Two-Sided Assortment Optimization

600

ABSTRACT

We study the two-sided assortment problem recently introduced by Rios, Saban and Zheng [17]. A platform must choose an assortment of profiles to show to each user on each side of the market in each stage. Users can either like/dislike as many profiles as they want, and a match occurs if two users see and like each other. The goal of the platform is to maximize the expected number of matches generated. Given the complexity of the problem, we focus on the one look-ahead version of the problem (i.e., two stages), and we provide performance guarantees for different variants of the problem. Specifically, we show that if users cannot see each other simultaneously there is an approximation guarantee of $1-1/e$ obtained using submodular optimization techniques. For the more challenging case where simultaneous shows are allowed, we provide an approximation guarantee of $(1-1/e)/24$. Finally, using data from our industry partner (a dating app in the US), we numerically show that our proposed heuristic outperforms several relevant benchmarks.

KEYWORDS

Matching, Assortment Optimization, Dating, Online platforms

ACM Reference Format:

600. 2022. Two-Sided Assortment Optimization. In *Proc.of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), Auckland, New Zealand, May 9–13, 2022*, IFAAMAS, 9 pages.

1 INTRODUCTION

A common feature of many marketplaces is their two-sided design, in which (1) both sides of the market have preferences, (2) both sides of the market can initiate an interaction with the other side, and (3) both users must mutually agree to generate a transaction. Examples include freelancing platforms such as TaskRabbit or UpWork, ride-sharing apps such as Blablacar, accommodation companies such as Airbnb, and dating platforms such as Hinge and Bumble, among others. In many of these platforms, the path towards a transaction starts with the platform eliciting preferences on both sides of the market. For instance, Airbnb requires guests to report the location and dates of their trips, and allows them to add filters regarding the price, the type of place, and so on. Similarly, Airbnb allows hosts to set the price and the minimum number of nights, as well as other preferences. After collecting this information, most of these platforms display a limited set of alternatives that agents can screen before starting their interaction with the other side of the market. Depending on the setting, this interaction can be one-directional, with one side of the market sending an initial request/message/like and the other side of the market either accepting or rejecting it; or it can be bi-directional, with both sides of the market being able to screen alternatives and start a transaction. Airbnb and Blablacar are examples of the former, while Hinge and Upwork are examples of the latter. Finally, in all these platforms a transaction takes place if, and only if, both sides of the market mutually accept/like each other.

As the previous discussion illustrates, one of the primary roles of the platform is to select the subset of alternatives—the *assortment*—to display, considering the preferences and characteristics of the agents on both sides of the market. We refer to this problem as a *two-sided assortment problem*. This is similar to the classic (one-sided) *assortment optimization problem*, in which a retailer must decide the subset of products to display to maximize the expected revenue obtained from a series of customers. However, the two-sidedness of the problem imposes some challenges. For instance, the two-sided nature of the problem implies that both users must mutually see and agree to each other. Hence, there is uncertainty on whether a transaction will occur even after the first user initiates a transaction. Moreover, the platform must carefully balance the trade-off between relevance—i.e., showing options that are likely to generate a transaction—and congestion, as the most popular users may get more request than what they are able to respond to.

The present paper aims to investigate how platforms should make this decision when both sides of the market can start a transaction. More specifically, the goal of this paper is twofold: (1) to present a two-sided assortment optimization model with bidirectional interactions, and (2) to provide performance guarantees on this problem.

We start describing a stylized model of a two-sided market mediated by a platform. The platform must choose which subset of profiles to show to each user. Then, users decide whether to accept or reject each of the options in their assortments. We extend previous work by allowing both sides of the market to start an interaction. As a result, matches can happen either (1) from simultaneous shows, where the two users see and agree to each other in the same stage; or (2) sequentially, where users see and agree to each other on different stages.

To facilitate the theoretical analysis of the problem, we assume a horizon of two stages. We show that popular algorithms, such as greedy and perfect matching, can perform arbitrarily poorly in this setting. Despite this negative result, we can obtain a approximation algorithms with constant factor guarantee. More specifically, when simultaneous matches are forbidden we get a guarantee of $1-1/e$. The problem is more challenging when simultaneous matches are allowed, but still we are able to provide a guarantee of $(1-1/e)/24$. Our results rely on submodular optimization techniques, specifically, finding different formulations to translate our problem into a submodular maximization problem over downward-closed constraints. Inspired by these formulations, we propose a new heuristic to solve the problem efficiently. This heuristic, called fix and match heuristic, uses a two-step procedure to find the assortments to show in the first stage and then optimally solves the second stage given the realization of the backlog.

Proc.of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), P. Faliszewski, V. Mascardi, C. Pelachaud, M.E. Taylor (eds.), May 9–13, 2022, Auckland, New Zealand.

Finally, we partnered with a major dating app[1] that provided us with real-data to numerically show the performance of our proposed algorithm. Dating is a direct application of our setting, and it is an exciting market on its own due to the massive growth it has experienced over the last decade.[2] Our numerical experiments show that the proposed algorithm can considerably increase the number of matches generated compared to the algorithm used by our industry partner and also compared to some other relevant benchmarks.

The reminder of the paper is organized as follows. In Section 2 we discuss the most related literature. In Section 3 we introduce our model, and in Section 4 we show that some natural approaches can perform arbitrarily poorly in this setting. In Section 5 we theoretically analyze the problem and provide our performance guarantees. Finally, in Section 6 we numerically compare our heuristic with other relevant benchmarks, and in Section 7 we conclude.

2 RELATED LITERATURE

Our paper is related to several strands of the literature. First, we contribute to the literature on assortment optimization. Most of this literature focuses on one-sided settings, where a retailer must choose the assortment of products to show in order to maximize the expected revenue obtained from a sequence of customers. This model, whose general version was introduced in [24], has been extended to include capacity constraints [20], different choice models [4, 7, 8, 21], search [27], learning [5, 20, 23], and also personalized assortments [3, 11]. We refer to [14] for an extensive review of the current state of the assortment planning literature in one-sided settings.

Over the last couple of years, a new strand of the assortment literature has focused on two-sided markets. [2] introduce a model where each customer chooses, simultaneously and independently, to either contact a supplier from the assortment offered to them by the platform or to remain unmatched. Then, the suppliers see the customers that chose to contact them and decide to either math with one of them or to remain unmatched. The goal of the platform is to select the assortment of suppliers to show to each customer in order to maximize the expected number of matches. The authors show that the problem is NP-hard, and they provide an algorithm that achieves a constant factor approximation. [25] study the same problem, and significantly improve the approximation factor obtained by [2]. Moreover, [25] provide the first guarantee to the problem with cardinality constraints. [1] extend this model and analyze the online setting, where arriving customers are shown an assortment of suppliers and choose at most one supplier to send a request, and the suppliers can choose to accept a match after some time on the platform. The authors show that when suppliers do not accept/reject requests immediately, then a simple greedy policy achieves a 1/2-factor approximation. [1] also propose balancing algorithms that perform relatively well under the Multinomial and Nested Logit models.

Notice that these papers consider a sequential two-sided matching market, where one side of the market initiates an interaction and the other side of the market decides whether to match one of the customers that contacted them in the first place. Moreover, customers are limited to contact only one supplier in their assortments. [17] introduce a model where a platform chooses an assortment for each user on both sides of the market, and users can like/dislike as many of the profiles shown to them as they want. Then, a match is formed if both users like each other, possibly on different time periods. The authors show that the problem is NP-hard, and they propose a family of algorithms that account for the negative effect that the matches obtained today may have on future periods. However, [17] provides no theoretical guarantees for the algorithms they propose.

The second stream of literature to which our paper is related is on matching platforms. Starting with the seminal work of [18], this literature has focused on participation, competition, and pricing, highlighting the role of cross-side externalities. In the dating context, [13] study how the search environment can impact users' welfare and the performance of the platform. They find that simple interventions, such as limiting what side of the market reaches out first or hiding quality information, can considerably improve the platform's outcomes. [12] show that two platforms can successfully coexist charging different prices by limiting the set of options offered to their users. They show that, depending on their outside option, users must balance two effects when choosing a larger platform: (1) a choice effect, whereby users are more likely to find a partner that exceeds their outside option; and (2) a competition effect, whereby agents on the other side of the market are less likely to accept a request as they have more options available. All these models consider a stylized matching market, where users interact with the other side of the market and leave the platform upon getting a match.

Finally, our paper is related to the submodular welfare problem [15], which belongs to the larger class of problems of submodular maximization under matroid constraints [9]. Within this literature, [26] shows that a $1 - 1/e$ approximation guarantee can be obtained for the problem of maximizing a monotone submodular function under partition constraints.

Our main contribution to the literature is to provide the first performance guarantee for the two-sided assortment optimization problem when users have no limit on the number of profiles they can contact/like in their assortment and can act either simultaneously or sequentially. Moreover, we show that many widely used algorithms, such as greedy-type policies or perfect matching, can perform arbitrarily bad due to the two-sided nature of the problem.

3 PROBLEM DEFINITION

In our two-sided assortment problem there are two sets of users I and J, each one representing a *side* of the market. Given their reported preferences and characteristics, the platform computes a set of potential partners $\mathcal{P}_\ell$—or simply *potentials*—for each user ℓ, which contains the subset of users that may be potentially shown to ℓ; i.e., $\mathcal{P}_\ell \subseteq J$ if $\ell \in I$, and $\mathcal{P}_\ell \subseteq I$ otherwise. For each stage $t \in T = \{1, 2\}$, the platform shows to each user $\ell \in I \cup J$ a subset

[1]We keep the name of the app undisclosed as part of our NDA.

[2]Since the launch of Tinder in 2012, hundreds of dating services have emerged worldwide, making this industry worth between \$2 and \$3 billion in the U.S. and \$12 billion worldwide [16]. As [19] discuss, online platforms have become the most common channel for couples to meet, replacing more traditional methods such as meeting through friends or co-workers. Indeed, using a nation-wide survey, they show that 39% of heterosexual couples that met in 2017 did so online, and this number raises to 65% for same-sex couples. Overall, approximately one out of five couples today met online.

S_ℓ^t of users from their set of potentials on the opposite side of the market, that is, $S_\ell^t \subseteq \mathcal{P}_\ell$. For each pair of users ℓ, ℓ' in different sides of the market, if we show to user ℓ agent ℓ' in stage t, that is $\ell' \in S_\ell^t$, there is a probability $\phi_{\ell\ell'}$ that ℓ likes ℓ', and we denote by $\Phi_{\ell\ell'}^t$ the binary random variable that indicates whether ℓ likes ℓ' in stage t or not. We remark that $\phi_{\ell\ell'}$ is not necessarily symmetric, namely $\phi_{\ell\ell'}$ can be different from $\phi_{\ell'\ell}$. Finally, let $\beta_{\ell\ell'} = \phi_{\ell\ell'} \cdot \phi_{\ell'\ell}$ be the probability that there is a match between users ℓ and ℓ' conditional that they see each other.

The *backlog* of a user ℓ, denoted by $\mathcal{B}_\ell \subseteq \mathcal{P}_\ell$, is the subset of users that liked ℓ in the first stage, that is, $\mathcal{B}_\ell = \{\ell' : \ell \in S_{\ell'}^1 \text{ and } \Phi_{\ell'\ell}^1 = 1\}$. We remark that the backlog is a random subset since it is defined over the random realization of the like events obtained from the first stage. During the second stage, each user can see profiles taken from their backlog, or they can also see profiles that were not shown to them in the first stage, that is, $S_\ell^2 \subseteq \mathcal{P}_\ell \setminus S_\ell^1$ for each $\ell \in I \cup J$. Finally, a *match* between users ℓ and ℓ' occurs if both users are shown to each other and they consequently like each other. Our main goal is to maximize the total expected number of matches.

Given that we are studying a two-stage model, then we can assume that the assortments that are shown in the second stage correspond to the *best response* with respect to the first stage decision. In other words, we only have to consider as a decision variable the assortments of the first stage. The decision of the first stage corresponds to a family of *feasible* assortments $S = (S_\ell^1)_{\ell \in I \cup J}$ such that, for a given budget $K \in \mathbb{Z}_+$, $S_\ell^1 \subseteq \mathcal{P}_\ell$ and $|S_\ell^1| \leq K$ for all $\ell \in I \cup J$.

Given a family of *feasible* assortments S^1, we denote by $\mathbf{M}_S$ the random variable that indicates the total number of matches achieved when S^1 is shown in the first stage. Now, we formally define problem our main optimization problem.

Problem 1. *Given a budget $K \in \mathbb{Z}_+$, the* two-sided assortment optimization *problem is the following*

$$\max \left\{ \mathbb{E}\left[\mathbf{M}_{S^1}\right] : S^1 = (S_\ell^1)_{\ell \in I \cup J},\ S_\ell^1 \subseteq \mathcal{P}_\ell,\ |S_\ell^1| \leq K,\ \forall \ell \in I \cup J \right\},$$

where the expectation is over the randomization of the agents' decision in the first stage.

In the remainder of this manuscript, we denote by OPT the optimal value of Problem 1. In what follows, we say that an algorithm is an *α-approximation* if it implements a feasible solution and the expected number of matches is at least an α fraction of OPT.

We denote by $E = \{\{\ell, \ell'\} : \ell \in I,\ \ell' \in J\}$ the set of all possible undirected edges between I and J. We denote by $\vec{E} = \vec{E}_I \cup \vec{E}_J$ the set where $\vec{E}_I = \{(\ell, \ell') : \ell \in I,\ \ell' \in J\}$ and $\vec{E}_J = \{(\ell', \ell) : \ell \in I,\ \ell' \in J\}$, that is, the sets of all directed arcs between I and J. For a given arc $a \in \vec{E}$, let $\overline{a}$ be the inverted arc, that is, $\overline{a} = (\ell', \ell)$ when $a = (\ell, \ell')$.

4 NATURAL APPROACHES

In this section, we some natural algorithms whose performance guarantees in the worst-case is asymptotically close to zero. Without loss of generality, consider the problem with $K = 1$. Then, the problem consist of deciding which profile to show to each user in each stage to maximize the number of matches, where a match is form either from simultaneous shows or from backlog queries. There are many natural algorithms that could solve the problem. For instance,

1. Greedily assign each user to the profile that maximizes the probability of a match.
2. Sequentially find a maximum weight perfect matching, where the weight of each edge is the probability of having a match between the users.

As we show now, both alternatives have a worst-case performance of zero.

Proposition 4.1. *Greedy has worst case performance equal to zero.*

Proof. Without loss of generality, suppose that there is a single stage, and that there are n users on each side of the market, i.e., $I = \{i_1, \ldots, i_n\}$ and $J = \{j_1, \ldots, j_n\}$. In addition, suppose that $\mathcal{P}_i = J$ for every $i \in I$, $\mathcal{P}_j = I$ for every $j \in J$, and that $\beta_{ij_1} = 1$ for all $i \in I$, while $\beta_{ij} = 1 - \epsilon$ for all i and $j \neq j_1$. In words, the estimated probability of matching with user j_1 is equal to 1 if users are shown to each other. In this setting, the Greedy policy will chose $S_i = \{j_1\}$ for every user i, and therefore only one match will take place in expectation. In contrast, an optimal solution is to assign $S_{i_k} = \{j_k\}$, which leads to $1 + (n-1) \cdot (1-\epsilon)$ matches in expectation. Then, the performance of Greedy is given by $1/(1 + (n-1)(1-\epsilon)) \to 0$, as $n \to \infty$ for ϵ sufficiently small. Hence, the performance of the Greedy algorithm can be arbitrarily bad. □

Proposition 4.2. *Sequential perfect match has worst case performance equal to zero.*

Proof. Suppose that $|I| = 2n$, $|J| = 2$, that $\mathcal{P}_i = J$ for every $i \in I$, $\mathcal{P}_j = I$ for every $j \in J$, and that $\phi_{ij} = p$ while $\phi_{ji} = q$ for all $i \in I$, $j \in J$. Then, it is easy to see that the sequential perfect match policy leads to $4pq$ matches in expectation. On the other hand, consider the policy where: i) In $t = 1$, $\{i_1, \ldots, i_n\}$ see j_1, $\{i_{n+1}, \ldots, i_{2n}\}$ see j_2, j_1 sees i_{2n} and j_2 sees i_1, and ii) in $t = 2$, $\{i_1, \ldots, i_n\}$ see j_2, $\{i_{n+1}, \ldots, i_{2n}\}$ see j_1, j_1 sees any profile that liked her in $t = 1$, and same for $j = 2$. Given this policy, the matches (i_{2n}, j_1) and (i_1, j_2) happen with probability pq each. On the other hand, j_1 matches with someone in $\{i_1, \ldots, i_n\}$ with probability $(1 - (1-p)^n) \cdot q$, and the same for j_2 matching with someone in $\{i_{n+1}, \ldots, i_{2n}\}$. Then, the overall expected number of matches is $2pq + 2q \cdot (1 - (1-p)^n)$. It is direct to check that this is the optimal expected number of matches. Then, the sequential perfect match policy achieves a ratio of $4pq/(2q(p + 1 - (1-p)^n)) \to 2/(1+n)$ when $p \to 0$, and $2/(1+n) \to 0$ when $n \to \infty$. □

5 APPROXIMATION ALGORITHMS

In this section, we devise approximation algorithms for the two-sided assortment problem. As shown in the previous section, natural greedy approaches might fail to provide approximation guarantees, and therefore we use techniques from integer programming and submodular optimization. Depending on the possibility of allowing simultaneous shows of users profiles in the platform, we split our analysis in three different cases that we detail in what follows.

5.1 No Simultaneous Shows

As a first step, we consider the case in which a feasible solution for the platform is further restricted to only show assortments to side I in the first stage, and to side J on the second stage. That is, $S_j^1 = \emptyset$ for each user $j \in J$ and $S_i^2 = \emptyset$ for each user $i \in I$. We refer to this setting as the model with *no simultaneous shows*. The main result of this section is the following:

Theorem 5.1. *There exists a $(1 - 1/e)$-approximation algorithm for Problem 1 with no simultaneous shows.*

We prove this result as follows: First, we focus on providing an explicit reformulation of Problem 1 when no simultaneous shows are allowed. With this new formulation, we use tools from submodular optimization to prove Theorem 5.1.

Given that in this section we focus on the setting with (1) no simultaneous shows, (2) first-stage assortments for users in I and (3) second-stage assortments for users in J, our decision variables will be determined only by the directed arcs in $\vec{E}_I$. To ease the explanation of our reformulation, in the remainder of this section, we avoid the subset notation for any assortment family S^1 and we use its corresponding indicator variable $x^1 \in \{0,1\}^{\vec{E}_I}$ defined over $\vec{E}_I$, where for any $a = (\ell, \ell')$ with $\ell \in I$ and $\ell' \in \mathcal{P}_\ell$, we have $x_a^1 = 1$ if $\ell' \in S_\ell^1$ and zero otherwise.

First, let us define the feasible region for the first-stage assortment. Since, we are considering cardinality constraints and first-stage assortments only for users in I, then we have the following feasible region:

$$P^1 = \left\{ x^1 \in \{0,1\}^{\vec{E}_I} : \sum_{\ell' \in J} x^1_{(\ell,\ell')} \leq K, \text{ for every } \ell \in I \right\}.$$

Since the second-stage assortments are only for users in J, then we write the random family of backlogs as $\mathcal{B} = (\mathcal{B}_\ell)_{\ell \in J}$, where $\mathcal{B}_\ell \subseteq \vec{E}_I$ for all $\ell \in J$, and a realization of it as $B = (B_\ell)_{\ell \in J}$. Then, given a first-stage decision $x^1 \in \{0,1\}^{\vec{E}_I}$, since users behave independently, the distribution of $\mathcal{B}$ is

$$\mathrm{P}_{x^1}(\mathcal{B} = B) = \prod_{\ell \in J} \left[\prod_{a \in B_\ell} \phi_a x_a^1 \prod_{a \notin B_\ell} (1 - \phi_a x_a^1) \right].$$

Now, let us focus on the second-stage objective, which will be the best-response to the first-stage assortment family. Recall that in the second-stage we only show assortments to agents in J. Given an scenario $\mathcal{B} = B$, we have

$$f(B) = \max_{y \in P_B^2} \left\{ \sum_{a \in \vec{E}_J} \phi_a y_a \right\}, \tag{1}$$

where the feasible region is given by

$$P_B^2 = \Big\{ y \in \{0,1\}^{\vec{E}_J} : \sum_{\ell' \in I} y_{(\ell',\ell)} \leq K, \text{ for every } \ell \in J,$$
$$y_a \leq 1_{\{\overline{a} \in B\}}, \text{ for every } a \in \vec{E}_J \Big\},$$

and $1_{\{\overline{a} \in B\}}$ is the indicator variable that takes value one if $\overline{a} = (i,j)$, with $i \in I$ and $j \in J$, is such that $\overline{a} \in B_j$, and zero otherwise. We remark that f can be evaluated by using a greedy algorithm.

Finally, the objective for Problem 1 in the setting of no simultaneous shows can be written as

$$F(\phi x^1) = \sum_B f(B) \cdot \mathrm{P}_{x^1}(\mathcal{B} = B), \tag{2}$$

where ϕx^1 is the vector given by $\phi_a x_a^1$ for all $a \in \vec{E}_I$ and the sum is over all possible backlog families. As we mentioned previously, to prove Theorem 5.1, we use tools from submodular optimization. First, let us define some key concepts.

Definition 5.2. Consider a ground set of elements $\mathcal{E}$. A non-negative set function $g : 2^{\mathcal{E}} \to \mathbb{R}_+$ is submodular if for every $X, Y \subseteq \mathcal{E}$ such that $X \subseteq Y$ and for every $x \notin Y$ we have $g(X \cup \{x\}) - g(X) \geq g(Y \cup \{x\}) - g(Y)$. The function g is monotone if for every $X, Y \subseteq \mathcal{E}$ with $X \subseteq Y$ we have $g(X) \leq g(Y)$.

We say that a subset $K \subseteq [0,1]^{\vec{E}_J}$ is *downward closed* if the following holds: when $q \in K$ and $w \leq k$ component-wise, then $w \in K$. We can obtain the following result.

Lemma 5.3. *The following holds:*

(a) The function f defined in (1) *is monotone and submodular.*
(b) For every B, the feasible region P_B^2 is downward closed.

Proof. For the first part, note that f can be written in the same form as the function in Proposition 3.1 of [9]. The second part easily follows by noting that the cardinality constraints are downward closed and also the second type of constraints. □

Definition 5.4. Given a ground set of elements $\mathcal{E}$ and a set function $g : 2^{\mathcal{E}} \to \mathbb{R}_+$, the multilinear extension of g, denoted by $G : [0,1]^{\mathcal{E}} \to \mathbb{R}_+$, is defined for any $y \in [0,1]^{\mathcal{E}}$ as the expected value of $g(O_y)$, where O_y is the random set generated by drawing independently each element $s \in \mathcal{E}$ with probability y_s. That is,

$$G(y) = \sum_{O \subseteq 2^{\mathcal{E}}} g(O) \prod_{s \in O} y_s \prod_{s \in \mathcal{E} \setminus O} (1 - y_s).$$

Vondrak [26] prove the following:

Theorem 5.5 ([26]). *Given a set function $g : 2^{\mathcal{E}} \to \mathbb{R}_+$ and its multilinear extension $G : [0,1]^{\mathcal{E}} \to \mathbb{R}_+$. Then, the following holds:*

(a) If g is monotone, then $\nabla G \geq 0$ component-wise.
(b) If g is submodular, then the Hessian matrix of G is non-positive component-wise.

Furthermore, there exists an efficient algorithm such that given a downward closed polyhedron R, it computes $y \in R$ such that $G(y) \geq (1 - 1/e) \cdot \max_{z \in R} G(z)$.

Note that function F defined in (2) is the multilinear extension of f defined in (1) evaluated in the vector $y_a = \phi_a x_a^1$ for all $a \in \vec{E}_I$. Therefore, Problem 1 in the setting of no simultaneous shows can be written as

$$\max \left\{ F(\phi x^1) : x^1 \in P^1 \right\}. \tag{3}$$

In what follows we focus on devising an approximation algorithm for Problem (3). Consider the optimization problem given by

$$\max \quad F(z) \tag{4a}$$

$$\text{s.t.} \quad \sum_{\ell' \in J : \phi_{\ell\ell'} > 0} \frac{z_{\ell\ell'}}{\phi_{\ell\ell'}} \leq K \qquad \text{for every } i \in I, \tag{4b}$$

$$0 \le z_a \le \phi_a \quad \text{for every } a \in \vec{E}_I. \tag{4c}$$

Lemma 5.6. *The optimal value of problem* (4a) *is an upper bound on the optimal value of Problem* (3).

Proof. Consider a feasible solution $x^1 \in \{0,1\}^{\vec{E}_I}$ for the optimization problem (3), that is, $\sum_{\ell' \in J} x^1_{(\ell,\ell')} \le K$ for every $\ell \in I$. Let $z = \phi_{\vec{E}_I} x^1$. Observe that

$$\sum_{j \in J : \phi_{ij} > 0} z_{ij}/\phi_{ij} = \sum_{j \in J : \phi_{ij} > 0} x^1_{(i,j)} \le K$$

for every $i \in I$, and $z_a = x^1_a \phi_a \le \phi_a$ for every $a \in \vec{E}_I$, and therefore z is a feasible solution for the problem (4a)-(4c). Since the objective value of x^1 in (3) is equal to the objective of z in (4a)-(4c), we conclude the lemma. □

In what follows we describe the algorithm used to prove Theorem 5.1. Our algorithm is based mainly on two different tools: the continuous greedy algorithm for submodular optimization [26] and the dependent randomized rounding algorithm by Gandhi et al. [10]. We formally state our method in Algorithm 1.

Algorithm 1 No Simultaneous Shows

Input: An instance for a two-sided assortment problem
Output: A feasible assortment

1. Compute a solution $\overline{z}$ for the problem (4a)-(4c) using the continuous greedy algorithm.
2. For each $i \in I$, let $\mu^i \in [0,1]^J$ be the vector such that the j-th entry of x_i is equal to $\overline{z}_{ij}/\phi_{ij}$ when $\phi_{ij} > 0$ and zero otherwise.
3. Independently for each user $i \in I$, run the dependent randomized rounding algorithm [10] to compute an integral random vector $\hat{x}_i \in \{0,1\}^J$.
4. For each $i \in I$ return $S^1_i = \{j \in J : \hat{x}_{ij} = 1\}$.

First, Theorem 5.5 says that there exists an efficient algorithm that computes a solution $\overline{z}$ feasible for the problem (4a)-(4c) such that

$$F(\overline{z}) \ge \left(1 - \frac{1}{e}\right) \cdot \mathrm{OPT}', \tag{5}$$

where OPT′ is the optimal value of (4a)-(4c).

Second, our algorithm utilizes the dependent randomized rounding algorithm of Gandhi et al. [10]. Given the fractional solution $\overline{z}$ satisfying (5), let $\mu_i \in [0,1]^J$ be the vector such that the j-th entry of μ_i is equal to $\overline{z}_{ij}/\phi_{ij}$ when $\phi_{ij} > 0$ and zero otherwise. Observe that thanks to constraint (4c) we have $\mu_i \in [0,1]^J$ for each $i \in I$. Then, independently for each user $i \in I$, by the algorithm of Gandhi et al. [10] it is possible to efficiently compute an integral random vector $\hat{x}_i \in \{0,1\}^J$ satisfying the following conditions:

(1) $\sum_{j \in J} \hat{x}_{ij} \le \lceil \sum_{j \in J} \mu_{ij} \rceil$, and
(2) $\mathbb{E}[\hat{x}_{ij}] = \mu_{ij}$ for each $i \in I$ and $j \in J$.

We are ready to prove our main result.

Proof of Theorem 5.1. Thanks to condition 1 of the randomized rounding algorithm, for each $i \in I$ we have

$$\sum_{j \in J} \hat{x}_{ij} \le \left\lceil \sum_{j \in J} x_{ij} \right\rceil = \left\lceil \sum_{j \in J : \phi_{ij} > 0} \frac{\overline{z}_{ij}}{\phi_{ij}} \right\rceil \le K,$$

where the last inequality holds since K is integral and $\overline{z}$ satisfies constraint (4b). Therefore, our algorithm gives a feasible solution for the two-sided assortment problem. We now analyze the approximation guarantee.

$$\begin{aligned}
\mathbb{E}_{\hat{x}}[F(\hat{x}^1)] &= \sum_B f(B) \cdot \mathbb{E}_{\hat{x}}\left(\prod_{\ell \in J} \left[\prod_{a \in B_\ell} \phi_a \hat{x}^1_a \prod_{a \notin B_\ell} (1 - \phi_a \hat{x}^1_a) \right] \right) \\
&= \sum_B f(B) \cdot \prod_{\ell \in J} \left[\prod_{a \in B_\ell} \phi_a \mathbb{E}_{\hat{x}}[\hat{x}^1_a] \prod_{a \notin B_\ell} (1 - \phi_a \mathbb{E}_{\hat{x}}[\hat{x}^1_a]) \right] \\
&= \sum_B f(B) \cdot \prod_{\ell \in J} \left[\prod_{a \in B_\ell} \phi_a \cdot \frac{\overline{z}_a}{\phi_a} \prod_{a \notin B_\ell} (1 - \phi_a \cdot \frac{\overline{z}_a}{\phi_a}) \right] \\
&= F(\overline{z}),
\end{aligned}$$

where the second equality comes from the fact that $\hat{x}_i$ is independent from $\hat{x}_\ell$ for every $i, \ell \in I$ with $i \ne \ell$ and the third equality comes from condition 2 of the randomized rounding procedure. Finally, using inequality (5) and Lemma 5.6 we conclude the proof. □

5.2 Simultaneous Shows in the First Stage

One of the limitations of the usual approaches in sequential two-sided assortment optimization is the inability of users to evaluate profiles in each stage. To accommodate this feature, we now consider the possibility that users see and like each other in the same stage. For simplicity, in this section, we will focus on allowing simultaneous shows in the first stage. A match between users $i \in I$ and $j \in J$ given by simultaneous shows in the first stage happens with probability $\phi_{ij}\phi_{ji}x^1_{ij}x^1_{ji}$. Therefore, an extra term has to be added to the optimization problem (3). Formally, Problem 1 with simultaneous shows only in the first stage can be formulated as follows

$$\max \left\{ \sum_{a \in \vec{E}} \phi_a \phi_{\overline{a}} x^1_a x^1_{\overline{a}} + F(\phi x^1) : x^1 \in P^1 \right\},$$

where P^1 and F are the same region and function, respectively, defined in the previous section. However, the objective function in this formulation is not necessarily submodular, so we cannot simply follow the same approach described in the previous section. To overcome this, we reformulate the objective function to make it submodular by adding extra variables and constraints. For every arc $a \in \vec{E}$, let x^1_a be the binary variable defined as in the previous section. Let us consider an extra binary variable z^1_e for every $e \in E$ (undirected edges) that takes the value one if $i \in I$ and $j \in J$ with $e = \{i, j\}$ simultaneously see each other in the first stage. Notice that for a given edge $e = \{i, j\}$, and arc $a = (i, j)$, only one of the variables z^1_e, x^1_a and $x^1_{\overline{a}}$ can be equal to one. For a given edge $e \in E$, denote by a_e and $\overline{a_e}$ the corresponding directed arcs. Our feasible region is therefore the following:

$$P^1 = \Big\{(x^1, z^1) : z_e^1 + x_{a_e}^1 + x_{\overline{a_e}}^1 \le 1, \text{ for every } e \in E,$$
$$\sum_{\ell' \in J} \left(x^1_{(\ell,\ell')} + z^1_{\{\ell,\ell'\}}\right) \le K, \text{ for every } \ell \in I,$$
$$\sum_{\ell' \in I} \left(x^1_{(\ell,\ell')} + z^1_{\{\ell,\ell'\}}\right) \le K, \text{ for every } \ell \in J,$$
$$x^1 \in \{0,1\}^{\vec{E}},\ z^1 \in \{0,1\}^E \Big\}.$$

We can interpret that each edge $e \in E$ has three types of states: without direction (simultaneous shows) or either of both directed versions (non-simultaneous shows). In the first stage, we can only choose one of these states for each edge $e \in E$, which corresponds to the first family of constraints in P^1. Each user $\ell \in I \cup J$ can only be shown at most K users in the opposite side. This can happen in two types of states: a directed arc or an edge. Therefore, for each user we have to choose at most K of these states, which corresponds to the second and third family of constraints.

Now, we can express the objective function of our problem in the following way:

$$\max \left\{ \sum_{e \in E} \phi_{a_e} \phi_{\overline{a_e}} z_e^1 + F(\phi x^1) :\ x^1 \in P^1 \right\}. \tag{6}$$

Observe that the objective corresponds to the sum of a linear function plus a submodular function, which gives a submodular function. Moreover, the new feasible region P^1 is downward closed. The following is the main result in this section.

Theorem 5.7. *There exists a $(1 - 1/e)/24$-approximation algorithm for Problem 1 with simultaneous shows only in the first stage.*

To obtain this result, we construct a similar upper bound than in the previous section and obtain a fractional solution (x^G, z^G) via continuous greedy whose value is at least $(1 - 1/e)\mathrm{OPT}'$, where OPT' is the optimal value of the upper bound, which we know is at least OPT. Since P^1 corresponds to packing constraints, we can use the technique of contention resolution schemes introduced in [6]. The system of inequalities is 3-sparse, i.e., every variable is in at most 3 constraints. Therefore, [6] gives a $(b, 1 - 6b)$-balanced contention resolution scheme and by using their methodology we can obtain a solution $(\hat{x}^1, \hat{z}^1) \in P^1$ such that its objective value is at least $b(1 - 6b)$ fraction of the objective value of (x^G, z^G), where $b \in (0, 1/6)$. Note that $b(1 - 6b)$ is maximized in $b = 1/2$ which gives a $1/24$-factor and concludes the proof.

5.3 Simultaneous Shows in Both Stages

In this section, we allow simultaneous shows in the first and second stages. This creates an extra challenge, since we cannot show an edge (or corresponding arc) more than once among stages. Recall that F is defined as

$$F(\phi x^1) = \sum_B f(B) \cdot \mathrm{P}_{x^1}(\mathcal{B} = B),$$

where f is a maximization problem defined over a region P_B^2. When simultaneous shows are allowed in both stages, the feasible region P_B^2 of the second stage depends on x^1 and z^1. In other words, instead of $f(B)$, we would have $f(B, x^1, z^1)$, which results in another layer of complexity since $\mathrm{P}_{x^1}(\mathcal{B} = B)$ already depends on x^1. To address this challenge, we will reformulate feasible regions P^1 and P_B^2, as well as the objective function of the second stage, in such a way that the distribution of the backlogs does not depend on x^1.

First, we add a binary variable in the first stage y_e^1 for every edge $e \in E$ that takes the value one if we decide to consider edge e in the second stage, and zero otherwise. Then, each edge has now four types of states: with no direction (simultaneous shows), both directions and left to the second stage. Since we can only choose one of those states for each edge, we have the following:

$$P^1 = \Big\{(x^1, z^1, y^1) : y_e^1 + z_e^1 + x_{a_e}^1 + x_{\overline{a_e}}^1 \le 1, \text{ for every } e \in E,$$
$$\sum_{\ell' \in J} \left(x^1_{(\ell,\ell')} + z^1_{\{\ell,\ell'\}}\right) \le K, \text{ for every } \ell \in I$$
$$\sum_{\ell' \in I} \left(x^1_{(\ell,\ell')} + z^1_{\{\ell,\ell'\}}\right) \le K, \text{ for every } \ell \in J$$
$$x^1 \in \{0,1\}^{\vec{E}},\ y^1, z^1 \in \{0,1\}^E \Big\}.$$

Note that the cardinality constraints are not affected by y^1. Now, let us focus on the second stage feasible region. Let us add the binary variable z_e^2 for every edge $e \in E$ with $e = \{i, j\}$ that takes value one if we simultaneously show j to i and i to j, and zero otherwise. Consider the following region:

$$P^2 = \Big\{(x^2, z^2) : \quad x_{\overline{a}}^2 \le x_a^1, \text{ for every } a \in \vec{E},$$
$$z_e^2 \le y_e^1, \text{ for every } e \in E,$$
$$\sum_{\ell' \in J} \left(x^2_{(\ell,\ell')} + z^2_{\{\ell,\ell'\}}\right) \le K, \text{ for every } \ell \in I$$
$$\sum_{\ell' \in I} \left(x^2_{(\ell,\ell')} + z^2_{\{\ell,\ell'\}}\right) \le K, \text{ for every } \ell \in J$$
$$x^2 \in \{0,1\}^{\vec{E}},\ z^2 \in \{0,1\}^E \Big\}.$$

Observe that given the constraints in P^1, we do not need to add the set of constraints $x_a^2 + x_{\overline{a}}^2 + x_e^2 \le 1$. We can rely on the two new families of constraints. The second family of constraints says that we can only simultaneously show an edge in the second stage if we decided to leave available that edge (this decision is represented by variable y^1). The first family of constraints ensures that we show the inverted arc $\overline{a}$ only if we showed a in the first stage; if not, then it should not be available in the second stage. The reason to add this set of constraints is to avoid the dependency of the distribution of the backlogs on x^1.

Formally, the distribution that we consider now is the following: We will sample all directed edges, without taking into account if they were selected in the first stage. In other words,

$$\mathrm{P}(\mathcal{B} = B) = \prod_{\ell \in J} \left[\prod_{a \in B_\ell} \phi_a \prod_{a \notin B_\ell} (1 - \phi_a) \right].$$

Next, we have to guarantee that arcs that were sampled, but whose inverted version was not selected in the first stage, then its contribution to the objective should be zero. Formally, given an scenario of a family of backlogs $\mathcal{B} = B$, we define the best response of the second stage as

$$f_B(x^1, y^1) = \max \left\{ \sum_{a \in \bigcup_\ell B_\ell} \phi_a x_a^2 + \sum_{e \in E} \phi_{a_e} \phi_{\overline{a_e}} z_e^2 : (x^2, z^2) \in P^2 \right\}.$$

Note in the first term of the sum, if $x_{\overline{a}}^1 = 0$, then $x_a^2 = 0$ and its contribution to the objective is zero (even if it was sampled as part of the scenario B). This formulation ensures that the distribution of the family of backlogs does not depend on first-stage variables.

Lemma 5.8. *Given any family of backlogs B, the following holds:*

(1) If x^1 is fixed, then f_B is monotone submodular on y^1.
(2) If y^1 is fixed, then f_B is monotone submodular on x^1.

Moreover, since the sets of elements that define x^1 and y^1 are disjoint, then f_B can be re-interpreted as unidimensional submodular function.

Proof. We use Proposition 3.1 in [9] to prove parts 1. and 2. Finally, since x^1 and y^1 are disjoint, [22] shows that f_B can be re-interpreted as unidimensional submodular function. □

Now, we can formulate Problem 1 with simultaneous shows in both stages in the following way:

$$\max \left\{ \sum_{e \in E} \phi_{a_e} \phi_{\overline{a_e}} z_e^1 + F(x^1, y^1) : \ (x^1, z^1, y^1) \in P^1 \right\}, \qquad (7)$$

where $F(x^1, y^1) = \sum_B f_B(x^1, y^1) \cdot \mathrm{P}(\mathcal{B} = B)$.

Given the submodularity of F (average of submodular functions), linearity of the first sum and that P^1 is downward closed. The following is the main result of this section.

Theorem 5.9. *There exists a $(1 - 1/e)/24$-approximation algorithm for Problem 1 with simultaneous shows in both stages.*

The proof of the theorem follows the same lines of Theorem 5.9.

5.4 Proposed Algorithm

Given the large number of possible backlogs that result from the assortments chosen for the first stage, it is not possible to efficiently solve (7). However, we know that for any set of profiles shown in the first stage, $\{x_{\ell\ell'}^1\}_{\ell' \in \mathcal{P}_\ell, \, \ell \in I \cup J}$ and for any given set of backlogs, $\vec{\mathcal{B}} = \{\mathcal{B}_\ell\}_{\ell \in I \cup J}$, the solution to the second stage problem can be efficiently obtained solving:

$$\begin{aligned}
\max \quad & \sum_{\ell \in I \cup J} \sum_{\ell' \in \mathcal{P}_\ell} y_{\ell\ell'} \cdot \phi_{\ell\ell'} + z_{\ell\ell'}^2 \cdot \phi_{\ell\ell'} \cdot \phi_{\ell'\ell} \\
st. \quad & y_{\ell\ell'} \leq 1_{\{\ell' \in \mathcal{B}_\ell\}}, \ \forall \ell' \in \mathcal{P}_\ell, \ell \in I \cup J \\
& x_{\ell\ell'}^2 + y_{\ell\ell'} \leq 1 - x_{\ell\ell'}^1, \ \forall \ell' \in \mathcal{P}_\ell, \ell \in I \cup J \\
& \sum_{\ell' \in \mathcal{P}_\ell} x_{\ell\ell'}^2 + y_{\ell\ell'} \leq K, \ \forall \ell \in I \cup J \\
& z_{\ell\ell'}^2 \leq x_{\ell\ell'}^2, \ z_{\ell\ell'}^2 \leq x_{\ell'\ell}^2, \ z_{\ell\ell'}^2 = z_{\ell'\ell}^2, \ \forall \ell' \in \mathcal{P}'_\ell, \ell \in I \cup J \\
& x_{\ell\ell'}^2, \ y_{\ell\ell'}, z_{\ell\ell'}^2 \in [0, 1], \ \forall \ell' \in \mathcal{P}_\ell, \ \ell \in I \cup J
\end{aligned} \qquad (8)$$

Hence, the challenge is to find a good solution for the first stage. As shown in [17], an upper bound for the problem with T stages can be obtained from solving the following LP:

$$\begin{aligned}
\max \quad & \sum_{t=1}^{T} \sum_{\ell \in I \cup J} \sum_{\ell' \in \mathcal{P}_\ell} y_{\ell\ell'}^t \cdot \phi_{\ell\ell'} + z_{\ell\ell'}^t \cdot \phi_{\ell\ell'} \cdot \phi_{\ell'\ell} \\
st. \quad & \sum_{\tau=1}^{t} y_{\ell\ell'}^\tau \leq 1_{\{\ell' \in \mathcal{B}_\ell\}} + \sum_{\tau=1}^{t-1} (x_{\ell'\ell}^\tau - z_{\ell\ell'}^\tau) \cdot \phi_{\ell'\ell}, \ \forall \ell' \in \mathcal{P}_\ell, \\
& \qquad \ell \in I \cup J, t \in [T] \\
& \sum_{\tau=1}^{T} x_{\ell\ell'}^\tau + y_{\ell\ell'}^\tau \leq 1, \ \forall \ell' \in \mathcal{P}_\ell, \ell \in I \cup J \\
& \sum_{j} x_{\ell\ell'}^t + y_{\ell\ell'}^t \leq K, \ \forall \ell \in I \cup J, t \in [T] \\
& z_{\ell\ell'}^t \leq x_{\ell\ell'}^t, \ z_{\ell\ell'}^t \leq x_{\ell'\ell}^t, \ z_{\ell\ell'}^t = z_{\ell'\ell}^t, \ \forall \ell' \in \mathcal{P}_\ell, \ell \in I \cup J, t \in [T] \\
& x_{\ell\ell'}^t, \ y_{\ell\ell'}^t, z_{\ell\ell'}^t \in [0, 1], \ \forall \ell' \in \mathcal{P}_\ell, \ell \in I \cup J, t \in [T]
\end{aligned} \qquad (9)$$

Notice that, when $T = 1$, this problem is similar to that in (8). The only difference lies in the first two constraints: in the first case, we consider the expected backlog instead of the realized one, and in the second one, we incorporate the decisions in the first stage to avoid showing a given profile more than once to the same user. Thus, when extended to the two stages, (9) weights in the expected backlog and the optimal decisions in the second stage.

The idea behind our heuristic—called fix and match heuristic (FMH)—is to use (9) to obtain a good first stage solution. As this solution may be fractional, the second step is to obtain an integer solution for the first stage. We accomplish this by fixing the first stage variables obtained from the linear relaxation and solving the integer version of (9). As the magnitude of the problem reduces considerably after fixing the variables with integer values, this step can be done efficiently, resulting in an integer solution for the first stage. Finally, the heuristic finds the optimal optimal decisions for the second stage after observing the realized backlogs by solving (8). This heuristic is formalized in Algorithm 2.

Algorithm 2 Fix and Match Heuristic (FMH)

Input: $\mathcal{P}_\ell, \ \ell \in I \cup J$, and $\phi_{\ell\ell'}^t, \ \ell' \in \mathcal{P}_\ell, \ \ell \in I \cup J$
Output: An assortment S_ℓ^t for each user $\ell \in \mathcal{I}$ and stage t

1: Solve 1 and keep $x^{*,1}$.
2: Solve 1 considering as additional constraints $x_{\ell\ell'}^1 = x_{\ell\ell'}^{*,1}$ whenever $x_{\ell\ell'}^{*,1} \in \{0, 1\}$, and $x_{\ell\ell'}^1 \in \{0, 1\}$ otherwise. Define $S_\ell^1 = \{\ell' : \bar{x}_{\ell\ell'}^1 = 1\}$, where $\bar{x}_{\ell\ell'}^1$ is the optimal value of x^1.
3: Given the assortments in stage 1, S^1, and the realized backlogs $\mathcal{B}$, solve 2 for the second stage.

---

6 EXPERIMENTS

In this section we evaluate the performance of the proposed algorithm using real data.

Table 1: Descriptives of Instance

	N	Score	Potentials	Like Rate
Women	173	5.330	69.491	0.270
		(2.423)	(15.978)	(0.174)
Men	113	2.763	94.035	0.571
		(1.504)	(22.701)	(0.132)

6.1 Data

To perform our experiments, we use a dataset obtained from our industry partner. This dataset includes a sample of heterosexual users from Houston, TX. For each user ℓ and for each profile ℓ' in their set of potentials, we compute the probability that ℓ likes ℓ' using the same approach described in [17]. In Table 1 we report several summary statistics of the instance (standard deviations in parenthesis), including the number of users, their score, the number of potentials available, and their like probabilities.

6.2 Benchmarks

To assess the performance of the proposed algorithm we compare it with several benchmarks:

(1) Naive: for each user, select the subset of profiles that maximizes the expected number of likes, i.e.,

$$S_\ell^t = \underset{S \subseteq \mathcal{P}_\ell \setminus \bigcup_{\tau=1}^{t-1} S_\ell^\tau : |S| \leq K}{\arg\max} \left\{ \sum_{j \in S} \phi_{\ell\ell'} \right\}$$

(2) Random: for each user, select a random subset of potentials of size K.

(3) Partner: for each user, select a subset of profiles based on the algorithm used by our industry partner.

(4) Sequential Perfect Matching (PM): for each user, select the profiles obtained from solving the problem

$$\begin{aligned} \max \quad & \sum_{t=1}^{T} \sum_{\ell \in I \cup J} \sum_{\ell' \in \mathcal{P}_\ell} z_{\ell\ell'}^t \cdot \phi_{\ell\ell'} \cdot \phi_{\ell'\ell} \\ st. \quad & \sum_{\tau=1}^{T} x_{\ell\ell'}^\tau \leq 1, \ \forall \ell' \in \mathcal{P}_\ell, \ell \in I \cup J \\ & \sum_{j} x_{\ell\ell'}^t \leq K, \ \forall \ell \in I \cup J, t \in [T] \\ & z_{\ell\ell'}^t \leq x_{\ell\ell'}^t, \ z_{\ell\ell'}^t \leq x_{\ell'\ell}^t, \ z_{\ell\ell'}^t = z_{\ell'\ell}^t, \ \forall \ell' \in \mathcal{P}_\ell, \ell \in I \cup J, t \in [T] \\ & x_{\ell\ell'}^t, z_{\ell\ell'}^t \in [0,1], \ \forall\forall \ell' \in \mathcal{P}_\ell, \ell \in I \cup J, t \in [T], \end{aligned} \tag{10}$$

i.e., $S_\ell^t = \left\{ j : x_{\ell\ell'}^t = 1 \right\}$ for each $\ell \in I \cup J$ and $t \in [T]]$.

(5) Greedy: for each user, select the subset of profiles that maximizes the expected number of matches, i.e.,

$$S_\ell^t = \underset{S \subseteq \mathcal{P}_\ell \setminus \bigcup_{\tau=1}^{t-1} S_\ell^\tau : |S| \leq K}{\arg\max} \left\{ \sum_{\ell' \in S} \phi_{\ell\ell'} \cdot \left(1_{\{\ell' \in \mathcal{B}_\ell^t\}} + \phi_{\ell'\ell} \cdot 1_{\{\ell' \notin \mathcal{B}_\ell^t\}} \right) \right\}$$

(6) DH: for each user, select the subset of profiles that results from the Dating Heuristic (DH) described in [17]. For each stage $t \in [T]$, DH consists of two steps: (i) optimization, where (9) considering one stage of look-ahead and the updated backlog for each user; and (ii) rounding, where first the profiles ℓ' for which $y_{\ell\ell'}^t > 0$ are added to S_ℓ^t first, and then if there is space left in the assortment, the profiles for which $x_{\ell\ell'}^t > 0$ are added to S_ℓ^t.

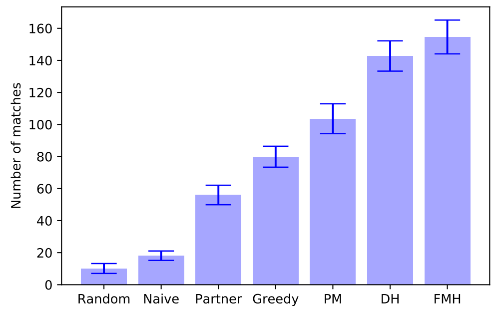

Figure 1: Matches by Benchmark

6.3 Results

For each benchmark, we perform 100 simulations where, in each stage, (1) we choose the assortment to show to each user, (2) we simulate the decisions of the users based on their like probabilities, and (3) we update the state of the system before moving on to the next stage.

Our main experimental results are summarized in Figure 1, which reports the average number of matches generated by each benchmark. First, we observe that the performance of PM and Greedy are considerably better than their worst-case performance. Indeed, we find that these two policies perform well compared to our industry partner's algorithm. Second, consistent with the results in[17], we observe that DH outperforms both Greedy and PM. Finally, we observe that the proposed heuristic outperforms all the other benchmarks. Moreover, we observe it achieves over 94% relative to the upper bound, equal to 163.63 for this instance.

7 CONCLUSIONS

We theoretically study a two-period version of the two-sided assortment optimization problem introduced in [17]. This problem departs from previous work as it allows for both sides of the market to start an interaction, as it is the case in many matching markets including dating, car-sharing, etc. Using tools from submodular optimization, we provide the first performance guarantees for different variants of the problem, including the cases where simultaneous shows are not allowed, and also considering the general when this constraint is removed. In addition, we provide a novel heuristic that is inspired by our theoretical analysis of the problem, and we show that this heuristic outperforms all the other benchmarks considered.

We believe that there are many interesting directions for future research. First, we believe that our results can be extended to the case with multiple periods, and also to the case when the like probabilities depend on the assortment shown. In addition, our model can be easily extended to the case where users log-in with some known probability. Another direction is to formally study the performance of the proposed FMH.

REFERENCES

[1] Ali Aouad and Daniela Saban. 2021. Online Assortment Optimization for Two-sided Matching Platforms. (2021).

[2] Itai Ashlagi, Anilesh Krishnaswamy, Rahul Makhijani, Daniela Saban, and Kirankumar Shiragur. [n.d.]. Assortment planning for two-sided sequential matching markets. ([n. d.]).

[3] Gerardo Berbeglia and Gwenaël Joret. 2015. Assortment Optimisation Under a General Discrete Choice Model: A Tight Analysis of Revenue-Ordered Assortments. (2015), 36 pages.

[4] Jose Blanchet, Guillermo Gallego, and Vineet Goyal. 2016. A Markov chain approximation to choice modeling. *Operations Research* 64, 4 (2016), 886–905.

[5] Felipe Caro and Jérémie Gallien. 2007. Dynamic assortment with demand learning for seasonal consumer goods. *Management Science* 53, 2 (2007), 276–292.

[6] Chandra Chekuri, Jan Vondrák, and Rico Zenklusen. 2014. Submodular function maximization via the multilinear relaxation and contention resolution schemes. *SIAM J. Comput.* 43, 6 (2014), 1831–1879.

[7] James M Davis, Guillermo Gallego, and Huseyin Topaloglu. 2014. Assortment Optimization Under Variants of the Nested Logit Model. *Operations Research* 62, 2 (2014), 250–273.

[8] Vivek F Farias, Srikanth Jagabathula, and Devavrat Shah. 2013. A Nonparametric Approach to Modeling Choice with Limited Data. *Management Science* 59, 2 (2013), 305–322.

[9] Marshall L Fisher, George L Nemhauser, and Laurence Wolsey. 1978. An analysis of approximations for maximizing submodular set functions–II. *Polyhedral combinatorics* (1978).

[10] Rajiv Gandhi, Samir Khuller, Srinivasan Parthasarathy, and Aravind Srinivasan. 2006. Dependent rounding and its applications to approximation algorithms. *Journal of the ACM (JACM)* 53, 3 (2006), 324–360.

[11] Negin Golrezaei, Hamid Nazerzadeh, and Paat Rusmevichientong. 2014. Real-Time Optimization of Personalized Assortments. *Management Science* 60, 6 (2014), 1532–1551.

[12] Hanna Halaburda, Mikolaj Jan Piskorski, and Pinar Yildirim. 2018. Competing by Restricting Choice: The Case of Search Platforms. *Management Science* 64, 8 (2018), 3574–3594.

[13] Yash Kanoria and Daniela Saban. 2017. Facilitating the Search for Partners on Matching Platforms: Restricting Agents' Actions. (2017).

[14] A. Kök, M. Fisher, and R. Vaidyanathan. 2015. *Retail Supply Chain Management: Quantitative Models and Empirical Studies.* Springer US, 175–236.

[15] Benny Lehmann, Daniel Lehmann, and Noam Nisan. 2006. Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behavior* 55, 2 (2006), 270–296.

[16] Melissa Lin. 2018. Online Dating Industry: The Business of Love. Link. Accessed November 11, 2019.

[17] Ignacio Rios, Daniela Saban, and Fanyin Zheng. 2021. Improving match rates in dating markets through assortment optimization. (2021).

[18] Jean-Charles Rochet and Jean Tirole. 2003. Two-Sided Markets. *Journal of the European Economic Association* (2003), 990–1029.

[19] Michael Rosenfeld, Reuben J Thomas, and Sonia Hausen. 2019. Disintermediating your friends. (2019), 15 pages.

[20] Paat Rusmevichientong, Zuo-Jun Max Shen, and David B. Shmoys. 2010. Dynamic Assortment Optimization with a Multinomial Logit Choice Model and Capacity Constraint. *Operations Research* 58, 6 (2010), 1666–1680.

[21] Paat Rusmevichientong, David Shmoys, Chaoxu Tong, and Huseyin Topaloglu. 2014. Assortment optimization under the multinomial logit model with random choice parameters. *Production and Operations Management* 23, 11 (2014), 2023–2039.

[22] Richard Santiago and F Bruce Shepherd. 2019. Multivariate submodular optimization. In *International Conference on Machine Learning*. PMLR, 5599–5609.

[23] Denis Sauré and Assaf Zeevi. 2013. Optimal Dynamic Assortment Planning with Demand Learning. *Manufacturing & Service Operations Management* 15, 3 (2013), 387–404.

[24] Kalyan Talluri and Garrett van Ryzin. 2004. Revenue Management Under a General Discrete Choice Model of Consumer Behavior. *Management Science* 50, 1 (2004), 15–33.

[25] Alfredo Torrico, Margarida Carvalho, and Andrea Lodi. [n.d.]. Multi-agent assortment optimization in sequential matching markets. ([n. d.]).

[26] Jan Vondrák. 2008. Optimal approximation for the submodular welfare problem in the value oracle model. In *Proceedings of the 40th Annual ACM Symposium on the Theory of Computing (STOC)*. 67–74.

[27] Ruxian Wang and Ozge Sahin. 2018. The Impact of Consumer Search Cost on Assortment Planning and Pricing. *Management Science* 64, 8 (2018), 3649–3666.