Improving Match Rates in Dating Markets through Assortment Optimization

Ignacio Rios

Naveen Jindal School of Management, University of Texas at Dallas, ignacio.riosuribe@utdallas.edu

Daniela Saban

Graduate School of Business, Stanford University

Fanyin Zheng

Graduate School of Business, Columbia University

Problem definition: We study how online platforms can leverage the behavioral considerations of their users to improve their assortment decisions. Motivated by our collaboration with a dating company, we study how a platform should select the assortments to show to each user in each period to maximize the expected number of matches in a time horizon, considering that a match is formed if two users like each other, possibly on different periods.

Academic/Practical Relevance: Increasing match rates is one of the most common objectives among many online platforms. We provide insights on how to leverage users' behavior towards this end.

Methodology: We model the platform's problem and we use econometric tools to estimate the main inputs of our model, namely, the like and log in probabilities, using our partner's data. We exploit a change in our partner's algorithm to estimate the causal effect of previous matches on the like behavior of users. Based on this finding, we propose a family of heuristics to solve for the platform's problem, and we use simulations and a field experiment to assess the benefits of our algorithm.

Results: First, we find that the number of matches obtained in the recent past has a negative effect on the like behavior of users. Leveraging this finding, we propose a family of heuristics that decide the assortment to show to each user on each day. Finally, using simulations and a field experiment we show that our algorithm can yield 40% more matches relative to our partner's algorithm.

Managerial Implications: Our results highlight the importance of correctly accounting for the behavior of users on both ends of a transaction to improve the operational efficiency of matching platforms. In addition, we identify and measure the effect of previous matches in the users' preferences, which is also leveraged by our algorithm. Our methodology can also be applied to online matching platforms in other settings.

Key words: assortment optimization, online platforms, matching, dating markets, behavioral operations.

1. Introduction

In the past two decades, hundreds of dating services have emerged worldwide, making dating an almost \$3-billion industry in the U.S. and a \$12-billion industry worldwide (Lin 2018). Moreover, online dating platforms have become the most common channel for couples to meet, replacing traditional channels such as meeting through friends or co-workers. Indeed, 39% of heterosexual couples and 65% of same-sex couples that met in the U.S. in 2017 did so online (Rosenfeld et al. 2019). Overall, approximately one out of five couples today met online.[1]

A common feature across many dating platforms is that they display a limited number of profiles of potential partners (or simply profiles) to each user in each day. Some platforms, like Tinder and Bumble, implement this by adding *swipe* limits; other platforms impose a limit on the number of likes (e.g., Hinge), while others explicitly limit the number of profiles displayed each day (e.g., Coffee Meets Bagel). As described in Bumble's website, platforms do so to "[...] *help foster more genuine, quality connections for our users and encourage more intentional swiping* [...]". As a result, one of the primary roles of dating platforms is to select the subset of profiles—the *assortment*—to display to each user in each day, considering the preferences and characteristics of the users involved.

The aforementioned problem resembles the classic *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, there are distinctive features in the dating context that make this problem particularly novel. First, both users must mutually agree—by liking each other—to generate a "match". The fact that both parties must agree to a match considerably affects the probability that a transaction occurs, so platforms should consider the behavior of the agents on both ends of a potential match when making the assortment decisions. Second, users may interact asynchronously, i.e., users need not see each other in the same period. Thus, platforms must carefully manage the timing of these interactions. Notice that these features are not exclusive to dating platforms, and may be relevant in other online platforms including those for freelancing (e.g., TaskRabbit or UpWork), ride-sharing (e.g., Blablacar), and accommodation (e.g., Airbnb). Indeed, Airbnb was able to improve the booking conversion by almost 4% by accounting for the preferences of the hosts when deciding which subset of listings to show to guests (Ifrach 2015).

The size and relevance of the dating market highlight the need to make these platforms more efficient. To contribute towards this goal, in September 2018 we partnered with a major dating app to help them optimize over the assortments to be shown to their users.[2] Our partner's platform offers a limited subset of profiles (ranging from 3 to 9) to each user in each day, and their primary objective is to maximize the number of matches generated. In addition, the assortments offered by their platform must satisfy a series of additional constraints: users can only be shown to each other if they find each other acceptable, no user can see a profile more than once, additional constraints on the composition of the assortment, among others.

[1] For example, the Pew Research Center reports that 15% of adults in the U.S. used online dating sites in 2016.

[2] We keep the name of our partner undisclosed as per the terms of our NDA.

Contributions. Our paper makes several contributions. First, we propose a model of a dynamic matching market mediated by a platform that captures the key elements of our industry partner's problem. Second, we introduce a class of algorithms that incorporates the findings on user behavior that were obtained through a rigorous analysis of our partner's data. Finally, we test the efficiency of our algorithms in a field experiment and find a significant and substantial increase in the number of matches. We now describe these contributions in more detail.

Problem formulation. To capture our partner's problem, we introduce a stylized model of a dynamic matching market mediated by a platform. The platform hosts a set of users and must decide, in each period, what subset of profiles to show to each user to maximize the overall expected number of matches. In our model, users log in each period with some time-dependent probability and, conditional on logging in, they observe a subset of profiles—an assortment—that satisfies the constraints imposed by the platform. Then, users decide whether to like or not like each profile in their assortment based on their preferences. If two users mutually like each other, possibly in different periods, a match is generated. Our goal is to find an algorithm to maximize the total expected number of matches generated by the platform over an entire time horizon. We show that the platform's problem is NP-hard, and we highlight that our model is general enough to capture a broad array of dating markets and, more broadly, other matching markets.

Estimating users' behavior from the data. To understand what drives users' behavior and guide the design of our algorithms, we use our partner's data to estimate both like decisions and log in probabilities. Using observational data we find that the probability of liking new profiles is negatively correlated with the number of matches obtained in the recent past. This result suggests that there exists a *history effect*, by which users are less likely to like other profiles when they have recently succeeded in matching. In order to address the potential endogeneity problem of these estimates, we use a quasi-experiment that introduced an exogenous variation in the probability of getting a match for some users. Our estimates show that each additional match reduces the probability of a new like by at least 3%. Finally, we also estimate the log in probabilities, and we find that there is a strong effect of the day of the week.

Proposed algorithms. Based on our previous findings—namely, that the platform's problem is computationally hard and that the history effect on the like probabilities is negative and significant—we establish an upper bound for the platform's problem, which can be obtained by solving a linear program. This linear program also serves as a building block for our family of algorithms, which we call *Dating Heuristics* (DH). Using simulations on real data, we show that the proposed heuristics outperform relevant benchmarks, improving the overall match rate by at least 40% relative to our partner's current algorithm.

Field experiment. The simulation results convinced our industry partner to test our algorithm in practice. In collaboration with the company, we designed and implemented a field experiment where we compare the number of matches in a treatment market that uses our algorithm with the number of matches attained in a set of control markets that use our industry partner's algorithm. The results of this field experiment confirm the benefits shown in our simulations. Since then, our algorithm has been launched in at least one other major market, and we are collaborating with the company to expand its use to other markets.

Managerial Implications. Our results provide valuable insights for platforms seeking to improve their search and recommendations systems. We show that designing algorithms to account for the fact that users must mutually like each other can lead to substantial improvement. Hence, platforms should focus on understanding the main drivers of users' preferences to make better recommendations. Second, a major finding is the existence of the history effect, by which users who have recently matched with others are less likely to like new potential partners. Our results show that this finding can be leveraged to generate more matches. Finally, although we focus on a dating market, we believe that these characteristics may also be present in other applications such as online labor markets and accommodation platforms, so our algorithmic framework may prove useful in those settings as well.

The remainder of this paper is organized as follows. Section 2 reviews the closest relevant literature. Section 3 describes our model, and Section 4 describes how the most relevant inputs to our model are estimated. Section 5 introduces our algorithms, and Section 6 reports the simulation results. Finally, Section 7 presents the results of a field experiment that tested our main algorithm in our partner's platform, and Section 8 concludes.

2. Literature

Our work sits at the intersection of several streams of literature. First, our paper contributes to the large literature on assortment optimization. Most of this literature assumes that incoming customers make independent purchasing choices, and that a decision maker must decide which subset of products to offer in order to maximize the expected profit. Talluri and van Ryzin (2004) introduce a general version of this problem, and more recent papers have extended this model to include capacity constraints (Rusmevichientong et al. 2010), different choice models (Davis et al. 2014, Rusmevichientong et al. 2014, Blanchet et al. 2016), search (Wang and Sahin 2018), learning (Caro and Gallien 2007, Rusmevichientong et al. 2010, Sauré and Zeevi 2013), and online selection of personalized assortments (Berbeglia and Joret 2015, Golrezaei et al. 2014). We refer the reader to Kök et al. (2015) for an extensive review the assortment planning literature. The setting we consider in this paper differs from the traditional assortment problem in several ways. First, our paper is one of the first to analyze an assortment problem where transactions (matches) are among users and occur only if users mutually see and like each other. Second, users in our setting have repeated interactions with the platform and can like as many alternatives as they want from their daily assortment; by contrast, most of the assortment optimization literature focuses on settings where consumers are short-lived and limited to one choice. This introduces several complications. For instance, the set of feasible assortments must be updated dynamically depending on users' past decisions. Moreover, due to the existence of the history effect, the probability that a user likes a profile is endogenous to the platform's previous choices, as it depends on the number of recent matches obtained, which in turn depends on the assortments seen by *all users* in the past.

Our paper is also related to the literature on matching platforms and, specifically, on dating platforms. In this context, Kanoria and Saban (2017) study how simple interventions, such as limiting what side of the market reaches out first or hiding quality information, can considerably improve the platform's outcomes. Halaburda et al. (2018) show that two platforms can successfully coexist charging different prices by limiting the set of options offered to their users. These stylized models assume that users leave the platform after matching. We contribute to this literature by modeling more closely how some dating platforms work, as users can accumulate matches and do not necessarily leave the platform once they get matched. Also related to our paper is the empirical literature on understanding user's preferences and behavior in dating markets. Previous papers show that preferences may differ across genders (Fisman et al. 2006, 2008), that there is no evidence that users behave strategically (Hitsch et al. 2010), and that there exist strong assortative patterns (Hitsch et al. 2013). Other papers empirically show the impact of design decisions and information on matching outcomes. For instance, Lee and Niederle (2014) show that dating platforms can increase the number of matches they generate by allowing users to signal their preferences, while Yu (2018) shows that users' beliefs about the market size affect their behavior. We contribute to this literature by empirically showing that the history of past success affects the like behavior of users, and we propose a dynamic assortment optimization algorithm that leverages this finding.

Our paper is also related to the behavioral economics and operations literature on context-dependent preferences (see Tversky and Simonson (1993)), and more specifically to the literature in satiation (see McAlister (1982)). This literature establishes that the history of consumption and interactions can affect the way that choices are made. We contribute to this literature by empirically showing that the context—through the history—can shape users' behavior. In addition, our paper contributes to the practically non-existent literature that analyzes how behavioral aspects can affect optimal assortment decisions (see Ovchinnikov (2018)).[3] To the best of our knowledge, the only exception is Wang (2018), who studies the effect of incorporating prospect theory in consumer choice models. More broadly, our paper contributes to the literature on behavioral operations management by taking one of the classic problems in the field—the assortment problem—and carefully studying how users make their decisions, taking behavioral aspects into account to design solutions and improve the platform's operational efficiency.

Finally, our paper also contributes to the literature on field experiments in online platforms. Most platforms constantly evaluate potential design changes using carefully crafted experiments, and these can be used by researchers to test hypotheses, measure the impact of new algorithms or interventions, among others. Recent examples in the operations management community include Gallino and Moreno (2018), Cui et al. (2019) in e-commerce, Cohen et al. (2018), Singh et al. (2019) in taxi or ride-sharing, Zhang et al. (2017) in education, among others.

3. Model

We now propose a model to capture the problem faced by our industry partner. Consider a discrete set of users, denoted by $\mathcal{I} = \{1, \ldots, I\}$, and a discrete set of periods, $\mathcal{T} = \{1, \ldots, T\}$. Each user $i \in \mathcal{I}$ is associated with a vector of time-invariant characteristics X_i. This vector may include information regarding user i's gender, race, religion, education, among others, and also her declared preferences regarding these characteristics in potential partners. For example, users can declare a preferred age and/or height range, as well as their preferences over other characteristics of their potential partners. In a slight abuse of notation, we denote by X_{ij} the vector that includes X_i, X_j, and also the interaction between i and j's characteristics and preferences for each pair $(i,j) \in \mathcal{I} \times \mathcal{I}$. For instance, this vector may include the age difference between the users, whether they share the same race or religion, or whether the users mutually satisfy their preferences, among others. Using this information, and for each user $i \in \mathcal{I}$ and each period t, the platform computes a set of *potential partners* that includes all users $j \in \mathcal{I}$ such that i and j mutually satisfy their preferences. We denote this set by $\mathscr{P}_i^t$, and we refer to it as the set of *potentials* of user i in period t. (Later on, we refine the definition of the set of potentials to accommodate for additional constraints.)

The sequence of events, summarized in Figure 1, can be described as follows. At each time period $t \in \mathcal{T}$, users can log in and use the platform, in which case we say that they are *active* in that period. Let Υ_i^t denote the random variable representing whether user i is active in period t. We

[3] This is in sharp contrast with other classical problems in operations management, including auction design (see Elmaghraby and Katok (2017) for a comprehensive review), procurement (Engelbrecht-Wiggans and Katok (2006), Tunca and Zenios (2006), Wan and Beil (2009), Beer et al. (2020), among others), and pricing (Katok et al. (2014), Özer and Zheng (2016), Baucells et al. (2017)). See Part III in Donohue et al. (2017) for a general overview of applications of behavioral operations.

Figure 1 **Timeline of Model**

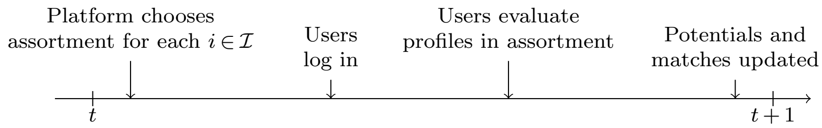

assume that $\Upsilon_i^t \sim Bernoulli(\upsilon_i^t)$ for some exogenous parameter υ_i^t. Moreover, we assume that these log-in random variables are independent across users and periods, and that the parameters υ_i^t are known by the platform for every user $i \in \mathcal{I}$ and every period $t \in \mathcal{T}$.[4] In every period and for each user, the platform selects a limited number of profiles—an *assortment*—taken from the user's set of potentials. Let $S_i^t \subseteq \mathscr{P}_i^t$ be the assortment selected by the platform to be offered to user i in period t. If the user logs in that period, she observes the assortment previously chosen by the platform. Upon being presented with an assortment $S_i^t = S$, user i decides whether to like/not like each profile $j \in S$.[5] We assume that user i makes her decision based on the (random) utility that she gets from getting matched with user j, $U_{ijt} = U(X_{ij}, M_i^t)$, which depends on the time-invariant characteristics of users i and j, X_{ij}, the number of matches that user i had in the recent past, and a random error.[6] To capture dependence on past matches, we define M_i^t to be the number of matches obtained by user i since their last session before period t. Then, user i decides to like j in period t if and only if $U_{ijt} \geq u_{i0}$, where u_{i0} is user i's outside option. In the next section we provide a concrete functional form for $U(X_{ij}, M_i^t)$, which will be validated using our partner's data.

Let $\vec{\Phi}_i^t = \left\{ \Phi_{ij}^t \, : \, j \in S \right\}$ denote the vector of random variables representing whether user i like/not like each profile in the assortment obtained in period t, where

$$\Phi_{ij}^t = \begin{cases} 1 \text{ if } i \text{ likes } j \text{ in period } t. \\ 0 \text{ otherwise.} \end{cases}$$

Then, given $M_i^t = M$ and $S_i^t = S$, we have that

$$\begin{aligned} \mathrm{P}\left(\Phi_{ij}^t = 1 \mid X_{ij}, S_i^t = S,\, M_i^t = M\right) &= \mathrm{P}\left(\Phi_{ij}^t = 1 \mid X_{ij},\, S_i^t = S,\, M_i^t = M,\, \Upsilon_i^t = 1\right) \cdot \mathrm{P}\left(\Upsilon_i^t = 1\right) \\ &= \mathrm{P}\left(U_{ijt} \geq u_{i0} \mid X_{ij},\, S_i^t = S,\, M_i^t = M,\, \Upsilon_i^t = 1\right) \cdot \mathrm{P}\left(\Upsilon_i^t = 1\right) \quad (1) \\ &= \phi_{ij}(S, M) \cdot \upsilon_i^t, \end{aligned}$$

[4] The results in Section 4.3 show that there is a significant effect of the day of the week on the log in probabilities, which implies that these probabilities are time dependent. Moreover, we find that the effect of the match history on log in probabilities is relatively small compared to the day of the week effect, and thus we assume that υ_i^t is exogenous, i.e., independent of the matching history so far.

[5] For simplicity we assume that users can either like or not like a profile, ruling out the skip option that is part of our partner's platform. This is without major loss of generality, since less than 5% of profiles were skipped whenever there was least one evaluation (like or not like) in the assortment.

[6] In Section 4 we extend this and allow for user-time specific unobservables.

where we define $\phi_{ij}(S,M) := \mathrm{P}\left(\Phi_{ij}^t = 1 \mid X_{ij},\, S_i^t = S,\, M_i^t = M,\, \Upsilon_i^t = 1\right)$ to be the probability that user i likes j conditional on logging in, observing assortment S, and having received M matches since their last session.

Following common practice, we assume that users can only evaluate profiles displayed to them (i.e., profiles in their assortment), and thus $\phi_{ij}(S,M) = 0$ whenever $j \notin S$. In addition, we assume that this function is known by the platform for each pair $(i,j) \in \mathcal{I} \times \mathcal{I}$. This assumption is standard in the literature, and it is likely to hold in practice as platforms collect large volumes of data that allow them to estimate these functions very accurately. Finally, we make the following assumption that we keep throughout:

Assumption 1. *For all $(i,j) \in \mathcal{I} \times \mathcal{I}$, and for any two periods $t, t' \in \mathcal{T}$, the decisions Φ_{ij}^t and $\Phi_{ji}^{t'}$ are independent conditional on the vector of time invariant-characteristics X_{ij}, the assortments $S_i^t, S_j^{t'}$, and the corresponding number of matches M_i^t and $M_j^{t'}$.*

Assumption 1 is likely to hold in practice, as users cannot signal their decisions and thus they do not know whether they were already evaluated by the other user and what the outcome was.

A match between users i and j takes place if both users mutually like each other at some point during the entire time horizon. Recall that users need not see each other simultaneously, i.e., user i may see j in one period, and j may see i several periods after that. However, we will assume that users see other users at most once. Let μ_{ij}^t be the random variable denoting whether a match between users i and j takes place in period t. Then, a match between users i and j takes place in period t if and only if one of the following disjoint events occur

$$\left\{\Phi_{ij}^t = 1 \text{ and } \Phi_{ji}^t = 1\right\} \text{ or } \cup_{\tau<t}\left\{\Phi_{ij}^t = 1 \text{ and } \Phi_{ji}^\tau = 1\right\} \text{ or } \cup_{\tau<t}\left\{\Phi_{ij}^\tau = 1 \text{ and } \Phi_{ji}^t = 1\right\}.$$

The first event corresponds to users i and j mutually liking each other in period t. The second event implies that user i likes j in period t, and that j did so in some prior period $\tau < t$. The third event captures the opposite case. Then, the number of matches obtained by user i in period $t+1$ since their last session can be expressed as

$$M_i^{t+1} = \sum_{j \in \mathscr{P}_i} \mu_{ij}^t + (1 - \Upsilon_i^t) \cdot M_i^t. \tag{2}$$

Let $H^t = \left\{\vec{\Upsilon}^\tau, \vec{\Phi}^\tau, \vec{S}^\tau\right\}_{\tau \in \mathcal{T},\, \tau<t}$ be the history of the system at the beginning of period t, where $\vec{\Upsilon}^\tau = \{\Upsilon_i^\tau\}_{i \in \mathcal{I}}$, $\vec{\Phi}^\tau = \left\{\vec{\Phi}_i^\tau\right\}_{i \in \mathcal{I}}$ and $\vec{S}^\tau = \{S_i^\tau\}_{i \in \mathcal{I}}$, and let $\mathcal{H}^t$ be the set of possible histories up to period t. Following our partner's practice, we assume that the assortments must satisfy pre-established constraints that may depend on the history of the system; we describe in detail the constraints considered in Section 5.1. More formally, we let $\mathcal{S}(H^t)$ denote the space of feasible

assortments at time t given the history of the system up to the beginning of period t, H^t. When clear from the context, we will remove the dependence from the history and we will use $\mathcal{S}^t$ to refer to the set of feasible assortments in period t.

An instance of the problem can be fully described in terms of the sets of users, $\mathcal{I}$, the initial sets of potentials, $\{\mathscr{P}_i^1\}_{i\in\mathcal{I}}$, the like probability functions, $\{\phi_{ij}(S,M)\}_{i\in\mathcal{I},\, j\in\mathcal{I}}$, and the log in probabilities, $\{v_i^t\}_{i\in\mathcal{I},t\in\mathcal{T}}$. The objective of the platform is to design a dynamic algorithm that selects a feasible *assortment* to show to each user in each period, in order to maximize the total expected number of matches through the entire horizon, i.e.,

$$\max_{\left\{\vec{S}^t\in\mathcal{S}^t\,:\,t\in\mathcal{T}\right\}} \mathbb{E}\left[\sum_{t\in\mathcal{T}}\sum_{i\in\mathcal{I}}\sum_{j\in\mathscr{P}_i^1}\mu_{ij}^t\right], \tag{3}$$

where the expectation is taken with respect to log in and like realizations. We will refer to the platform's problem described above as the *dynamic two-sided assortment problem*. Notice that the problem is *dynamic* since the platform observes the past history of the system (including like/not like decisions) before deciding which assortments to show in the next period.

Our first result establishes that the dynamic two-sided assortment problem is computationally hard. To show this result, we focus on the following decision-theoretic formulation of our problem: given a number of matches M, are there assortments that result in an expected number of matches of M or more for problem (3)?

Proposition 1. *The decision-theoretic formulation of the dynamic two-sided assortment problem is NP-hard.*

To conclude this section, we would like to highlight that our problem departs from other standard dynamic matching problems in several meaningful ways. First, in contrast with the traditional online matching problems and the ad display problem, users remain in the platform and interact with each other throughout the entire time horizon. Therefore, our problem does not consider uncertainty over future arrivals. Second, in our setting, users will be offered an assortment in each period they log in and will make like/not like decisions in all such periods. Therefore, the dynamics described above will result in users matching multiple times. This is in contrast with most matching problems, which assume that each user can be matched at most once. Third, current period decisions depend on a users' own past decisions and, crucially, on other users' past decisions, as they are a function of the number of matches obtained by a user since their last log-in, and users can be matched with others asynchronously. This introduces complicated market-level dynamics that are typically absent in other matching settings. Finally, it is worth noting that while our focus is on dating markets and our model is intended to closely resemble such interactions, some of the aforementioned characteristics may also be featured in other contexts, such as online labor platforms (e.g., Task Rabbit), hospitality platforms (e.g., Airbnb), car-sharing (e.g. Blablacar), among others. Hence, some of our findings may also translate into such settings.

4. Estimation

To understand how users make their decisions and tailor the heuristics accordingly, in this section we estimate the two main elements of the model described in Section 3. First, we estimate the like probabilities $\phi_{ij}(S, M)$. We focus on estimating the causal effect of the number of matches in the past on users' like decisions. This estimation task is particularly challenging since the number of matches and like decisions can be both correlated with unobserved user characteristics, introducing bias in the estimation. As we later explain in Section 4.2.2, it is also difficult and costly to conduct randomized experiments without interference in this market. Instead, we use a quasi-experimental design to address these challenges and construct an estimator for the history effect without bias. The second element in the model we estimate is the log in probabilities υ_i^t. We first describe our data in Section 4.1 and then provide the details of the estimation of the like probabilities and log in probabilities in Sections 4.2 and 4.3, respectively.

4.1. Data

The data we use in the analysis consists of two parts:

1. User characteristics: for each user, we have their profile information including age, height, education, religion, and race, and their attractiveness measures, which depend on the number of likes and evaluations received in the past.

2. User decisions and backlog queries: for each user in each period, we observe whether the user logged in, and if she did, we observe all the profiles shown to her and also her evaluations. In addition, using this data we compute a set of usage metrics in the recent past, including the number of days active, the number of matches obtained, the number of likes and not likes given and received, among others. We compute each of these metrics for different time windows, including the last session (i.e., the last day in which the user logged in), the last day, the last week, the last month, among others. Finally, we also use this information to determine the *backlog* of each user in each period. The backlog of user i in period t is the set of all users that have liked i in the past (i.e., in any period $\tau < t$) and that i has not seen yet. The backlog is especially relevant because it allows users to see each other asynchronously, generating a match immediately if a user likes a profile from their backlog. For this reason, the algorithm used by our industry partner prioritizes (to some extent) profiles from the backlog. From now on we will refer to profiles taken from the backlog and placed in an assortment as *backlog queries*.

A unique feature of this data is that it allows us to observe the exact assortment offered to each user in each period, including all the characteristics of the profiles involved. In addition, we have access to the full history of interactions between each user and the platform, so we can describe the complete history of each user in each period and include this information in the estimation.

We note that the samples we use for estimating the like probabilities and the log in probabilities are different. For the former, we use data across all markets for a short time window to leverage a quasi-experimental variation that enables us to properly estimate the history effect without bias. In contrast, to estimate log in probabilities we use data from a single market—Houston, TX—which is also used as the treatment market in the field experiment described in Section 7. Finally, throughout the rest of the paper we focus on heterosexual users for simplicity, i.e., users who declared a gender and that they are only interested in users of the opposite gender. Therefore, in the rest of this section we assume that the market has two different sides (one per gender).[7] This is without major loss of generality, as such users represent over 93.74% of the total number of users in the markets where we focus our analysis.

4.2. Like Probabilities

As discussed in Section 3, we assume that user i decides whether to like/not like user j based on the indirect utility U_{ijt} that she perceives from getting matched with j in period t. This utility is not directly observed, and thus needs to be estimated from the data. In particular, we can model the indirect utilities as

$$U_{ijt} = X_{ij}'\beta + M_i^t\gamma + \xi_{it} + \epsilon_{ijt}, \tag{4}$$

where X_{ij} and M_i^t are as defined in Section 3, i.e., M_i^t represents the number of matches obtained by i since her last session, and X_{ij} encodes a set of time-invariant observable characteristics of users i and j. More specifically, we include in X_{ij} three groups of covariates. First, we include time-invariant characteristics of users i and j that are directly observed from the data, including their age,[8] height, race, religion, education level and whether the users have been using the platform for more than 50 days. Second, X_{ij} includes measures of the attractiveness of users, given by the ratio of likes and evaluations received during their time on the platform, and the quintile of attractiveness compared to all users of their same gender. These covariates allow us to account for the absolute and relative attractiveness of users. The last group of covariates in X_{ij} are interactions between the characteristics of users i and j, which allow us to capture preferences that depend on the combination of characteristics, e.g., users looking for partners from the same race, religion,

[7] We do not have enough data to estimate models for both sides separately, and thus we pool the data and control for gender differences.

[8] Since we consider a short horizon, we assume that age is time-invariant.

or of a similar age. Specifically, following Hitsch et al. (2013), we include in X_{ij} the squared positive difference $|x_j - x_i|_+^2$ and the squared negative difference $|x_j - x_i|_-^2$ for each pair of numerical variables x_i and x_j, and the interaction $\mathbb{1}\{d_{ik} = 1,\, d_{jl} = 1\}$ for each pair of categorical variables d_{ik} and d_{jl}. ξ_{it} are user-time specific unobserved characteristics, which can include how many people the user is currently dating (in or outside the platform), whether there is someone with whom the user feels especially connected, whether the user has had positive or negative experiences in the recent past, among other unobservable features that may affect the user's decision. Finally, ϵ_{ijt} are idiosyncratic error terms that are independent and identically distributed and follow extreme value distribution.

We first estimate the utility function using panel logit regressions including user and time fixed effects. The results are presented in Appendix B.1. From this analysis we observe that the number of matches in the recent past has a negative and significant effect on the like probabilities, and that the magnitude of the effect is smaller for older matches. This provides initial evidence supporting the existence of the *history effect*.[9]

4.2.1. Endogeneity. Although the fixed effect regressions account for unobserved user-specific time-invariant confounders, they omit user-time specific unobservables, ξ_{it}, which could lead to biased estimates. As previously discussed, ξ_{it} captures user-time specific unobservables such as user i's recent dating experience both on and outside the platform, which may affect i's willingness to like or not like the profiles seen in period t. In particular, ξ_{it} is likely to be correlated with the number of past matches, leading to biased estimate of the history effect.[10]

Formally, let X, M and ξ be the matrices $\{X_{ij}\}_{(i,j)\in\mathcal{I}\times\mathcal{I}}$, $\{M_i^t\}_{i\in\mathcal{I},t\in\mathcal{T}}$ and $\{\xi_{it}\}_{i\in\mathcal{I},t\in\mathcal{T}}$, respectively. To obtain consistent estimators of β and γ, one necessary condition is that $\mathbb{E}[\xi \cdot X] = \mathbb{E}[\xi \cdot M] = 0$. However, given ξ_i^t is likely to be correlated with M_i^t, this introduces an endogeneity problem that leads to biased estimates of γ.

4.2.2. Quasi-Experiment. Estimating the history effect without bias in the dating market is particularly challenging. In an ideal world, one would randomly assign users to treatment and control groups, and for users in the treatment group would exogenously generate a random number of new matches while keeping the behavior of users on the other side of the market unchanged. Then, comparing the like decisions of users in the following session between the treated and control

[9] We also find a negative and significant effect of the average attractiveness of the other profiles in the assortment, which we refer to as the *assortment effect*. However, the magnitude of this effect is very small and, crucially, considerably smaller than that of the *history effect*, so we decided to focus on the latter.

[10] The number of matches that a user gets also depends on the like behavior of users on the other side of the market and the frequency by which i logs in. The latter not only affects how many profiles user i evaluates—and potentially gets matched with—but also influences how frequently the platform shows his profile on the other side of the market.

groups, we would be able to measure the history effect. Unfortunately, it is difficult to implement this ideal experiment in a dating market, as users must like each other in order to generate a match. In other words, one would need to modify the behavior of the users on the other side of the market to generate matches, which is difficult to implement in practice and can introduce interference to the randomized experiment design. To address this challenge, in this section we describe a quasi-experiment that exploits an unexpected change in the platform's algorithm, which exogenously changed the probability of getting new matches for some users, and we use it to estimate the causal effect of an extra match in users' like behavior.

Background. As we previously discussed, backlog queries are particularly important for generating matches. Each user gets up to three backlog queries each day depending on the size of their backlog and the attractiveness of the users in their backlog. Before May 17th, 2019, backlog queries only included active users, i.e., users whose last log in was within 45 days before the creation of the assortment. Starting on May 17th, 2019, this constraint was removed, so all inactive users in the backlogs were eligible to be used as part of these queries. As a result, some users experienced a change in the number of backlog queries they received relative to the previous sessions, which increased their probability of getting new matches. Since users did not know that this change in the algorithm was implemented or whether they happened to have inactive profiles in their backlog, and since users cannot differentiate a profile that is active from one that is not, the change in the number of new matches due to the change in the algorithm is not correlated with users' behavior or unobserved characteristics conditional on their observed characteristics. Hence, we use whether a user is affected by the change in the algorithm as a quasi-experiment to estimate the history effect.

We argue that using this change in the algorithm as a quasi-experiment provides a good approximation to the ideal experiment described above for several reasons. First, users were not notified that this change in the algorithm was implemented, and users cannot differentiate a profile that is active from one that is not. In addition, whether users happened to have inactive users in their backlog is random conditional on the their observed characteristics such as their attractiveness and how long they have been active on the platform. Thus, the extra matches obtained from inactive profiles can be thought as exogenous to users' behavior conditional on their observed characteristics. Second, the fact that the extra matches are generated from inactive profiles guarantees that the change in the algorithm does not have an impact on the behavior of users on the other side of the market. Finally, since the additional matches involve inactive users, the users getting these extra matches are likely to remain active in the platform in the short run.

Figure 2 Timeline of Quasi-Experiment

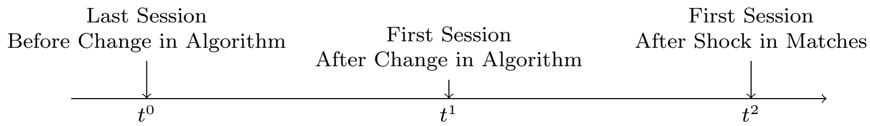

Quasi-Experiment. The set of potentially treated users are those who observed inactive profiles within a short time window after the change in the algorithm, increasing their chances of getting a match. Formally, we define a user as treated if the following two conditions are satisfied:

1. The user received at least one inactive profile as a backlog query.
2. The composition of the assortment shown in terms of the backlog queries is different from what the user would have seen with no change in the algorithm.[11]

The first condition is necessary for the user to be affected by the change in the algorithm. The second condition ensures that the user would not have obtained the same assortment composition of backlog queries if the change in the algorithm was not implemented. In other words, the second condition is necessary for the user's probability of obtaining a match to be affected by the change in the algorithm. We are able to assess the second condition because we have access to our partner's algorithm, and thus we can compute the counterfactual assortment—and its composition in terms of backlog queries—that users would obtain if no change in the algorithm was implemented.

Similarly, we define the set of users in the control group as those who experienced no change in the composition of the assortments due to the change in the algorithm. Notice that this definition allows users in the control group to see inactive profiles taken from their backlogs, as long as the composition of the assortment in terms of backlog queries did not change relative to the counterfactual assortment with no change in the algorithm. We further restrict the control group to be those who had at least one backlog query prior to the quasi-experiment period to make the treatment and control groups more comparable.

The timeline of the quasi-experiment is as follows (see Figure 2). Let t_i^0 be the last time user i logged in before the change in the algorithm, which is the pre-treatment period. t_i^1 and t_i^2 are the first and second time user i logged in after the change in the algorithm, respectively. In other words, t_i^1 is the treatment period, and t_i^2 is the post-treatment period where the outcome is measured. In order to avoid possible responses that might interfere with the quasi-experiment, we restrict the analysis to a short time window around the change in the algorithm. In particular, we consider no more than 3 days around t_i^1 in order to isolate the history effect and avoid the interference

[11] There are different types of backlog queries that depend on characteristics of the users. Hence, a change in composition implies that the number of backlog queries of each type included in the assortment changed.

Table 1 Summary Statistics by Treatment and Period

			Num. Profiles Viewed		Num. Likes		Num. Matches	
Period	Treatment	N	Mean	Std.	Mean	Std.	Mean	Std.
t^0	0	10059	3.615	0.560	1.404	1.182	0.431	0.721
	1	3177	3.528	0.553	1.336	1.152	0.333	0.610
t^1	0	10059	3.588	0.560	1.439	1.189	0.483	0.764
	1	3177	3.558	0.542	1.285	1.177	0.648	0.870
t^2	0	10059	3.541	0.561	1.386	1.175	0.389	0.678
	1	3177	3.510	0.541	1.294	1.118	0.387	0.680

from the response of the other side of the market.[12] As t_i^1 is either May 17th or May 18th for most users, we add the constraints that t_i^0 is no sooner than May 14th and t_i^2 is no later than May 21st. Since the variation in t_i^τ, $\tau \in \{0, 1, 2\}$ across users is small, we drop the i subscript for simplicity, and we exclude from the analysis all users that were not active in at least one of the three periods $\{t^0, t^1, t^2\}$. Given these definitions, our treatment and control groups have 3,177 and 10,059 users, respectively. In Table 1 we present the mean and standard deviation of the number of profiles viewed, the number of likes, and the number of matches per user in period t^0, t^1, and t^2. We find that users in the treatment group received a slightly higher number of matches on average in period t^1 compared to the users in the control group. However, the difference is not statistically significant.

Estimation. Our estimation procedure includes three steps. First, based on the user's observed characteristics and usage metrics in period t^0, we estimate the propensity that the user receives the treatment in the quasi-experiment. Using the estimated propensity scores, we construct weights to be applied in the second and the third steps of the estimation. The second and third steps are similar to the standard two-stage least squares (2SLS) method. In the second step, we use the treatment indicator as an instrumental variable (IV) for the number of matches that users receive in period t^1, and estimate its impact on the like probabilities in period t^2. In this section, we first discuss the validity of the IV, and then describe the details of the three-step estimation procedure.

For our treatment variable to be a valid IV we need the relevance condition and the exclusion restriction to be satisfied. The relevance condition requires that the change in the algorithm led to a higher number of matches for the treated users in period t^1 conditional on the observed user characteristics. This is the second step of our estimation (first-stage of 2SLS), and we present the estimation results later in this section. For the exclusion restriction to be satisfied, we need the

[12] We calculate the number of days between a like by a user (say, like of user i to j) and the appearance of the corresponding backlog query on the other side of the market (in j's assortment) using the data before the change in the algorithm. We find that the difference between the two events is more than 3 days for 80% of the users. We also conduct the analysis using a time window of two days and four days around the change in the algorithm and results remain unchanged.

treatment in the quasi-experiment to be uncorrelated with unobserved user characteristics ξ_{it}. As discussed in the description of the quasi-experiment, we argue that this is a reasonable assumption for the following reasons. First, the change in the algorithm was not announced, so users' like decisions are not correlated with the change in the algorithm. Second, users cannot distinguish between active and inactive profiles, so it is unlikely that they noticed that there was a change in the algorithm in the short-term. Third, conditioning on users' observed characteristics and the propensity score, the treatment is close to randomly assigned. Although more attractive users are more likely to have larger backlogs which potentially include more inactive profiles, and thus may be more likely to be part of the treatment group, we observe a rich set of user characteristics and usage metrics pre-treatment which we use to estimate the propensity score and construct appropriate weights for the 2SLS regressions.

Step 1: Propensity Scores (PS) Estimation. Using data from period t^0, we estimate a logistic model considering the following specification:

$$e_i(X_i, M_i^0) = X_i\beta_0 + \gamma M_i^0 + \epsilon_i,$$

where e_i is the propensity score of user i. In other words, $e_i(X_i, M_i^0) = Pr(W_i = 1|X_i, M_i^0)$, where W_i is the binary treatment indicator. X_i is a matrix of pre-treatment characteristics of user i that includes age, height, education, race, region, attractiveness score, quintile of attractiveness, the number of profiles of different types in the backlog, and the number of backlog profiles the user observed in period t^0. We provide the detailed description of the full set of variables in Appendix B.2. M_i^0 is a vector that contains the number of matches obtained by user i in period t^0, in the previous day, in the previous week, and so on. Using the estimated coefficients, for each user we compute the estimated propensity score, which we denote by $\hat{e}_i$. Finally, to increase the degree of overlap between the distributions of propensity scores among the two groups, we conduct symmetric trimming of $\hat{e}_i$ at the 10% level (see Figure 6 in Appendix B.2).

Step 2: First-Stage Regression of 2SLS. We compute weights ω_i using the estimated propensity scores $\hat{e}_i$ following Hirano, Imbens, and Ridder (2003), i.e.,

$$\omega_i = \begin{cases} \left(\sum_j W_j\right) \cdot \hat{e}_i^{-1} / \sum_{j:W_j=1} \hat{e}_j^{-1} & \text{if } W_i = 1 \\ \left(\sum_j (1-W_j)\right) \cdot (1-\hat{e}_i)^{-1} / \sum_{j:W_j=0} (1-\hat{e}_j)^{-1} & \text{if } W_i = 0 \end{cases}.$$

Using these weights, we estimate the following model,

$$M_i^{t^1} = \theta W_i + X_i\beta_1 + Z_i^{t^1}\delta_1 + \varepsilon_i$$

where $M_i^{t^1}$ represents the number of matches obtained by user i in period t^1, and $Z_i^{t^1}$ is a matrix of observed characteristics of the profiles viewed by user i in period t^1, including their average age, height, education, score, and the fraction of profiles in the assortment that share the same race and religion with the user. Using the estimated parameters from this model, we compute the predicted number of matches for each user in period t^1, denoted by $\hat{M}_i^{t^1}$.

Step 3: Second-Stage Regression of 2SLS. Using the estimated weights in Step 1, we estimate the following model:

$$\Phi_{ij}^{t^2} = \gamma \hat{M}_i^{t^1} + X_{ij}\beta_2 + Z_i^{t^1}\delta_2 + \varepsilon_{ij},$$

where $\Phi_{ij}^{t^2}$ is a binary variable that is equal to 1 if user i liked j in period t^2 and 0 otherwise; X_{ij} is as previously defined in Section 4.2; and $Z_i^{t^1}$ is the matrix of characteristics of the assortment observed in period t^1.

Results. In Appendix B.2 we report the results of the propensity score estimation using different methods, including logit, probit, LASSO-logit, random forests, neural nets, and logit with the optimal specification following Imbens and Rubin (2015). For our main analysis below, we use the estimates from the logit and LASSO-logit (LASSO) models due to their simplicity, and also because they provide more conservative estimates for the history effect. Nevertheless, the results are qualitatively the same for the other methods.

In Table 2 we report the first stage results of 2SLS by estimating a Poisson regression. We use the Poisson regression because the dependent variable takes on discrete values $\{0, 1, 2, \ldots\}$; the results are qualitatively the same if a linear model is used instead. The first column reports the results obtained using the propensity scores estimated from the logit model, and the second column reports the results obtained using the propensity scores estimated from the LASSO model. We observe that, in both cases, the treatment variable is positive and significant, which confirms that our instrument satisfies the relevance condition. In addition, we observe that the estimates are very similar, which suggests that the method used to estimate propensity scores does not play an important role. Finally, we observe that the average marginal effects corresponding to the treatment variable are 0.250 and 0.249 for the logit and LASSO models, respectively. This suggests that users in the treatment group received around 0.25 more matches than those in the control group on average.

In Table 3 we compare the second stage of 2SLS results obtained from using the logit and LASSO models for estimating $\hat{e}_i$. In addition, the columns labeled as *w/o Outliers* report the results obtained after removing users for whom the predicted number of matches from the first stage of 2SLS is above the 99^{th} percentile.

We observe that the coefficient corresponding to $\hat{M}_i^{t^1}$ is negative and significant for all the specifications considered. In addition, depending on the specification, we find that the average marginal effect of $\hat{M}_i^{t^1}$ ranges from -0.036 to -0.044. In other words, an extra match in period t^1

Table 2 First Stage of 2SLS

	Dependent variable: $M_i^{t^1}$	
	Logit	LASSO
Treated	0.639***	0.616***
	(0.033)	(0.032)
Observations	10,931	11,944
Log Likelihood	−7,792.923	−8,717.577
Akaike Inf. Crit.	15,729.850	17,583.150
Note:	*p<0.1; **p<0.05; ***p<0.01	

Table 3 Second Stage of 2SLS

	Dependent variable: Liked (binary)			
	Logit		LASSO	
	Full	w/o Outliers	Full	w/o Outliers
$\hat{M}_i^{t^1}$	−0.170***	−0.170***	−0.162**	−0.239**
	(0.081)	(0.096)	(0.080)	(0.095)
$\bar{X}_{-j}$ Constant	−2.193***	−2.256***	10.222	−1.949***
	(0.627)	(0.630)	(99.674)	(0.599)
Observations	38,203	37,811	41,907	41,477
Log Likelihood	−19,815.710	−19,643.090	−21,318.160	−21,126.570
Akaike Inf. Crit.	39,807.410	39,462.180	42,816.310	42,429.140
Note:	*p<0.1; **p<0.05; ***p<0.01			

reduces the like probability in period t^2 by between -3.6% and -4.4%. These results confirm the existence of the history effect found using the fixed effect regressions presented in Table 7 (see Appendix B.1), but the magnitude of the effect is larger according to this three-step procedure.

4.3. Log In Probabilities

To estimate the log in probabilities, we construct a panel where we include all heterosexual users from Houston, TX, who logged in at least once between February 18th and August 17th, 2020. For each user in the sample, and for each day between March 18th and August 17th, 2020, we record whether or not the user was active. We assume that a user is active if they logged in and they evaluated at least one profile in their assortment. This excludes users who opened the app without the intention of evaluating new profiles, e.g., to continue a conversation with a previous match. We compute a set of usage metrics for different time windows in the recent past, including the number of days the user was active, the number of likes given, the number of matches obtained, among others. In addition, for each user we include time-invariant observable characteristics, including age, height, education, race, attractiveness score, among others.

Since our goal is to obtain a precise prediction of the log in probabilities, in Figure 7 in Appendix B.3 we report the results of five different models, namely, Logit, Probit, Lasso-logit, Random Forest and Neural Nets. We observe that all these models lead to similar results in the test sample, so from now on we focus on the logit model for simplicity. To estimate this model we consider the following specification:

$$\Upsilon_i^t = X_i\alpha_1 + Z_i^t\alpha_2 + \epsilon_i^t, \tag{5}$$

where Υ_i^t and X_i are as defined in Section 3, i.e., X_i are observable characteristics of user i and $\Upsilon_i^t = 1$ if user i is active in period t and $\Upsilon_i^t = 0$ otherwise. Z_i^t is a vector of time-varying observable characteristics, including the day of the week, the number of days since the last active session, the number of days active, the number of likes, the number of evaluations, and the number of matches for different time windows in the past; and ϵ_i^t is an idiosyncratic shock.

In Table 4 we report the estimation results considering all users that logged in at least once in the first month of our sample (i.e., between February 18th and March 18th, 2020).[13] In column (1) we only control for time-invariant characteristics, while in column (2) we add usage metrics. In column (3) we add user fixed effects, while in column (4) we control for both user fixed effects and additional usage metrics. First, we observe that women are significantly less likely to log in compared to men. However, after controlling for the level of activity, the magnitude of the difference becomes smaller (see column (2)). In addition, we observe significant day of the week effects, with Friday and Saturday being the days with the lowest log in rates. We also observe that the coefficient corresponding to the variable *Days since Last Session* is negative and significant for all the models considered, suggesting that users that have been active in previous days are more likely to log in than users that have not been recently active. We leave the micro-foundation of the log in decisions for future work, and we assume for simplicity that the log in probabilities υ_i^t are exogenously given conditional on the user observables. Finally, we also observe that the number of matches since the last session has an unclear effect. Since the magnitude of the effect is small in all specifications, it is without major loss to assume that log in probabilities do not depend on the number of matches in the recent past.

Table 4 Log In Model

	Dependent variable: Log In			
	(1)	(2)	(3)	(4)
Female	-0.617^{***}	-0.109^{***}	-	-
	(0.020)	(0.024)	-	-
Days since Last Session	-0.245^{***}	-0.041^{***}	-0.125^{***}	-0.044^{***}
	(0.002)	(0.001)	(0.001)	(0.001)
Matches since Last Session	-0.063^{***}	0.015	0.038^{***}	0.017
	(0.010)	(0.013)	(0.012)	(0.014)
Monday	0.175^{***}	0.234^{***}	0.211^{***}	0.234^{***}
	(0.018)	(0.021)	(0.020)	(0.021)
Tuesday	0.195^{***}	0.269^{***}	0.238^{***}	0.267^{***}
	(0.018)	(0.021)	(0.020)	(0.021)
Wednesday	0.166^{***}	0.244^{***}	0.218^{***}	0.246^{***}
	(0.018)	(0.021)	(0.020)	(0.021)
Thursday	0.070^{***}	0.112^{***}	0.093^{***}	0.110^{***}
	(0.017)	(0.021)	(0.020)	(0.021)
Saturday	0.027	0.018	0.020	0.017
	(0.018)	(0.021)	(0.020)	(0.021)
Sunday	0.168^{***}	0.206^{***}	0.194^{***}	0.207^{***}
	(0.018)	(0.021)	(0.020)	(0.021)
User Fixed Effects	No	No	Yes	Yes
Additional Usage Metrics	No	Yes	No	Yes
Observations	290,083	290,083	290,232	290,232
McFadden R^2	0.341	0.475	0.516	0.545
Note:		*p<0.1; **p<0.05; ***p<0.01		

[13] We leave out the first month of data in the estimation because we use it to compute the usage metrics.

5. Heuristics

The goal of this section is to introduce a family of algorithms that leverage the main findings from our empirical analysis in Section 4. To this end, we start in Section 5.1 by describing the set of business requirements that define a feasible assortment. In Section 5.2 we provide an upper bound for the optimal value of the platform's problem (3), and in Section 5.3 we present a family of heuristics that can be parametrized via a market-level penalty function. Finally, in Section 5.4 we discuss one specific approach to choose this penalty function.

5.1. Incorporating Business Constraints

Following our partner's practice, we limit the number of potential partners that a user can see in each given day by restricting the assortments to be of (at most) a fixed size K, i.e., $|S_i^t| \leq K$ for some $K \ll I$. In addition, we require that each user sees a potential partner at most once, i.e., $S_i^\tau \cap S_i^t = \emptyset$ for every user $i \in \mathcal{I}$ and every two periods $\tau, t \in \mathcal{T}$, $\tau < t$. Moreover, users are not allowed to see profiles from users who have rejected them in the past.

To express the above constraints, we denote by $\mathscr{P}_i^t$ the set of potentials available for user i at the beginning of period t, and by $\mathscr{B}_i^t$ the subset of potentials that have liked user i in the past, i.e.,

$$\mathscr{B}_i^t = \left\{ j \in \mathscr{P}_i^t \,:\, \Upsilon_j^\tau \cdot \Phi_{ji}^\tau = 1 \text{ for some } \tau < t \right\}.$$

Then, we can write the last two constraints simply as $S_i^t \subseteq \mathscr{P}_i^t$, where we assume that the set of potentials $\mathscr{P}_i^t$ is updated every period by removing both the profiles that were shown to the user in the last period (if any) and also removing all users that disliked user i in the last period. Finally, our partner also imposes additional constraints on which profiles can be part of an assortment; these constraints are additively separable and affine in the set of profiles that are part of the assortment.[14] For instance, they may require that a minimum number of profiles are shown from each user's backlog, or that a minimum number of profiles with some level of attractiveness are included in the assortment. We refer to these constraints collectively as *business constraints*, and we denote the set of business constraints by $\mathcal{L} = \{1, \ldots, L\}$. In Section 5.3 we provide more details on how to incorporate these constraints into the solution of our problem.

5.2. An Upper Bound on the Expected Number of Matches

A major implication of our empirical findings is that the probability that each user i likes a profile j is upper bounded by the like probability when only j is shown in the assortment and when user i has had no matches in the recent past, i.e., $\phi_{ij}(S, M) \leq \phi_{ij}(\{j\}, 0)$. To ease notation, throughout the rest of the paper we use ϕ_{ij} to denote $\phi_{ij}(\{j\}, 0)$.

[14] A set function $f(S)$ is additively separable and affine if $f(S) = \sum_{j \in S} f_j(\{j\})$ and each $f_j(\{j\}) = a_j + b_j \cdot \mathbb{1}_{\{j \in S\}}$.

Following this observation, we propose the following linear program that, as we will establish in Proposition 2 below, can be used to obtain an upper bound to the platform's problem in (3), and that will also play a fundamental role in constructing our heuristics in Section 5.3:

$$
\begin{aligned}
\tilde{\pi} \quad := \quad \max_{x,y,z} \quad & \sum_{t\in\mathcal{T}}\sum_{i\in\mathcal{I}}\sum_{j\in\mathscr{P}_i^1} \upsilon_i^t\cdot\phi_{ij}\cdot y_{ij}^t+\frac{1}{2}\cdot\upsilon_i^t\cdot\upsilon_j^t\cdot\phi_{ij}\cdot\phi_{ji}\cdot z_{ij}^t \\
st. \quad & y_{ij}^t \le \mathbb{1}_{\left\{j\in\mathscr{B}_i^1\right\}}-\sum_{\tau=1}^{t-1} y_{ij}^\tau\cdot\upsilon_i^\tau+\sum_{\tau=1}^{t-1}\left(x_{ji}^\tau-z_{ji}^\tau\right)\cdot\upsilon_j^\tau\cdot\phi_{ji},\ \forall i\in\mathcal{I},\ j\in\mathscr{P}_i^1,\ t\in\mathcal{T} \\
& \sum_{t=1}^{T}\left(x_{ij}^t+y_{ij}^t\right)\cdot\upsilon_i^t\le 1,\ \forall i\in\mathcal{I},\ t\in\mathcal{T} \\
& \sum_{j\in\mathscr{P}_i^1} x_{ij}^t+y_{ij}^t\le K,\ \forall i\in\mathcal{I},\ j\in\mathscr{P}_i^1,\ t\in\mathcal{T} \\
& z_{ij}^t\le x_{ij}^t,\ \forall i\in\mathcal{I},\ j\in\mathscr{P}_i^1,\ t\in\mathcal{T} \\
& z_{ij}^t\le x_{ji}^t,\ \forall i\in\mathcal{I},\ j\in\mathscr{P}_i^1,\ t\in\mathcal{T} \\
& z_{ij}^t= z_{ji}^t,\ \forall i\in\mathcal{I},\ j\in\mathscr{P}_i^1,\ t\in\mathcal{T} \\
& x_{ij}^t, y_{ij}^t\in[0,1],\ \forall i\in\mathcal{I},\ j\in\mathscr{P}_i^1,\ t\in\mathcal{T}.
\end{aligned}
\tag{6}
$$

The decision variables x_{ij}^t and y_{ij}^t may be interpreted as the probability that user i gets a profile j in period t before and after j evaluates i's profile, respectively. In other words, y_{ij}^t represents the probability that i sees j as part of a *backlog query* (that is, j has previously seen i and liked her). By contrast, x_{ij}^t is the probability that i sees j before being evaluated by the latter. In addition, z_{ij}^t represents the probability that users i and j see each other simultaneously in period t. The first constraint captures the definition of y_{ij}^t and the evolution of the (expected) backlog. Specifically, for any period t, the term $\sum_{\tau=1}^{t-1} y_{ij}^\tau\cdot\upsilon_i^\tau$ captures the probability that i saw j as part of a backlog query in the past, and the term $\sum_{\tau=1}^{t-1}\left(x_{ji}^\tau-z_{ji}^\tau\right)\cdot\upsilon_j^\tau\cdot\phi_{ji}$ represents the probability that j liked i in some prior period $\tau<t$ without j being shown to i in the same period τ. Hence, if j is not in the backlog of i at the beginning of the horizon, the probability that j is included in the assortment of user i in period t as a backlog query is limited by the probability that j sees and likes i prior to t and that i has not seen j before period t. On the other hand, if $j\in\mathscr{B}_i^1$, then y_{ij}^t may be one starting from the first period. The second constraint guarantees that each profile is seen with probability at most one in the entire horizon. The third constraint ensures that at most K profiles are shown in expectation to each user in each period. The next three constraints define the variable z_{ij}^t, while the last constraint ensures that x_{ij}^t and y_{ij}^t are valid probabilities. In a slight abuse of notation, we denote by $\mathcal{S}^{UB}$ the polytope defined by all constraints in (6) given the initial state of the system, which can be fully described in terms of the initial sets of potentials $\{\mathscr{P}_i^1\}_{i\in\mathcal{I}}$, the initial backlogs

$\{\mathscr{B}_i^1\}_{i\in\mathcal{I}}$, the initial number of matches $\{M_i^1\}_{i\in\mathcal{I}}$, the initial log in probabilities $\{\upsilon_i^1\}_{i\in\mathcal{I}}$, the like probabilities $\{\phi_{ij}\}_{i,j\in\mathcal{I}}$, and the horizon $\mathcal{T}$. Hence, (6) can be re-written as

$$\begin{aligned} \max_{x,y,z} \quad & \sum_{t\in\mathcal{T}}\sum_{i\in\mathcal{I}}\sum_{j\in\mathscr{P}_i^1} \upsilon_i^t\cdot\phi_{ij}\cdot y_{ij}^t+\frac{1}{2}\cdot\upsilon_i^t\cdot\upsilon_j^t\cdot\phi_{ij}\cdot\phi_{ji}\cdot z_{ij}^t \\ st. \quad & \{\vec{x}^{\,t},\vec{y}^{\,t},\vec{z}^{\,t}\}_{t\in\mathcal{T}}\in\mathcal{S}^{UB}. \end{aligned}$$

Proposition 2. *Let π^* be the optimal value of the platform's problem introduced in* (3)*, and let $\tilde{\pi}$ be an optimal solution to* (6)*. Then, $\pi^*\leq\tilde{\pi}$.*

Notice that this upper bound is not likely to be tight, as it does not take into account neither the history effect nor the business constraints.

5.3. The Dating Heuristics

We present a family of algorithms called Dating Heuristics (DH), which take into account the effect that the assortments chosen in the current period will generate in the future. Each algorithm is defined by a penalty function that aims to capture this effect and, as described in Algorithm 1, each algorithm works in two steps: (i) optimization, and (ii) rounding.

Optimization. The first step is to solve an optimization problem similar to that in (6), but we modify it in four important ways. First, we consider as input the realized state of the system up to the beginning of period t, which can be fully described in terms of the set of potentials $\{\mathscr{P}_i^t\}_{i\in\mathcal{I}}$, the backlogs $\{\mathscr{B}_i^t\}_{i\in\mathcal{I}}$, the number of matches $\{M_i^t\}_{i\in\mathcal{I}}$, and the log in probabilities $\{\upsilon_i^t\}_{i\in\mathcal{I}}$. Second, instead of considering the full horizon, we consider only one period of look-ahead (i.e., considering $\tau\in\{t,t+1\}$), and we replace the second constraint in (6) by $\sum_{\tau=t}^{t+1}x_{ij}^\tau+y_{ij}^\tau\leq 1$.[15] We denote by $\mathcal{S}^t$ the polytope resulting from these changes. Third, we update the objective function to incorporate a penalty to account for the history effect. We provide more details on this penalty in Section 5.4. Finally, we also include the business constraints to satisfy the requirements of our industry partner. As these constraints are additively separable and affine in the profiles that are part of the assortments, we can linearize them as

$$\sum_{j\in\mathscr{P}_i^t} x_{ij}^\tau\cdot a_{ijl}^{\tau,x}+y_{ij}^\tau\cdot a_{ijl}^{\tau,y}\geq b_{il}^\tau,\ \forall i\in\mathcal{I},\ \tau\in\{t,t+1\}\,,\ l\in\mathcal{L},$$

where $a_{ijl}^{\tau,x},a_{ijl}^{\tau,y}$ and b_{il}^τ are constants that depend on the state of the system, namely, the sets of potentials and the backlogs in period t, $\{\mathscr{P}_i^t,\mathscr{B}_i^t\}_{i\in\mathcal{I}}$. The resulting optimization problem can be found in (7) in Algorithm 1.

[15] We chose not to include υ_i^t to increase the chances of obtaining an integer solution.

Rounding. After solving the optimization problem described above, we obtain a solution (x^*, y^*, z^*) that may be fractional. Hence, to decide the assortments to show in period t, the second step in Algorithm 1 is to round the solution obtained for the first period in the horizon, i.e., $x^{*,t}, y^{*,t}$. To do so, we first include profiles in decreasing order of $y_{ij}^{*,t}$ and, if there is space left in the assortment, we add profiles in decreasing order of $x_{ij}^{*,t}$ until the assortment reaches the desired size. By considering this rounding procedure we guarantee that the main business constraints used by our industry partner are satisfied.

Algorithm 1 Dating Heuristic (DH)

Input: $\mathscr{P}_i^t, \mathscr{B}_i^t, M_i^t, \upsilon_i^t, \phi_{ij}$ for each user $i \in \mathcal{I}$, $j \in \mathscr{P}_i^t$, and $\bar{\xi} \leq 0$

Output: An assortment S_i^t for each user $i \in \mathcal{I}$

Step 1. Optimization: solve

$$\begin{aligned}
x^*, y^*, z^* = \underset{x,y,z}{\arg\max} \quad & \sum_{\tau=t}^{t+1} \sum_{i \in \mathcal{I}} \sum_{j \in \mathscr{P}_i^t} \upsilon_i^\tau \cdot \phi_{ij}^t \cdot y_{ij}^\tau + \frac{1}{2} \cdot \upsilon_i^\tau \cdot \upsilon_j^\tau \cdot \phi_{ij}^t \cdot \phi_{ji}^t \cdot z_{ij}^\tau \\
& + \bar{\xi} \cdot \left(\Psi^t \left(\vec{x}^{\,t}, \vec{y}^{\,t}, \vec{z}^{\,t}, \vec{M}^t \right) + \Psi^{t+1} \left(\vec{x}^{\,t}, \vec{y}^{\,t}, \vec{z}^{\,t}, \vec{x}^{\,t+1}, \vec{y}^{\,t+1}, \vec{z}^{\,t+1}, \vec{M}^t \right) \right) \\
st. \quad & \sum_{j \in \mathscr{P}_i^t} x_{ij}^\tau \cdot a_{ijl}^{\tau,x} + y_{ij}^\tau \cdot a_{ijl}^{\tau,y} \geq b_{il}^\tau, \; \forall i \in \mathcal{I}, \; \tau \in \{t, t+1\}, \; l \in \mathcal{L}, \\
& \{\vec{x}^{\,\tau}, \vec{y}^{\,\tau}, \vec{z}^{\,\tau}\}_{\tau \in \{t,t+1\}} \in \mathcal{S}^\tau
\end{aligned} \tag{7}$$

Keep $(x^{*,t}, y^{*,t})$, discard the rest of the solution, and re-define $x^* = x^{*,t}, y^* = y^{*,t}$.

Step 2. Rounding: for each $i \in \mathcal{I}$,

Step 2.a. Set $S_i^t = \underset{\substack{S \in \mathscr{P}_i^t, \, |S| \leq K \\ y_{ij}^* > 0, \, \forall j \in S}}{\arg\max} \left\{ \sum_{j \in S} y_{ij}^* \right\}$

Step 2.b. If $|S_i^t| < K$, then update $S_i^t = S_i^t \cup \underset{\substack{S \in \mathscr{P}_i^t \setminus S_i^t, \, |S| \leq K - \left|S_i^t\right| \\ x_{ij}^* > 0, \, \forall j \in S}}{\arg\max} \left\{ \sum_{j \in S} x_{ij}^* \right\}$

5.4. Penalty

As previously discussed, to decide which assortments to show in period t, our heuristic uses a penalty function that accounts for the negative effect that the matches in each period $\tau \in \{t, t+1\}$ have in future ones. Specifically, our penalty function is designed to capture the effect that matches in each period will have on individual like probabilities in the future periods. We now provide an informal discussion to motivate our choice of penalty function.

First, the number of matches generated in each period depends on the state of the system and the assortment decisions for that period. Therefore, for the optimization problem in Algorithm 1 to remain linear, the penalty function $\Psi(\cdot)$ must be a linear function of these decision variables. To accomplish this, we use the idea behind a first order Taylor expansion to approximate the change in the like probabilities, i.e., for any two values of matches since the last session M_i^τ and $M_i^{\tau+1}$,

$$\phi_{ij}(M_i^{\tau+1}) - \phi_{ij}(M_i^\tau) \approx \left(M_i^{\tau+1} - M_i^\tau\right) \cdot \gamma_{ij}(M_i^\tau), \tag{8}$$

where $\gamma_{ij}(M_i^\tau)$ is the local marginal effect of an extra match in the probability that user i likes profile j when the former has M_i^τ matches. If instead of using the local marginal effect we use the average marginal effect $\bar{\gamma}$, the expected change in the like probabilities in period $t+1$, conditional on the decisions made and the state of the system up to period t, can be approximated by

$$\mathbb{E}\left[\phi_{ij}(M_i^{t+1}) - \phi_{ij}(M_i^t)\middle| \left\{\vec{x}^{\,t}, \vec{y}^{\,t}, \vec{z}^{\,t}\right\}, \vec{M}^t\right] \approx \left(\mathbb{E}\left[M_i^{t+1}\middle| \left\{\vec{x}^{\,t}, \vec{y}^{\,t}, \vec{z}^{\,t}\right\}, \vec{M}^t\right] - M_i^t\right) \cdot \bar{\gamma},$$

where the expectations are taken over the users' decisions in period t given the algorithm (probability over) decisions in t and the number of matches by t. Observe that:

$$\begin{aligned}
\mathbb{E}\left[M_i^{t+1}\middle| \{\vec{x}^{\,t}, \vec{y}^{\,t}, \vec{z}^{\,t}\}, \vec{M}^t\right] &= \mathbb{E}\left[\mathbb{E}\left[M_i^{t+1}\middle| \{\vec{x}^{\,t}, \vec{y}^{\,t}, \vec{z}^{\,t}\}, \vec{M}^t, \Upsilon_i^t\right]\right] \\
&= \upsilon_i^t \left(\sum_{j \in \mathscr{P}_i^t} \left(\phi_{ij}^t \cdot y_{ij}^t + \upsilon_j^t \cdot \phi_{ij}^t \cdot \phi_{ji}^t \cdot z_{ij}^t\right) + \sum_{j \in \mathcal{I}} \upsilon_j^t \cdot \phi_{ji}^t \cdot y_{ji}^t \right) + (1 - \upsilon_i^t) \left(M_i^t + \sum_{j \in \mathcal{I}} \upsilon_j^t \cdot \phi_{ji}^t \cdot y_{ji}^t \right),
\end{aligned}$$

where $\phi_{ij}^t = \phi_{ij}(M_i^t)$ is the probability that i likes j in period t given that i got M_i^t matches since the last session, and the analogous definition holds for ϕ_{ji}^t. The second equality is obtained as follows. The first term corresponds to the case wheres user i logs in period t. Then, the matches since the last session for period $t+1$, M_i^{t+1}, are only those formed in period t; these matches can be obtained by one of the three ways described before. By contrast, if i does not log in (second term), then the number of matches in $t+1$ will be the ones at the beginning of period t, M_i^t, plus those obtained in period t. Note that the latter can only be obtained when another user j, who was previously liked by i, logs in and likes i back in period t. Therefore, we define the penalty function for period t as

$$\begin{aligned}
\Psi^t\left(\vec{x}^{\,t}, \vec{y}^{\,t}, \vec{z}^{\,t}, \vec{M}^t\right) &= \sum_{i \in \mathcal{I}} \mathbb{E}\left[M_i^{t+1}\middle| \{\vec{x}^{\,t}, \vec{y}^{\,t}, \vec{z}^{\,t}\}, \vec{M}^t\right] - M_i^t \\
&= \sum_{i \in \mathcal{I}} \upsilon_i^t \left(\sum_{j \in \mathscr{P}_i^t} \left(\phi_{ij}^t \cdot y_{ij}^t + \upsilon_j^t \cdot \phi_{ij}^t \cdot \phi_{ji}^t \cdot z_{ij}^t\right) + \sum_{j \in \mathcal{I}} \upsilon_j^t \cdot \phi_{ji}^t \cdot y_{ji}^t \right) + (1 - \upsilon_i^t) \left(M_i^t + \sum_{j \in \mathcal{I}} \upsilon_j^t \cdot \phi_{ji}^t \cdot y_{ji}^t \right) - M_i^t.
\end{aligned}$$

Using a similar reasoning, we define

$$\begin{aligned}
\Psi^{t+1}\left(\vec{x}^{\,t}, \vec{y}^{\,t}, \vec{z}^{\,t}, \vec{x}^{\,t+1}, \vec{y}^{\,t+1}, \vec{z}^{\,t+1}, \vec{M}^t\right) &= \sum_{i \in \mathcal{I}} \mathbb{E}\left[M_i^{t+2} - M_i^{t+1} \,|\, \{\vec{x}^{\,\tau}, \vec{y}^{\,\tau}, \vec{z}^{\,\tau}\}_{\tau \in \{t, t+1\}}, M_i^t\right] \\
&\approx \sum_{i \in \mathcal{I}} \sum_{j \in \mathscr{P}_i^t} \upsilon_i^{t+1} \cdot \left(y_{ij}^{t+1} \cdot \phi_{ij}^t + z_{ij}^{t+1} \cdot \upsilon_j^{t+1} \cdot \phi_{ij}^t \cdot \phi_{ji}^t\right) + \sum_{j \in \mathcal{I}} y_{ji}^{t+1} \cdot \upsilon_j^{t+1} \cdot \phi_{ji}^t \\
&\quad - \sum_{i \in \mathcal{I}} \upsilon_i^{t+1} \cdot \left(\sum_{j \in \mathscr{P}_i^t} \upsilon_i^t \cdot \left(y_{ij}^t \cdot \phi_{ij}^t + z_{ij}^t \cdot \upsilon_j^t \cdot \phi_{ij}^t \cdot \phi_{ji}^t\right) + \sum_{j \in \mathcal{I}} y_{ji}^t \cdot \upsilon_j^t \cdot \phi_{ji}^t + M_i^t \cdot (1 - \upsilon_i^t) \right).
\end{aligned}$$

Finally, notice that we use $\bar{\xi}$ as an input to our algorithm. By multiplying the penalty by $\bar{\xi}$, this input allows us to control the relative magnitude of the penalty.

6. Simulations

In this section we numerically evaluate the performance of our algorithm and compare it with relevant benchmarks.

6.1. Setup

Data and Simulation Setting. We use a dataset similar to that described in Section 4, which includes all heterosexual users in Houston, TX, that observed at least 100 potentials between September 1st, 2019 and April 1st, 2020, and that logged in at least once between March 1st and April 1st, 2020. For each of these users, we assume that their initial set of of potentials is formed by the profiles they saw between September 1st, 2019 and April 1st, 2020, and we assume that they had no backlog nor previous matches. As a result, we end up with a market with 852 women and 865 men, who have on average 180.28 and 167.01 potentials available, respectively. We also ran simulations considering different markets, including Austin and Dallas, and different initial conditions for the set of potentials, number of matches since the last session, etc. The results are qualitatively similar to those that will be reported in Section 6.2.

Having defined the market, the next step is to define the like and log in probabilities. To compute the former, we use real data on the characteristics of the users in the sample, and we use the parameters reported in the first column of Table 3 to compute the probabilities. For the latter, we use for simplicity the observed mean values of log in rates in the sample, which are approximately 0.372 for women and 0.537 for men. Finally, for each policy we consider a fixed assortment size of $K = 3$, a time horizon of a week, i.e., $T = 7$, and we consider as business constraints the two most relevant ones according to our industry partner.[16]

Each simulation can be summarized as follows. In each period we start by choosing the assortments that will be shown to each user who logs in. Then, the subset of users who are active in that period is realized, and each of these users makes like/not like decisions about the profiles shown in their assortment. Based on these decisions, we compute the number of matches generated, and we also update the sets of potentials, the backlogs, and the number of matches since the last session for each user. Finally, having updated the state of the system, we move on to the next period and repeat this process until the end of the horizon.

Benchmarks. We compare the performance of our heuristic against the following relevant benchmarks:

1. Partner: implementation of our partner's current algorithm.

[16] We cannot disclose what are these business constraints due to the terms in our NDA.

2. Naive: this benchmark selects, for each user i, the assortment that maximizes the *expected number of likes* in the current period, without considering the probability of being liked back, nor the log in probabilities, nor the history effect on like probabilities.

3. Greedy: this benchmark selects, for each user i, the assortment that maximizes the *expected number of matches* in the current period, without considering the history effect on the like probabilities.

For Naive, Partner and Greedy we consider a constraint on the number of times that each profile is shown in a given period (equal to 10) and on the backlog size of a user to be eligible to be shown (equal to 5). These values are the ones leading to the maximum number of matches for these policies (see Appendix B.4 for further details).

6.2. Results

Our main simulation results are summarized in Figure 3. Figure 3a reports the average number of matches generated over 100 simulations by each policy, while Figure 3b reports boxplots with the improvement of each heuristic relative to Partner and Greedy. In both cases we consider $\bar{\xi} = -0.05$ for DH. We observe that our heuristic considerably outperforms the other benchmarks. Indeed, the improvement of DH is 42.30% relative to our Partner's algorithm, and 13.52% relative to Greedy. In addition, we observe that the improvement of DH relative to Partner and Greedy is always positive and quite substantial. In Appendix B.5.1 we show that these differences remain for different values of the history effect, and in Appendix B.5.2 we show that similar results are obtained in other markets.

To identify the sources of improvement and quantify how much is due to (1) finding better matches and to (2) considering a penalty in the objective, in Figure 3c we plot the average number of matches obtained by Greedy and by DH from 100 simulations (same setup as before) considering different values for $\bar{\xi}$. In addition, we plot the results obtained by DH if we do not take into account the history effect, i.e., if we consider $\bar{\xi} = 0$ and $\phi_{ij} = \phi_{ij}(0)$ for all $i \in \mathcal{I}$ and $j \in \mathscr{P}_i$. The idea of including this additional benchmark, which we label as No History (NH), is to separate the improvement due to better matches from that obtained from including the history effect.

First, we observe that our heuristic with no history effect (NH) considerably outperforms the Greedy policy. This improvement is solely based on finding better matches, as in both cases we do not consider the history effect. In addition, we observe that when $\bar{\xi} = 0$, the DH heuristic generates 718.15 matches on average, which represents an improvement of 0.22% relative to the NH heuristic (which leads to 716.53 matches on average). This improvement is fully explained by taking into account the history effect in the like probabilities, i.e., by using $\phi_{ij}(M_i^t)$ instead of $\phi_{ij}(0)$ in each period. Finally, the improvement obtained with a lower value of $\bar{\xi}$ relative to the case $\bar{\xi} = 0$ is the result of including the penalty in the objective function of our heuristic. In the best case, when $\bar{\xi} = -0.15$, the improvement is 5.76% (759.57 matches on average) relative to the case when $\bar{\xi} = 0$.

Figure 3 **Simulation Results**

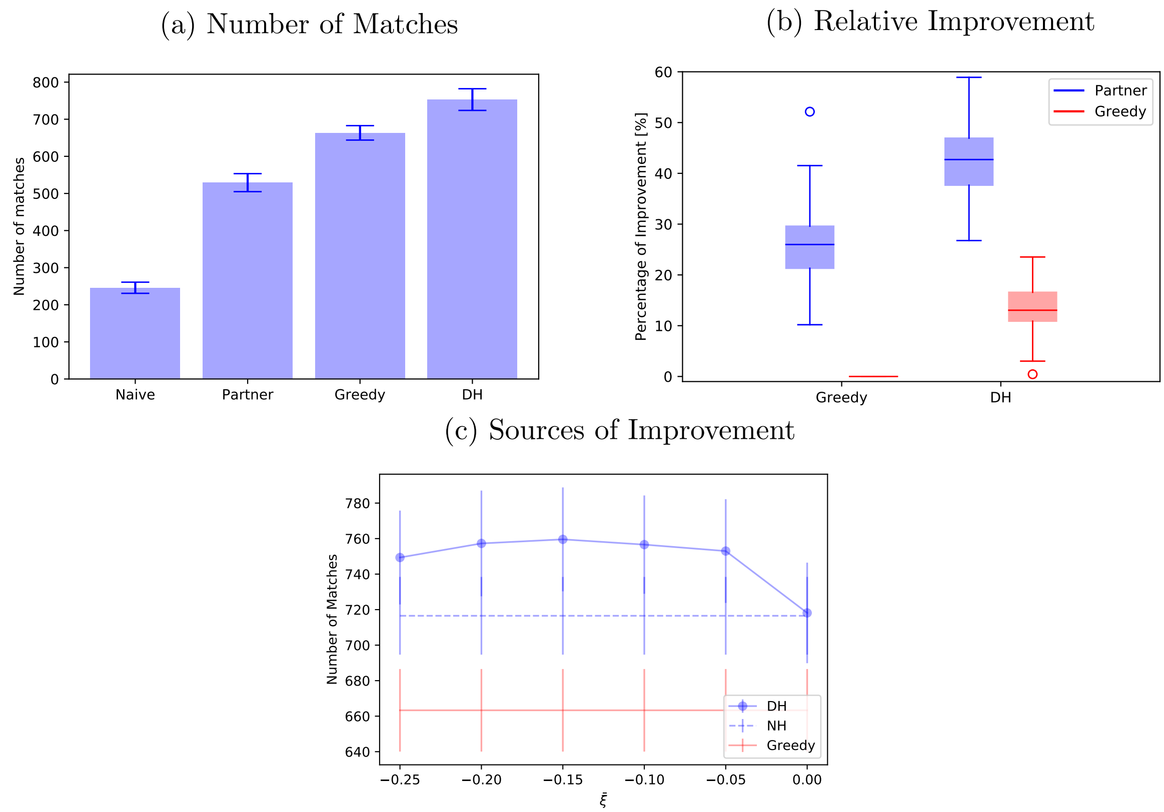

7. Field Experiment

In this section, we describe the results of a field experiment aiming to test if the improvements of our proposed heuristics translate to practice.

7.1. Setup

A field experiment to measure the impact of our heuristic would ideally assign identical markets into treatment and control groups. The experimenter would then offer assortments obtained with our proposed algorithm to users in each of the treatment markets, while keeping the default algorithm in the control markets. Under such a field experiment, a simple comparison of the average number of matches generated in the treatment and control markets would provide an estimate of the causal effect of our proposed algorithm on the number of matches.

In practice, however, there are no two identical markets. As a result, the field experiment we performed approximates this ideal field experiment using a difference-in-differences (DID) design. The DID design allows us to remove biases generated from the differences across markets and across time periods if the parallel trends assumption is satisfied. To ensure the latter, we consider the three largest markets in the state of Texas, namely, Dallas-Fort Worth, Houston, and Austin, and we randomly chose one of these markets—Houston—to be assigned to the treatment group, while

Table 5 Summary Statistics - Treatment and Control Markets

			Age		Score		Login Rate		Like Rate		Match rate	
		N	Mean	Std.	Mean	Std.	Mean	Std.	Mean	Std.	Mean	Std.
Austin	Women	942	32.150	7.438	4.714	2.185	0.459	0.267	0.270	0.251	0.086	0.151
	Men	902	32.776	7.023	2.656	1.569	0.601	0.257	0.482	0.293	0.064	0.116
Dallas-Ft. Worth	Women	1,486	31.632	7.373	4.678	2.230	0.450	0.265	0.247	0.255	0.079	0.146
	Men	1,254	33.199	7.655	2.517	1.493	0.619	0.262	0.471	0.284	0.059	0.108
Houston	Women	1,057	32.503	7.304	4.297	2.230	0.458	0.270	0.246	0.257	0.067	0.142
	Men	950	33.092	7.131	2.481	1.442	0.600	0.260	0.486	0.287	0.059	0.111

the two other markets were assigned to the control group. We chose these three markets due to their geographic proximity and their similarity in the distribution of the main variables of interest. In Table 5, we report the mean and standard deviation of the age, the attractiveness score, and the log in, like, and match rates of users that were active at least once between August 5th and August 19th, 2020,[17] by market and gender. We find that there is no statistically significant differences for all these variables across the three markets. In addition, in Figure 10 (see Appendix B.6.1) we show that the number of matches generated before the intervention are relatively stable for these three markets, suggesting that the parallel trends assumption is likely to hold. A formal test confirming that this assumption holds is also presented in Appendix B.6.1.

7.2. Implementation

Between August 19th and August 25th, 2020, the users in the treatment market received assortments chosen with our heuristic, while the markets in the control group kept using the default algorithm provided by the platform. As an input for our heuristic we use the parameters in the first column of Table 3 to estimate the like probabilities, and the results of the first model with fixed effects (third column in Table 4) to predict the log in probabilities.We chose these parameters because they were the simplest to implement. Moreover, since the estimates are very similar across models, we believe that this choice had no major impact on the results of the experiment.

As a pilot experiment, instead of showing users all the profiles obtained by our heuristic, we randomized a subset of them to be shown, while the remaining profiles were chosen by the default algorithm. In Table 9 (see Appendix B.6.1) we report the total number of profiles shown in the treatment market and the number of profiles shown that were selected by our algorithm each day. On average, 40.91% of the profiles shown during the time window of the experiment were profiles chosen by our algorithm. We expect that increasing the fraction of profiles chosen by our heuristic would increase the size of the improvement. In other words, the results reported in this section provide a lower bound on the effect of the intervention.

[17] We restrict the analysis to heterosexual users that started using the platform before August 5th, 2020.

7.3. Estimation

As discussed above, to estimate the effect of our heuristic we follow a DID approach. Formally, let $W_m = 1$ if market m received the treatment and $W_m = 0$ otherwise, i.e., $W_m = 1$ for Houston, and $W_m = 0$ for Austin and Dallas-Fort Worth. In addition, let $Z_t = 1$ if period t is after the beginning of the experiment (i.e., t is August 19, 2020 or later), and $Z_t = 0$ otherwise. Finally, let M_{mt} be the number of matches generated in market m in period t. Then, the DID estimator can be obtained from estimating the following models:

$$M_{mt} = \alpha_m + Z_t \cdot \gamma + W_m \cdot Z_t \cdot \delta + \epsilon_{mt}, \quad \text{and} \quad M_{mt} = \alpha_m + \lambda_t + W_m \cdot Z_t \cdot \delta + \epsilon_{mt},$$

where α_m are market-specific fixed effects that account for permanent differences between the treated and control markets, γ (λ_t) captures the potential trends affecting both treated and control markets, and δ is the parameter of interest, which captures the treatment effect of the intervention. Notice that the second model allows for more flexible period-specific effects.

7.4. Results

In Figure 4 we plot the number of matches generated in each market in Houston, Austin and Dallas-Fort Worth, between August 12th and August 25th, 2020. The vertical line marks August 19th, 2020, the day when the experiment started in the treated market.

Figure 4 **Matches within Region**

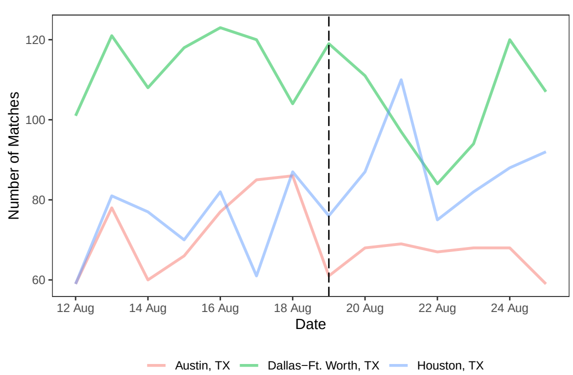

Table 6 **Difference-in-Differences Results**

	Dependent variable: Num. of Matches	
	(1)	(2)
Post	-8.143^{**}	-
	(3.923)	-
Treated × Post	21.429^{***}	21.429^{***}
	(6.795)	(6.267)
Austin	73.429^{***}	59.619^{***}
	(3.397)	(6.000)
Dallas-Ft. Worth	113.143^{***}	99.333^{***}
	(3.397)	(6.000)
Houston	73.857^{***}	60.048^{***}
	(3.923)	(6.267)
Period Fixed Effects	No	Yes
Observations	42	42
R^2	0.988	0.993
Note:		*p<0.1; **p<0.05; ***p<0.01

We observe that, starting from August 19th, the number of matches generated in the treated market considerably increased relative to the previous days. In addition, we find that the increase in the number of matches persisted for all days after the start of the experiment. These results suggest that our algorithm significantly increased the number of matches generated in Houston. In addition, Figure 4 graphically illustrates that the parallel trends assumption holds, as there is no clear difference in the trends experienced by each market before the start of the experiment.

To test if the aforementioned effect is statistically significant, in Table 6 we report the estimation results. The first column provides the results using the first model, while the second column provides the results with fixed effects for each time period. We observe that the variable of interest (Treated×Post) is positive and significant for both models. Moreover, we observe that the DID estimator is equal to 21.429, which provides an estimate of the average number of extra matches that our algorithm produced in the treated market. Comparing this value with the fixed effect corresponding to the Houston market, we observe that our algorithm improved by at least 29.01% the number of matches generated in that market.

We note that the improvement obtained in the field experiment is less than that obtained from the simulations. The main reason is that less than half of the total number of profiles shown during the time of the experiment were chosen by our algorithm. We expect that, once our algorithm is fully implemented in the market, we would obtain an improvement similar to that described in Section 6.2. Nevertheless, these results show that the implementation of our algorithm substantially increased the number of matches generated.

8. Conclusions

Motivated by our collaboration with a dating company, we study how matching platforms should decide on the assortments to show to their users. To accomplish this, we introduce a model of a dynamic matching market mediated by a platform, where users can repeatedly interact with the platform and must mutually like each other to generate a match. Using data from our industry partner, we estimate the main parameters of the model, and we find that the number of matches in the recent past reduces the probability that a user likes other profiles. Based on this finding, we propose a family of algorithms to optimize the set of assortments offered by the platform, and we show through simulations that the proposed algorithms considerably outperform relevant benchmarks. Finally, the results of a field experiment confirm that the aforementioned improvements translate into our partner's platform. This motivated our industry partner to run a second field experiment in another major market and, given the positive results, we are collaborating to deploy our heuristic as their main algorithm to choose which assortments to show in other markets.

Overall, our results showcase that platforms should leverage their knowledge on the drivers of its users' behavior to improve their assortment decisions. Moreover, we believe that insights derived from this work are applicable in other matching contexts. Specifically, it would be interesting to analyze whether similar history effects are present in other online matching markets, such as those for accommodation or labor, and whether similar algorithms to the ones we propose can also prove to be effective in those settings.

Acknowledgments

We would like to thank the data science team of our industry partner for their invaluable help during this project. We are also grateful to Mohsen Bayati, Fuhito Kojima, Tomas Larroucau, Greg Macnamara, Paulo Somaini, Stefan Wager, Xavier Warnes, Gabriel Weintraub, and attendants to the Stanford OIT seminar for their helpful comments.

References

Autor DH (2003) Outsourcing at will: The contribution of unjust dismissal doctrine to the growth of employment outsourcing. *Journal of Labor Economics* 21(1):1–42.

Baucells M, Osadchiy N, Ovchinnikov A (2017) Behavioral anomalies in consumer wait-or-buy decisions and their implications for markdown management. *Operations Research* 65(2):357–378.

Beer R, Rios I, Saban D (2020) Increased transparency in procurement: The role of peer effects.

Berbeglia G, Joret G (2015) Assortment Optimisation Under a General Discrete Choice Model: A Tight Analysis of Revenue-Ordered Assortments.

Blanchet J, Gallego G, Goyal V (2016) A Markov chain approximation to choice modeling. *Operations Research* 64(4):886–905.

Caro F, Gallien J (2007) Dynamic assortment with demand learning for seasonal consumer goods. *Management Science* 53(2):276–292.

Cohen M, Fiszer MD, Kim BJ (2018) Frustration-based promotions: Field experiments in ride-sharing. *SSRN Electronic Journal* .

Cui R, Zhang DJ, Bassamboo A (2019) Learning from inventory availability information: Evidence from field experiments on amazon. *Management Science* 65(3):1216–1235.

Davis JM, Gallego G, Topaloglu H (2014) Assortment Optimization Under Variants of the Nested Logit Model. *Operations Research* 62(2):250–273.

Donohue K, Katok E, Leider S (2017) *The Handbook of Behavioral Operations* (Wiley).

Elmaghraby WJ, Katok E (2017) Behavioral research in competitive bidding and auction design. Donohue K, Katok E, Leider S, eds., *The Handbook of Behavioral Operations* (Wiley).

Engelbrecht-Wiggans R, Katok E (2006) E-sourcing in Procurement: Theory and Behavior in Reverse Auctions with Noncompetitive Contracts. *Management Science* 52(4):581–596.

Fisman R, Iyengar SS, Kamenica E, Simonson I (2006) Gender differences in mate selection: Evidence from a speed dating experiment. *Quarterly Journal of Economics* 121(2):673–697.

Fisman R, Iyengar SS, Kamenica E, Simonson I (2008) Racial preferences in dating. *Review of Economic Studies* 75(1):117–132.

Gallino S, Moreno A (2018) The value of fit information in online retail: Evidence from a randomized field experiment. *Manufacturing & Service Operations Management* 20(4):767–787.

Golrezaei N, Nazerzadeh H, Rusmevichientong P (2014) Real-Time Optimization of Personalized Assortments. *Management Science* 60(6):1532–1551.

Halaburda H, Piskorski MJ, Yildirim P (2018) Competing by Restricting Choice: The Case of Search Platforms. *Management Science* 64(8):3574–3594.

Hitsch G, Hortaçsu A, Ariely D (2013) Online Dating Matching and Sorting in. *American Economic Association* 100(1):130–163.

Hitsch GJ, Hortaçsu A, Ariely D (2010) What makes you click?-mate preferences in online dating. *Quantitative Marketing and Economics* 8(4):393–427.

Ifrach B (2015) How airbnb uses machine learning to detect host preferences. Link, accessed November 11, 2019.

Imbens GW, Rubin DB (2015) *Causal inference: For statistics, social, and biomedical sciences an introduction* (Cambridge University Press).

Kanoria Y, Saban D (2017) Facilitating the Search for Partners on Matching Platforms: Restricting Agents' Actions.

Katok E, Olsen T, Pavlov V (2014) Wholesale pricing under mild and privately known concerns for fairness. *Production and Operations Management* 23(2):285–302.

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

Lee S, Niederle M (2014) Propose with a rose? Signaling in internet dating markets. *Experimental Economics* 18(4):731–755.

Lin M (2018) Online dating industry: The business of love. Link, accessed November 11, 2019.

McAlister L (1982) A Dynamic Attribute Satiation Model of Variety-Seeking Behavior. *Journal of Consumer Research* 9(2):141.

Ovchinnikov A (2018) *Incorporating Customer Behavior into Operational Decisions*, chapter 17, 587–617 (John Wiley and Sons, Ltd).

Özer Ö, Zheng Y (2016) Markdown or everyday low price? the role of behavioral motives. *Management Science* 62(2):326–346.

Rosenfeld M, Thomas RJ, Hausen S (2019) Disintermediating your friends.

Rusmevichientong P, Shen ZJM, Shmoys DB (2010) Dynamic Assortment Optimization with a Multinomial Logit Choice Model and Capacity Constraint. *Operations Research* 58(6):1666–1680.

Rusmevichientong P, Shmoys D, Tong C, Topaloglu H (2014) Assortment optimization under the multinomial logit model with random choice parameters. *Production and Operations Management* 23(11):2023–2039.

Sauré D, Zeevi A (2013) Optimal Dynamic Assortment Planning with Demand Learning. *Manufacturing & Service Operations Management* 15(3):387–404.

Singh J, Teng N, Netessine S (2019) Philanthropic campaigns and customer behavior: Field experiments on an online taxi booking platform. *Management Science* 65(2):913–932.

Talluri K, van Ryzin G (2004) Revenue Management Under a General Discrete Choice Model of Consumer Behavior. *Management Science* 50(1):15–33.

Tunca TI, Zenios SA (2006) Supply Auctions and Relational Contracts for Procurement. *Manufacturing & Service Operations Management* 8(1):43–67.

Tversky A, Simonson I (1993) Context-dependent preferences. *Management Science* 39(10):1179–1189.

Wan Z, Beil DR (2009) RFQ Auctions with Supplier Qualification Screening. *Operations Research* 57(4):934–949.

Wang R (2018) When Prospect Theory Meets Consumer Choice Models: Assortment and Pricing Management with Reference Prices. *Manufacturing & Service Operations Management* 20(3):583–600.

Wang R, Sahin O (2018) The Impact of Consumer Search Cost on Assortment Planning and Pricing. *Management Science* 64(8):3649–3666.

Yu J (2018) Search, Selectivity, and Market Thickness in Two-Sided Markets. Technical report.

Zhang DJ, Allon G, Mieghem JAV (2017) Does social interaction improve learning outcomes? evidence from field experiments on massive open online courses. *Manufacturing & Service Operations Management* 19(3):347–367.

Appendix A: Proofs

A.1. Proof of Proposition 1

To show that the *dynamic two-sided assortment problem* (DTSAP) is NP-hard, we show that we can reduce the exact cover problem (ECP) to DTSAP. An instance of ECP consists of a set of elements $\mathcal{I}=\{1,\ldots,I\}$ and a collection of subsets $\mathcal{S}$ such that $S\subseteq\mathcal{I}$ for all $S\in\mathcal{S}$. The decision problem is whether there is exists an exact cover, i.e., a subset $\mathbf{S}\subseteq\mathcal{S}$ such that $\mathcal{I}\subseteq\cup_{S\in\mathbf{S}}S$ and $S\cap S'=\emptyset$ for all $S,S'\in\mathbf{S}$.

Given an instance $(\mathcal{I},\mathcal{S})$ of ECP we construct the following instance of the DTSAP. There are $I+1$ users, one per each element in $\mathcal{I}$ and an extra user i which will be referred as the focal user. Users that correspond to elements in $\mathcal{I}$ will like user i with probability 1, and dislike all other users with probability 1. User i likes all other users with probability 1. These probabilities are assumed to be constant; specifically, they do not depend on the number of matches.

Let $\bar{K}=\max\{|S|:S\in\mathcal{S}\}$ be the cardinality of the largest subset in $\mathcal{S}$; we let $K=\bar{K}$ and we require that every assortment shown in our problem has cardinality at most K.[18] Let $\underline{K}=\min\{|S|:S\in\mathcal{S}\}$ the cardinality of the smallest subset in $\mathcal{S}$, and let $T=\lceil I/\underline{K}\rceil$ be the time horizon. Define the set of potentials as $\mathscr{P}_i=\mathcal{I}$, for the focal user i and $\mathscr{P}_j=\{i\}$ for all other users j. We require that each user sees a potential partner at most once in the entire time horizon, and we impose the additional constraint that, for every period t, the assortment S_i^t that is be presented to user i must be an element of $\mathcal{S}$.

We show that an exact cover exists for the instance $(\mathcal{I},\mathcal{S})$ if and only if the number of matches produced in the instance of DTSAP we just defined is at least I.

$\Rightarrow$ Suppose that there exists an exact cover $\mathbf{S}=\{S_1,\ldots,S_n\}\subseteq\mathcal{S}$. Then, we construct a feasible solution to the DTSAP that achieves I matches as follows. Our focal user i will be shown assortments $\mathbf{S}'=\{S_1',\ldots,S_n'\}$ in the first n periods, and an empty assortment in the remaining ones. All other users will be shown assortment $\{i\}$ in the first period, and an empty assortment thereafter. By definition, we have that $\mathcal{I}\subseteq\bigcup_{t=1}^n S_t'$, $|S_t'|\le\bar{K}$, and that $S_{t_1}'\cap S_{t_2}'=\emptyset$ for all $t_1\neq t_2$. In addition, we know that $n\le T$ (otherwise $\{S_1,\ldots,S_n\}$ would not be an exact cover). Thus, the constructed solution is feasible. Since $\mathbf{S}'$ also covers $\mathcal{I}$ and the focal user likes all other users with probability 1, and he is shown to and will be liked back by all users corresponding to elements in $\mathcal{I}$, we conclude that the expected number of matches must be greater than or equal to I.

$\Leftarrow$ Suppose that there exists a solution $\bar{\mathbf{S}}'=\left\{\bar{S}_1',\ldots,\bar{S}_T'\right\}$ to the DTSAP that generates at least I matches. Recall that every user corresponding to an element in $\mathcal{I}$ likes user i with probability 1, and likes all other users with probability 0. This implies that the total number of matches must be equal to the number of matches obtained by the focal user i. Moreover, by definition, user i likes all other users with probability 1. Therefore, we have that the total number of matches must be equal to the total number of profiles in $\cup_{t=1}^T S_t'^i$, where $S_t'^i\in\bar{S}_t$ is the assortment shown to the focal user i in period t. Since $\bar{\mathbf{S}}'\subseteq\bar{\mathcal{S}}$ is feasible, we know that the sets $\{S_t'^i\}_{t\le T}$ are disjoint and that $S_t'^i\subset\mathcal{S}$ for all t. Therefore, $|\cup_{t=1}^T S_t'^i|\ge I$ together with the feasibility of the solution allows us to conclude that $\mathbf{S}$ is an exact cover of $\mathcal{I}$.

[18] We could also require that assortments have cardinality exactly K by adding the corresponding “dummy” users to our problem.

A.2. Proof of Proposition 2

To prove Proposition 2 it is enough to focus on the case where $M_i^t = 0$ for all $i \in \mathcal{I}$ and $t \in \mathcal{T}$, since $\phi_{ij}(S,M) \leq \phi_{ij}(\{j\},0) := \phi_{ij}$ and thus we directly obtain an upper bound.

Let the random variables $\{\mathcal{X}_{ij}^t\}_{i,j,t}$ denote whether, under the optimal policy, user i gets a non-backlog profile j in period t. Also, under the optimal policy, let $\mathcal{Y}_{ij}^t = 1$ be the random variable denoting whether user i gets a backlog profile j in period t, and let $\mathcal{Z}_{ij}^t$ be the random variable representing whether users i and j see each other simultaneously in period t be defined as $\mathcal{Z}_{ij}^t = \mathcal{X}_{ij}^t \cdot \mathcal{X}_{ji}^t$. Since the optimal policy is non-anticipating and decides the assortments for period t before observing whether the user logs in, it is direct that these decisions are independent of log in decision in period t, i.e., $\mathcal{X}_{ij}^t, \mathcal{Y}_{ij}^t, \mathcal{Z}_{ij}^t \perp \Upsilon_i^t$ for all $i \in \mathcal{I}$, $t \in \mathcal{T}$, and $j \in \mathscr{P}_i^t$. Notice that $\mathcal{X}_{ij}^t$, $\mathcal{Y}_{ij}^t$, $\mathcal{Z}_{ij}^t$, Υ_i^t and Φ_{ij}^t are random variables, and by independence of users' evaluations we also know that

$$\mathrm{P}\left(\Phi_{ij}^t = 1 \mid \mathcal{X}_{ij}^t = 1, \Upsilon_i^t = 1\right) = \mathrm{P}\left(\Phi_{ij}^t = 1 \mid \mathcal{Y}_{ij}^t = 1, \Upsilon_i^t = 1\right) = \mathrm{P}\left(\Phi_{ij}^t = 1 \mid \mathcal{Z}_{ij}^t = 1, \Upsilon_i^t = 1\right) = \phi_{ij}.$$

By the feasibility of the optimal policy, we know that users see at most K potential partners per period, i.e.,

$$\sum_{j \in \mathscr{P}_i^1} \left(\mathcal{X}_{ij}^t + \mathcal{Y}_{ij}^t\right) \leq K, \; \forall i \in \mathcal{I}, \; t \in \mathcal{T}$$

and users see a given profile at most once in the entire horizon, i.e.,

$$\sum_{t=1}^{T} \left(\mathcal{X}_{ij}^t + \mathcal{Y}_{ij}^t\right) \cdot \Upsilon_i^t \leq 1, \; \forall j \in \mathscr{P}_i^1, \; i \in \mathcal{I}.$$

Also, by definition of $\mathcal{X}_{ij}^t$ and $\mathcal{Z}_{ij}^t$, we know that

$$\mathcal{Z}_{ij}^t \leq \mathcal{X}_{ij}^t, \; \mathcal{Z}_{ij}^t \leq \mathcal{X}_{ji}^t, \; \mathcal{Z}_{ij}^t = \mathcal{Z}_{ji}^t, \; \forall j \in \mathscr{P}_i^1, \; i \in \mathcal{I}, \; t \in \mathcal{T},$$

where the first two inequalities guarantee that $\mathcal{Z}_{ij}^t = 1$ only if both $\mathcal{X}_{ij}^t = 1$ and $\mathcal{X}_{ji}^t = 1$, while the last equality ensures the consistency of $\mathcal{Z}$.

Moreover, we know that j can only be added to the assortment of user i in period t as part of a backlog query if i has not seen j in previous periods and if j has liked i in the past, i.e.,

$$\mathcal{Y}_{ij}^t \leq \mathbb{1}_{\left\{j \in \mathscr{B}_i^1\right\}} - \sum_{\tau=1}^{t-1} \mathcal{Y}_{ij}^\tau \cdot \Upsilon_i^\tau + \sum_{\tau=1}^{t-1} \Upsilon_j^\tau \cdot \Phi_{ji}^\tau \cdot (\mathcal{X}_{ji}^t - \mathcal{Z}_{ji}^t), \; \forall t \in \mathcal{T}.$$

Taking expectation of both sides of all previous inequalities we obtain the following system of inequalities,

$$\begin{aligned}
&\sum_{j \in \mathscr{P}_i^t} \mathbb{E}\left[\left(\mathcal{X}_{ij}^t + \mathcal{Y}_{ij}^t\right)\right] \leq K, \; \forall i \in \mathcal{I}, \; t \in \mathcal{T} \\
&\sum_{t=1}^{T} \mathbb{E}\left[\left(\mathcal{X}_{ij}^t + \mathcal{Y}_{ij}^t\right) \cdot \Upsilon_i^t\right] \leq 1, \; \forall i \in \mathcal{I}, \; j \in \mathscr{P}_i^1 \\
&\mathbb{E}\left[\mathcal{Z}_{ij}^t\right] \leq \mathbb{E}\left[\mathcal{X}_{ij}^t\right], \; \forall i \in \mathcal{I}, \; j \in \mathscr{P}_i^1, \; t \in \mathcal{T} \\
&\mathbb{E}\left[\mathcal{Z}_{ij}^t\right] \leq \mathbb{E}\left[\mathcal{X}_{ji}^t\right], \; \forall i \in \mathcal{I}, \; j \in \mathscr{P}_i^1, \; t \in \mathcal{T} \\
&\mathbb{E}\left[\mathcal{Z}_{ij}^t\right] = \mathbb{E}\left[\mathcal{Z}_{ji}^t\right], \; \forall i \in \mathcal{I}, \; j \in \mathscr{P}_i^1, \; t \in \mathcal{T} \\
&\mathbb{E}\left[\mathcal{Y}_{ij}^t\right] \leq \mathbb{1}_{\left\{j \in \mathscr{B}_i^1\right\}} - \sum_{\tau=1}^{t-1} \mathbb{E}\left[\mathcal{Y}_{ij}^\tau \cdot \Upsilon_i^\tau\right] + \sum_{\tau=1}^{t-1} \mathbb{E}\left[\Upsilon_j^\tau \cdot \Phi_{ji}^\tau \cdot (\mathcal{X}_{ji}^\tau - \mathcal{Z}_{ji}^\tau)\right], \; \forall i \in \mathcal{I}, \; j \in \mathscr{P}_i^1, \; t \in \mathcal{T}
\end{aligned}$$

By as the optimal policy is non-anticipating, we have that $\mathbb{E}\left[\mathcal{X}_{ij}^t \cdot \Upsilon_i^t\right] = \mathbb{E}\left[\mathcal{X}_{ij}^t\right] \cdot \mathbb{E}\left[\Upsilon_i^t\right] = \upsilon_i^t \cdot \mathbb{E}\left[\mathcal{X}_{ij}^t\right]$, and similarly $\mathbb{E}\left[\mathcal{Y}_{ij}^t \cdot \Upsilon_i^t\right] = \upsilon_i^t \cdot \mathbb{E}\left[\mathcal{Y}_{ij}^t\right]$. On the other hand, we can further simplify the last inequality noticing that:

$$\begin{aligned}
\mathbb{E}\left[\Upsilon_j^t \cdot \Phi_{ji}^t \cdot (\mathcal{X}_{ji}^t - \mathcal{Z}_{ji}^t)\right] &= \mathbb{E}\left[\Phi_{ji}^t \cdot (\mathcal{X}_{ji}^t - \mathcal{Z}_{ji}^t) \mid \mathcal{X}_{ji}^t = 1, \Upsilon_j^t = 1\right] \cdot \mathrm{P}\left(\mathcal{X}_{ji}^t = 1, \Upsilon_j^t = 1\right) \\
&= \mathbb{E}\left[\Phi_{ji}^t \cdot \mathcal{X}_{ji}^t \cdot (1 - \mathcal{X}_{ij}^t) \mid \mathcal{X}_{ji}^t = 1, \Upsilon_j^t = 1\right] \cdot \mathrm{P}\left(\mathcal{X}_{ji}^t = 1, \Upsilon_j^t = 1\right) \\
&= \mathbb{E}\left[\Phi_{ji}^t \cdot (1 - \mathcal{X}_{ij}^t) \mid \mathcal{X}_{ji}^t = 1, \Upsilon_j^t = 1\right] \cdot \mathrm{P}\left(\mathcal{X}_{ji}^t = 1, \Upsilon_j^t = 1\right) \\
&= \left(\mathbb{E}\left[\Phi_{ji}^t \mid \mathcal{X}_{ji}^t = 1\right] \cdot \mathbb{E}\left[(1 - \mathcal{X}_{ij}^t) \mid \mathcal{X}_{ji}^t = 1, \Upsilon_j^t = 1\right]\right) \cdot \mathrm{P}\left(\mathcal{X}_{ji}^t = 1, \Upsilon_j^t = 1\right) \\
&= \phi_{ji} \cdot \left(\mathbb{E}\left[(1 - \mathcal{X}_{ij}^t) \mid \mathcal{X}_{ji}^t = 1, \Upsilon_j^t = 1\right]\right) \cdot \mathrm{P}\left(\mathcal{X}_{ji}^t = 1, \Upsilon_j^t = 1\right) \\
&= \phi_{ji} \cdot \mathbb{E}\left[(\mathcal{X}_{ji}^t - \mathcal{Z}_{ji}^t) \mid \mathcal{X}_{ji}^t = 1, \Upsilon_j^t = 1\right] \cdot \mathrm{P}\left(\mathcal{X}_{ji}^t = 1, \Upsilon_j^t = 1\right) \\
&= \phi_{ji} \cdot \mathbb{E}\left[(\mathcal{X}_{ji}^t - \mathcal{Z}_{ji}^t) \mid \mathcal{X}_{ji}^t = 1, \Upsilon_j^t = 1\right] \cdot \mathrm{P}\left(\mathcal{X}_{ji}^t = 1\right) \cdot \mathrm{P}\left(\Upsilon_j^t = 1\right) \\
&= \upsilon_j^t \cdot \phi_{ji} \cdot \left(\mathbb{E}\left[\mathcal{X}_{ji}^t\right] - \mathbb{E}\left[\mathcal{Z}_{ji}^t\right]\right),
\end{aligned}$$

where the second equality follows by the definition of $\mathcal{X}$ and $\mathcal{Z}$, the third follows from the independence of Φ_{ji}^t and $\mathcal{X}_{ij}^t$ conditional on $\mathcal{X}_{ji}^t$, the fourth comes by definition of Φ_{ji}^t, and the second to last equation follows from the non-anticipativity of the optimal policy.

Hence, if we define $\bar{x}_{ij}^t = \mathbb{E}\left[\mathcal{X}_{ij}^t\right]$, $\bar{y}_{ij}^t = \mathbb{E}\left[\mathcal{Y}_{ij}^t\right]$ and $\bar{z}_{ij}^t = \mathbb{E}\left[\mathcal{Z}_{ij}^t\right]$, we observe that $\left\{(\bar{x}_{ij}^t, \bar{y}_{ij}^t, \bar{z}_{ij}^t)\right\}_{j \in \mathscr{P}_i, i \in \mathcal{I}, t \in \mathcal{T}}$ is a feasible solution to our optimization problem.

Furthermore, we know that a match between users i and j takes place in period t if and only if one of three mutually exclusive events occurs:

- $\mathcal{Y}_{ij}^t = 1$, $\Upsilon_i^t = 1$ and $\Phi_{ij}^t = 1$,
- $\mathcal{Y}_{ji}^t = 1$, $\Upsilon_j^t = 1$ and $\Phi_{ji}^t = 1$,
- $\mathcal{Z}_{ij}^t = \mathcal{Z}_{ji}^t = 1$, $\Upsilon_i^t = \Upsilon_j^t = 1$ and $\Phi_{ij}^t = \Phi_{ji}^t = 1$.

The first event captures the case when user j belongs to i's backlog and i gets their profile and likes it. The second event covers the exact opposite case, i.e., i is in j's backlog and the latter likes the former in period t. In the last case, both users see each other simultaneously, and thus a match happens if both users mutually like each other.

$$\begin{aligned}
\mathbb{E}\left[\mu_{ij}^t\right] &= \mathrm{P}\left(\Phi_{ji}^t = 1, \Upsilon_j^t = 1, \mathcal{Y}_{ji}^t = 1\right) + \mathrm{P}\left(\Phi_{ij}^t = 1, \Upsilon_i^t = 1, \mathcal{Y}_{ij}^t = 1\right) + \mathrm{P}\left(\Phi_{ji}^t = \Phi_{ij}^t = 1, \Upsilon_i^t = \Upsilon_j^t = 1, \mathcal{Z}_{ij}^t = \mathcal{Z}_{ji}^t = 1\right) \\
&= \mathrm{P}\left(\Phi_{ji}^t = 1 \mid \mathcal{Y}_{ji}^t = 1, \Upsilon_j^t = 1\right) \cdot \mathrm{P}\left(\mathcal{Y}_{ji}^t = 1, \Upsilon_j^t = 1\right) + \mathrm{P}\left(\Phi_{ij}^t = 1 \mid \mathcal{Y}_{ij}^t = 1, \Upsilon_i^t = 1\right) \cdot \mathrm{P}\left(\mathcal{Y}_{ij}^t = 1, \Upsilon_i^t = 1\right) \\
&\quad + \mathrm{P}\left(\Phi_{ji}^t = 1, \Phi_{ij}^t = 1 \mid \mathcal{Z}_{ij}^t = \mathcal{Z}_{ji}^t = 1, \Upsilon_i^t = \Upsilon_j^t = 1\right) \cdot \mathrm{P}\left(\mathcal{Z}_{ij}^t = \mathcal{Z}_{ji}^t = 1, \Upsilon_i^t = \Upsilon_j^t = 1\right) \\
&= \mathrm{P}\left(\Phi_{ji}^t = 1 \mid \mathcal{Y}_{ji}^t = 1, \Upsilon_j^t = 1\right) \cdot \mathrm{P}\left(\mathcal{Y}_{ji}^t = 1, \Upsilon_j^t = 1\right) + \mathrm{P}\left(\Phi_{ij}^t = 1 \mid \mathcal{Y}_{ij}^t = 1, \Upsilon_i^t = 1\right) \cdot \mathrm{P}\left(\mathcal{Y}_{ij}^t = 1, \Upsilon_i^t = 1\right) \\
&\quad + \mathrm{P}\left(\Phi_{ji}^t = 1 \mid \mathcal{Z}_{ij}^t = 1, \Upsilon_j^t = 1\right) \cdot \mathrm{P}\left(\Phi_{ij}^t = 1 \mid \mathcal{Z}_{ij}^t = 1, \Upsilon_i^t = 1\right) \cdot \mathrm{P}\left(\mathcal{Z}_{ij}^t = 1, \Upsilon_i^t = 1, \Upsilon_j^t = 1\right) \\
&= \upsilon_j^t \cdot \phi_{ji} \cdot \bar{y}_{ji}^t + \upsilon_i^t \cdot \phi_{ij} \cdot \bar{y}_{ij}^t + \upsilon_i^t \cdot \upsilon_j^t \cdot \phi_{ij} \cdot \phi_{ji} \cdot \bar{z}_{ij}^t
\end{aligned} \tag{9}$$

The first equality is by definition of a match and using the fact that the three events are disjoint. The third equality follows from Assumption 1 and the fact that $\mathcal{Z}_{ij}^t = \mathcal{Z}_{ji}^t$. Finally, the last equality is by definition of y_{ij}^t and z_{ij}^t. Hence, we have that

$$\mathbb{E}\left[\sum_{t\in\mathcal{T}}\sum_{i\in\mathcal{I}}\sum_{j\in\mathscr{P}_i}\mu_{ij}^t\right] = \sum_{t\in\mathcal{T}}\sum_{i\in\mathcal{I}}\sum_{j\in\mathscr{P}_i} \upsilon_j^t\cdot\phi_{ji}\cdot\bar{y}_{ji}^t + \upsilon_i^t\cdot\phi_{ij}\cdot\bar{y}_{ij}^t + \upsilon_i^t\cdot\upsilon_j^t\cdot\phi_{ij}\cdot\phi_{ji}\cdot\bar{z}_{ij}^t$$
$$= \sum_{t\in\mathcal{T}}\sum_{i\in\mathcal{I}}\sum_{j\in\mathscr{P}_i} \upsilon_j^t\cdot\phi_{ji}\cdot\bar{y}_{ji}^t + \frac{1}{2}\cdot\upsilon_i^t\cdot\upsilon_j^t\cdot\phi_{ij}\cdot\phi_{ji}\cdot\bar{z}_{ij}^t,$$

where the last equality follows by changing the terms in the summation and avoiding to count matches twice. Notice that this is equivalent to the objective value that the solution $\left\{(\bar{x}_{ij}^t, \bar{y}_{ij}^t, \bar{z}_{ij}^t)\right\}_{j\in\mathscr{P}_i, i\in\mathcal{I}, t\in\mathcal{T}}$ generates for our problem. So, there exists a feasible solution of the LP that leads to an objective value equal to that resulting from (3), and thus we conclude that the optimal objective value of the aforementioned problem provides an upper bound for (3).

Appendix B: Additional Results

B.1. Estimation

Table 7 Logit Model with Fixed Effects

	Dependent variable: Liked		
	(1)	(2)	(3)
Matches $[t-7, t-1]$	-0.118^{***}	-0.119^{***}	-0.118^{***}
	(0.016)	(0.017)	(0.016)
Matches $[t-14, t-8]$	-0.080^{***}	-0.080^{***}	-0.080^{***}
	(0.017)	(0.017)	(0.017)
Matches $[t-21, t-15]$	-0.067^{***}	-0.066^{***}	-0.066^{***}
	(0.018)	(0.018)	(0.018)
Matches $[t-28, t-22]$	-0.049^{***}	-0.047^{***}	-0.049^{***}
	(0.017)	(0.017)	(0.017)
User Fixed Effects	Yes	Yes	Yes
Date Fixed Effects	No	Yes	No
Day of Week Fixed Effects	No	No	Yes
Observations	37,468	37,468	37,468
Note:			*p<0.1; **p<0.05; ***p<0.01

B.2. Estimation of Propensity Scores.

For all the models we consider as covariates the observable characteristics of the user, including gender, region, age, height, education, race, days in the platform, and usage metrics including the number of days active and the number of matches obtained for different time windows in the past.[19] In addition, we include covariates for the backlog composition, i.e., for the number of active and inactive users of each type in the backlog.[20] The results in Figure 5 show that Logit-OPT leads to the best fit. Nevertheless, we observe that all estimates are relatively close to each other. Moreover, Logit and Lasso lead to the most conservative estimates for the history effect in the last stage of our estimation procedure, so using these estimates provides a lower bound on the impact of our heuristics that leverage the history effect.

[19] We include these variables for the day before, $t-1$, and also for the intervals $[t-7, t-2], [t-14, t-8], [t-21, t-15], [t-28, t-22]$, i.e., for each week in the last month.

[20] For LASSO-logit we include second order terms for each of these variables, and also the interaction with gender and region. We also tested using all interactions and the results are relatively similar, although the running time is considerably larger.

Figure 5 Propensity Score Estimation

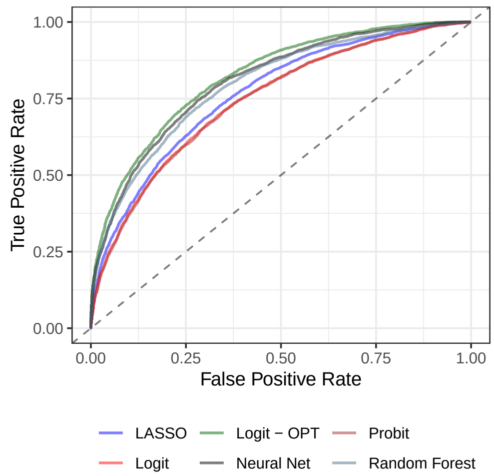

Figure 6 Distribution of Propensity Scores

(a) Original (b) Trimmed

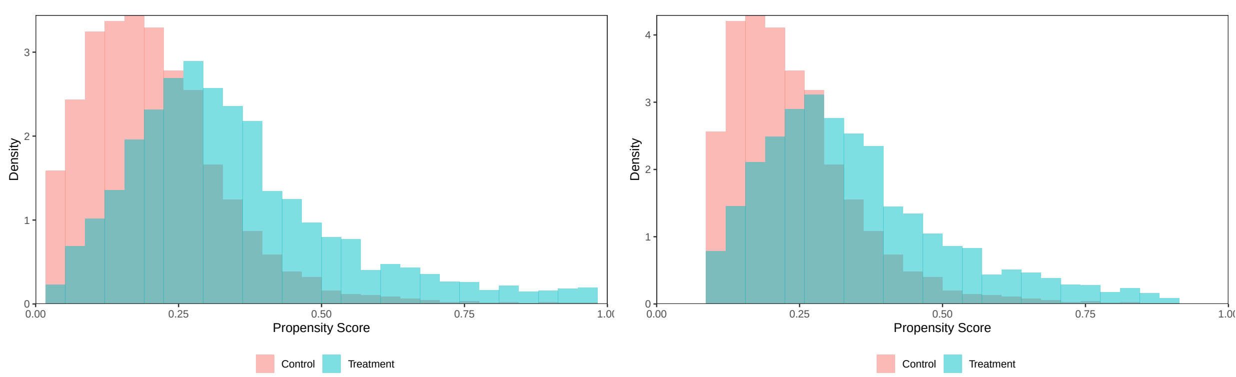

B.3. Log In Probabilties

Figure 7 Log In Models - ROC Curves

(a) Train (b) Test

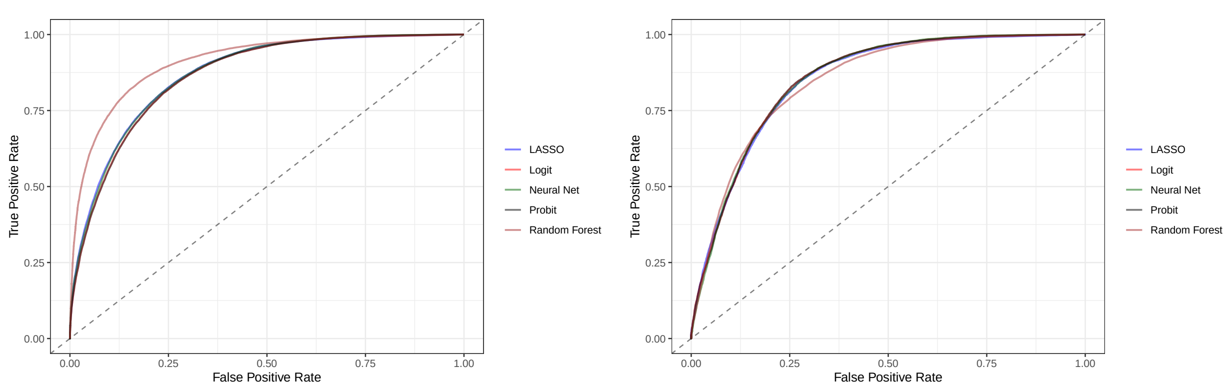

B.4. Sensitivity to Bounds - Greedy and Partner

Our actual implementation of Greedy and our Partner's algorithm considers the addition of bounds on the number of times that each profile is shown in a given period n_s, and also on the maximum backlog size of a user to be eligible to be shown (n_b). We add these bounds in order to alleviate congestions and obtain better results for these benchmarks.

To evaluate the effect of these bounds, we simulate the number of matches obtained by each heuristic varying n_s and n_b. We consider the same data and simulation setup described in Section 6.

In Figure 8 we report the average from 100 simulations considering $m = 3$, $T = 7$, $K = 3$, and no business constraints.

Figure 8 **Sensitivity to Bounds - Partner and Greedy**

(a) Partner (b) Greedy

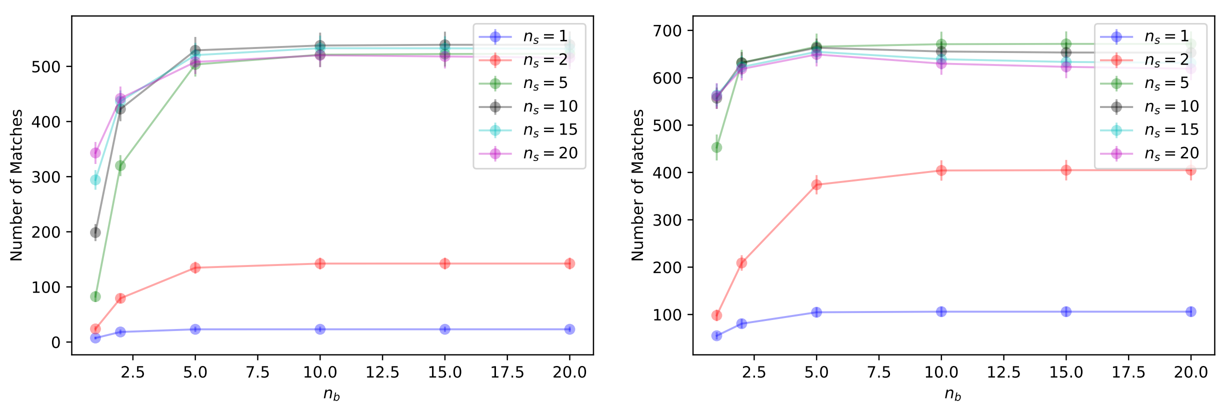

We observe that, for Partner, the highest number of matches is obtained when $(n_s) = 10$ and $n_b \geq 5$, while the maximum for Greedy is achieved when $(n_s, n_b) = (10, 5)$. For this reason, we use the latter set of parameters to run our simulations.

B.5. Simulations

B.5.1. Sensitivity to History Effect. To assess the sensitivity of the results to changes in the magnitude of the history effect, we perform simulations changing the coefficient of the history effect (keeping the rest of the setup fixed), which is multiplied by a factor in $\{0, 0.5, 1, 1.5, 2\}$ (x-axis on the plots). In Figure 9a we plot the average number of matches obtained for each policy considered. As expected, the number of matches is decreasing on the multiplier of the history effect. Following our previous example, we observe that having a history effect twice as large as the one estimated leads to a reduction of 25.4 matches on average (from 752.98 to 724.99), which represents a relative change of 3.86%. In addition, we observe that DH outperforms all the other benchmarks considered for all values of the history effect. Finally, we observe that the difference between the heuristics remains relatively constant as we change the magnitude of the history effect. This is confirmed by Figure 9b, where we report boxplots with the percentage of improvement relative to Greedy and our Partner's algorithm.

Figure 9 Sensitivity History - Matches

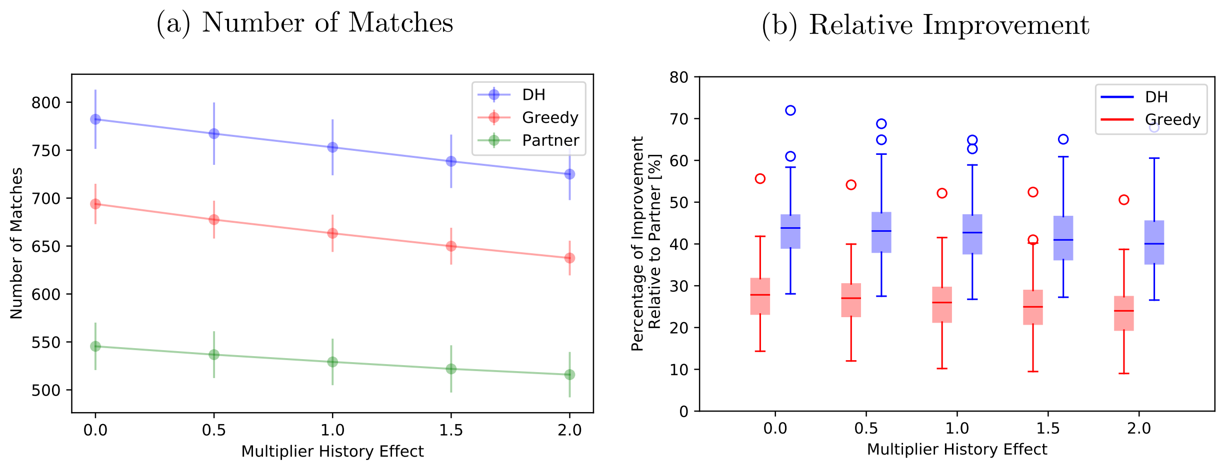

B.5.2. Sensitivity Region. To test the sensitivity of our results to the region considered, in Table 8 we compare the number of matches obtained for different regions. In each region we construct a market following the guidelines described in Section 6. Moreover, we consider the same parameters to compute like probabilities, and we assume that log in rates to be fixed and equal to 0.372 for women and 0.537 for men.

As before, to perform each of these simulations we consider a time horizon of one week, i.e., $T = 7$, and assortments of size $K = 3$. In addition, we consider as login rates and the same business constraints as before. Finally, for DH we consider $\bar{\xi} = -0.05$.

Table 8 Comparison Regions

	Market Size		Heuristics			
	N	N	Partner	Greedy	DH	UB
Austin	754	791	493.26	624.55	689.99	1166.87
			(21.44)	(21.97)	(26.28)	
Houston	852	865	529.14	663.30	752.98	1256.84
			(24.27)	(19.47)	(29.20)	
Dallas	1162	1143	741.64	950.55	1058.08	1766.08
			(27.11)	(23.54)	(35.17)	

We observe that the improvement in the number of matches generated between DH and our partner's algorithm is 39.63% in Austin, 42.30% in Houston, and 42.66% in Dallas. Therefore, we conclude that there are no significant differences in the improvements obtained across different markets.

B.6. Field Experiment

B.6.1. Validity of Difference-in-Difference For the difference-in-difference estimator to be interpreted as causal we need to check the *parallel trends* assumption, i.e., that the difference between the treatment and control groups is constant over time absent of a treatment.

One approach to test if this assumption holds is to check if a simple model of trends fits the data before the experiment. More specifically, we estimate the following model

$$M_{it} = \alpha_i + \beta \cdot t + \epsilon_{it},$$

where α_i are market fixed effects and β represents the time trend. Then, if the predicted values of this model fit the data reasonably well for all markets, then there is evidence supporting that the assumption holds. The solid lines in Figure 10 represent the number of matches obtained for each market between August 5th and August 18th, 2020 (i.e., in the two weeks before the start of the experiment) for each market. On the other hand, the dashed lines represent the predicted values from the previous model. We observe that the predicted lines from the model fit the data of each market reasonable well, suggesting that there are no significant differences between markets in terms of their trends before the experiment.

Figure 10 Pre-Trend Test

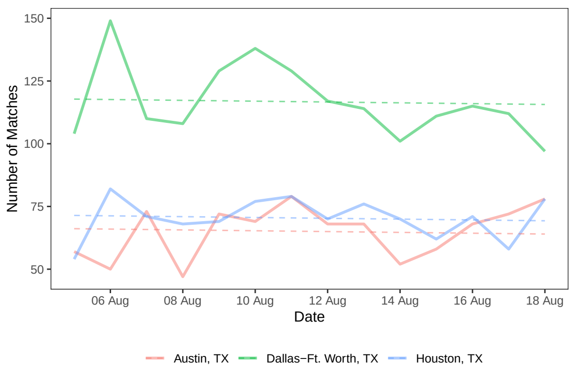

Table 9 Share of Algorithm Potentials Shown

Date	Queries Shown	Total Queries	Percentage
8/19/20	2337	4452	0.525
8/20/20	1286	4911	0.262
8/21/20	1557	4195	0.371
8/22/20	1892	4067	0.465
8/23/20	1836	4216	0.435
8/24/20	2028	4939	0.411
8/25/20	1716	4343	0.395

Another approach to test the parallel trends assumption (following Autor (2003)) is to estimate the following model:

$$M_{it} = \alpha_i + \lambda_t + \sum_{\tau=-p}^{q} W_i \cdot \mathbb{1}_{\{t=t_0+\tau\}} \delta_\tau + \epsilon_{it},$$

where t_0 is the period at which the treatment is being switched on in Houston, TX, p, q are the number of *leads* and *lags* to be included, respectively, δ_τ captures the treatment effect in the τ-th lead or lag. Then, a test for the parallel trends assumption is $\delta_\tau = 0, \forall \tau < 0$. This test can be performed by estimating the aforementioned model and comparing it with a model where $\delta_\tau = 0, \forall \tau < 0$; if the fit provided by the former model is significantly higher, then the parallel trends assumption would not hold. In our case, after estimating the two variants of the model and performing an ANOVA test we obtain a p-value equal to 0.3958, so we cannot reject the hypothesis that the full model provides a better fit. Therefore, we conclude that the parallel trends assumptions holds in our data.