

On the other hand, let (S_1, S_2, \dots, S_T) be a feasible solution for [APV](#). Because, $\ell_2 = T$, we have to include product 2 in every assortment S_t . Adding product 1 to an assortment only increases the revenue because $R(\{1, 2\}) = \frac{1}{2+M} > R(\{2\}) = 0$. Therefore, it is optimal to offer product 1 and 2 in every assortment S_t in an optimal solution of [APV](#). Therefore, $S_t^* = \{1, 2\}$ for all $t = 1, \dots, T$. Now, consider the ratio

$$\frac{TR(S^*)}{\sum_{t=1}^T R(S_t^*)} = \frac{T \frac{1}{2}}{T \frac{1}{2+M}} = \frac{M}{2} + 1,$$

which goes to infinity as M increases. Hence the gap can be arbitrarily large. While in this pathological example, we see that enforcing products with very low price and high weight can drive the revenue very low, our numerical experiments in [Section 5.3](#) will illustrate how more common distributions of the prices and weight react to the enforcement of visibility constraints.

5.2 Sharing the loss

In this section, we explore a scenario where each product within our universe is associated with a specific vendor. As mentioned earlier, vendors can impose visibility constraints on their products within the platform. These constraints can be established through service level agreements or product sponsorships. However, it is important to note that enforcing these constraints may result in a decrease in the platform's revenue. To address this issue, the platform can implement a fee structure based on the vendors' contributions to the revenue loss. Let S^* be an optimal solution for [AP](#) and let $(S_1^*, S_2^*, \dots, S_T^*)$ be an optimal solution for [APV](#). We denote the revenue

²We refer the reader to [Appendix A](#) for further discussion of assortment optimization under MNL.

loss due to the visibility constraints as

$$\Delta := T \cdot R(S^*) - \sum_{t=1}^T R(S_t^*). \quad (12)$$

A first naive approach. One approach is to allocate the loss based solely on the parameters ℓ_i . In this case, the proportion of the loss assigned to the vendor of product i would be determined by $\frac{\ell_i}{\sum_{j=1}^n \ell_j}$. However, this distribution would not be equitable in the sense that we should not impose any charges on a product that already belongs to the optimal set S^* , even if it satisfies the constraint $\ell_i > 0$. Moreover, this allocation fails to consider the impact of each product on the overall loss. For example, if there is a product with exceptionally high preference weight v_i but a significantly lower price p_i , while other products have higher prices and lower preference weights, enforcing the visibility of the first product would drive the revenue down, whereas the others would have a lesser impact. In this scenario, the former product should be responsible for covering almost the entire revenue loss.

Our approach. Let S_1^*, \dots, S_T^* be an optimal solution to APV. First observe that $R(S_t^*) = \frac{\sum_{i \in S_t^*} p_i v_i}{1 + \sum_{i \in S_t^*} v_i}$, implies that

$$R(S_t^*) = \sum_{i \in S_t^*} (p_i - R(S_t^*)) v_i$$

for every set S_t^* . This decomposition of the revenue gives us which products drive the revenue down (the ones with $p_i < R(S_t^*)$) and which products increase the revenue ($p_i > R(S_t^*)$). It also shows that the contribution of product i to the revenue is proportional to the difference between the price of product i and the actual revenue, as well as proportional to the preference weight v_i . We can then rewrite the total revenue of the assortments $(S_1^*, S_2^*, \dots, S_T^*)$ as

$$\sum_{t=1}^T R(S_t^*) = \sum_{t=1}^T \sum_{i \in S_t^*} (p_i - R(S_t^*)) v_i = \sum_{i=1}^n \sum_{t=1}^T \mathbb{1}(i \in S_t^*) (p_i - R(S_t^*)) v_i.$$

Therefore, we view the contribution of product $i \in \mathcal{N}$ to the total revenue as

$$C_i := \sum_{t=1}^T \mathbb{1}(i \in S_t^*) (p_i - R(S_t^*)) v_i.$$

Pricing the loss. For each product $i \in \mathcal{N}$, we propose to charge its vendor the fraction of the loss corresponding to the negative contribution of this product, divided by the sum of the negative contributions of all the products:

$$\Gamma_i := \frac{C_i^-}{\sum_{j=1}^n C_j^-} \cdot \Delta \quad (13)$$

where $x^- = \max(-x, 0)$ is the negative part of x , and recalling that Δ is the total loss in revenue due to enforcing visibility constraints.

Next, we discuss three important properties of this strategy:

(a) Fair distribution of fees. First, it is worth noting that the loss in revenue is exactly shared between the products whose contribution to the total revenue is negative as on one hand we have $\sum_{i \in \mathcal{N}} \Gamma_i = \Delta$, and on the other hand we have $\Gamma_i > 0$ if and only if $C_i^- > 0$, i.e., the product's contribution to the revenue is negative. Moreover, the fee Γ_i for each product takes into account its actual impact on the revenue: the first observation is that the products in S^* all have a nonnegative contribution, and their vendors are consequently exempt from a fee as expected. The second observation is that for any product with a negative contribution, its impact on the revenue is magnified by a lower price and a greater preference weight. It turns out that our policy does take this observation into account, as the lower/greater the price/weight of such a product, the greater the loss it incurs to the total revenue, and thus the greater the fee imposed on the vendor. The third observation is that even if a products $i \notin S^*$ and/or if there exists some $t \in [T]$ such that $i \in S_t^*$ and $p_i \leq R(S_t^*)$, such a product can still have a positive contribution, in which case $\Gamma_i = 0$, concurring with a product whose impact overall is positive. The final observation is that this pricing strategy guarantees a fair treatment of equal products. Indeed, if two distinct products i and j have identical prices and preference weights, then imposing similar visibility constraints for the two products implies a similar fee for the vendors.

(b) Monotonicity of the fee. For every product i , the fee Γ_i is nondecreasing as a function of ℓ_i . In other words, if a vendor wishes to display her product more often to customers, her fee gets higher. Let us show this intuitive result mathematically. Recall that an optimal solution to APV is given by $S_t^* = \overline{L_t \cup \dots \cup L_T}$ for $t \in [T]$, as shown in Theorem 3.7. Similarly, let \tilde{S}_t be the optimal solution given by Equations (6) when ℓ_i is increased by 1. Note that a direct consequence of the formula giving these optimal solution, we have $\tilde{S}_t = S_t^*$ for all $t \neq \ell_i + 1$, and $\tilde{S}_{\ell_i+1} = \overline{L_t \cup \dots \cup L_{\ell_i+1} \cup \{i\}}$. First, if $i \in S_{\ell_i+1}^*$, then nothing changes and we have $\tilde{C}_i = C_i$, where $\tilde{C}_i := \sum_{t=1}^T \mathbb{1}(i \in \tilde{S}_t)(p_i - R(\tilde{S}_t))v_i$ is product i 's contribution to the loss after increasing ℓ_i . Otherwise, $C_i \geq \tilde{C}_i$, and therefore $C_i^- \leq \tilde{C}_i^-$. Furthermore, for every $j \neq i$, since $R(S_{\ell_i+1}^*) \geq R(\tilde{S}_{\ell_i+1})$, we have $\tilde{C}_j \geq C_j$, since the term associated with $t = \ell_i + 1$ in the definition of C_j can only increase when ℓ_i is increased, while all the other terms remain unchanged. Letting $\tilde{\Delta}$ and $\tilde{\Gamma}_i$ be the total loss incurred and the fee imposed on vendor i respectively, after ℓ_i is increased by 1, we have

$$\begin{aligned} \tilde{\Gamma}_i &= \frac{\tilde{C}_i}{\tilde{C}_i + \sum_{j \neq i} \tilde{C}_j} \cdot \tilde{\Delta} \\ &\geq \frac{\tilde{C}_i}{\tilde{C}_i + \sum_{j \neq i} C_j} \cdot \Delta \\ &\geq \frac{C_i}{C_i + \sum_{j \neq i} C_j} \cdot \Delta \\ &= \Gamma_i. \end{aligned}$$

The third line follows from the fact that the map $x \mapsto x/c + x$ is nondecreasing on $[0, +\infty)$, for any $c > 0$.

(c) Computational tractability. The fees charged are easy to compute: each of them can be computed in polynomial time $O(nT)$ since they only require to solve APV. Additionally,

consider the situation where a vendor is interested in knowing the fee that she needs to pay in order to increase the visibility of the product by one customer. This corresponds to the situation where we change ℓ_i to $\ell_i + 1$ (without changing any other ℓ_j). In this case, we only have to recompute $S_{\ell_i+1}^*$, because the others assortments S_t^* stay the same, as explained in the previous paragraph. Therefore, for a fixed product i , can compute efficiently the value of the fee Γ_i for all the values of $\ell_i \in [T] \cup \{0\}$.

5.3 Numerical Study

To illustrate our theoretical contributions, we perform some numerical experiments on randomly generated data.

Our experimental setup. We fix the number of products equal to $n = 20$, and fix the number of customers to $T = 30$. We generate the weights v_i , prices p_i and visibility constraints ℓ_i from several distributions. Namely, to generate the weights v_i and prices p_i we used uniform and exponential distributions with varying parameters. For the visibility constraints ℓ_i , we ran our experiments using three distributions: we used a standard integer uniform $\mathcal{U}(\{0\} \cup [T])$, then an integer uniform $\ell_i \sim \mathcal{U}(\{\lfloor \frac{i}{n}T \rfloor, \dots, T\})$, and finally the constant $\ell_i = T$, for all $i \in \mathcal{N}$. Recalling that products are ordered by non-increasing price values, the intuition behind the second distribution is that products with lower prices are likelier to receive higher visibility constraints.

For each set of parameters, we generate 1000 samples, then we compute ratio of the optimal solution of **APV** divided by $T \cdot R(S^*)$, i.e., the unconstrained optimal expected revenue (without any visibility constraints), namely

$$\eta := \frac{\sum_{t=1}^T R(S_t^*)}{T \times R(S^*)}.$$

This captures the fraction of the revenue that we conserve once visibility constraints are enforced. We compute statistics on ratio η : mean, standard deviation, minimum and maximum values. ³

Overall results and analysis. Our results are gathered in Table 1.

As expected, the ratio η decreases when we give higher visibility ℓ_i to the products with smaller prices p_i . Furthermore, we observe that the higher the range of fluctuation of the weights v_i and prices p_i , the lower the revenue becomes once we enforce visibility constraints. This is rather intuitive, because we saw previously that the cases in which the visibility constraints can severely harm the revenue are when there are some products with extreme value of the price or weight compared to the others. However, we observe that extreme events when the revenue of **APV** becomes very low are really rare since the standard deviation is rather small compared to the range of fluctuation of η . Based on our results, it seems that in most cases, as long as there are no abnormally extreme values of the weights and prices, we can always hope to keep a fraction close to half of the revenue after enforcing visibility constraints.

³Initially, we made T vary, but we observed that the results were almost independent of T compared to the randomness of the generation of the parameters, so we kept only $T = 30$. Indeed, varying T just extends the stream of customers, and if the visibility constraints are extended by the same factor, it seems natural that the ratio of conserved revenue remains constant.

Table 1: For each distribution of p_i, v_i and ℓ_i , we compute the ratio η between the optimal assortment of **APV** and the optimal assortment of the unconstrained problem **AP** on 1000 samples. We provide its mean, standard deviation, minimum and maximum values.

Distribution of v_i and p_i	Distribution of ℓ_i	η (in %)			
		Mean	Standard deviation	Min	Max
$v_i \sim \mathcal{U}(0, 1), \quad p_i \sim \mathcal{U}(0, 1)$	$\mathcal{U}(\{0\} \cup [T])$	82.9	5.7	58.8	96.8
	$\mathcal{U}(\{\lfloor \frac{i}{n} T \rfloor, \dots, T\})$	73.7	6.4	51.9	91.3
	$\ell_i = T$	71.6	6.9	47.4	91.6
$v_i \sim \mathcal{U}(0, 10), \quad p_i \sim \mathcal{U}(0, 10)$	$\mathcal{U}(\{0\} \cup [T])$	71.0	7.3	43.9	89.2
	$\mathcal{U}(\{\lfloor \frac{i}{n} T \rfloor, \dots, T\})$	60.3	7.1	39.5	81.6
	$\ell_i = T$	58.1	7.6	33.0	81.0
$v_i \sim \mathcal{E}(1), \quad p_i \sim \mathcal{E}(1)$	$\mathcal{U}(\{0\} \cup [T])$	64.3	9.4	34.5	88.1
	$\mathcal{U}(\{\lfloor \frac{i}{n} T \rfloor, \dots, T\})$	50.9	9.4	24.8	87.6
	$\ell_i = T$	48.2	10.1	20.0	89.1
$v_i \sim \mathcal{E}(1/10), \quad p_i \sim \mathcal{E}(1/10)$	$\mathcal{U}(\{0\} \cup [T])$	49.4	11.2	16.1	79.1
	$\mathcal{U}(\{\lfloor \frac{i}{n} T \rfloor, \dots, T\})$	37.2	9.8	13.4	63.7
	$\ell_i = T$	34.7	9.4	11.3	67.1




Figure 1: Comparison between the ratios of revenue decrease and of sales increase as ℓ goes from 0 to T and $\ell_i = \ell, \forall i \in \mathcal{N}$.

Revenue vs market shares. To assess the influence of the visibility constraints on the market share, we compute the optimal value of **APV** for all visibility constraints $(\ell_i)_{1 \leq i \leq n}$ equal to a same value ℓ going from 0 to T (for $v_i \sim \mathcal{U}(0, 1), p_i \sim \mathcal{U}(0, 1)$). We then compute the expected sales associated to this expected revenue, and plot in Figure 1 the ratio of revenue decrease versus the ratio of sales increase by dividing them by their unconstrained value, for normalization.

It is interesting to see that, while the revenue decreases as we enforce more visibility constraints on products, on the contrary, the sales increase. Indeed, based on the structure of the optimal nested solution we identified, $S_t^* = \overline{L_t \cup \dots \cup L_T}$, we see that the sizes of the S_t^* increase when the ℓ_i 's increase. Furthermore, if we take a look at the sales optimization problem

$$\max_{S_1, \dots, S_T} \sum_{t=1}^T \frac{V(S_t)}{1 + V(S_t)},$$

that corresponds to $p_i = 1$ for all $i \in \mathcal{N}$, we can see that the more products we add to each S_t , the higher the sales since the function $x \mapsto \frac{x}{1+x}$ is increasing with respect to x .

Therefore, we can see that while visibility constraints decrease the revenue, they increase the sales, which can be interesting if the objective of the platform is to capture market shares.

Marginal study of one product. Finally, it is interesting to illustrate the marginal effect of the visibility constraints on one product. For this purpose, we take an instance of our generated data with $v_i \sim \mathcal{E}(1)$, $p_i \sim \mathcal{E}(1)$, $\ell_i \sim \mathcal{U}([0, T])$. Then, we select a particular $i \in \mathcal{N}$, and vary ℓ_i from 0 to T while the others ℓ_j stay fixed at their generated value ℓ_j . In Figure 2, we plot the variation of the revenue loss with respect to ℓ_i




Figure 2: Expected revenue loss as ℓ_i goes from 0 to T for a specific product i .

We remark that the loss is convex with respect to ℓ_i . Indeed, let $(S_t^*)_{1 \leq t \leq T} = (\overline{L_T \cup \dots \cup L_t})_{1 \leq t \leq T}$ be the nested solution for constraint ℓ_i , and $\ell_i \mapsto \Delta(\ell_i)$ the revenue loss function, as defined in (12), and considered as a function of constraint ℓ_i when all the other visibility constraints $(\ell_j)_{j \neq i}$ are fixed. We have

$$\Delta(\ell_i + 1) - \Delta(\ell_i) - (\Delta(\ell_i) - \Delta(\ell_i - 1)) = \overline{R}(S_{t+1}^* \cup \{i\}) + \overline{R}(S_t^*) - \overline{R}(S_t^* \cup \{i\}) - \overline{R}(S_{t+1}^*) \geq 0,$$

where the inequality holds by supermodularity and $S_t^* \subseteq S_{t+1}^*$, which shows $\ell_i \mapsto \Delta(\ell_i)$ is convex.

We can then, in the case where all the revenue from the sales of a product comes to its vendor, compute the expected revenue of product i versus the fee its vendor will pay as the visibility constraint ℓ_i increases. The expected revenue of product i is given by: $\sum_{t=1}^T \mathbb{1}(i \in S_t^*) p_i \frac{v_i}{1+V(S_t^*)}$, while the fee charged Γ_i is defined in (13). This is depicted in Figure 3.

For this particular product, we can see that enforcing a visibility constraint ℓ_i is profitable until $\ell_i = 17$ (with respect to $T = 30$), and becomes non profitable after. The profit associated with this constraint is the difference between the revenue and the fee, and is maximized for the value $\ell_i = 7$ (see Figure 3). This is an example of a product that does not belong to the optimal unconstrained set S^* (since the revenue is 0 when $\ell_i = 0$, it means the product was not included), but for which adding a visibility constraint allows to boost sales and the revenue of the vendor, while paying a small fee to compensate the platform's revenue loss.

One interesting observation here is that the fee paid by the vendor of product i (0.8 when $\ell_i = T = 30$ on Graph 3) is higher than the loss caused by the introduction of $\ell_i = T = 30$ while the others products were enforced previously (loss of 0.5 as we can read on Graph 2). However,




Figure 3: Expected revenue of product i versus fee paid by product i as ℓ_i goes from 0 to T .

this can be explained: as the expanded revenue function is supermodular and decreasing, the revenue decreases most when the first products are enforced, and then decreases less. Therefore, enforcing visibility constraints on product i decreases less the revenue if we already have some visibility constraints on the other products than if i is the first product on which the constraints are enforced. However, our pricing policy Γ_i does not take into account such sequential order of arrival, and prices each product globally, so that two products with the same price and weight pay the same fee for identical visibility constraints.

6 Conclusions and future directions

In this paper, we introduced the problem of assortment optimization with visibility constraints [APV](#), motivated by situations in e-commerce and online advertising in which a platform aims at ensuring a minimal exposure for some or all products. We prove that this problem can be solved in polynomial time, and devise an efficient algorithm to compute an optimal solution. Our algorithm leverages the supermodularity of an altered version of the expected revenue function, that allows us to identify the nested structure of an optimal solution. We also consider an extension of the problem with cardinality constraints on the assortments offered. We prove that the problem becomes strongly NP-hard even under uniform prices, and that in particular, it admits no FPTAS, unless $P = NP$. We then devise a PTAS for the special case of equal prices. Finally, we evaluate the revenue loss caused by the visibility constraints enforced, and propose a fair pricing strategy to charge each vendor a fee proportional to the contribution of its product to the revenue loss. Finally,

Future directions. A promising future research direction involves developing an approximation algorithm for the assortment optimization problem with visibility and cardinality constraints, considering general prices. Furthermore, exploring the assortment problem with visibility using alternative choice models such as the Markov Chain choice model or Nested Logit represents another fruitful avenue for investigation.

References

- Aydın Alptekinoglu and John H Semple. The exponential choice model: A new alternative for assortment and price optimization. *Operations Research*, 64(1):79–93, 2016. 5
- Qinyi Chen, Negin Golrezaei, Fransisca Susan, and Edy Baskoro. Fair assortment planning. *arXiv preprint arXiv:2208.07341*, 2022. 4
- Wang Chi Cheung and David Simchi-Levi. Thompson sampling for online personalized assortment optimization problems with multinomial logit choice models. *Available at SSRN 3075658*, 2017. 5
- James Davis, Guillermo Gallego, and Huseyin Topaloglu. Assortment planning under the multinomial logit model with totally unimodular constraint structures. *Technical report, Columbia University*, 2013. 2, 4
- James M Davis, Huseyin Topaloglu, and David P Williamson. Assortment optimization over time. *Operations Research Letters*, 43(6):608–611, 2015. 5
- Antoine Désir, Vineet Goyal, Danny Segev, and Chun Ye. Constrained assortment optimization under the markov chain-based choice model. *Management Science*, 66(2):698–721, 2020. 2
- Antoine Désir, Vineet Goyal, and Jiawei Zhang. Capacitated assortment optimization: Hardness and approximation. *Operations Research*, 70(2):893–904, 2022. 4
- Benjamin Doerr. Probabilistic tools for the analysis of randomized optimization heuristics. *Theory of evolutionary computation: Recent developments in discrete optimization*, pages 1–87, 2020. 20
- Omar El Housni and Huseyin Topaloglu. Joint assortment optimization and customization under a mixture of multinomial logit models: On the value of personalized assortments. *Operations Research*, 2022. 4
- Omar El Housni, Marouane Ibn Brahim, and Danny Segev. Maximum load assortment optimization: Approximation algorithms and adaptivity gaps. *arXiv preprint arXiv:2309.01772*, 2023. 4
- Guillermo Gallego and Huseyin Topaloglu. Constrained assortment optimization for the nested logit model. *Management Science*, 60(10):2583–2601, 2014. 2
- Guillermo Gallego, Richard Ratliff, and Sergey Shebalov. A general attraction model and sales-based linear program for network revenue management under customer choice. *Operations Research*, 63(1):212–232, 2015. 4, 12
- Rajiv Gandhi, Samir Khuller, Srinivasan Parthasarathy, and Aravind Srinivasan. Dependent rounding in bipartite graphs. In *The 43rd Annual IEEE Symposium on Foundations of Computer Science, 2002. Proceedings.*, pages 323–332. IEEE, 2002. 17, 20

- Pin Gao, Yuhang Ma, Ningyuan Chen, Guillermo Gallego, Anran Li, Paat Rusmevichientong, and Huseyin Topaloglu. Assortment optimization and pricing under the multinomial logit model with impatient customers: Sequential recommendation and selection. *Operations Research*, 69(5):1509–1532, 2021. [4](#)
- Michael R Garey and David S. Johnson. Complexity results for multiprocessor scheduling under resource constraints. *SIAM journal on Computing*, 4(4):397–411, 1975. [36](#)
- Michael R Garey and David S Johnson. “strong” np-completeness results: Motivation, examples, and implications. *Journal of the ACM (JACM)*, 25(3):499–508, 1978. [36](#)
- Jerry Hausman and Daniel McFadden. Specification tests for the multinomial logit model. *Econometrica*, 52(5):1219–40, 1984. [4](#)
- Zhaolin Li. A single-period assortment optimization model. *Production and Operations Management*, 16:369–380, 2009. [5](#)
- Wentao Lu, Ozge Sahin, and Ruxian Wang. A simple way towards fair assortment planning: Algorithms and welfare implications. *Available at SSRN 4514495*, 2023. [4](#)
- R Duncan Luce. Individual choice behavior. *John Wiley*, 1959. [4](#)
- Siddharth Mahajan and Garrett Van Ryzin. Stocking retail assortments under dynamic consumer substitution. *Operations Research*, 49(3):334–351, 2001. [4](#)
- Daniel McFadden. *Conditional logit analysis of qualitative choice behavior*. Institute of Urban and Regional Development, University of California, 1973. [4](#)
- Sentao Miao and Xiuli Chao. Dynamic joint assortment and pricing optimization with demand learning. *Manufacturing & Service Operations Management*, 23(2):525–545, 2021. [5](#)
- Paat Rusmevichientong, Zuo-Jun Max Shen, and David B Shmoys. Dynamic assortment optimization with a multinomial logit choice model and capacity constraint. *Operations research*, 58(6):1666–1680, 2010. [4](#)
- Paat Rusmevichientong, David Shmoys, Chaoxu Tong, and Huseyin Topaloglu. Assortment optimization under the multinomial logit model with random choice parameters. *Production and Operations Management*, 23(11):2023–2039, 2014. [4](#)
- Danny Segev and Sahil Singla. Efficient approximation schemes for stochastic probing and prophet problems. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pages 793–794, 2021. [19](#)
- Mika Sumida, Guillermo Gallego, Paat Rusmevichientong, Huseyin Topaloglu, and James Davis. Revenue-utility tradeoff in assortment optimization under the multinomial logit model with totally unimodular constraints. *Management Science*, 67(5):2845–2869, 2021. [4](#)
- Kalyan Talluri and Garrett Van Ryzin. Revenue management under a general discrete choice model of consumer behavior. *Management Science*, 50(1):15–33, 2004. [1](#), [4](#), [26](#), [35](#)

Chenhao Wang, Yao Wang, and Shaojie Tang. When advertising meets assortment planning: Joint advertising and assortment optimization under multinomial logit model. *Available at SSRN 3908616*, 2021. 4

Ruxian Wang. Capacitated assortment and price optimization under the multinomial logit model. *Operations Research Letters*, 40(6):492–497, 2012. 5

A Revenue Maximization Under MNL

In this appendix, we present for completeness two standard known results about assortment optimization under MNL. Lemma A.1, shows that the optimal assortment for AP is revenue-ordered. This is a standard result in assortment optimization under MNL (Proposition 6 Talluri and Van Ryzin, 2004, 25). Lemma A.2 provides a necessary and sufficient conditions under which adding a given product to a given assortment increases the overall revenue. Consider the classic unconstrained revenue maximization problem under MNL, defined as follows

$$\max_{S \subseteq \mathcal{N}} R(S). \quad (\text{AP})$$

LEMMA A.1. *Let $R^* = \max_{S \subseteq \mathcal{N}} R(S)$ be the optimal value of AP. There exists a revenue ordered optimal solution to AP, given by $S^* = \{i \in \mathcal{N}, p_i \geq R^*\}$.*

Before proving Lemma A.1, we state and prove the following more general lemma, which shows that adding a product i to any given assortment S increases the objective if and only if the price of the added product i is greater than or equal to the original revenue $R(S)$.

LEMMA A.2. *For any assortment $S \subseteq \mathcal{N}$ and any product $k \in \mathcal{N} \setminus S$, the three following propositions are equivalent:*

$$(i) R(S \cup \{k\}) \geq R(S), \quad (ii) p_k \geq R(S), \quad (iii) p_k \geq R(S \cup \{k\}).$$

Proof of Lemma A.2. We simply show that $R(S \cup \{k\})$ is a convex combination of $R(S)$ and p_k . In fact, simple algebra gives

$$R(S \cup \{k\}) = \sum_{i \in S} p_i \cdot \phi(i, S \cup \{k\}) + p_k \cdot \phi(k, S \cup \{k\}) = R(S) \cdot \alpha + p_k \cdot (1 - \alpha),$$

where $\alpha = (1 + V(S))/(1 + V(S \cup \{k\}))$. It is easy to see that $\alpha \in (0, 1)$, which proves that $R(S \cup \{k\})$ is a convex combination of $R(S)$ and p_k . Consequently, $R(S \cup \{k\})$ belongs to the closed interval whose extremities are $R(S)$ and p_k . In particular, we have $R(S \cup \{k\}) \geq R(S)$ if and only if $p_k \geq R(S)$, and if and only if $p_k \geq R(S \cup \{k\})$. ■

Proof of Lemma A.1. Let \tilde{S} be the optimal assortment of AP with maximal cardinality. In the case of ties, let \tilde{S} be any arbitrary such assortment. In the following, we show that $S^* = \tilde{S}$ (which in particular implies that there can be no ties). Let $i \in S^*$, and assume by contradiction that $i \notin \tilde{S}$. On one hand, we know by optimality of \tilde{S} that $R(\tilde{S}) = R^*$. On the other hand,

we know that $p_i \geq R^*$ by definition of S^* . Therefore, $p_i \geq R(\tilde{S})$ and we can apply lemma A.2, which implies that $R(\tilde{S} \cup \{i\}) \geq R(\tilde{S})$, and hence that $\tilde{S} \cup \{i\}$ is also an optimal assortment. Since $\tilde{S} \subsetneq \tilde{S} \cup \{i\}$. This contradicts the definition of \tilde{S} as the optimal assortment with maximal cardinality, and therefore shows that $i \in \tilde{S}$. Now, let $i \in \tilde{S}$. Assume by contradiction that $i \notin S^*$, i.e., $p_i < R^* = R(\tilde{S})$. In particular, this implies using Lemma A.2 that $R(\tilde{S} \setminus \{i\}) > R(\tilde{S})$, which contradicts the optimality of \tilde{S} , and therefore shows that $i \in S^*$. This proves the second inclusion, and thereby completes the proof of the lemma. ■

B Additional Proofs from Section 4

B.1 Proof of Theorem 4.1

In order to show the result, we reduce the 3-partition problem to APVC. In the 3-PARTITION problem, the input consists on an integer T , and a set of $3T$ integers $\mathcal{A} = \{a_1, \dots, a_{3T}\}$ such that $\sum_{i=1}^{3T} a_i = BT$ for some $B \geq 0$, and seeks to determine whether \mathcal{A} can be partitioned into T triplets A_1, \dots, A_T , and such that the sum of elements of each triplet is B . This problem is known to be strongly NP-hard (see Garey and Johnson (1975) for example, 3-PARTITION is referred to there as $P[3, 1]$).

Given an arbitrary instance of 3-PARTITION, we construct the following instance of APVC. We consider $3T$ products, simply denoted by the universe $\mathcal{N} = \{1, \dots, 3T\}$, and such that for all $i \in \mathcal{N}$, $v_i = a_i$, $p_i = 1$ and $\ell_i = 1$. In particular, note that $V(\mathcal{N}) = BT$. Finally, let $k = 3$ in our instance, the upper bound on the cardinality of any offered assortment. A first important observation is that any feasible solution for this instance of APVC shows each product to exactly one customer. In other words, if S_1, \dots, S_T is a feasible solution for this instance of APVC, then for any product $i \in \mathcal{N}$, there exists a unique $t \in [T]$ such that $i \in S_t$. Indeed, on one hand, the visibility constraints impose that each product be shown at least once, and on the other hand, since each assortment can contain at most $k = 3$ products and there are $3T$ products in our universe, this implies that each product must be shown at most once, which shows the observation. A natural consequence of this observation is that for any feasible solution S_1, \dots, S_T of our instance of APVC, we have $\sum_{t=1}^T V(S_t) = V(\mathcal{N}) = BT$. We denote by $S^* = (S_1^*, \dots, S_T^*)$ an optimal sequence of assortments for this instance of APVC. In order to show the strong NP-hardness, we introduce the following claim.

CLAIM B.1. *The answer to the 3-PARTITION is yes if and only if $\sum_{t=1}^T R(S_t^*) = T \cdot B / (1 + B)$.*

Before proving the claim, note that it directly implies the strong NP-hardness of APVC, even in the special case of equal prices and integer preference-weights, as it shows that 3-PARTITION reduces to this special case of APVC (see (Observation 5 Garey and Johnson, 1978, 504)). Let us now prove the claim.

Proof of Claim B.1.

Direct implication. Suppose that a valid 3-PARTITION exists, which we denote by A_1, \dots, A_T ,

and let $S_t = \{i : a_i \in A_t\}$ for $t \in [T]$. On one hand, we have

$$\sum_{t=1}^T R(S_t^*) \geq \sum_{t=1}^T R(S_t) = T \cdot \frac{B}{1+B},$$

and on the other hand, we have using Jensen's inequality and the concavity of the map $x \mapsto x/(1+x)$,

$$\sum_{t=1}^T R(S_t^*) = T \sum_{t=1}^T \frac{1}{T} \cdot \frac{V(S_t^*)}{1+V(S_t^*)} \leq T \cdot \frac{\sum_{t=1}^T \frac{V(S_t^*)}{T}}{1 + \sum_{t=1}^T \frac{V(S_t^*)}{T}} = T \cdot \frac{B}{1+B}.$$

This implies that $\sum_{t=1}^T R(S_t^*) = T \cdot B/(1+B)$ and proves the direct implication.

Indirect implication. Assume that

$$\sum_{t=1}^T R(S_t^*) = T \cdot \frac{B}{1+B}.$$

On one hand, we have according to Cauchy-Schwarz inequality,

$$\sum_{t=1}^T (1+V(S_t^*)) \cdot \sum_{t=1}^T \frac{1}{1+V(S_t^*)} \geq T^2,$$

with equality if and only if there exists some λ such that for all $t \in [T]$, $1+V(S_t^*) = \lambda/(1+V(S_t^*))$, which in turn is equivalent to $V(S_1^*) = V(S_2^*) = \dots = V(S_T^*) = \sqrt{\lambda} - 1$.

On the other hand, it is easy to see that

$$\sum_{t=1}^T R(S_t^*) = T - \sum_{t=1}^T \frac{1}{1+V(S_t^*)}.$$

Therefore,

$$\begin{aligned} \sum_{t=1}^T (1+V(S_t^*)) \cdot \sum_{t=1}^T \frac{1}{1+V(S_t^*)} &= (T + V(\mathcal{N})) \cdot \left(T - \sum_{t=1}^T R(S_t^*) \right) \\ &= (T + TB) \cdot \left(T - T \cdot \frac{B}{1+B} \right) \\ &= T^2, \end{aligned}$$

This proves the equality case in Cauchy-Schwarz inequality, and consequently that $V(S_1^*) = V(S_2^*) = \dots = V(S_T^*)$. Finally, taking $A_t = \{a_i : i \in S_t^*\}$ for $t \in [T]$, it is easy to see that A_1, \dots, A_T is a valid 3-PARTITION, which concludes the proof of the claim. \blacksquare

In the remainder of this proof, we show that the problem does not admit an FPTAS, relying on the following claim.

CLAIM B.2. *If the answer to the 3-PARTITION is no, then*

$$\sum_{t=1}^T R(S_t^*) \leq (1 - \delta) \cdot T \cdot \frac{B}{1 + B},$$

where $\delta = \frac{2}{TB^2(B+2)}$.

Finally, the non-existence of an FPTAS for this problem follows naturally. Indeed, assume for the sake of contradiction that an FPTAS exists for APV. Then by taking $\epsilon < \delta$, say $\epsilon = \delta/2$ for example, the FPTAS would return a solution (S_1^*, \dots, S_T^*) with an objective greater than or equal to $(1 - \epsilon) \cdot TB/(1 + B)$ if and only if a valid partition exists. This yields an algorithm that solves the 3-PARTITION problem in pseudo-polynomial time, thereby contradicting its strong NP-hardness.

Conclusion. APV is strongly NP-hard, and admits no FPTAS.

Proof of Claim B.2. Suppose the answer to the 3-PARTITION problem is *no*. Consider the following integer program:

$$\begin{aligned} \max_{\mathbf{b} \in \mathbb{N}^T} \quad & \sum_{t=1}^T \frac{b_t}{1 + b_t} \\ \text{s.t.} \quad & \sum_{t=1}^T b_t = BT \\ & \mathbf{b} \neq (B, B, \dots, B). \end{aligned} \tag{14}$$

Intuitively, this program's objective is to partition a budget of BT into T bins, under the constraint that not all the bins receive the same fraction of budget, and such that each bin $t \in [T]$ yields a return of $b_t/(1 + b_t)$. Let \mathbf{b}^* be the optimal solution of this problem. The idea of the proof is to show that:

$$\sum_{t=1}^T R(S_t^*) \leq \sum_{t=1}^T \frac{b_t^*}{1 + b_t^*} = (1 - \delta) \cdot \frac{TB}{1 + B}. \tag{15}$$

The inequality. Let $\mathbf{b} = (V(S_1^*), \dots, V(S_T^*))$. We have $\sum_{t=1}^T V(S_t^*) = BT$, and since a valid partition does not exist, then $\mathbf{b} \neq (B, B, \dots, B)$ (see Claim B.1). Therefore \mathbf{b} is feasible for (14), and we have:

$$\sum_{t=1}^T R(S_t^*) = \sum_{t=1}^T \frac{b_t}{1 + b_t} \leq \sum_{t=1}^T \frac{b_t^*}{1 + b_t^*},$$

by the optimality of \mathbf{b}^* .

The equality. We show that $(B+1, B-1, B, \dots, B)$ is an optimal solution for (14). Towards this purpose, we start by showing that the number of indices u for which $b_u^* \neq B$ is at least two, before showing by contradiction that it is exactly 2, then we show that the entries of these two indices in \mathbf{b}^* are exactly $B+1$ and $B-1$. Since the indices are all interchangeable, this shows that $(B+1, B-1, B, \dots, B)$ is an optimal solution for (14).

Let $\mathcal{U} \subseteq [T]$ be the set of indices u such that $b_u^* \neq B$, i.e., $\mathcal{U} = \{u \in [T] : b_u^* \neq B\}$. First, $\mathcal{U} \neq \emptyset$, by the second constraint of (14). Also, if we suppose by contradiction that $|\mathcal{U}| = 1$, say $\mathcal{U} = \{1\}$ without loss of generality, then $\sum_{t=1}^T b_t^* = b_1^* + (T-1)B \neq BT$, which contradicts the first constraint of (14). Therefore, $|\mathcal{U}| \geq 2$.

Suppose by contradiction that $|\mathcal{U}| > 2$. Since the elements of \mathcal{U} cannot all be smaller, or all be greater than B (otherwise the second constraint in (14) does not hold), there exists $u_1, u_2 \in [T]$ such that $b_{u_1}^* \geq B+1$ and $b_{u_2}^* \leq B-1$. Also, since $|\mathcal{U}| > 2$, there exists some $u_3 \neq u_1, u_2$, such that $b_{u_3}^* \neq B$. For simplicity, we assume without loss of generality that $u_1 = 1$, $u_2 = 2$ and $u_3 = 3$. Let $\tilde{\mathbf{b}} = (b_1^* - 1, b_2^* + 1, b_3^*, \dots, b_T^*)$, in other words, we transfer a unit of the budget from b_1^* to b_2^* . First, $\sum_{t=1}^T \tilde{b}_t = \sum_{t=1}^T b_t^* = BT$. Moreover, since $3 \in \mathcal{U}$, $\tilde{b}_3 = b_3^* \neq B$, which implies that $\tilde{\mathbf{b}}$ is feasible for (14). Finally, we have

$$\begin{aligned} \sum_{t=1}^T \frac{\tilde{b}_t}{1 + \tilde{b}_t} &= \frac{b_1^* - 1}{1 + b_1^* - 1} + \frac{b_2^* + 1}{1 + b_2^* + 1} + \sum_{t=3}^T \frac{b_t^*}{1 + b_t^*} \\ &> \frac{b_1^*}{1 + b_1^*} + \frac{b_2^*}{1 + b_2^*} + \sum_{t=3}^T \frac{b_t^*}{1 + b_t^*} \\ &= \sum_{t=1}^T \frac{b_t^*}{1 + b_t^*}, \end{aligned}$$

where the inequality in the second line is due to the fact that the map

$$f : x \mapsto \frac{b_1^* - x}{1 + b_1^* - x} + \frac{b_2^* + x}{1 + b_2^* + x}$$

is monotonically (strictly) increasing on $[0, 1]$ (which we show at the end of the proof for completeness). This contradicts the optimality of \mathbf{b}^* , and thereby shows that $|\mathcal{U}| = 2$.

Assume without loss of generality that $\mathcal{U} = \{1, 2\}$, and that $b_1^* = B + \alpha$, and $b_2^* = B - \alpha$ for some $\alpha \in \mathbb{N} \setminus \{0\}$. Assume by contradiction that $\alpha \geq 2$, then using the same argument as the contradiction above, the map

$$x \mapsto \frac{b_1^* - x}{1 + b_1^* - x} + \frac{b_2^* + x}{1 + b_2^* + x}$$

is increasing, and we can therefore transfer a unit of the budget from 1 to 2, and strictly increase the objective of the obtained solution. Moreover, since $\alpha \geq 2$, $b_1^* - 1 \neq B$, and the second constraint in (14) is valid, which proves the feasibility of the obtained solution, and contradicts the optimality of \mathbf{b}^* . Therefore, $\alpha = 1$ and $(B+1, B-1, B, \dots, B)$ is an optimal solution for (14).

Going back to the sought equality (15), we have

$$\begin{aligned} \sum_{t=1}^T \frac{b_t^*}{1 + b_t^*} &= \frac{B-1}{1+B-1} + \frac{B+1}{1+B+1} + (T-2) \cdot \frac{B}{1+B} \\ &= \frac{TB}{1+B} \cdot (1 - \delta). \end{aligned}$$

The second equality follows from basic algebra operations. This proves Equation (15), which in turn proves the claim.

Leftover: Proof of the strict monotonicity of f . For all $x \in [0, 1]$, recall that

$$f(x) = \frac{b_1^* - x}{1 + b_1^* - x} + \frac{b_2^* + x}{1 + b_2^* + x}.$$

We differentiate f , and we have for all $x \in [0, 1]$

$$f'(x) = \frac{1}{(1 + b_2^* + x)^2} - \frac{1}{(1 + b_1^* + x)^2}$$

Therefore, it is easy to see that $f'(x) > 0$ for all $x < (b_1^* - b_2^*)/2$. Since $b_1^* \geq B+1$ and $b_2^* \leq B-1$, then $(b_1^* - b_2^*)/2 \leq 1$, which proves that f is monotonically strictly increasing on $[0, 1]$. ■

B.2 Proof of Lemma 4.4

In this proof, we show separate inequalities for light, medium, and heavy customers. Combining these inequalities then gives the desired result.

Heavy customers: If $t \in [T]$ is heavy then by the fifth constraint of [ILP](#), we have

$$\sum_{i \in \mathcal{N}} v_i x_{it}^* \geq \frac{1}{\epsilon}.$$

Therefore,

$$\frac{\sum_{i \in \mathcal{N}} v_i x_{it}^*}{1 + \sum_{i \in \mathcal{N}} v_i x_{it}^*} \geq \frac{1}{1 + \epsilon} \geq (1 - \epsilon) \cdot \frac{v(S_t^*)}{1 + v(S_t^*)}. \quad (16)$$

Medium customers: If $t \in G_\ell$ for some $\ell \in \{1, \dots, L\}$, then we have

$$\sum_{i \in \mathcal{N}} v_i x_{it} \geq \epsilon \cdot (1 + \epsilon)^{\ell-1} \geq \frac{1}{1 + \epsilon} v(S_t^*),$$

where the first inequality follows from the fifth constraint of [ILP](#), and the second follows from the fact that $v(S_t^*) \leq \epsilon \cdot (1 + \epsilon)^\ell$ by definition of customers of type ℓ . Therefore,

$$\frac{\sum_{i \in \mathcal{N}} v_i x_{it}^*}{1 + \sum_{i \in \mathcal{N}} v_i x_{it}^*} \geq \frac{1}{1 + \epsilon} \cdot \frac{v(S_t^*)}{1 + v(S_t^*)} = (1 - \epsilon) \cdot \frac{v(S_t^*)}{1 + v(S_t^*)}. \quad (17)$$

Light customers: Let \mathbf{X}^* be the solution defined as follows, for every $t \in [T]$, and every $i \in \mathcal{N}$, $X_{it}^* = \mathbb{1}\{i \in S_t^*\}$. It is easy to see that by construction of [ILP](#), the solution \mathbf{X}^* is feasible for [ILP](#), and therefore, it is also feasible for its linear relaxation. Therefore, since \mathbf{x}^* is optimal for [ILP](#), we have

$$\sum_{t \in G_{light}} \sum_{i \in \mathcal{N}} v_i x_{it}^* \geq \sum_{t \in G_{light}} \sum_{i \in \mathcal{N}} v_i X_{it}^* = \sum_{t \in G_{light}} v(S_t^*).$$

We also have for all $t \in G_{light}$, $\sum_{i \in \mathcal{N}} v_i x_{it}^* \leq \epsilon$ by the sixth constraint of [ILP](#). Therefore,

$$\begin{aligned}
\sum_{t \in G_{light}} \frac{\sum_{i \in \mathcal{N}} v_i x_{it}^*}{1 + \sum_{i \in \mathcal{N}} v_i x_{it}^*} &\geq \frac{1}{1 + \epsilon} \cdot \sum_{t \in G_{light}} \sum_{i \in \mathcal{N}} v_i x_{it}^* \\
&\geq \sum_{t \in G_{light}} v(S_t^*) \\
&\geq \sum_{t \in G_{light}} \frac{v(S_t^*)}{1 + v(S_t^*)}.
\end{aligned} \tag{18}$$

We conclude the proof by combining Equations [\(16\)](#), [\(17\)](#) and [\(18\)](#), as we have

$$\sum_{t=1}^T \frac{\sum_{i \in \mathcal{N}} v_i x_{it}^*}{1 + \sum_{i \in \mathcal{N}} v_i x_{it}^*} \geq (1 - \epsilon) \cdot \sum_{t=1}^T \frac{v(S_t^*)}{1 + v(S_t^*)}.$$

B.3 Proof of Claim [4.7](#)

We have

$$\begin{aligned}
W_t^l &= \sum_{q=1}^Q \sum_{i \in D_q} v_i X_{it} \\
&= \sum_{q=1}^Q \epsilon^5 \cdot (1 + \epsilon)^{q-1} \sum_{i \in D_q} X_{it} \\
&= \sum_{q=1}^Q \epsilon^5 \cdot (1 + \epsilon)^{q-1} \sum_{i \in D_q} x_{it}^* \\
&= \sum_{q=1}^Q \sum_{i \in D_q} v_i x_{it}^* \\
&= w_t^l,
\end{aligned}$$

where the second equality follows from the definition of D_q , and the third inequality follows from the degree preservation property [\(P2\)](#).

B.4 Proof of Claim [4.8](#)

First, note that the proof of Claim [4.7](#) holds even for the case where t is a light customer. Therefore, we can decompose $v(S_t)$ into a sum of a random variable W_t^s representing the contribution of products in D_0 , and a deterministic random variable $W_t^l \stackrel{a.s.}{=} w_t^l$ (which will be treated as a constant from this point on), i.e.,

$$v(S_t) = W_t^s + w_t^l.$$

For simplicity of notation, and only in this proof, we use the notation $A := W_t^s$ and $b := w_t^l$. This proof consists on two steps. In the first step, we lower bound the expectation of $v(S_t)/(1+v(S_t))$ with $\sum_{i \in \mathcal{N}} v_i x_{it}^*$, up to a constant $\alpha < 1$, mainly using a convexity argument. In other words,

we provide a constant α such that

$$v(S_t)/(1 + v(S_t)) \geq \alpha \cdot \sum_{i \in \mathcal{N}} v_i x_{it}^*.$$

Then, in the second step, we show that $\alpha = (1 - O(\epsilon))$.

Step 1: Let \mathcal{A} be the support of the random variable A , i.e.,

$$\mathcal{A} = \{a \in \mathbb{R} : \mathbb{P}[A = a] > 0\}.$$

We have

$$\begin{aligned} \mathbb{E} \left[\frac{v(S_t)}{1 + v(S_t)} \right] &= \mathbb{E} \left[\frac{A + b}{1 + A + b} \right] \\ &= \sum_{a \in \mathcal{A}} \frac{a + b}{1 + a + b} \mathbb{P}[A = a] \\ &= \left(\sum_{a \in \mathcal{A}} \frac{1}{1 + a + b} \cdot \frac{(a + b) \cdot \mathbb{P}[A = a]}{b + \mathbb{E}[A]} \right) \cdot (b + \mathbb{E}[A]). \end{aligned}$$

For each $a \in \mathcal{A}$, let $f(a) = \frac{1}{1+a+b}$ and let $z_a = \frac{(a+b) \cdot \mathbb{P}[A=a]}{b + \mathbb{E}[A]}$. It is easy to see that $\sum_{a \in \mathcal{A}} z_a = 1$, and that f is a convex function on $(0, +\infty)$. Therefore, using Jensen's inequality, and recalling that $\mathbb{E}[A] + b = \mathbb{E}[v(S_t)]$ we have

$$\mathbb{E} \left[\frac{v(S_t)}{1 + v(S_t)} \right] \geq f \left(\sum_{a \in \mathcal{A}} z_a \cdot y \right) \cdot \mathbb{E}[v(S_t)] = \alpha \cdot \mathbb{E}[v(S_t)],$$

where $\alpha = f \left(\sum_{a \in \mathcal{A}} z_a \cdot a \right)$.

Step 2. Let us now show that $\alpha = 1 - O(\epsilon)$. We have

$$\begin{aligned} \alpha &= \frac{1}{1 + b + \frac{\sum_{a \in \mathcal{A}} (a+b) \cdot \mathbb{P}[A=a] \cdot a}{b + \mathbb{E}[A]}} \\ &= \frac{1}{1 + b + \frac{b\mathbb{E}[A] + \mathbb{E}[A^2]}{b + \mathbb{E}[A]}} \\ &= \frac{1}{1 + b + \mathbb{E}[A] + \frac{\text{Var}[A]}{b + \mathbb{E}[A]}}. \end{aligned} \tag{19}$$

In order to show the desired property, we need to bound the denominator in Equation (19). To this purpose we have

$$\begin{aligned} \text{Var}[A] &= \text{Var} \left[\sum_{i \in D_0} v_i X_{it} \right] \\ &\leq \sum_{i \in D_0} v_i^2 \cdot \text{Var}[X_{it}] \end{aligned}$$

$$\begin{aligned}
&\leq \epsilon^5 \cdot \sum_{i \in D_0} v_i \cdot \mathbb{E}[X_{it}] \\
&= \epsilon^5 \cdot \mathbb{E}[A].
\end{aligned}$$

Therefore, by replacing in Equation (19), we have

$$\alpha \geq \frac{1}{1 + \epsilon + \epsilon^5} \geq \frac{1}{1 + 2\epsilon} \geq 1 - 2\epsilon.$$

Finally, combining steps 1 and 2, we have

$$\frac{v(S_t)}{1 + v(S_t)} \geq \alpha \cdot \mathbb{E}[v(S_t)] \geq (1 - 2\epsilon) \cdot \sum_{i \in \mathcal{N}} v_i x_{it}^* \geq (1 - 2\epsilon) \cdot \frac{\sum_{i \in \mathcal{N}} v_i x_{it}^*}{1 + \sum_{i \in \mathcal{N}} v_i x_{it}^*}.$$