Submitted to *Manufacturing & Service Operations Management*
manuscript

Improving Match Rates in Dating Markets through Assortment Optimization

(Authors' names blinded for peer review)

Problem Definition: Motivated by our collaboration with an online dating company, we study how a platform should dynamically select the set of potential partners to show to each user in each period in order to maximize the expected number of matches in a time horizon, where a match is formed only after two users like each other, possibly in different periods.

Academic/Practical Relevance: Increasing match rates is a prevalent objective of online platforms. We provide insights into how to leverage users' preferences and behavior towards this end. Our proposed algorithm was piloted by our collaborator, a major online dating company in the US.

Methodology: Our work combines several methodologies. We model the platform's problem as a dynamic optimization problem. We use econometric tools and exploit a change in the company's algorithm in order to estimate the users' preferences and the causal effect of previous matches on the like behavior of users, as well as other parameters of interest. Leveraging our data findings, we propose a family of heuristics to solve the platform's problem, and use simulations and field experiments to assess their benefits.

Results: We find that the number of matches obtained in the recent past has a negative effect on the like behavior of users. We propose a family of heuristics to decide the profiles to show to each user on each day that accounts for this finding. Two field experiments show that our algorithm yields at least 27% more matches relative to our industry partner's algorithm.

Managerial Implications: Our results highlight the importance of correctly accounting for the preferences, behavior, and activity metrics of users on both ends of a transaction to improve the operational efficiency of matching platforms. In addition, we propose a novel identification strategy to measure the effect of previous matches on the users' preferences in a two-sided matching market, the result of which is leveraged by our algorithm. Our methodology may also be applied to online matching platforms in other domains.

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

---

1. Introduction

In the past two decades, hundreds of dating services have emerged worldwide, making dating a $12-billion industry worldwide (Lin 2018). Moreover, online dating platforms have become one of the most common channels for couples to meet: 39% of heterosexual couples and 65% of same-sex couples that met in the US in 2017 did so online (Rosenfeld et al. 2019).

A common feature across many dating platforms is that they display a limited number of potential partners' profiles (or simply profiles) to each user on each day. Some platforms, like Tinder and Bumble, implement this by imposing *swipe* limits, others put in place a limit on the number of likes (e.g., Hinge), and still others explicitly limit the number of profiles displayed on each day (e.g., Coffee Meets Bagel). As described on 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 set of profiles—the *assortment*—to display to each user on each day, based on the preferences and characteristics of the users involved. This is the problem we study in this paper.

The aforementioned problem resembles the classic *assortment optimization problem*, where a retailer must decide the set of products to display in order to maximize the expected revenue obtained from a series of customers. However, distinctive features from the dating context make our problem particularly novel. First, both users must mutually agree—by liking each other—to generate a "match," which considerably affects the probability that a transaction occurs. Thus, platforms should consider the preferences and behavior of the users on both ends of a potential match when making assortment decisions. Second, users interact often and repeatedly with the platform, with those living in the same geographical area being part of the same "market." Importantly, users may interact sequentially; i.e., users need not see each other's profile (henceforth, see each other) in the same period. Thus, platforms must carefully manage the timing of these interactions. Notice that some of these features are not exclusive to dating platforms, and may be relevant in other online platforms, including freelancing (e.g., UpWork), ride-sharing (e.g., Blablacar), and accommodation platforms (e.g., Airbnb).[1]

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 the assortments to be shown to their users.[2] Our partner's platform offers a limited set of profiles (ranging from 3 to 9) to each user on each day, and their primary objective is to maximize the number of matches generated. In addition, the assortments offered by the platform must satisfy a series of business constraints; e.g., users can be shown to each other only if they find each other acceptable, no user can see a profile more than once, etc.

Contributions. Our paper combines a variety of methodologies and 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 estimate users' preferences and behavior on the platform, using our partner's data. In particular, we identify an effect of past matches on users' current behavior, and propose a novel identification strategy to estimate it

[1] Indeed, Airbnb improved the booking conversion by 4% by accounting for the preferences of the hosts when deciding which listings to show to guests (Ifrach 2015).

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

without bias. Third, we introduce a class of algorithms that incorporates our estimation findings. Finally, we test the efficiency of our algorithms in two field experiments and find a significant increase in the number of matches. We now describe these contributions in more detail.

Problem formulation. To capture our industry 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, which subset of profiles to show to each user in order to maximize the overall expected number of matches over some horizon. In our model, users log in each period with some time-dependent probability and, conditional on logging in, they observe a set 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. A novel component of our model is that we allow the like decisions to depend on users' past experiences in the platform. If two users 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 computationally hard. Finally, we highlight that our model is general enough to capture a broad array of dating markets and other matching markets.

Estimating users' preferences and behavior from the data. To understand what drives users' behavior on the platform, and to guide the design of our algorithms, we use our industry partner's data to estimate users' preferences and like decisions. 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* on users' behavior, by which users are less likely to like other profiles when they have recently succeeded in obtaining more matches. In order to address the potential endogeneity problem in the estimation, we use a quasi-experiment that introduces exogenous variation in the number matches obtained by some users. Our estimates show that each additional match reduces the probability of a new like by at least 3%. Our identification strategy also provides an important example of using quasi-experiments in matching markets without interference, which is an exciting new area of research.

Proposed algorithms. Based on our previous findings—namely, the estimated like probabilities, the fact 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). Our algorithms differ from current practice by (i) using improved personalized estimates for the like probabilities, (ii) explicitly accounting for the probability that a profile is liked back, and (iii) accounting for the history effect and for the fact that the frequency with which users log in may vary. Using simulations on real data, we show that the proposed heuristics outperform relevant benchmarks, improving the overall match rate by 20% to 45% relative to our partner's current algorithm. Roughly, 80% of the improvement comes from finding better matches (via (i) and (ii) above), and the remaining 20% of the improvement is due to accounting for the history effect.

Field experiments. The simulation results convinced our industry partner to test our algorithm in practice. In collaboration with the company, we designed and implemented two field experiments to 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 partner's algorithm. The results of the field experiments show that the number of matches increased by at least 27%, confirming the benefits shown in our simulations. Given these positive results, we are collaborating with the company to expand the use of our algorithm to other markets.

Managerial implications. Our results provide valuable insights into platforms seeking to improve their search and recommendation systems. Our approach shows that having an algorithm that (i) uses personalized estimates for the like probabilities to account for idiosyncratic differences in taste across users, (ii) accounts for the probability that a profile is liked back, which allows to optimize the utilization of a scarce resource (slots to display profiles) more efficiently, and (iii) accounts for time-dependent user behavior such as the history effect and the varying log-in rates, leads to substantial improvement. Our simulations (in Section 6) attempt to quantify how much each of these features contributes to the improvement, and our field experiments (in Section 7) further validate our approach. Finally, although we focus on a dating market, some of the aforementioned characteristics may also be present in other markets, such as online labor markets, and so our algorithmic framework may prove useful in those settings as well.

The remainder of this paper is organized as follows. Section 1.1 reviews the closest literature. Section 2 describes our partner's platform and the data. Section 3 introduces our model, and Section 4 describes how the like probabilities are estimated. Section 5 presents our algorithms. Section 6 numerically evaluates the performance of our algorithms. Section 7 presents the results of our field experiments, and Section 8 concludes.

1.1. Related Literature

Our work lies 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 of 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 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 example, 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 despite charging different prices by limiting the set of options offered to their users. We contribute to this literature by modeling more closely how some dating platforms work, as users can accumulate matches and do not leave the platform once they obtain a match. Also related to our paper is the empirical literature on understanding users' 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. 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 using a novel identification strategy to empirically show that the history of past success affects the like behavior of users, and by proposing a dynamic algorithm that leverages this finding.

Our paper is also related to the behavioral economics and operations literature on context-dependent preferences (Tversky and Simonson 1993), and, more specifically, to the literature on satiation (McAlister 1982). These literatures establish that the history of consumption and interactions can affect the way that choices are made. We contribute to these literatures by empirically showing that the context—through the history—can shape users' behavior. In addition, our paper contributes to the nascent literature that analyzes how behavioral aspects can affect optimal assortment decisions (Ovchinnikov 2018).[3] To the best of our knowledge, the only paper in this literature is Wang (2018), who studies the effect of incorporating prospect theory into 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 studying how users make their decisions, taking behavioral aspects into account in order to design solutions that improve the platform's operational efficiency.

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

2. Description of the Dating Platform and the Data

Our partner's platform has roughly 800,000 active users in more than 150 geographical markets, and uses the same algorithm in all markets. We now briefly explain how the platform works and describe the data used for our empirical analysis.

2.1. How the Dating Platform Works

When users sign up to use the dating platform, they report some personal information including their age, gender, height, race, religion, education, location, etc. They also declare preferences regarding these characteristics in potential partners. For example, users can declare a preferred age range, height range, a maximum distance from their location, etc. Using this information, the platform computes a set of potential partners (*potentials*, for short) for each user i that includes all users j such that i and j satisfy each other's preferences.

On each day and for each user, the platform selects a limited number of profiles—an *assortment*—taken from the user's set of potentials. If a user logs in during that day, they observe the assortment previously chosen by the platform. Each assortment contains between 3 and 9 profiles (the median is 3; the average is 3.53 with a standard deviation of 0.67). Upon being presented with the assortment,

[3] This is in sharp contrast to other classic 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), 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.

the user decides whether to like or not like each profile in the assortment.[4] A match between two users occurs if both users like each other. When a match is formed, both users are notified.[5]

Importantly, users need not see each other in the same period; i.e., if today's assortment for user j contains user i, this does not imply that j will be included in i's assortment on that same day. In fact, user j may be included in i's assortment in the future (or never). As a result, there are two mechanisms by which matches can be formed. The first mechanism is *simultaneous shows*; i.e., both users see and like each other on the same day. The second one is what we call the backlog. Suppose that user j sees user i, and i has not yet seen j. Then, if j likes i, j is added to i's backlog. Formally, the *backlog* of user i on day t is the set of all users that have liked i in the past (i.e., on any day $\tau < t$) and that i has not yet seen. The backlog is particularly relevant because it allows users to see each other sequentially, generating a match immediately whenever a user likes a profile from their backlog. We will refer to the profiles in the backlog as *backlog profiles*, and to the backlog profiles shown to a user in an assortment as *backlog queries*. Backlogs will play a crucial role in our empirical analysis in Section 4 and in our proposed algorithm.

We have access to the algorithm used by our partner to select the assortments to show to each user on each day, which we will use in our analysis. In addition to guaranteeing that users see profiles only from their set of potentials, the assortments offered by the platform must satisfy a series of additional constraints, which we refer to as *business constraints*. Examples include that users find each other acceptable, that no user can see the same profile more than once, and additional constraints on the composition of the assortment. It is worth noting that the algorithm we will develop also satisfies these constraints.

2.2. Data

We have access to data from all markets in which the platform operates. The data we use in the analysis consists of two parts.

1. *User characteristics*: For each user, we observe their profile information, including their age, height, location, education, religion, race, as well as their attractiveness score, which depends on their evaluations (likes/not likes) received in the past.

2. *User decisions and backlog queries*: For each user on each day, we observe whether the user logged in, and if they did, we observe all the profiles shown to them and their evaluations. Using this data we compute a set of usage metrics for the recent past, including the number of days

[4] For simplicity, we focus on the cases where 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 loss of generality, as less than 5% of profiles were skipped whenever there was at least one evaluation (like or not like) in the assortment.

[5] By opening the notification, users can directly see their new match and start a conversation. Moreover, the app allows users to directly access their matches and conversations *without* observing the newly selected profiles.

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 on which the user logged in), the last day, the last week, the last month, among others. Finally, we also use this data to determine the *backlog* of each user on each day.

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 characteristics of the profiles involved. In addition, we have access to the full history of interactions between each user and the platform, and so we can describe the complete history of each user in each period and include this information in the estimation.

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 (days),[6] $\mathcal{T} = \{1, \ldots, T\}$. Each user $i \in \mathcal{I}$ is associated with a vector of time-invariant characteristics X_i. This vector includes i's personal information (e.g., height, location), and also their declared preferences regarding these characteristics in potential partners (e.g., preferred height range, maximum distance radius). In a slight abuse of notation, we denote by X_{ij} the vector that includes X_i, X_j, and also the interactions between i's and j's characteristics and preferences for each pair $(i,j) \in \mathcal{I} \times \mathcal{I}$. The latter may include the age difference between the users, or whether they share the same race or religion.

Using this information, and for each user $i \in \mathcal{I}$ and each period t, the platform computes an initial set of potential partners that includes all users $j \in \mathcal{I}$ such that i and j satisfy each other's preferences and such that i has not seen j before. The latter is to ensure, as is common in dating platforms, that users see each profile at most once. We denote the initial set of *potentials* of user i by $\mathscr{P}_i^1$, and use $\mathscr{P}_i^t$ to denote the set of potentials of i in period t. Moreover, the platform also knows the *initial backlog* for each user i, i.e., the subset of i's potentials that have liked them in the past. We denote it by $\mathscr{B}_i^1$, and we use $\mathscr{B}_i^t \subseteq \mathscr{P}_i^t$ to denote the backlog of user i at the beginning of period t, i.e., those users that have liked i before t and whose profiles i has not yet seen.

The sequence of events, summarized in Figure 1, can be described as follows. In 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 assume that Υ_i^t follows a Bernoulli distribution with an exogenous time-dependent parameter υ_i^t.

[6] To capture our industry partner's problem and simplify exposition, we focus on a short-time horizon throughout the paper. As a result, we make several assumptions that are reasonable in this context, but that would not hold if we focused on the long run. For example, the fact that the set of users is fixed or that login probabilities do not depend on the history. The short-term assumption is standard in the literature, and it is consistent with our empirical analysis in Section 4, the algorithm introduced in Section 5, and our industry partner's objective. We believe it captures the main features of our problem.

Figure 1 Timeline of the Within-Period Dynamics of the Model

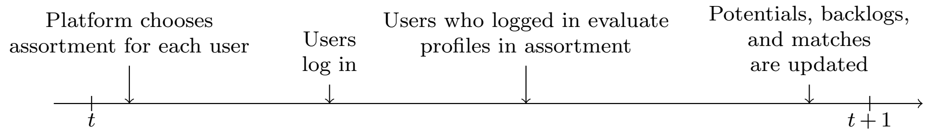

We assume that these variables are independent across users and periods, and that the parameters v_i^t can be estimated acurately by the platform for every user $i \in \mathcal{I}$ and every period $t \in \mathcal{T}$.

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 user i logs in during period t, they observe the assortment previously chosen by the platform, S_i^t, and decide whether to like/not like each profile $j \in S_i^t$. Based on the resulting evaluations and before the period ends, the platform (i) computes the new resulting matches and notifies the corresponding users, (ii) for each user i, it updates their backlog by adding those users that i has not yet seen and who liked i in this period, and (iii) following the constraint that a user can see each profile at most once, it updates the potentials and backlogs of the users who logged in by removing the profiles seen by those users in this period.

We assume that user i makes their like/not like decision for each user j in their assortment based on the (random) utility that they get from matching with user j in period t, U_{ijt}. This utility will naturally depend on the time-invariant characteristics of users i and j, which we denote by X_{ij}. Additionally, a novel component of our model is that we allow U_{ijt} to depend on user i's past experience on the platform; this is consistent with the behavioral observation that the history of interactions can affect the way in which choices are made (e.g., McAlister 1982). For concreteness, we assume that the utility depends on M_i^t, which is defined as the number of matches obtained by user i since their last session before period t; in general, M_i^t can be defined using any measure of previous activities.[7] Examples of such a dependence might include that a user who got more recent matches might become more picky when evaluating the new assortment as they may have less bandwidth to pursue new conversations, or they may be more optimistic about the prospect of landing a date soon.

Based on the previous discussion, we assume that U_{ijt} depends on the characteristics of users i and j, i's number of matches since their last session, and a random error. Hence, we write it as[8]

[7] However, our empirical analysis shows that the current definition is the most relevant measure, as it is the only measure that has a statistically significant effect of a meaningful magnitude.

[8] One may also conjecture that the utility user i gets from being matched with user j can depend on the other users that are shown together with j in the assortment. We empirically tested this and found a negative and significant effect of the average attractiveness of the other profiles in the assortment on the like probabilities. However, the magnitude of this effect is very small and, crucially, considerably smaller than that of the history effect, and so we decided to focus on the latter.

$U_{ijt} = U(X_{ij}, M_i^t)$. Then, user i decides to like j in period t if and only if $U(X_{ij}, M_i^t) \geq u_{i0}$, where u_{i0} is user i's outside option. In Section 4, we provide an expression for $U(X_{ij}, M_i^t)$, which will be validated using our partner's data.

Let $\vec{\Phi}_i^t = \left\{\Phi_{ij}^t : j \in \mathcal{I}\right\}$ denote the vector of random variables representing whether user i liked each profile in period t, i.e., $\Phi_{ij}^t = 1$ if i likes j in period t and 0 otherwise. Following common practice, users can evaluate only profiles displayed to them (i.e., i may like user j in period t only if i logs in during t and j is in i's assortment, i.e., $j \in S_i^t$). Thus, we assume that

$$\Phi_{ij}^t = \begin{cases} 0 & \text{if } \Upsilon_i^t = 0 \text{ or } j \notin S_i^t, \\ 1 \text{ w.p. } \phi_{ij}(M_i^t) & \text{otherwise,} \end{cases}$$

where we define $\phi_{ij}(M)$ to be the probability that user i likes j conditional on logging in, observing an assortment containing j, and having received M matches since their last session, i.e.,

$$\phi_{ij}(M) = \mathrm{P}\left(U_{ijt} \geq u_{i0} \mid X_{ij},\, j \in S_i^t,\, M_i^t = M,\, \Upsilon_i^t = 1\right). \tag{1}$$

We assume that the function $\phi_{ij}(\cdot)$ can be estimated 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 the rest of the section.

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 have already been evaluated by the other user. When estimating the like probabilities in Section 4, we relax Assumption 1 to allow for user–time-specific unobservables.

A match between users i and j takes place if both users 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. Moreover, users can see each other's profile 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. To ease exposition, define $\Phi_{ji}^0 = 1$ for every $j \in \mathscr{B}_i^1$, i.e., every j in i's initial backlog. Then, a match between users i and j takes place in period t if and only if one of the following disjoint events occurs:

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

The first event corresponds to users i and j liking each other in period t. The second event implies that user i likes j in period t, and that j liked i 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

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

An instance of the problem, which we name the *dynamic two-sided assortment problem*, can be fully described in terms of the set of users, $\mathcal{I}$, the initial sets of potentials and backlogs, $\{\mathscr{P}_i^1\}_{i \in \mathcal{I}}$ and $\{\mathscr{B}_i^1\}_{i \in \mathcal{I}}$, the like probability functions, $\{\phi_{ij}(\cdot)\}_{i \in \mathcal{I},\, j \in \mathcal{I}}$, 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 throughout the entire horizon. An algorithm π for the dynamic two-sided assortment problem describes a (possibly randomized) sequence of assortments $\{\vec{S}^{t,\pi}\}_{t=1}^T = \{\{S_i^{t,\pi}\}_{i \in \mathcal{I}}\}_{t=1}^T$ to show to each user in each period, where the choice of assortments for period t may depend on the past history of the system (including which users logged in, the assortments that were shown, and the resulting like/not like decisions) up to the start of period t.

Following our partner's practice, we assume that the assortments must satisfy additional business constraints that may depend on the history of the system; we describe these in more detail in Section 5.1. Formally, 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 it is clear from the context, we will remove the dependence from the history and use $\mathcal{S}^t$ to refer to the set of feasible assortments in period t. We denote by Π the set of all admissible algorithms. The platform's objective is to maximize the total expected number of successful matches over the time horizon,

$$\sup_{\pi \in \Pi} \mathbb{E}\left[\sum_{t \in \mathcal{T}} \sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{I}: j < i} \mu_{ij}^{t,\pi}\right], \tag{3}$$

where the expectation is taken with respect to log-in realizations, like decisions, and possibly random selections of the algorithm, if the latter is not deterministic.

It is worth noting that the above problem could be modeled as a Markov decision problem (MDP), where the state space is given by the set of potentials and the backlogs.[9] However, as described in Section 5, our focus is on generating algorithms that can be easily implemented by our industry partner and that run relatively fast.

To conclude this section, we note that our problem departs from other well-studied dynamic matching problems in several meaningful ways. First, in contrast to traditional online matching

[9] However, one can readily observe that our problem suffers from the so-called "curse of dimensionality." Moreover, consider the following decision-theoretic formulation of our problem: given a number of matches M, are there assortments that result in an expected number of M or more matches for problem (3)? Using a reduction from the exact cover problem, we establish in the appendix that the aforementioned problem is NP-complete.

and Adwords problems, users remain in the platform and interact with each other throughout the entire time horizon. Hence, our problem does not consider uncertainty over future arrivals. Second, in our setting, users get an assortment in each period they log in and make like/not like decisions in all such periods. Thus, the dynamics described above result in users matching multiple times, unlike in many matching problems where users can be matched at most a fixed number of times. Third, current-period decisions depend on a user's 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 sequentially. This introduces complicated market-level dynamics that are typically absent in other matching settings. Finally, while our focus is on dating markets and our model is intended to closely resemble such interactions, some of the above characteristics may be present in other matching markets and our approach may be useful in those settings.

4. Estimation

To understand how users make their decisions and design our algorithms accordingly, we estimate the like probability functions $\phi_{ij}(\cdot)$ introduced in Section 3, which depend on detailed pairwise characteristics of the users. One challenge in the estimation is to recover the causal effect of the number of recent matches on the user's (current) like decision without bias, as both the number of matches and the like decisions may be correlated with unobserved user characteristics. We use a quasi-experimental design to address this challenge and construct an estimator for the history effect. We provide the details of the estimation strategy in Section 4.1, and present the results in Section 4.2.

Before describing our estimation procedure, we note two special features about our empirical setting. First, we use the data described in Section 2.2 as well as knowledge about our partner's algorithm to estimate the history effect without bias. In particular, we have complete knowledge about the set of user characteristics that the platform uses to make the assortment decisions; in all subsequent estimation results, we control for these characteristics. Second, we focus our analysis on heterosexual users, i.e., users who declared a gender and who are only interested in users of the opposite gender, as such users represent 93.7% of the total number of users in the markets where we conduct our analysis. Thus, in the rest of this section, we assume that the market has two different sides (one per gender).[10]

[10] We do not have enough data to estimate a model for users who declared other preferences or to estimate different models for the two sides separately, and thus we pool the data and control for gender differences.

4.1. Like Probability Estimation

As discussed in Section 3, we assume that user i decides whether to like or not like user j based on the utility $U_{ijt} = U(X_{ij}, M_i^t)$ that i derives 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 model the U_{ijt} 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 their 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, including their height, race, religion, education level, etc. Second, we include measures of the attractiveness of users i and j, including their attractiveness score (ratio of likes to evaluations received during their time on the platform) and their quintile of attractiveness compared to all users of the same gender. These covariates provide a measure of the absolute and relative attractiveness of the users. The last group of covariates in X_{ij} are interactions between the characteristics of users i and j. For example, user i might be looking for partners from the same race, religion, or age group. Following Hitsch et al. (2013), for each numerical variable x_k included in X_{ij}, we include the squared positive difference $|x_{jk} - x_{ik}|_+^2$ and the squared negative difference $|x_{jk} - x_{ik}|_-^2$ and, for each categorical variable d_k and each pair of values l, l' included in X_{ij}, we include the interaction $\mathbb{1}\{d_{ikl} = 1,\ d_{jkl'} = 1\}$. The rich set of user i's and user j's characteristics and their interactions included in the utility function allow the like probabilities to be ij-specific, which we incorporate into our algorithms to generate personalized assortments. ξ_{it} are unobserved user–time-specific characteristics, which can include the quality and attractiveness of a recently uploaded picture, how many people the user is currently dating, whether there is someone to whom the user feels especially connected, whether the user has had positive or negative dating experiences in the recent past, among other unobservable features that may affect the user's decision. Finally, ϵ_{ijt} are independent and identically distributed error terms that follow an extreme value distribution.

We first estimate the like probabilities using panel logit regressions, including user and time fixed effects. The results are presented in Appendix B.1. We observe that the number of matches in the recent past has a negative and significant effect on the like probabilities. This provides suggestive evidence supporting the existence of the *history effect*.

4.1.1. Challenge in Estimation Although the fixed effect regressions account for user-specific and time-specific unobservables, they potentially omit endogeneity issues caused by the user–time-specific unobservables, ξ_{it}, which can lead to biased estimates. As previously discussed,

ξ_{it} captures user–time-specific unobservables such as user i's recent dating experiences both on and outside the platform, which are potentially correlated with the number of matches user i obtained since their last session, M_i^t. These user–time-specific unobservables may also affect i's willingness to like or not like the profiles seen in period t. For example, if user i had a positive dating experience outside the platform in the recent past, they might have become pickier when evaluating profiles on the platform in recent periods, which may have led to fewer matches since their last session at the beginning of period t, M_i^t. This effect and the picky attitude could last for more than one period, and user i might also like fewer profiles shown to them in period t. As a result, directly estimating Equation (4) would underestimate the magnitude of the history effect because ξ_{it} is positively correlated with M_i^t. In other words, M_i^t is endogenous.

Estimating the history effect without bias in the dating market is particularly challenging. In an ideal world, one could run an experiment that randomly assigns users to treatment and control groups, where users in the treatment group would receive extra matches compared to those in the control group. Then, comparing between the treatment and control groups in terms of the like decisions of the users in the following session, one would be able to measure the history effect. Unfortunately, one cannot implement this randomized experiment in a dating market, as users must like each other to generate a match. In other words, to generate matches for a user in the market, one would need to manipulate the like decisions of other users, which is unfeasible to implement in practice. More importantly, changing the like decisions of those users might introduce interference in the randomized experimental design: some of these users might be in the control group, changing these users' like decisions about the treated users might affect their like decisions about other users in the control group, etc. This is a particularly challenging problem and, to the best of our knowledge, it has not been studied much in the literature.

To address this challenge, we utilize a quasi-experiment generated by a change in the platform's algorithm that exogenously changed the probability of getting new matches for some users, and that had a limited impact on other users. We use this quasi-experiment to estimate the causal effect of an extra match on users' subsequent like decisions. Our analysis also provides an important example that shows that, while randomized experiments and quasi-experiments in general suffer from interferences in two-sided market settings, properly designed experiments can still be used to estimate treatment effect without bias.

4.1.2. Quasi-Experiment As we discussed in Section 2.1, backlog queries are particularly important for generating matches. Most users get up to three backlog queries each day depending on the size and composition of their backlog. Our partner's algorithm ranks the backlog profiles in terms of the attractiveness score, and uses this metric to decide which backlog profiles to show

Figure 2 Timeline of Quasi-Experiment

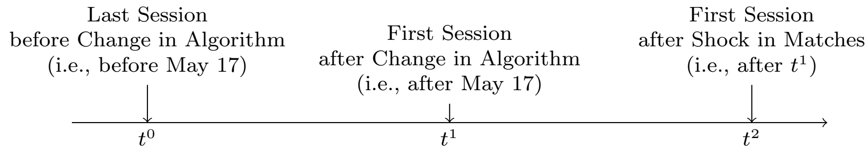

(more attractive profiles are shown first). Before May 17, 2019, backlog profiles *eligible to be shown* only included *active users*, i.e., users whose last log in was within 45 days of the creation of the assortment. We also use "active" to refer to users who log in in the model. Starting on May 17, 2019, this constraint was removed, and so many *inactive users* in the backlogs became eligible to be shown.

As a result, some users whose backlog contained inactive profiles experienced a change in the assortment and, possibly, in the number of backlog queries they received. This occurred when either (i) the user had no active backlog profiles with high enough priority (relative to other business constraints) to be shown, but had an inactive backlog profile that was now eligible to be shown; or (ii) the attractiveness of any of those inactive backlog profiles was above the attractiveness of the active backlog queries that would have been shown before the change in the algorithm. This change, in turn, increased those users' probability of getting new matches instantaneously (and thus before their next session), if they liked any of those inactive backlog queries. For some users, however, although they had inactive profiles in their backlog after the change in the algorithm, they did not receive them in their assortments as those inactive backlog profiles remained ineligible to be shown due to other reasons. Importantly, as we will later explain in detail, some of the reasons for the ineligibility are uncorrelated with the focal users' characteristics or behavior on and outside the platform. We use the change in the algorithm and the eligibility of the inactive backlog profiles as a quasi-experiment to estimate the history effect. Notice that, as the additional matches are formed with inactive users, they do not introduce interference with the other active users in the market in a short period of time. This is key in our quasi-experimental design.

Next, we describe the quasi-experiment in detail. The timeline of the quasi-experiment is as follows (see Figure 2). Let t_i^0 be the last time user i logged in to see a new assortment before the change in the algorithm on May 17, which is the pre-treatment period. t_i^1 and t_i^2 are the first and second times user i logged in to see a new assortment 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. To avoid potential interference in the quasi-experiment, we restrict the analysis to a short time window around the change in the algorithm on May 17. Specifically, we consider no more than three days around t_i^1 to estimate the history effect without bias and, in particular, to avoid the interference from the other users responding to the change in the focal users' like decisions in the quasi-experiment.[11] As t_i^1 is either May 17 or May 18 for most users, we add the constraints that t_i^0 is no sooner than May 14 and that t_i^2 is no later than May 21. As the variation in t_i^τ, $\tau \in \{0,1,2\}$ across users is small, we drop the subscript i for brevity, and we exclude from the analysis all users that were not active in at least one of the three periods $\{t^0, t^1, t^2\}$. We also restrict our sample to users who did not change their preferences, relationship status, or geographic region (as described in Section 2) three months prior to t^0.

We define the *treated users* as those who were shown at least one inactive backlog profile in period t^1. Recall that our partner's algorithm ranks each user's backlog profiles in terms of their attractiveness score and only uses this metric to select (up to) the top three backlog profiles to show for most users. Some treated users were shown more backlog profiles than they would have seen before the change in the algorithm (including those that had no active users in their backlogs). For those treated users who would have seen backlog profiles with no change in the algorithm, these inactive backlog profiles may be more attractive than some of the backlog profiles they would have seen. These two elements combined increased the treated users' probability of getting new matches as liking these additional inactive backlog queries instantaneously generates a match.

We define the set of users in the *control group* as those whose observed characteristics and attractiveness scores from active and inactive users in their backlog are similar to those of the treated users, but who were not shown any inactive backlog profiles in period t^1 because the inactive backlog profiles were not eligible to be shown because the (inactive) users: 1) disabled their profiles on the platform; 2) changed their relationship status to "not single"; 3) changed their preferences; or 4) moved to a different region. We emphasize that these reasons are all uncorrelated with the unobserved characteristics or dating activities of the focal user, captured by ξ_{it} in the model. Moreover, we have access through our partner to all inactive backlog profiles, their eligibility to be shown, and the reason for their ineligibility. In other words, the control group users satisfy the following four conditions:

I. The user had inactive backlog profiles after the change in the algorithm,

II. The user's characteristics and the number and attractiveness of their active and inactive backlog profiles are similar to those in the treatment group,

III. The user was not shown any inactive backlog profiles in period t^1,

IV. At least one of the inactive backlog profiles is more attractive than the least attractive active backlog profile that would have been shown in period t^1 (if any) and is ineligible to be shown for one of the four reasons listed above.

[11] As we discuss later in the section, the short window also allows us to guarantee that the treated users are not able to tell that their additional matches are with an inactive user.

The first two conditions ensure that the control group users are similar to those in the treatment group in terms of their characteristics and the composition and attractiveness of their backlogs. The last two conditions establish that the user had inactive backlog profiles in period t^1 but was not shown any because of the ineligibility of these inactive profiles due to the exogenous reasons listed above. In other words, the probability of getting a match in t^1 is not affected by the change in the algorithm for these users for reasons uncorrelated with the unobserved characteristics or dating activities of these users.

Conditions I and III are straightforward to implement. For II, we first estimate the users' propensity to be treated as a function of the full set of user characteristics and the number and attractiveness distribution measures (mean, standard deviation, min, max) of both their active and inactive backlog profiles. Then, we use the estimated propensity scores to construct weights or matches in our analysis to balance the treatment and control groups. That is, conditional on the propensity score or the observables of each user, whether the user is in the treatment or control group is determined by whether an inactive backlog profile happens to be eligible to be shown after the change in the algorithm. We note that condition IV implies that there might be potential differences in the attractiveness scores between the backlog queries of the treatment and control users. Thus, we control for the attractiveness score of every profile shown to all users in the analysis.

Compared to the control users, the treated users have a higher probability of obtaining matches in period t^1. Moreover, since the treated users' additional matches are with inactive users, and the active/inactive state of their matches remains unknown in a short period of time, the impact of these "inactive" matches on the treated users' subsequent like decisions is the same as that of the matches with active users. In Appendix B.3, we provide statistics on the amount of time between when a match is formed and the completion of one round of messaging between the users (one message from each user) and show that users could not tell the difference between an active and an inactive match within three days of obtaining the match. Using the inactive matches and a short time window helps us avoid the interference common to experimental designs in our two-sided market setting.

4.1.3. Estimation Procedure Our estimation procedure includes three steps. First, based on users' observed characteristics and usage metrics from the pre-treatment period t^0, we estimate the propensity that users are treated in the quasi-experiment. Using the estimated propensity scores, we construct weights to be applied in the second and 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 for the number of matches that users receive between periods t^1 and t^2. In the third step, we estimate the impact of the estimated number of matches on the like probabilities in period t^2 [12]. Next, we describe the details of the three-step estimation procedure.

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

$$e_i(X_i, M_i^{t^1}) = X_i\beta_0 + \gamma M_i^{t^1} + \epsilon_i,$$

where e_i is the propensity score of user i. In other words, $e_i(X_i, M_i^{t^1}) = Pr(W_i = 1|X_i, M_i^{t^1})$, 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, number of backlog profiles shown to the user observed in period t^0, number of profiles liked by i in period t^0, number of active and inactive profiles in i's backlog in period t^0, and summary statistics of the distribution of both active and inactive backlog profiles' attractiveness scores, including the mean, standard deviation, minimum, and maximum. We provide the detailed description of the full set of variables in Appendix B.2. $M_i^{t^1}$ represents the number of matches obtained by user i since their last session in the pre-treatment period, i.e., between t^0 and t^1. 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 of 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. Following Hirano et al. (2003), we compute weights ω_i using the estimated propensity scores $\hat{e}_i$, 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:

$$M_i^{t^2} = \theta W_i + X_i\beta_1 + Z_i^{t^1}\delta_1 + \varepsilon_i,$$

where $M_i^{t^2}$ represents the number of matches obtained by user i between periods t^1 and t^2, 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, attractiveness score, and the fraction of profiles in the assortment that share the same race and religion with the user. We also control for the same covariates used in the propensity score estimation. Using the estimated parameters from this model, we compute the predicted number of matches since their last session before t^2 for each user i, i.e., $\hat{M}_i^{t^2}$.

[12] As robustness checks, we conduct the estimation using two alternative specifications: without propensity score weighting and on a matched sample based on the estimated propensity scores. The results are similar and are omitted due to lack of space. Note that the propensity score weighting is not necessary, but we include it to provide a fairer comparison between the treated and control groups.

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

$$\phi_{ij}^{t^2} = \mathrm{P}\left(\gamma \hat{M}_i^{t^2} + X_{ij}\beta_2 + Z_i^{t^1}\delta_2 + \varepsilon_{ij}\right),$$

where $\phi_{ij}^{t^2}$ is the probability that user i liked j in period t^2, $\mathrm{P}(\cdot)$ is a cumulative distribution function, X_{ij} is as previously defined in Section 4.1, and $Z_i^{t^1}$ is the matrix of characteristics of the assortment observed in period t^1. In an alternative specification, we also control for quality of matches obtained since the last session.

4.2. Estimation Results

As a result of our treatment and control definitions, our quasi-experiment consists of 8,398 control and 6,412 treated users. In Appendix B.2 we report the results of the propensity score estimation and we also include an extensive set of statistics about the users' characteristics and activity measures for the treated and the control group after propensity score weighting. We find that there is no statistically significant difference between the two groups across all these statistics.

In the upper panel of Table 1 we report the first-stage results of 2SLS estimated using an ordinary least squares (OLS) regression and a negative binomial regression. We include the latter because the dependent variable takes discrete values. We observe that, in both models, the coefficient of the treatment variable is positive and significant, which confirms that our instrument satisfies the relevance condition. Moreover, we observe that the estimated marginal effects are very similar, which suggests that the choice of the first-stage model does not play an important role.

In the lower panel of Table 1, we present the second stage results of the 2SLS procedure as described in Step 3. We address the forbidden regression problem of the negative binomial first-stage following Angrist and Pischke (2008).

Table 1 Quasi-Experiment Estimation Results

	OLS	Neg. Bin.
	First Stage	
Treated	0.228***	0.555***
	(0.011)	(0.027)
	Second Stage	
$\hat{M}_i^{t_2}$	-0.423***	-0.222**
	(0.147)	(0.112)
Constant	-3.020***	-2.961***
	(1.148)	(1.141)
Observations	51,561	50,533
Pseudo-R^2	0.323	0.322

Note: The first-stage is estimated using OLS and negative binomial regressions. The second stage for the negative binomial first-stage regressions addresses the forbidden regression problem. Standard errors are clustered at the user level. Significance reported: *p<0.1; **p<0.05; ***p<0.01.

We observe that the coefficient corresponding to $\hat{M}_i^{t^2}$ is negative and statistically significant for both specifications. We calculate the average marginal effect of $\hat{M}_i^{t^2}$ and find that an extra match in period t^1 reduces the like probability in period t^2 by 3.2% to 6.3%. These results provide an estimate of the history effect and of the utility function, which we will use as an input for our proposed algorithm in the simulations and in the field experiment. In Appendix B.5, we show that these results are robust to controlling for the attractiveness of the matches obtained since the last session. More specifically, we control for the mean, the standard deviation, the minimum, and the maximum attractiveness scores of the matches received since the last session. We also control for the positive and negative differences relative to the score of the focal user. In all alternative specifications, we find that the estimated history effect is larger in magnitude than those reported in Table 1; i.e., these results provide a conservative measure on the magnitude of the history effect.

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.

5.1. Incorporating Business Constraints

Following our partner's practice, we limit the number of potential partners that a user can see on each given day by restricting the assortments to be of (at most) a fixed size, i.e., $|S_i^t| \leq K_i^t$ for some $K_i^t \ll |\mathcal{I}|$. 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 of users who have rejected them in the past.

To express the above constraints, we use the notation introduced in Section 3. 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 "minimum requirement" (covering) constraints, and examples include: (1) if a user's backlog is not empty, show at least one profile from the backlog; (2) a minimum number of profiles with some level of attractiveness should be included, etc. Importantly, these constraints are assigned an order and they need to be satisfied in that order. We refer to these constraints collectively as *business constraints*. 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 user i has no matches in the recent past, i.e., $\phi_{ij}(M) \le \phi_{ij}(0)$. To ease notation, throughout the rest of the paper we use ϕ_{ij}^0 to denote $\phi_{ij}(0)$.

Following this observation, we propose the following linear program that, as we establish in Proposition 1 below, can be used to obtain an upper bound for the platform's problem in (3), and that also plays a fundamental role in constructing our heuristics in Section 5.3:

$$
\begin{aligned}
\tilde{\pi} \quad := \quad \max_{x,y,z} \quad & \sum_{t\in\mathcal{T}}\sum_{i\in\mathcal{I}}\sum_{j\in\mathscr{P}_i^1} \upsilon_i^t\cdot\phi_{ij}^0\cdot y_{ij}^t + \frac{1}{2}\cdot\upsilon_i^t\cdot\upsilon_j^t\cdot\phi_{ij}^0\cdot\phi_{ji}^0\cdot z_{ij}^t \\
\text{s.t.} \quad & y_{ij}^t \le \mathbb{1}_{\left\{j\in\mathscr{B}_i^1\right\}} - \sum_{\tau=1}^{t-1} y_{ij}^\tau\cdot\upsilon_i^\tau + \sum_{\tau=1}^{t-1}\left(x_{ji}^\tau - \upsilon_i^\tau\cdot z_{ji}^\tau\right)\cdot\upsilon_j^\tau\cdot\phi_{ji}^0, \quad \forall i\in\mathcal{I},\, j\in\mathscr{P}_i^1,\, t\in\mathcal{T} \\
& \sum_{t=1}^{T}\left(x_{ij}^t + y_{ij}^t\right)\cdot\upsilon_i^t \le 1, \quad \forall i\in\mathcal{I},\, j\in\mathscr{P}_i^1,\, t\in\mathcal{T} \\
& \sum_{j\in\mathscr{P}_i^1} x_{ij}^t + y_{ij}^t \le K_i^t, \quad \forall i\in\mathcal{I},\, j\in\mathscr{P}_i^1,\, t\in\mathcal{T} \\
& z_{ij}^t \le x_{ij}^t, \quad z_{ij}^t \le x_{ji}^t, \quad z_{ij}^t = z_{ji}^t, \quad x_{ij}^t, y_{ij}^t, z_{ij}^t \in [0,1], \quad \forall i\in\mathcal{I},\, j\in\mathscr{P}_i^1,\, t\in\mathcal{T}
\end{aligned}
\tag{5}
$$

The decision variables x_{ij}^t and y_{ij}^t may be interpreted as follows. y_{ij}^t represents the probability that j is included in i's assortment as a backlog query (i.e., j has previously seen and liked i). By contrast, x_{ij}^t is the probability that j is included in i's assortment as a non-backlog query (i.e., i has not been evaluated by j yet). 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 $\sum_{\tau=1}^{t-1} y_{ij}^\tau\cdot\upsilon_i^\tau$ denotes the probability that i saw j as part of a backlog query in the past, and the term $\sum_{\tau=1}^{t-1}\left(x_{ji}^\tau - \upsilon_i^\tau z_{ji}^\tau\right)\cdot\upsilon_j^\tau\cdot\phi_{ji}^0$ denotes 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 take a value equal to one starting from the first period. The second constraint guarantees that each profile is seen at most once throughout the entire horizon in expectation. The third constraint ensures that at most K_i^t profiles are shown in expectation to each user in the corresponding period. The next constraints define the variable z_{ij}^t and ensure 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 (5) 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 horizon $\mathcal{T}$, the log-in probabilities $\{\upsilon_i^t\}_{i\in\mathcal{I},t\in[\mathcal{T}]}$, and the like probabilities $\{\phi_{ij}^0\}_{i,j\in\mathcal{I}}$.

PROPOSITION 1. *Let π^* be the optimal value of the platform's problem introduced in* (3)*, and let $\tilde{\pi}$ be an optimal solution to* (5)*. Then, $\pi^* \leq \tilde{\pi}$.*

We conclude by noting that this upper bound is not likely to be tight, as it does not take into account either the business constraints or the history effect.

5.3. The Dating Heuristics

Our finding that the like probabilities depend on the number of recent matches introduces significant challenges from an optimization perspective. Typically, one would like to treat the like probabilities as (time-invariant) parameters to our algorithm; however, due to the history effect, these are endogenous to the choice of the algorithm. Specifically, the algorithm decides the assortments for today, which have an effect on the likes and thus the matches formed today, which in turn affect the like probabilities tomorrow. To address this challenge, we next 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 (5), 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 $\{M_i^t\}_{i\in\mathcal{I}}$, the log-in probabilities $\{\upsilon_i^t\}_{i\in\mathcal{I}}$, and the like probabilities $\left\{\phi_{ij}^t \in \mathcal{I}\right\}$, where $\phi_{ij}^t = \phi_{ij}\left(M_i^t\right)$. Second, instead of considering the full horizon, we consider only one period of look-ahead (i.e., $\tau \in \{t, t+1\}$). We denote by $\mathcal{S}^t$ the polytope resulting from these changes. Third, we update the objective function by incorporating a penalty in order to account for the effect that the initial number of matches together with the decisions in these periods can have in future ones. While our main focus moving forward will be to design this penalty function to capture the history effect, it is worth highlighting that this penalty can also be used to capture other considerations that may impact future matches, such as exhausting all good options for a picky user. Finally, we also include the business constraints to satisfy the requirements of our industry partner. Recall from our earlier discussion in Section 5.1 that the business constraints are of the form such that "a minimum number of profiles satisfying criteria X should be included if possible," and that these constraints must be satisfied (if possible) in some pre-defined order. We denote the ordered set of business constraints by $\mathcal{L} = \{1, \ldots, L\}$ and observe that, in each period t, business constraint l can be expressed as $\sum_{j\in\mathscr{P}_i^t} x_{ij}^t \cdot a_{ijl}^{t,x} + y_{ij}^t \cdot a_{ijl}^{t,y} \geq b_{il}^t$ where $a_{ijl}^{t,x} \in \{0,1\}$, $a_{ijl}^{t,y} \in \{0,1\}$ and $b_{il}^t \in \mathbb{N}_0$ are constants that may depend on the state of the system, namely, the sets of potentials and the backlogs in period t, i.e., $\{\mathscr{P}_i^t, \mathscr{B}_i^t\}_{i\in\mathcal{I}}$. Importantly, only constraints that can be satisfied will be added to the formulation (e.g., if a user has an empty backlog, no constraint on their backlog will be added). The resulting optimization problem can be found in (6) in Algorithm 1.

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 construct feasible solutions by satisfying the business constraints sequentially, in the order that mimics the one followed by our industry partner. We first include backlog profiles in decreasing order of $y_{ij}^{*,t}$ followed by profiles in decreasing order of $x_{ij}^{*,t}$ until the first business constraint is satisfied or the assortment is full. We then proceed to the second constraint, and so on. Once all the constraints are satisfied, we complete the assortment by including profiles in decreasing order of $y_{ij}^{*,t}$ and, if there is space left in the assortment, we add profiles in decreasing order of $x_{ij}^{*,t}$. Observe that this rounding technique will prioritize showing backlog profiles.[13]

Penalty. 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 periods. As matches today affect matches tomorrow through the like probabilities, our penalty function uses a first-order approximation to capture the effect that matches in each period will have on individual like probabilities in 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 on the assortment decisions in 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 $M_i^{\tau+1}$,

$$\phi_{ij}(M_i^{\tau+1}) - \phi_{ij}(M_i^\tau) \approx \left(M_i^{\tau+1} - M_i^\tau\right) \cdot \gamma_{ij}(M_i^\tau), \tag{7}$$

where $\gamma_{ij}(M_i^\tau)$ is the local marginal effect of an extra match on 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

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

[13] We tested the performance of other rounding rules in simulations; the results are omitted for the sake of space.

Algorithm 1 Dating Heuristic (DH)

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

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

Step 1. Optimization: solve

$$\begin{aligned} x^*, y^*, z^* = \arg\max_{x,y,z} \quad & \sum_{\tau=t}^{t+1} \sum_{i\in\mathcal{I}} \sum_{j\in\mathscr{P}_i^t} \upsilon_i^\tau \cdot \phi_{ij}^t \cdot y_{ij}^\tau + \frac{1}{2} \cdot \upsilon_i^\tau \cdot \upsilon_j^\tau \cdot \phi_{ij}^t \cdot \phi_{ji}^t \cdot z_{ij}^\tau \\ & + \bar{\xi} \cdot \left(\Psi^t \left(\vec{x}^{\,t}, \vec{y}^{\,t}, \vec{z}^{\,t}, \vec{M}^t \right) + \Psi^{t+1} \left(\vec{x}^{\,t}, \vec{y}^{\,t}, \vec{z}^{\,t}, \vec{x}^{\,t+1}, \vec{y}^{\,t+1}, \vec{z}^{\,t+1}, \vec{M}^t \right) \right) \\ \text{s.t.} \quad & \sum_{j\in\mathscr{P}_i^t} x_{ij}^\tau \cdot a_{ijl}^{\tau,x} + y_{ij}^\tau \cdot a_{ijl}^{\tau,y} \ge b_{il}^\tau, \ \forall i \in \mathcal{I}, \ \tau \in \{t, t+1\}, \ l \in \mathcal{L}, \\ & (x, y, z) \in \mathcal{S}^\tau. \end{aligned} \tag{6}$$

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 $i \in \mathcal{I}$, set $S_i^t = \emptyset$.

For $l = 1, \ldots, \mathcal{L}$;

If constraint l is not satisfied by S_i^t:

Let $\mathscr{P}_i^{t,y}(l)$ be the subset of potentials for which $a_{ijl}^{\tau,y} = 1$, $y_{ij}^* > 0$ and $j \notin S_i^t$.

Define $\mathscr{P}_i^{t,x}(l)$ as the subset of potentials for which $a_{ijl}^{\tau,x} = 1$, $x_{ij}^* > 0$ and $j \notin S_i^t \cup \mathscr{P}_i^{t,y}(l)$.

Greedily add profiles in $\mathscr{P}_i^{t,y}(l)$ in decreasing order of y_{ij}^* to S_i^t until the constraint is satisfied;

If no profiles are left in $\mathscr{P}_i^{t,y}(l)$ and the constraint is still not satisfied:

Greedily add profiles in $\mathscr{P}_i^{t,x}(l)$ in decreasing order of x_{ij}^* until the constraint is satisfied.

If $|S_i^t| < K$, complete the assortment by adding profiles in decreasing order of y_{ij}^* (not included thus far) and, if there is still space, add profiles in decreasing order of x_{ij}^* that have not been included so far.

where the expectations are taken over the users' decisions in period t given the algorithm's (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 $\phi_{ij}^t = \phi_{ij}(M_i^t)$. The second equality is obtained as follows. The first term corresponds to the case where user i logs in 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 above. 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 be obtained only 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

$$\Psi^t \left(\vec{x}^{\,t}, \vec{y}^{\,t}, \vec{z}^{\,t}, \vec{M}^t \right) = \sum_{i\in\mathcal{I}} \mathbb{E}\left[M_i^{t+1} \middle| \{\vec{x}^{\,t}, \vec{y}^{\,t}, \vec{z}^{\,t}\}, \vec{M}^t \right] - M_i^t$$

$$= \sum_{i\in\mathcal{I}} \upsilon_i^t \left(\sum_{j\in\mathscr{P}_i^t} \left(\phi_{ij}^t \cdot y_{ij}^t + \upsilon_j^t \cdot \phi_{ij}^t \cdot \phi_{ji}^t \cdot z_{ij}^t\right) + \sum_{j\in\mathcal{I}} \upsilon_j^t \cdot \phi_{ji}^t \cdot y_{ji}^t \right) + (1-\upsilon_i^t) \left(M_i^t + \sum_{j\in\mathcal{I}} \upsilon_j^t \cdot \phi_{ji}^t \cdot y_{ji}^t \right) - M_i^t.$$

Using a similar reasoning, we define

$$\Psi^{t+1}\left(\vec{x}^{\,t}, \vec{y}^{\,t}, \vec{z}^{\,t}, \vec{x}^{\,t+1}, \vec{y}^{\,t+1}, \vec{z}^{\,t+1}, \vec{M}^t\right) = \sum_{i\in\mathcal{I}} \mathbb{E}\left[M_i^{t+2} - M_i^{t+1} \mid \{\vec{x}^{\,\tau}, \vec{y}^{\,\tau}, \vec{z}^{\,\tau}\}_{\tau\in\{t,t+1\}}, M_i^t\right]$$

$$\approx \sum_{i\in\mathcal{I}} \sum_{j\in\mathscr{P}_i^t} \upsilon_i^{t+1} \cdot \left(y_{ij}^{t+1} \cdot \phi_{ij}^t + z_{ij}^{t+1} \cdot \upsilon_j^{t+1} \cdot \phi_{ij}^t \cdot \phi_{ji}^t\right) + \sum_{j\in\mathcal{I}} y_{ji}^{t+1} \cdot \upsilon_j^{t+1} \cdot \phi_{ji}^t$$

$$- \sum_{i\in\mathcal{I}} \upsilon_i^{t+1} \cdot \left(\sum_{j\in\mathscr{P}_i^t} \upsilon_i^t \cdot \left(y_{ij}^t \cdot \phi_{ij}^t + z_{ij}^t \cdot \upsilon_j^t \cdot \phi_{ij}^t \cdot \phi_{ji}^t\right) + \sum_{j\in\mathcal{I}} y_{ji}^t \cdot \upsilon_j^t \cdot \phi_{ji}^t + M_i^t \cdot (1-\upsilon_i^t) \right).$$

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 against relevant benchmarks.

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 profiles from their potentials between September 1, 2019 and April 1, 2020, and that logged in at least once between March 1 and April 1, 2020. For each of these users, we assume that their initial set of of potentials is comprised of the profiles they saw between September 1, 2019 and April 1, 2020, and we assume that they had no backlog or 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 for different markets, including Austin and Dallas, and different initial conditions for the set of potentials, the number of matches since the last session, etc. The results are qualitatively similar to those that will be reported next.

Having defined the market, we next 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 second column of Table 1 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 $K = 3$ for all users, 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.[14]

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

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

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 (by removing from each user i's potentials all users j who saw and disliked user i, and also the profiles evaluated by i (if any) in that period), the backlogs (by adding to each user i's backlog all users j who saw and liked user i, and also by removing the backlog profiles evaluated by i (if any) in that period), and the number of matches obtained since the last session for each user. Finally, having updated the state of the system, we proceed to the next period and repeat this process until the end of the horizon.

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

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

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, the log-in probabilities, and the history effect on the 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. In other words, it does not update the current-period like probabilities by accounting for the matches obtained since the last session, and it does not have a look-ahead period or a penalty term.

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 in order for the profile 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 C.2 for further details).

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. Moreover, the improvement of DH relative to Partner and Greedy is always positive and quite substantial. In Appendices C.1 and C.1.1 we respectively show that these differences remain for different values of the history effect and that similar results are obtained in other markets.

To identify the sources of improvement, we aim to quantify how much is due to (1) finding better matches (by using improved personalized estimates for the like probabilities and explicitly accounting for the probability that a profile is liked back) and (2) considering a penalty in the objective that accounts for the history effect. In Figure 3c we plot the average number of matches

Figure 3 Simulation Results

(a) Number of Matches (b) Relative Improvement (c) Sources of Improvement

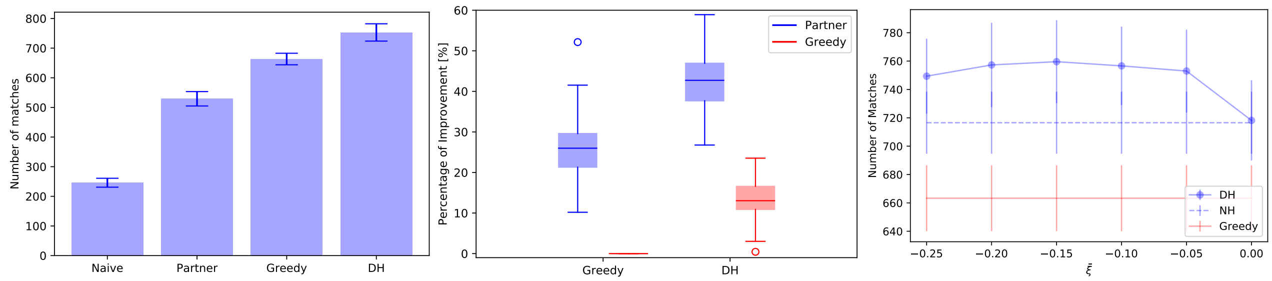

obtained by Greedy and DH from 100 simulations (same setup as before) for different values of $\bar{\xi}$. In addition, we plot the results obtained by DH if we do not take into account the history effect, i.e., if we consider $\bar{\xi} = 0$ and $\phi_{ij} = \phi_{ij}(0)$ for all $i \in \mathcal{I}$ 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 the improvement due to including the history effect.

First, we observe that our no-history (NH) heuristic considerably outperforms the Greedy policy. This improvement is solely based on finding better matches by using the one-period look-ahead policy, 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 generates 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.[15] Second, 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. When $\bar{\xi} = -0.15$, the absolute improvement is 5.76% (759.57 matches on average) relative to the case $\bar{\xi} = 0$. Overall, from these results we conclude that 80% of the improvement (or roughly, 185 extra matches) comes from finding better matches and the remaining 20% of the improvement is due to accounting for the history effect.

Finally, an additional set of simulations reported in Appendix C.3 show that the improvement obtained by DH is also meaningful and stable over a longer time horizon. We also observe that DH aims to keep the like rate stable (neutralize the history effect) by means of two mechanisms: (1) saving backlog profiles when the number of matches since the last session increases and the backlog size is not large, and (2) showing more attractive profiles as the number of matches since the last session increases. These mechanisms explain the improvement achieved by DH over NH.

[15] Alternatively, one could use the like probabilities obtained from estimating a model without the history effect. We note here that both the estimated coefficients and the predictive power of such probabilities when there are no matches since the last session are very similar to those of $\phi_{ij}(0)$.

7. Field Experiments

In this section, we describe the results of two field experiments aiming to test if and how 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 to treatment and control groups.[16] 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 identical markets. As a result, we perform our field experiments in similar markets, 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.

We considered 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 and Austin in the first and second field experiments, respectively—to be assigned to the treatment group, while the other two 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 13 (see Appendix D.1) we show that there are no significant differences in the main demographics across the three markets.

7.2. First Field Experiment

During the seven days between August 19 and August 25, 2020, the users in the treatment market (Houston) received assortments chosen with our heuristic, while the markets in the control group kept the default algorithm provided by the platform. As an input for our heuristic, we use the parameters in the first column of Table 1 to estimate the like probabilities. To predict the log-in probabilities, we use the estimation result of a model with user and time fixed effects and detailed activity measures; we describe the model in detail in Appendix B.4.[17] As the estimates are very similar across models, we believe that this choice does not affect the results of the experiment.

[16] Running the experiment at the user level in our matching-market setting would lead to interference for reasons similar to those explained in Section 4.1.2. Suppose that we apply our algorithm to a randomly selected treatment group of users and keep the default algorithm provided by the platform for the users in the control group. As the potentials of the treated users may contain control users, the selection of assortments by our algorithm also affects the potentials and backlogs of those users in the control group.

[17] The results in Appendix B.4 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 the log-in probabilities is either negative or statistically insignificant. Our result assumes that v_i^t is exogenous; i.e., the history effect is estimated by conditioning on the log-in probabilities. If the effect of the match history on the log-in probabilities is negative, our result provides a lower bound on the size of the history effect. Our field experiment results provide additional validation of this assumption.

Recall that our algorithm still needs to satisfy the business constraints imposed by our industry partner, i.e., the algorithm will first satisfy the business constraints and, if there is space remaining in the assortment, it will select additional profiles to add. As a result, 38.30% of the profiles shown during the time window of the experiment were chosen freely by DH.

7.2.1. Overview of Results. In Figure 4 we plot the number of matches per day in Houston, Austin, and Dallas–Fort Worth, between August 12 and August 25, 2020. The vertical line marks August 19, 2020, the day when the experiment started in the treated market. We observe that, starting from August 19, 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 on all days after the start of the experiment. These results suggest that our algorithm significantly increased the number of matches generated in Houston.

Figure 4 First Field-Experiment: Number of Matches and Active Users

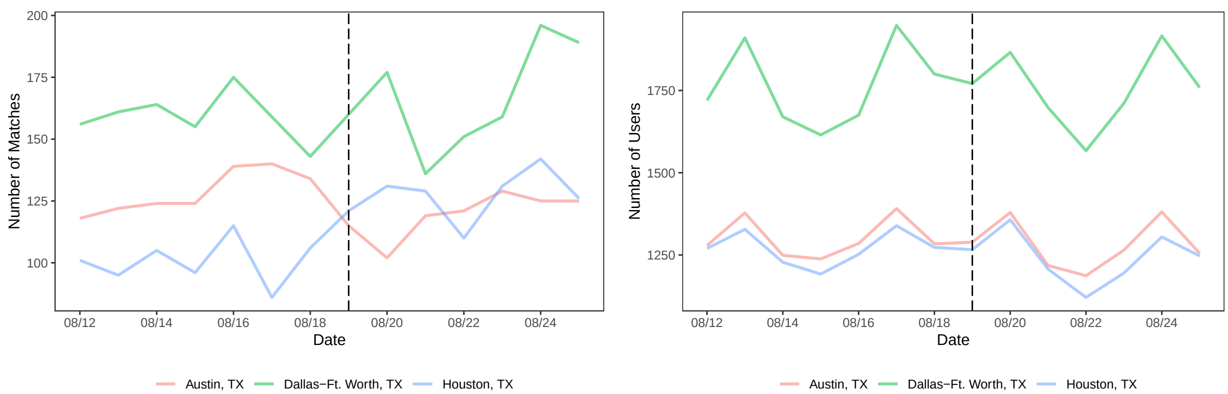

:

Note Number of matches and number of users who logged in between August 12 and August 25, 2020. The vertical line marks the day when the experiment started in the treated market (Houston).

Panel 1 in Table 2 summarizes these findings. The Before period corresponds to the week of August 12–18, 2020 (before the change in the algorithm), whereas the After period corresponds to the week of August 19–25, 2020. The Queries column includes the total number of profiles that were shown in each market and time period; the Likes column indicates how many of those profiles were liked. Finally, the Matches column counts the matches that were formed. We observe that the number of queries remained approximately constant. However, there is a significant increase in the number of matches obtained in Houston during the After period. We will formalize this finding in the next subsection. We conclude Section 7.2 by exploring what is driving the increase in matches.

Table 2 First Field-Experiment: Results

Panel 1: Summary of Results

Market	Period	Queries	Likes	Matches
Austin	After	27633	11243	836
Austin	Before	27580	11440	901
Dallas	After	38234	14796	1168
Dallas	Before	38931	15067	1113
Houston	After	27337	10091	890
Houston	Before	27796	11083	704

Panel 2: Difference-in-Differences Results

	Dependent Variable: Number of Matches	
	(1)	(2)
Post	-0.714	-
	(7.611)	-
Post×Treated	27.286***	27.286***
	(7.611)	(9.147)
Austin	124.429***	120.000***
	(4.129)	(7.283)
Dallas	163.286***	158.857***
	(4.129)	(7.283)
Houston	100.571***	96.143***
	(4.768)	(7.607)
Observations	42	42
R^2	0.296	0.596

Note: Panel 1 reports the overall number of queries, likes, and matches in the Before period (from August 12, to August 18, 2020) and the After period (from August 19 to August 25, 2020) of the field experiment. Panel 2 reports the estimation results. Column (1) includes a dummy for the periods after the start of the experiment, while column (2) considers date fixed effects.

7.2.2. Estimation. To estimate the effect of our heuristic we follow a DID approach. 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. 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:

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

where α_m are market-specific fixed effects that account for the 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.

To validate the DID approach, we tested and confirmed that there are no significant differences in the trends before the intervention, proving that the parallel-trends assumption holds in our setting. In the interest of space, we have relegated the details to Appendix D.2.

In Panel 2 of Table 2 we report the estimation results. The first column provides the results of the first model, while the second column provides the results with the fixed effects for each time period. We observe that the coefficient for the variable of interest (Post×Treated) is positive and significant in both models, and that the estimated average number of extra matches that our algorithm produced in the treated market is 27.286. Comparing this value with the estimated fixed effect corresponding to the Houston market, we observe that our algorithm improved the number of matches generated in that market by at least 27.13%.

To assess the robustness of our results, we performed a placebo test excluding data from Houston and assigning Austin to the treatment group. As the results in Table 15 in Appendix D.3 show, we find no significant effect in the variable of interest. We also estimated our DID model for different subsets of data (e.g., removing one control market at the time, removing the last day of the intervention to avoid the end-of-horizon effects, and removing the first two days of the intervention as these rely mostly on the backlog generated before the intervention) and obtained similar results (see Tables 16 and 17 in Appendix D.3). Finally, we compared the improvement in Houston against that in all other markets with at least 400 matches per week, and we find that it more than doubles the second-largest improvement (see Table 18 in Appendix D.3).

7.2.3. Discussion of the sources of the improvement. Recall that there are two mechanisms by which matches can be formed. The first mechanism is *simultaneous shows*; i.e., both users see and like each other in the same period. The second is *sequentially* through backlog queries. That is, user i first sees and likes user j and i is added to j's backlog; then, user i is shown to j and, if j likes i, a match is automatically formed. From Table 3 we observe that the vast majority of matches were formed through the latter mechanism. Specifically, 839 of the 890 matches in Houston during the treatment period were formed sequentially as a result of backlog queries.

Table 3 also shows that the number of backlog queries in the Houston After period is significantly larger than in the Houston Before period. This is not surprising: as explained in Section 5, our algorithm favors showing backlog profiles more than our partner's algorithm. However, this raises the following concern: is it the case that all the improvement comes from the fact that we are "depleting" the existing backlogs of the users that were generated by our partner's algorithm in the previous periods? If the latter is true, the improvement achieved by DH may not be sustainable in a longer time horizon.

To get a better understanding of what is driving the improvement, in Table 3 we provide a picture of the non-backlog queries. We focus our discussion on the Houston market. The total numbers of non-backlog queries in Houston in the Before and After periods are similar. However, the like rate during the After period is significantly lower. This is to be expected: as we showed in Section 6, our algorithm takes into account the probability that a match is formed (i.e., that both parties like each other) whereas our partner's algorithm is more biased towards maximizing likes. In total, of the 23,395 non-backlog queries during the After period, 9,252 were liked, resulting in a total of 9,252 new additions to the backlog, compared to the 10,401 additions to the backlog during the Before period. However, 1,946 of these 9,252 new backlog queries were shown during the

After period, resulting in 436 matches. This implies that *48.99% of the matches that were obtained within the experiment window were a result of the backlog generated by our own algorithm* within that same window. By contrast, 37.92% of the matches obtained in the period Before were a result of the backlog generated within that period. This provides evidence that our algorithm is obtaining matches not by depleting the backlog previously generated, but rather by exploiting the backlog generated by itself.

Table 3 First Field-Experiment: Backlog and Non-Backlog Queries

		Backlog			Non-Backlog					
Market	Period	Queries	Likes	Matches	Queries	Likes	LR	NB	Shown	Matches
Austin	After	3787	810	810	23846	10433	0.438	10433	922	269
Austin	Before	3712	869	869	23868	10571	0.443	10571	923	297
Dallas	After	4828	1115	1115	33406	13681	0.410	13681	1477	452
Dallas	Before	4936	1063	1063	33995	14004	0.412	14004	1646	439
Houston	After	3942	839	839	23395	9252	0.395	9252	1946	436
Houston	Before	2929	682	682	24867	10401	0.418	10401	954	267

Note: Number of backlog and non-backlog queries, and the resulting number of likes and matches. LR stands for like rate; NB stands for the new backlog generated within the corresponding period.

7.3. Second Field Experiment

As previously mentioned, our second field experiment considered Austin as the treated market; i.e., we used our DH algorithm to select the assortments, while Houston and Dallas were the control markets. The experiment ran for seven days, between September 9 and September 15, 2020, and all the parameters considered were the same as in the first field experiment.

We evaluate the effect of our intervention by applying the DID regression described in Section 7.2.2 (see Appendix D.5 for details).[18] First, we find that the variable of interest, Post×Treated, is positive and significant and that the number of matches increased by 37.95% (see Panel 1 of Table 20). This effect is more prominent than in the first experiment because the number of profiles shown during the time window of the experiment that were chosen freely by DH—after the business constraints were satisfied—increased: 49.57% in the second experiment compared to 38.30% in the first experiment. Second, we perform the placebo test by removing the treated market (Austin) data and labeling one of the control markets as treated (Houston). We observe that the variable of interest is not statistically significant in the placebo test, suggesting that the increase in the number of matches is due to the implementation of DH in Austin (see Panel 2 of Table 20)).

To assess the strength of these results, we perform the same set of robustness checks that we run for the first field experiment. First, we confirm that the parallel-trends assumption holds (see

[18] Specifically, we define $W_m = 1$ for Austin and $W_m = 0$ for Dallas and Houston. We also update Z_t so that $Z_t = 1$ if t is September 9, 2020 or later, and $Z_t = 0$ otherwise.

Table 21). Second, we estimate our results by excluding one control market at a time, and we find that our results are robust to the control markets considered (see Table 22). Third, we observe that the results are robust to excluding some days of the experiment, i.e., the last or the first two days (see Table 23). Finally, we also observe that the fraction of matches obtained from the backlog generated in the period after the experiment is higher than that in the period before the experiment (see Table 24), confirming that the improvement is not driven by exploiting the existing backlog at the beginning of the experiment.

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 like each other to generate a match. Using data from our industry partner, we estimate the main parameters of the model. Using a novel identification strategy, we find that matches in the recent past reduce 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 two field experiments confirm that the improvements translate into our partner's platform. Given the positive results, we are collaborating with our industry partner to deploy our heuristic as their main algorithm for choosing which assortments to show in other markets.

Overall, our results showcase how platforms can leverage their knowledge of the drivers of its users' behavior to improve their display decisions. Broadly speaking, we believe that it is of general importance to account for the effect that algorithmic decisions have on user input in dynamic settings and on how this impacts decisions in future periods. Our problem presents one specific setting where accounting for such an effect with a simple one-period look-ahead policy leads to at least a 5% improvement. Moreover, the results from the field experiment also provide additional evidence of the good performance of the one-period look-ahead policies in practice.

References

Angrist J, Pischke JS (2008) *Mostly Harmless Econometrics: An Empiricist's Companion* (Princeton University Press), 1 edition.

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.

Hirano K, Imbens GW, Ridder G (2003) Efficient estimation of average treatment effects using the estimated propensity score. *Econometrica* 71(4):1161–1189.

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.

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. Computational Hardness of the Dynamic Two-sided Assortment Problem

PROPOSITION 2. *The decision-theoretic formulation of the dynamic two-sided assortment problem is NP-hard.*

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 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[19] K. Let $\underline{K} = \min\{|S| : S \in \mathcal{S}\}$ be 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, in every period t, the assortment S_i^t that is 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 the focal user 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 $\bar{\mathbf{S}}' = \{\bar{S}'_1, \ldots, \bar{S}'_T\}$ 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 $\mathcal{I}$.

A.2. Proof of Proposition 1

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. (To ease notation, we do not explicitly include the dependency on the optimal policy π^*.) Also, under the optimal policy, let $\mathcal{Y}_{ij}^t$ be the random variable denoting whether user i gets a backlog profile j in period t. Since the optimal policy is non-anticipating and decides the assortments for period t *before* observing whether the user logs in, these decisions are independent of the log-in decisions in period t, i.e., $\mathcal{X}_{ij}^t, \mathcal{Y}_{ij}^t \perp \Upsilon_i^t$ for all $i \in \mathcal{I}$, $t \in \mathcal{T}$, and $j \in \mathscr{P}_i^t$. Let $\mathcal{Z}_{ij}^t$ be the random variable representing whether users i and j see each other simultaneously in period t and let it be defined as $\mathcal{Z}_{ij}^t = (\mathcal{X}_{ij}^t \cdot \Upsilon_i^t) \cdot (\mathcal{X}_{ji}^t \cdot \Upsilon_i^t)$. 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

$$\mathrm{P}\left(\Phi_{ij}^t = 1 \mid \mathcal{X}_{ij}^t = 1, \Upsilon_i^t = 1\right) = \mathrm{P}\left(\Phi_{ij}^t = 1 \mid \mathcal{Y}_{ij}^t = 1, \Upsilon_i^t = 1\right) = \mathrm{P}\left(\Phi_{ij}^t = 1 \mid \mathcal{Z}_{ij}^t = 1, \Upsilon_i^t = 1\right) \le \phi_{ij}^0.$$

By the feasibility of the optimal policy, we know that assortments contain at most K_i^t potential partners per period, i.e., with probability one,

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

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

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

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

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

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

and, even when $\Upsilon_i^t=1$ and $\Upsilon_j^t=1$, with probability one we have that

$$\mathcal{Z}_{ij}^t\le\mathcal{X}_{ij}^t,\quad \mathcal{Z}_{ij}^t\le\mathcal{X}_{ji}^t,\quad \forall j\in\mathscr{P}_i^1, i\in\mathcal{I},\, t\in\mathcal{T}.$$

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.,

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

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

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

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:

$$\begin{aligned}
\mathbb{E}\left[\Upsilon_j^t\cdot\Phi_{ji}^t\cdot(\mathcal{X}_{ji}^t-\mathcal{Z}_{ji}^t)\right]&=\mathbb{E}\left[\Phi_{ji}^t\mid\mathcal{X}_{ji}^t=1,\Upsilon_j^t=1\right]\cdot\mathrm{P}\left(\mathcal{X}_{ji}^t=1,\Upsilon_j^t=1\right)-\mathbb{E}\left[\Phi_{ji}^t\mid\mathcal{Z}_{ji}^t=1\right]\cdot\mathrm{P}\left(\mathcal{Z}_{ji}^t=1\right)\\
&=\mathbb{E}\left[\Phi_{ji}^t\mid\mathcal{X}_{ji}^t=1,\Upsilon_j^t=1\right]\cdot\mathrm{P}\left(\mathcal{X}_{ji}^t=1\right)\cdot\mathrm{P}\left(\Upsilon_j^t=1\right)\\
&\quad-\mathbb{E}\left[\Phi_{ji}^t\mid\mathcal{Z}_{ji}^t=1,\Upsilon_j^t=1\right]\cdot\mathrm{P}\left(\mathcal{Z}_{ji}^t=1\mid\Upsilon_j^t=1,\Upsilon_i^t=1\right)\cdot\mathrm{P}\left(\Upsilon_j^t=1\right)\cdot\mathrm{P}\left(\Upsilon_i^t=1\right)\\
&\le\phi_{ji}^0\cdot\mathrm{P}\left(\mathcal{X}_{ji}^t=1\right)\cdot\mathrm{P}\left(\Upsilon_j^t=1\right)-\phi_{ji}^0\cdot\mathrm{P}\left(\mathcal{Z}_{ji}^t=1\mid\Upsilon_j^t=1,\Upsilon_i^t=1\right)\cdot\mathrm{P}\left(\Upsilon_j^t=1\right)\cdot\mathrm{P}\left(\Upsilon_i^t=1\right)\\
&=\upsilon_j^t\cdot\phi_{ji}^0\cdot\left(\mathbb{E}\left[\mathcal{X}_{ji}^t\right]-\upsilon_i^t\cdot\mathbb{E}\left[\mathcal{Z}_{ji}^t\mid\Upsilon_j^t=1,\Upsilon_i^t=1\right]\right),
\end{aligned}$$

where the second equality follows from the independence and $\mathcal{X}_{ij}^t$ and Υ_j^t, the inequality follows by the definition of Φ_{ji}^t (i.e., the like probability conditional on *seeing* a profile) together with the fact that $\phi_{ji}(M)\le\phi_{ji}^0$ for all M, and the last equality uses the definition of υ_j^t.

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 $\bar{z}_{ij}^t=\mathbb{E}\left[\mathcal{Z}_{ij}^t\mid\Upsilon_j^t=1,\Upsilon_i^t=1\right]$, 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:

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

The first event captures the case where 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. Then,

$$\begin{aligned}
\mathbb{E}\left[\mu_{ij}^t\right] &= \mathrm{P}\left(\Phi_{ji}^t=1,\Upsilon_j^t=1,\mathcal{Y}_{ji}^t=1\right)+\mathrm{P}\left(\Phi_{ij}^t=1,\Upsilon_i^t=1,\mathcal{Y}_{ij}^t=1\right)+\mathrm{P}\left(\Phi_{ji}^t=1,\Phi_{ij}^t=1,\mathcal{Z}_{ij}^t=\mathcal{Z}_{ji}^t=1\right)\\
&= \mathrm{P}\left(\Phi_{ji}^t=1\mid\mathcal{Y}_{ji}^t=1,\Upsilon_j^t=1\right)\cdot\mathrm{P}\left(\mathcal{Y}_{ji}^t=1,\Upsilon_j^t=1\right)+\mathrm{P}\left(\Phi_{ij}^t=1\mid\mathcal{Y}_{ij}^t=1,\Upsilon_i^t=1\right)\cdot\mathrm{P}\left(\mathcal{Y}_{ij}^t=1,\Upsilon_i^t=1\right)\\
&\quad+\mathrm{P}\left(\Phi_{ji}^t=1,\Phi_{ij}^t=1\mid\mathcal{Z}_{ij}^t=\mathcal{Z}_{ji}^t=1\right)\cdot\mathrm{P}\left(\mathcal{Z}_{ij}^t=\mathcal{Z}_{ji}^t=1\right)\\
&\le \phi_{ji}^0\cdot\mathrm{P}\left(\mathcal{Y}_{ji}^t=1\right)\cdot\mathrm{P}\left(\Upsilon_j^t=1\right)+\phi_{ji}^0\cdot\mathrm{P}\left(\mathcal{Y}_{ij}^t=1\right)\cdot\mathrm{P}\left(\Upsilon_i^t=1\right)\\
&\quad+\mathrm{P}\left(\Phi_{ji}^t=1\mid\mathcal{Z}_{ij}^t=1\right)\cdot\mathrm{P}\left(\Phi_{ij}^t=1\mid\mathcal{Z}_{ij}^t=1\right)\cdot\mathrm{P}\left(\mathcal{Z}_{ij}^t=1\mid\Upsilon_i^t=1,\Upsilon_j^t=1\right)\cdot\mathrm{P}\left(\Upsilon_i^t=1\right)\cdot\mathrm{P}\left(\Upsilon_j^t=1\right)\\
&\le \upsilon_j^t\cdot\phi_{ji}^0\cdot\bar{y}_{ji}^t+\upsilon_i^t\cdot\phi_{ij}^0\cdot\bar{y}_{ij}^t+\upsilon_i^t\cdot\upsilon_j^t\cdot\phi_{ij}^0\cdot\phi_{ji}^0\cdot\bar{z}_{ij}^t.
\end{aligned}$$

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

$$\mathbb{E}\left[\sum_{t\in\mathcal{T}}\sum_{i\in\mathcal{I}}\sum_{\{j\in\mathscr{P}_i^1:j<i\}}\mu_{ij}^t\right]\quad\le\quad\sum_{t\in\mathcal{T}}\sum_{i\in\mathcal{I}}\sum_{j\in\mathscr{P}_i^j}\upsilon_j^t\cdot\phi_{ji}^0\cdot\bar{y}_{ji}^t+\frac{1}{2}\cdot\upsilon_i^t\cdot\upsilon_j^t\cdot\phi_{ij}^0\cdot\phi_{ji}^0\cdot\bar{z}_{ij}^t,$$

where in the first term we take the summation over $\{j\in\mathscr{P}_i^1:j<i\}$ to avoid counting matches twice, and equality follows by changing the terms in the summation and avoiding counting matches twice for the simultaneous shows. 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. Therefore, 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: Appendix to Section 4: Additional Estimation Results

B.1. Fixed Effects Model

In Table 4 we report the results of logit models with user fixed effects. We observe that the number of matches in the recent past has a negative and significant effect for all the specifications considered. In addition, we observe that the magnitude of the effect is decreasing in time; i.e., users are less impacted by older matches. These results support our choice to focus on matches obtained since the last active session in order to reduce the complexity of the problem.

Table 4 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
Pseudo-R^2	0.241	0.242	0.242
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, religion, days on the platform, and additional metrics including the number of days active and the number of likes and the number of profiles from the backlog observed in period t^0 (last day before the change in the algorithm), and also the number of matches obtained between period t^0 and t^1. In addition, we include covariates for the backlog composition in t^1, i.e., for the number of active and inactive users of each type in the backlog. Finally, we also include summary statistics of the attractiveness of the profiles in the user's backlog, including the minimum, maximum, and average and standard deviation of the score.

The results in Figure 5 show that neural nets achieve the best fit. Nevertheless, we observe that all estimates are relatively close to each other. Moreover, logit leads to the most conservative estimates for the history effect in the last stage of our estimation procedure, and so using these estimates provides a lower bound on the impact of our heuristics that leverage the history effect.

Figure 5 Propensity Score Estimation: ROC Curves

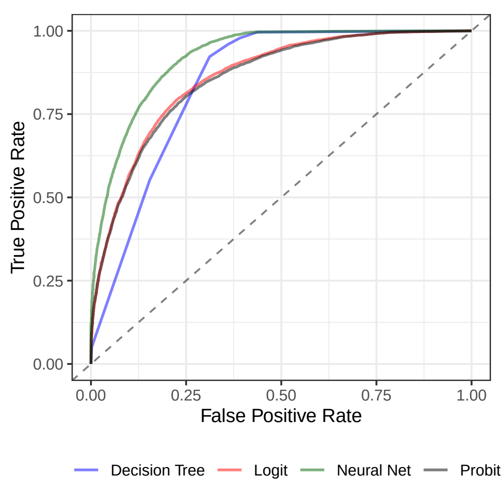

In Figure 6 we plot the distribution of propensity scores for the original and the trimmed sample (at 10%). We observe that there is overlap in the distributions for the treatment and control group, which ensures that we will not have to rely on extrapolation to compare these two groups.

Figure 6 Distribution of Propensity Scores

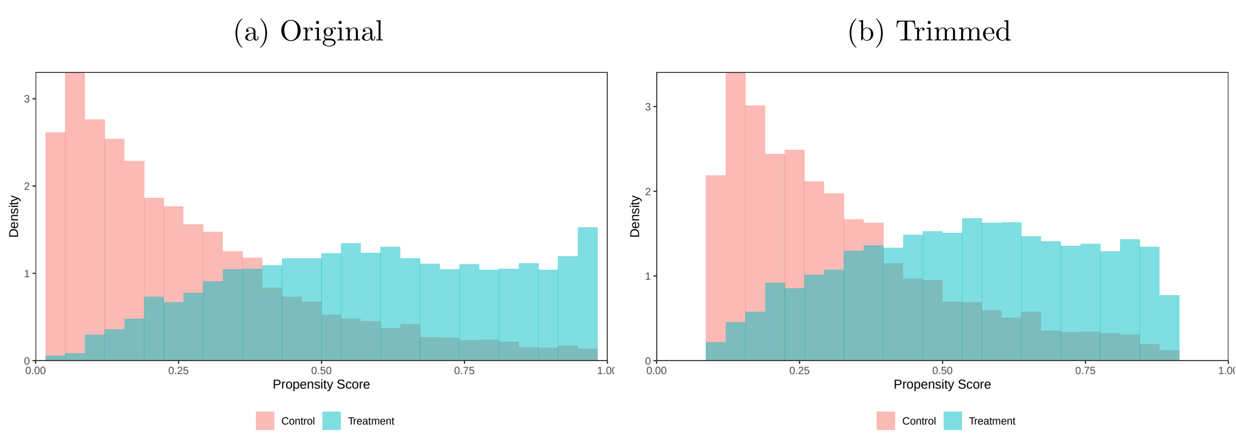

As discussed in Section 4, we use these propensity scores to compute weights that are used in the estimation of the history effect. To illustrate that the treatment and control groups in the sample considered for estimation are well balanced, in Table 5 we report several balance statistics for each of the most relevant variables. Notice that the normalized difference is lower than 0.1 for all the variables considered (except for one them, for which it is less than 0.2), which suggests that our sample is well balanced.

Table 5 Balance Measures

	Treated		Control		Balance Measures			
	Mean	Std.	Mean	Std.	Norm. Diff	Log Var. Ratio	$\pi_C^{0.05}$	$\pi_T^{0.05}$
is female	0.4	0.49	0.41	0.49	-0.02	0	0	0
age pid	30.73	6.58	30.7	6.48	0	0.02	0.06	0.04
height pid	68.87	3.97	68.84	4.05	0.01	-0.02	0.02	0.02
edu pid	2	0.83	2.01	0.84	-0.01	-0.01	0	0
score pid	3.85	1.88	3.81	1.84	0.02	0.02	0.04	0.05
quintalg pid	3.29	1.3	3.24	1.29	0.04	0.01	0	0
days in app	600.79	363.25	582.58	344.16	0.05	0.05	0.03	0.08
exp user	0.99	0.09	0.99	0.1	0.01	-0.06	0	0.01
asian	0.08	0.27	0.08	0.27	-0.01	-0.02	0	0
white	0.78	0.41	0.77	0.42	0.02	-0.01	0	0
black	0.05	0.22	0.05	0.22	0	-0.01	0	0
hispanic	0.05	0.22	0.05	0.22	0	-0.01	0	0
other	0.08	0.27	0.09	0.28	-0.01	-0.02	0	0
num from backlog t0	0.58	0.93	0.56	0.86	0.01	0.08	0	0.05
num matches t1 since last	0.34	0.71	0.33	0.68	0.02	0.04	0.01	0.02
num likes t0	1.4	1.21	1.38	1.19	0.02	0.02	0	0.01
backlog asq active cts	5.22	11.91	4.92	9.72	0.03	0.2	0.01	0.04
backlog bhq active cts	3.48	5.3	3.43	5.43	0.01	-0.02	0.03	0.02
backlog ahq active cts	4.42	8.2	3.94	7.27	0.06	0.12	0.02	0.03
backlog asq inactive cts	0.82	2.19	0.66	1.58	0.08	0.33	0	0.07
backlog bhq inactive cts	0.58	0.8	0.51	0.92	0.08	-0.13	0.01	0.03
backlog ahq inactive cts	0.71	1.39	0.47	1.16	0.19	0.18	0	0.08
backlog asq active avg	3.22	2.57	3.13	2.62	0.03	-0.02	0.02	0.03
backlog bhq active avg	1.88	1.74	1.86	1.79	0.01	-0.03	0.02	0.02
backlog ahq active avg	1.4	1.42	1.32	1.46	0.05	-0.03	0.02	0.03
backlog asq inactive avg	1.71	2.46	1.61	2.54	0.04	-0.03	0.02	0.03
backlog bhq inactive avg	1.08	1.49	1.04	1.69	0.03	-0.13	0.03	0.02
backlog ahq inactive avg	0.76	1.2	0.67	1.35	0.07	-0.12	0.02	0.03
backlog asq active sd	0.43	0.52	0.41	0.52	0.04	0.01	0.02	0.03
backlog bhq active sd	0.32	0.44	0.33	0.46	-0.01	-0.05	0.02	0.03
backlog ahq active sd	0.36	0.46	0.34	0.47	0.04	-0.01	0.02	0.03
backlog asq inactive sd	0.1	0.31	0.09	0.32	0.04	-0.02	0.01	0.06
backlog bhq inactive sd	0.06	0.26	0.05	0.24	0.02	0.07	0.01	0.05
backlog ahq inactive sd	0.08	0.29	0.06	0.26	0.09	0.11	0.01	0.08
backlog asq active min	2.69	2.29	2.63	2.32	0.03	-0.01	0.02	0.02
backlog bhq active min	1.47	1.57	1.45	1.59	0.01	-0.01	0.02	0.03
backlog ahq active min	0.94	1.18	0.9	1.19	0.03	-0.01	0.02	0.03
backlog asq inactive min	1.61	2.35	1.53	2.43	0.04	-0.04	0.02	0.03
backlog bhq inactive min	1.04	1.45	1	1.64	0.03	-0.12	0.03	0.02
backlog ahq inactive min	0.68	1.11	0.62	1.28	0.05	-0.14	0.02	0.03
backlog asq active max	3.85	3.08	3.73	3.12	0.04	-0.01	0.01	0.04
backlog bhq active max	2.27	2.04	2.25	2.12	0.01	-0.04	0.03	0.02
backlog ahq active max	1.86	1.82	1.75	1.87	0.06	-0.03	0.02	0.03
backlog asq inactive max	1.81	2.63	1.69	2.68	0.05	-0.02	0.01	0.05
backlog bhq inactive max	1.13	1.56	1.08	1.76	0.03	-0.12	0.02	0.03
backlog ahq inactive max	0.84	1.34	0.73	1.46	0.08	-0.09	0.02	0.04

Note: Data considers all users included in the regression analysis based on IPTW, i.e., after trimming. Reported statistics are weighted averages and weighted standard deviations for each of the listed variables. In addition, we include different balance measures for each variable, including the normalized difference, the log variation ratio, among others.

B.3. Conversations between Users

To assess whether users may realize that they got matched with inactive users, we analyze the conversations between users who got matched in period t^1, and we compare them with those conversations between users matched in period t^0. Specifically, we analyze the time between when a match is formed and the completion of one round of messaging between the users (one message from each user). In Figure 7a we compare the fraction of matches (formed in periods t^0 and t^1) that complete one round of messaging before and after the sender of the first message gets a new assortment. The *After* bar includes matches for which there was a first message and the first reply came after the first sender's next assortment. The *Before* bar matches in t^1 for which the first message and reply happened before the first sender's next assortment. The *NF* bar includes all matches for which there was no *first message* ever. Finally, the *NS* bar includes all matches for which there was no second message (i.e., no reply to the message from the first sender) ever. We find that 13.93% of the matches formed in t^0 completed one round of messaging before the first sender's next assortment, while this percentage is 11.84% for the matches formed in t^1. Since the difference is only 2% and statistically insignificant, we conclude that users could not tell the difference between an active match and an inactive match by period t^2. In Figure 7b we modify this analysis and classify matches based on whether the completion of the first round of messages happened within three days of the creation of the match. As before, we do not find significant differences between the matches formed in t^0 and t^1. Hence, we conclude that users cannot tell the difference between inactive and active users in the short run based on their conversation experience.

Figure 7 **Comparison of Conversation Patterns between Matches Formed in t^0 vs. t^1**

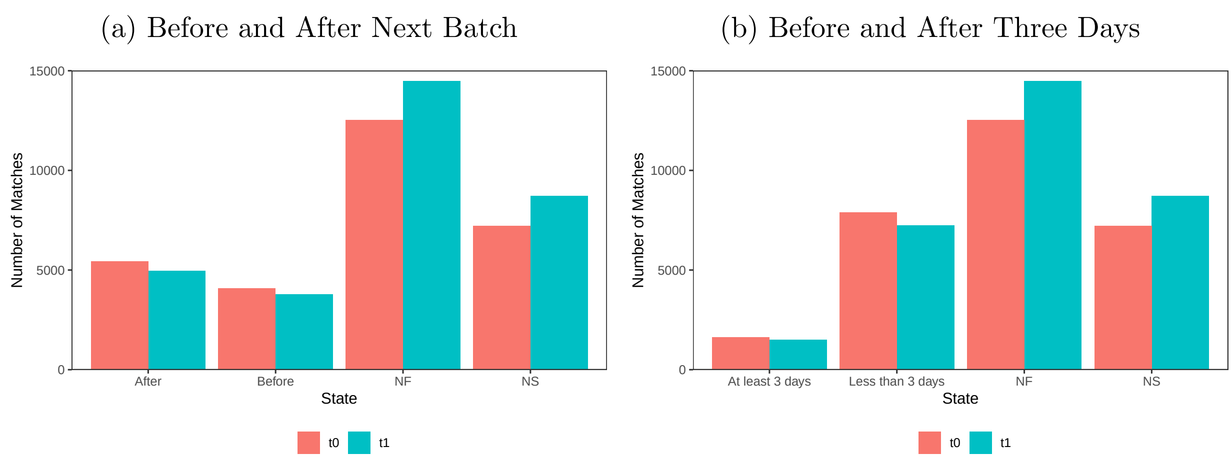

B.4. Log-in Probabilities

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 18 and August 17, 2020. For each user in the sample, and for each day between March 18 and August 17, 2020, we record whether or not the user was active. We assume that a user was 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, including the number of days since the user's last session, the number of likes given in their last session, and the number of matches obtained since the user's last session. In addition, for each user we include time-invariant observable characteristics, including age, height, education, race, attractiveness score, among others.

In Figure 8 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, and so from now on we focus on the logit model for simplicity.

To estimate this model we consider the following specification:

$$\Upsilon_i^t = X_i\alpha_1 + Z_i^t\alpha_2 + \epsilon_i^t, \tag{8}$$

Figure 8 Log-in Model: ROC Curves

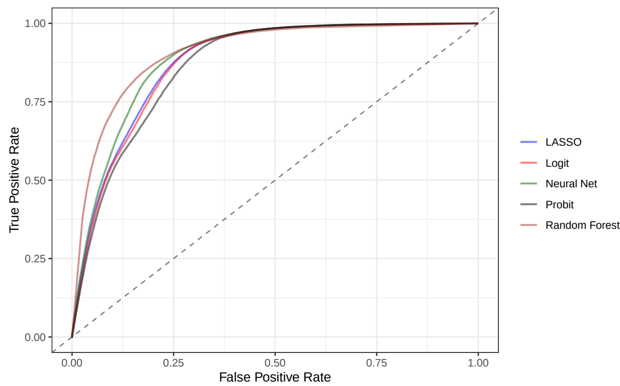

where Υ_i^t and X_i are as defined in Section 3; i.e., X_i are observable characteristics of user i and $\Upsilon_i^t = 1$ if user i is active in period t and $\Upsilon_i^t = 0$ otherwise. Z_i^t is a vector of time-varying observable characteristics, including the day of the week, the number of days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

In Table 6 we report the estimation results for all users that logged in at least once in the first month of our sample (i.e., between February 18 and March 18, 2020).[20] In columns (1) and (2) we report the estimation results obtained for users in Houston, TX. These are the parameters that were used as part of both field experiments (column (2)). Columns (3) and (4) report the results for users in Austin, TX. In addition, columns (2) and (4) control for fixed effects at the user level. As previously discussed, all these specifications include three usage metrics: (1) the number of days since the last session, (2) the number of likes in the last session, and (3) the number of matches since the last session. We also estimated the model with more usage metrics and the results are consistent.

First, we observe that the results obtained for Houston are relatively similar to those obtained for Austin. This suggests that using the parameters obtained for Houston in both field experiments is without major loss. Moreover, it suggests that the improvement obtained in our field experiments is not driven by the market chosen to estimate the parameters of the log-in model. Second, we observe significant day of the week effects, with Friday and Saturday being the days with the lowest log-in rates. Finaly we observe that the number of matches obtained since the last session is either not significant or negative. In other words, matches in the recent past either do not have any effect of reducing the probability that users will log in. Hence, our heuristic that penalizes for the effect that matches today have in future periods would also internalize the negative effect of matches on log ins. Nevertheless, we also observe that the effect of matches since the last session is relatively small compared to the effect of the day of the week. For this reason, assuming that log-in probabilities are exogenous and that these do not depend on the history effect is without major loss.

B.5. Additional Robustness Checks

B.5.1. Regression Including Quality of Matches and Difference Relative to Score of Focal User. We estimate our model including controls for several aspects of the attractiveness of the matches obtained by users, including the average, minimum, and maximum score of the matches obtained, and also including the positive and negative difference between the score of the focal user and the average score of the matches obtained. The first set of covariates allows us to control for the absolute attractiveness of the matches obtained, while the latter two covariates capture the effect relative to the attractiveness of the user obtaining these additional matches.

In Table 7 we present the results. Columns (1) and (2) control for the average attractiveness of the matches obtained, while columns (3) and (4) control for the positive and negative difference between the score of the focal user and the mean attractiveness of the matches obtained between t^1 and t^2. Odd columns estimate the first stage with OLS, while even columns do so using a negative binomial regression.[21]

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

[21] We compute the positive difference as $[s_i - \bar{s}_i]^+ = (s_i - \bar{s}_i) \cdot 1_{s_i \geq \bar{s}_i}$, where s_i is the score of the focal user and $\bar{s}_i$ is the average score of the matches obtained by user i since their last session. Similarly, we compute the negative difference as $[s_i - \bar{s}_i]^- = (\bar{s}_i - s_i) \cdot 1_{s_i \leq \bar{s}_i}$.

Table 6 Log-in Model: Houston vs. Austin

	Dependent Variable: Log-in			
	Houston		Austin	
	(1)	(2)	(3)	(4)
Female	-0.570***	-	-0.260***	-
	(0.094)	-	(0.090)	-
Experience	-0.184***	-0.825***	-0.174***	-1.050***
	(0.068)	(0.078)	(0.062)	(0.081)
Monday	0.175***	0.213***	0.182***	0.227***
	(0.017)	(0.021)	(0.017)	(0.020)
Saturday	0.025	0.019	0.037**	0.032*
	(0.016)	(0.020)	(0.015)	(0.018)
Sunday	0.167***	0.194***	0.190***	0.224***
	(0.017)	(0.021)	(0.017)	(0.021)
Thursday	0.070***	0.095***	0.077***	0.109***
	(0.015)	(0.019)	(0.015)	(0.018)
Tuesday	0.197***	0.242***	0.199***	0.250***
	(0.017)	(0.021)	(0.017)	(0.020)
Wednesday	0.166***	0.221***	0.192***	0.254***
	(0.016)	(0.020)	(0.016)	(0.020)
Matches since Last Session	-0.130***	-0.011	-0.129***	-0.038**
	(0.020)	(0.018)	(0.019)	(0.017)
Usage Metrics	Yes	Yes	Yes	Yes
User Fixed Effects	No	Yes	No	Yes
Observations	290,083	276,973	321,589	303,497
Pseudo R^2	0.343	0.441	0.357	0.438

Note: Columns (1) and (2) consider data from Houston, while (3) and (4) consider data from Austin. Usage metrics include days since last active session, likes in last active session, and number of matches since last session. Standard errors are clustered at the user level. Significance reported: *p<0.1; **p<0.05; ***p<0.01.

Table 7 Regression Results: Adding Score of Matches in t^1

	(1)	(2)	(3)	(4)
$\hat{M}_i^{t_2}$	−0.737***	−0.490***	−0.719***	−0.492***
	(0.150)	(0.117)	(0.150)	(0.117)
Constant	−3.050***	−2.992***	−2.984***	−2.926***
	(1.126)	(1.118)	(1.128)	(1.120)
	Score of Matches since t^1			
Mean	0.159	0.221	-	-
	(0.576)	(0.330)	-	-
Min.	−0.265	−0.304	−0.271	−0.311
	(0.290)	(0.290)	(0.286)	(0.285)
Max.	0.182	0.157	0.153	0.125
	(0.291)	(0.291)	(0.287)	(0.287)
Positive Diff.	-	-	−0.222	−0.291
	-	-	(0.569)	(0.568)
Negative Diff.	-	-	0.135	0.198
	-	-	(0.569)	(0.568)
Observations	51,561	51,561	51,561	51,561
Pseudo-R^2	0.327	0.327	0.328	0.327

Note: The first stage is estimated using OLS (columns (1) and (3)) and negative binomial (columns (2) and (4)) regressions. The second stage for the negative binomial first-stage regressions addresses the forbidden regression problem. Standard errors are clustered at the user level. Significance reported: *p<0.1; **p<0.05; ***p<0.01.

First, we observe that the coefficient of the number of matches since the last session is negative and significant for all specifications considered. In addition, we observe that the magnitude of the effect is larger than that obtained when no controls for the quality of matches between t^1 and t^2 are included. Second, we observe that none of the measures of the attractiveness of the matches obtained since t^1 are significant. These results suggest that the attractiveness of the matches obtained has no effect on the like behavior of users in period t^2, and that our results are robust to controlling for the quality of the matches obtained.

B.5.2. Effect of Matches since Last Session on Log-in Decisions during Quasi-Experiment. A potential concern regarding the results reported in Appendix B.4 is that this analysis may suffer from a similar endogeneity problem to that described in Section 4. To rule out this concern, we estimate the effect of the matches obtained since the last session on log-in decisions by means of the same quasi-experimental variation used in the analysis of the history effect.

Notice that, to accomplish this, we cannot directly use the sample and the treatment definition described in the analysis of the history effect (i.e., the effect of matches since the last session) on like behavior. The reason is that our history effect analysis is conditional on users logging in at least one time before and two times after the change in the algorithm (i.e., log-ins in periods t^0, t^1, and t^2). For this reason, we only condition on users logging in at least once before and after the change in the algorithm. The condition of logging in once after the change in the algorithm is to allow users to see their inactive backlog profiles and effectively get the treatment, while the condition of logging in once before the change in the algorithm is to make the periods before and after the change in the algorithm comparable.

Similar to the analysis of the history effect, we define a user as treated if they see an inactive backlog profile in their first session after the change in the algorithm, and we define the control group as users that have an inactive backlog but do not see it for exogenous reasons. Then, to assess the effect of the additional matches obtained due to the change in the algorithm on log-in decisions, we adapt the procedure described in Section 4 to estimate the effect of matches on log-in decisions. More specifically, we proceed as follows.

1. Propensity Score (PS) Estimation: estimate propensity scores for the probability of being treated, using as controls the same set of covariates as in the analysis of the history effect. We add to this set of controls two additional covariates that are relevant to understanding log-in decisions: (1) the number of times the user logged in in the ten days before the change in the algorithm, and (2) the number of days between the last and the first time the user logged in before and after the change in the algorithm. Based on the estimated propensity scores, we compute weights that we later use for trimming and also in the estimation of the first and second stages.
2. First-Stage Regression of 2SLS: using the weights estimated in the previous step, estimate a weighted model where the dependent variable is the number of matches obtained by the user since their last session, and the controls are the average characteristics of the profiles seen in period t^1 and all the controls included in the propensity score estimation. Using the estimated parameters, compute the predicted number of matches in the recent past. It is important to notice that, if the user logged in, we count the number of matches obtained between their last session and right before they evaluate new profiles. On the other hand, for users that did not log-in we count matches until the end of the day.[22]
3. Second-Stage Regression of 2SLS: using the weights estimated in the first step, estimate a model where the dependent variable is a dummy equal to 1 if the user logged in at time t, and 0 otherwise. The main variable of interest is the estimated number of matches since the last session at time t, obtained in the previous step.

We estimate this model for three subsets of data: (1) the day after t^1, i.e., $t = t^1 + 1$; (2) the first three days after t^1, i.e., $t \in \{t^1 + 1, \ldots, t^1 + 3\}$; and (3) the full week following t^1, i.e., $t \in \{t^1 + 1, \ldots, t^1 + 7\}$. By considering (1) we can estimate the direct effect that matches obtained due to the change in the algorithm had on log-in decisions on the very next day. By considering the next three days, we can assess whether the matches obtained had an effect on log-in decisions on the following days, including[23] t^2. Finally, we consider a full week as an additional robustness check.

In Table 8 we report the estimation results. As in the analysis of the history effect, we trim observations with a propensity score outside $[0.1, 0.9]$ to increase overlap.[24] Panel 1 reports the results of the first-stage

[22] We tried counting until the middle of the day and the results did not change.

[23] Recall that, in the analysis of the history effect, we restrict t^2 to be within three days from t^1, i.e., no later than May 2, 2019.

[24] The results remain unchanged if we consider the entire sample and do not trim observations.

regressions, while Panel 2 reports the results of the second stage. For the second stage, columns (2), (4), and (6) add controls for the quality of the matches obtained since the last session, including the mean, minimum, and maximum attractiveness score.

Table 8 Effect of Matches on Log-in Rate

	Dependent Variable: Log-in					
	One		Three		Seven	
	(1)	(2)	(3)	(4)	(5)	(6)
	Panel 1: First Stage					
Treated	0.330***		0.288***		0.220***	
	(0.018)		(0.012)		(0.008)	
	Panel 2: Second Stage					
Matches since Last Session	0.110	0.169	-0.016	0.023	0.011	0.089
	(0.097)	(0.105)	(0.079)	(0.083)	(0.069)	(0.077)
Constant	0.961	0.950	2.54***	2.56***	2.07***	2.07***
	(1.26)	(1.29)	(0.882)	(0.897)	(0.763)	(0.776)
Observations	23,488	23,488	70,464	70,464	164,416	164,416
Pseudo-R^2	0.259	0.260	0.268	0.268	0.286	0.286

Note: All columns consider the trimmed sample at 10%. Even columns add controls for the quality of the matches obtained since the last session, including the mean, minimum, and maximum attractiveness score. The first stage is estimated using negative binomial regressions. The second stage addresses the forbidden regression problem. Standard errors are clustered at the region level for columns (1) and (2), and at the user level for the other columns. Significance reported: $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

First, from Panel 1 we observe that the treatment dummy has a positive and significant effect on the number of matches obtained since the last session for all models considered. This result confirms that our treatment variable is a relevant instrument for the number of matches obtained since the last session, as previously discussed in the analysis of the history effect. In addition, we observe that the magnitude of the coefficient is decreasing as we move from the next day to the full week following t^1. This result is expected, because the change in the algorithm had a significant effect in the days right after the change in the algorithm, and its effect decayed as time progressed. Second, from Panel 2 we observe that the predicted number of matches since the last session has no significant effect on log-in decisions for any of the models considered. Finally, we observe no major differences if we add controls for the quality of the matches obtained since the last session.

These results are in line with our previous findings and confirm that the number of matches obtained since the last session has no significant effect on log-in behavior.

Appendix C: Appendix to Section 6: Additional Simulation Results

C.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 with different coefficients 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 in the multiplier of the history effect. Following our previous example, we observe that having a history effect twice as large as the estimated one 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 in Figure 9b, where we report boxplots with the percentage of improvement relative to Greedy and our partner's algorithm.

Figure 9 Sensitivity History: Matches

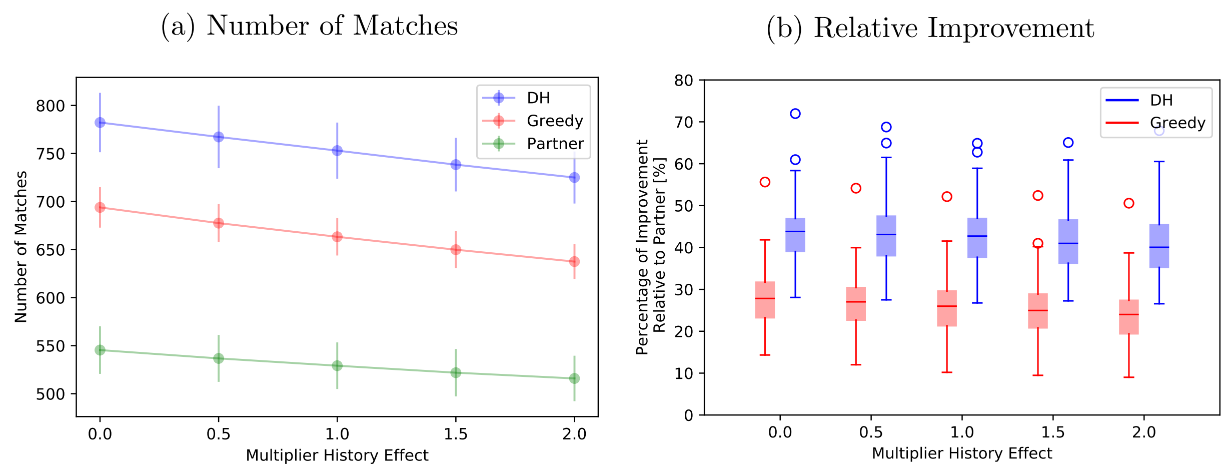

C.1.1. Sensitivity Region. To test the sensitivity of our results to the region considered, in Table 9 we compare the number of matches obtained for different regions. In each region we construct a market following the guidelines described above. Moreover, we consider the same parameters to compute like probabilities, and we assume that log-in rates are 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 the same log-in rates and business constraints as before. Finally, for DH we consider $\bar{\xi} = -0.05$.

Table 9 Comparison of 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.

C.2. 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 in order for their profile 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 for different n_s and n_b. We consider the same data and simulation setup described above.

In Figure 10 we report the average from 100 simulations for $m = 3$, $T = 7$, $K = 3$, and no business constraints.

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.

Figure 10 Sensitivity to Bounds: Partner and Greedy

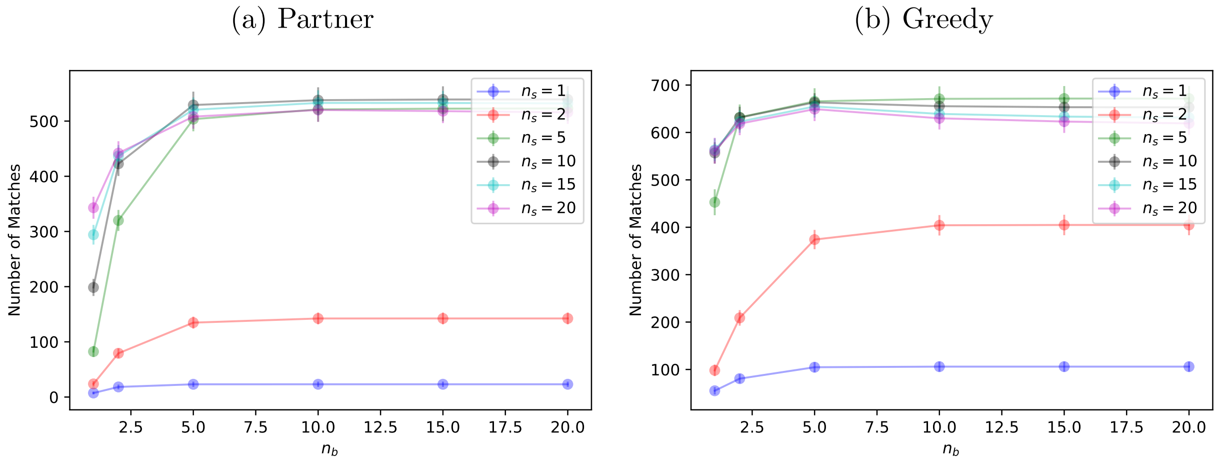

C.3. Sensitivity to Horizon and Characteristics of Profiles Shown

To analyze the effect of the proposed heuristic in a longer horizon, in this section we present simulation results involving a time horizon of two weeks. In these simulations we consider an instance created with data from Houston, right before the first field experiment. More specifically, we construct a submarket that includes users that have logged in at least once in the 28 days prior to the field experiment (i.e., between July 22 and August 18, 2019), and that have at least 30 potentials in the submarket. In addition, we consider the actual backlogs and initial number of matches that each user had on August 19.

To compute like probabilities we use the same parameters as in the field experiment, and we use fixed log-in rates to solve the optimization problem in the first step of DH. However, log-ins are realized according to the actual login data between August 19 and September 2, 2020. Finally, as in the simulations reported in Section 6, we consider assortment size $K = 3$ and scale the penalty term by $\xi = 0.1$. We simulated this market 100 times.

In Table 10 we report summary statistics of the instance considered, including the average scores, number of potentials, backlog sizes, and matches obtained since the last session.

Table 10 Summary Statistics of the Instance Used for the Simulations

	N	Score	Potentials	Backlog	Prev. Matches
Female	830	4.681	128.711	21.892	0.161
		2.356	89.832	41.024	0.453
Male	757	2.610	120.524	3.466	0.185
		1.514	82.167	8.717	0.485

In Table 11 we compare the performance of DH and the default algorithm. On the one hand, we observe that the default algorithm achieves a higher number of likes. This result is expected and in line with our field experiment results, as we know that our partner's algorithm prioritizes likes. On the other hand, we observe that DH achieves a higher number of matches. Indeed, we observe that the improvement of DH over Partner is 16.74%.

In Figure 11 we report the evolution of the number of queries and backlog queries shown, and also the number of likes and matches over the two weeks of the simulation horizon. First, we observe that DH shows a significantly higher number of backlog queries. Second, we observe that the default algorithm achieves a higher number of likes in all periods considered, but the number of matches generated in each period is higher for DH. Finally, we observe that the difference in the number of matches generated is relatively stable over time. These results confirm that the improvement observed is not due to exploiting the initial backlog.

To better understand the mechanism behind the improvement, in Table 12 we report summary statistics for the total number of queries, likes, and matches obtained from backlog and non-backlog queries. First, as

Table 11 Comparison of DH vs. Default

Algorithm	Queries	Likes	Matches
Default	23225.180	9542.060	1227.120
	(2.717)	(66.675)	(30.227)
DH	23254.100	8565.330	1432.600
	(2.928)	(54.871)	(30.877)

Note: Summary statistics for the number of queries, likes, and matches obtained over 100 simulations.

Figure 11 Evolution over Two Weeks

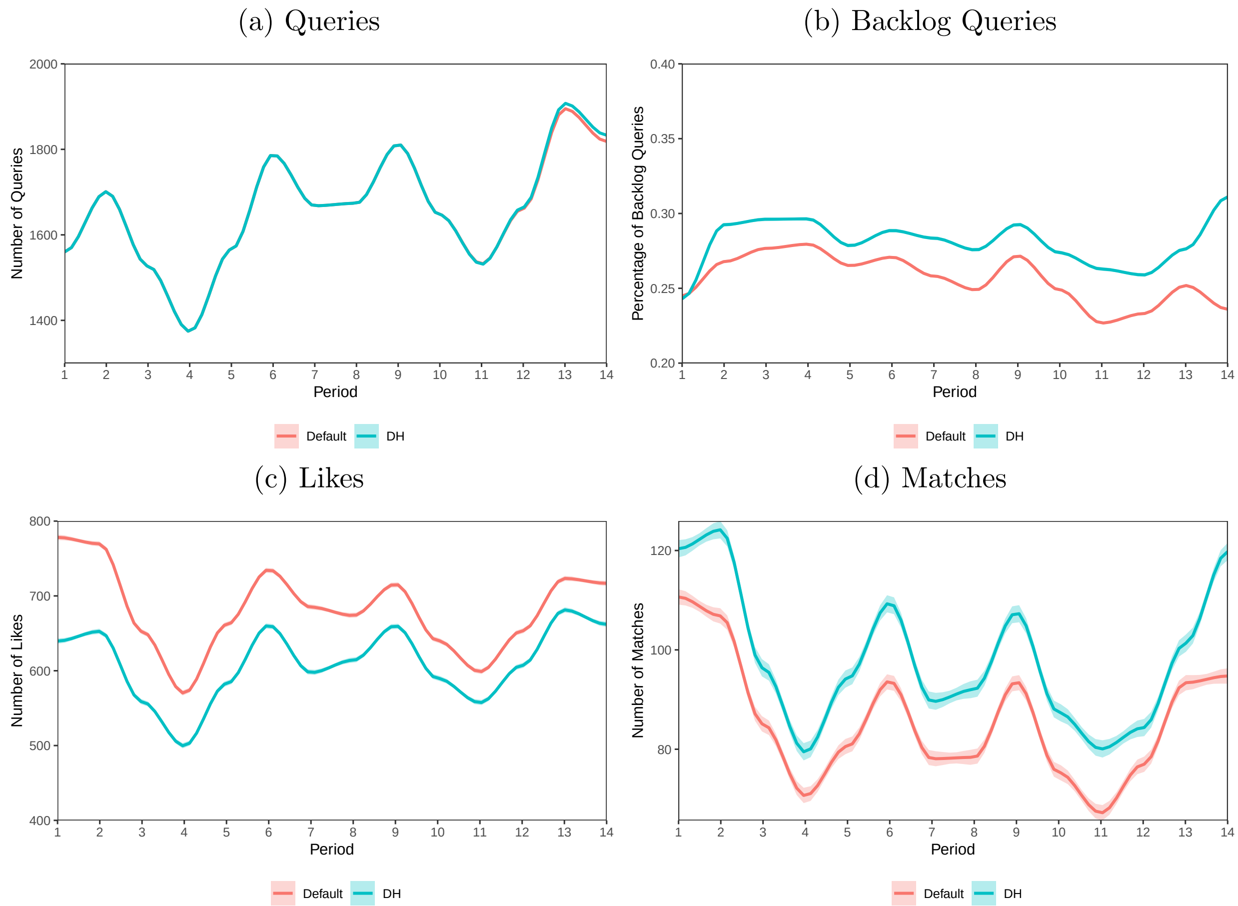

shown in Figure 11, we observe that the total number of backlog queries shown by DH is higher that that for the default algorithm. We also observe that the number of matches obtained from backlog queries generated by DH is higher (854.780 vs. 691.030). This suggests that DH selects non-backlog profiles that are more likely to result in a match. For this reason, we also observe that the number of initial non-backlog queries that end up being shown as backlog queries later is also higher for DH than for the default algorithm (4577.740 vs. 3747.980). This result is especially interesting considering that the number of non-backlog queries liked is considerably smaller for DH (7229.470 vs. 8364.460).

C.3.1. DH and the History Effect. To understand how DH accounts for the history effect, in this section we analyze in more detail how the profiles chosen by each algorithm change with the number of matches obtained by the user since their last session.

In Figure 12 we report the distribution of the number of matches since the last session and the backlog size for all user sessions. On the one hand, from Figure 12a we observe that most users have either 0, 1, or 2 matches since their last active session, and less than 1% of users have more than 2 matches since their last session. For this reason, in the remainder of the analysis we will focus on these three cases (matches since last

Table 12 Mechanism of Improvement

	Backlog			Non-Backlog				
Algorithm	Queries	Likes	Matches	Queries	Likes	Shown	Matches (SS)	Matches
Default	5957.740	1177.600	1177.600	17267.440	8364.460	3747.980	49.520	691.030
	(30.569)	(30.147)	(30.147)	(30.589)	(53.299)	(34.143)	(9.611)	(23.520)
DH	6529.840	1335.860	1335.860	16724.260	7229.470	4577.740	96.740	854.780
	(32.626)	(29.475)	(29.475)	(32.637)	(46.083)	(37.564)	(13.593)	(25.465)

Note: Reported statistics are the average and standard deviation (in parentheses) over 100 simulations for each algorithm. Shown represents non-backlog queries that were initially liked and that were shown as backlog queries within the horizon of the simulation. Matches (SS) represents the number of matches generated out of simultaneous shows, i.e., when two users see each other in the same period as part of non-backlog queries. Matches represents the number of non-backlog queries that resulted in a match within the horizon, i.e., that were initially liked, shown, and liked back as a backlog query, or that were part of a simultaneous show where both users liked each other.

session equal to 0, 1 or 2). On the other hand, from Figure 12b we observe that there is more heterogeneity in the backlog size, but most users have a backlog size with less than 10 profiles.

Figure 12 Distribution Sessions: Matches since Last Session and Backlog Size

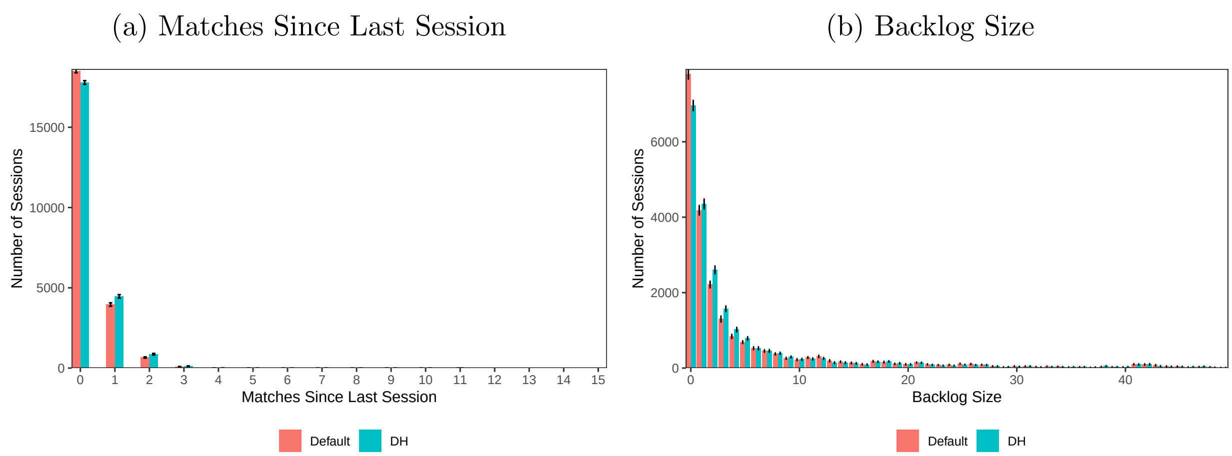

In Figure 13 we analyze the relation between the number of matches obtained since the last session and the fraction of backlog queries shown in the assortment. From Figure 13a we observe that, as the number of matches since the last session increases, the fraction of backlog queries in the assortment also increases. However, notice that when the number of matches since the last session is 2, this fraction is lower for DH. This result is counterintuitive considering that DH prioritizes backlog queries as previously discussed. To better understand this result, in Figures 13b and 13c analyze the fraction of backlog queries in the assortment as a function of the number of matches since the last session and of the backlog size. As expected, we observe that the fraction of backlog queries is increasing in the backlog size for both algorithms. However, conditional on the backlog size, we observe that this fraction is decreasing in the number of matches for DH, while there is no clear pattern for the default algorithm. Finally, we observe that when the backlog size increases, DH is less sensitive to the number of matches obtained since the last session.

One possible explanation for this pattern is that, when the number of matches since the last session is higher, DH tries to avoid showing more backlog profiles and "saves them" for a time when the number of matches since the last session is lower and thus the chances of generating a match out of those backlog queries is higher. However, DH's main goal is to generate matches, and so it will show more backlog profiles as the backlog size increases.

Figure 14 analyzes the attractiveness of the profiles shown by each algorithm. We observe that the average attractiveness of the profiles shown by DH is increasing in the number of matches obtained since the last session. In other words, when the history effect is high (due to a high number of matches since the last

Figure 13 Percentage of Queries from Backlog

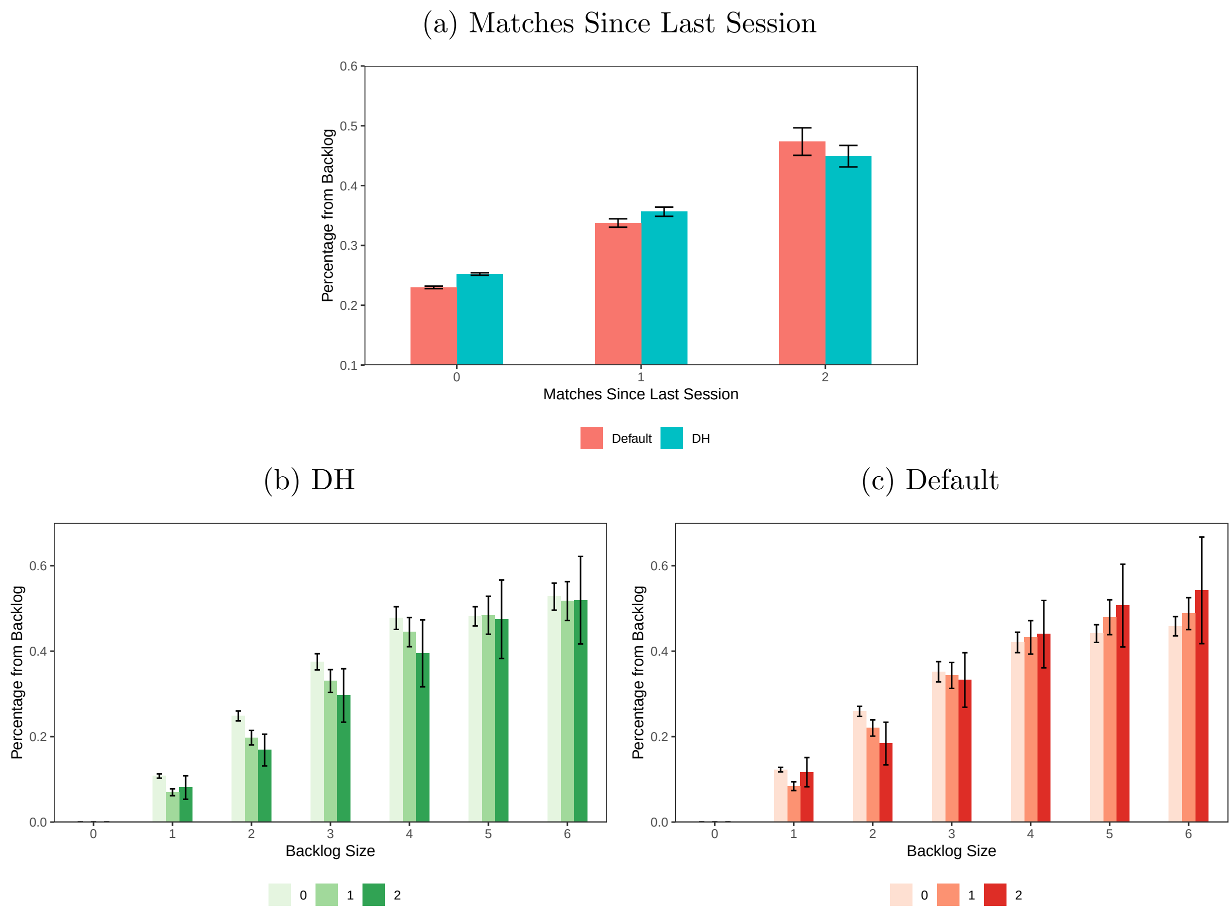

Note: The color scales represent the number of matches since the last session, with 0 matches being the lightest color and 2 being the darkest one.

session), DH compensates and chooses even more attractive profiles, saving less attractive ones that may be more likely to produce a match for later periods (when the history effect is smaller).

Appendix D: Appendix to Section 7 (Field Experiment)

D.1. Comparison of Markets

In Table 13, 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 5 and August 19, 2020, by market and gender.

Table 13 Summary Statistics: Treatment and Control Markets

			Age		Score		Log-in Rate		Like Rate		Match Rate	
		N	Mean	Std.	Mean	Std.	Mean	Std.	Mean	Std.	Mean	Std.
Austin	Women	1239	33.421	11.120	4.745	2.186	0.530	0.327	0.209	0.215	0.042	0.098
	Men	1223	33.795	10.454	2.708	1.543	0.709	0.389	0.425	0.283	0.039	0.094
Dallas	Women	1828	32.457	9.829	4.636	2.197	0.518	0.323	0.194	0.211	0.038	0.098
	Men	1604	34.500	12.272	2.620	1.590	0.704	0.376	0.416	0.279	0.034	0.084
Houston	Women	1285	33.371	10.121	4.280	2.242	0.528	0.337	0.200	0.223	0.037	0.106
	Men	1178	34.053	10.527	2.479	1.429	0.697	0.319	0.432	0.283	0.033	0.082

Figure 14 **Attractiveness of Profiles Shown**

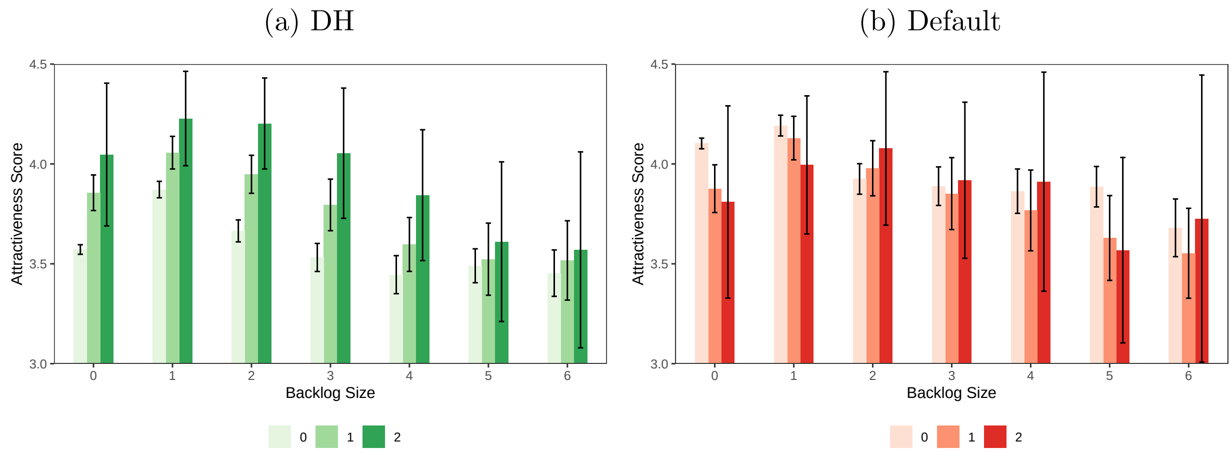

Note: The color scales represent the number of matches since the last session, with 0 matches being the lightest color and 2 being the darkest one.

D.2. Parallel-Trends

A critical check for the validity of our DID analysis is to satisfy the parallel-trends assumption. In words, there should be no significant differences in the trends of the markets before the start of the experiment. As discussed in the previous version of our manuscript, a common approach to testing the parallel-trends assumption (following (Autor 2003)) is to estimate the following model:

$$M_{mt} = \alpha_m + \lambda_t + \sum_{\tau=-p}^{q} W_i \cdot 1\{t = t_0 + \tau\}\, \delta_\tau + \epsilon_{mt}, \tag{9}$$

where t_0 is the period in which the treatment is implemented, p, q are the number of *leads* and *lags* to be included, respectively, and δ_τ is 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 checking whether $\delta_\tau = 0, \forall \tau < 0$. Alternatively, the test can be performed by estimating a reduced model where $\delta_\tau = 0, \forall \tau < 0$, and then checking whether there are any differences between the full and restricted models using an Anova test. In our case, we obtain that all these coefficients are not statistically different from zero, and our Anova test leads to p-values of 0.9095 and 0.8813 for one and two weeks prior to the experiment, respectively. Hence, we conclude that the parallel-trends assumption holds in both datasets.

An alternative approach to checking that the parallel-trends assumption holds is to estimate the pre-trend directly for each market. This can be achieved by estimating the following model:

$$M_{mt} = \alpha_m + \lambda_t + \rho_m \cdot P_t + \epsilon_{mt}, \tag{10}$$

where α_m is a market fixed effect, λ_t is a date fixed effect, and P_t represents how many days have elapsed since the beginning of the pre-treatment week (e.g., $P_t = 1$ for the first day in the pre-treatment week). Hence, ρ_m captures the slope of the pre-trend in each market m. Then, if there are no pre-trends, we would expect that $\rho_{AUS} = \rho_{HOU} = \rho_{DFW}$.

In Table 14 we show the results for both datasets. Column (1) reports the results for only one week before the start of each experiment, while column (2) reports the trends for two weeks before the start of the experiment.

The first two rows, Austin × Period and Dallas × Period, represent the coefficients ρ_{AUS} and ρ_{DFW}, respectively. Notice that both coefficients are not statistically significant. This result implies that there are no significant differences between these coefficients and ρ_{HOU}, which is the default for this specification. Hence, there are no significant differences in the trends before the experiments, confirming that the parallel-trends assumption holds. This result is confirmed by Figure 15, where we plot the daily number of matches in the week prior to the experiment along with the predicted pre-trends from the model.

D.3. Additional Results and Robustness Checks

D.3.1. Placebo Test. In Table 15 we report the results of a placebo test that removes data on Houston and labels Austin as the treated market. We observe that the variable of interest is not significant, which confirms that the effect we find is due to the intervention performed in Houston.

Table 14 First Field-Experiment: Pre-Trends Analysis (Equation (10))

	Dependent Variable: Number of Matches	
	(1)	(2)
Period	0.250	−0.466
	(1.665)	(0.773)
Period × Austin	3.286	1.470
	(2.355)	(1.093)
Period × Dallas	−1.393	−0.620
	(2.355)	(1.093)
Austin	89.821***	121.824***
	(18.619)	(6.581)
Dallas	171.571***	172.429***
	(18.619)	(6.581)
Houston	97.821***	107.209***
	(18.619)	(6.581)
Observations	21	42

Note: Columns (1) and (2) report the results for one week and two weeks of the pre-treatment period, respectively. Significance reported: *p<0.1; **p<0.05; ***p<0.01.

Figure 15 First Field-Experiment: Pre-Trends

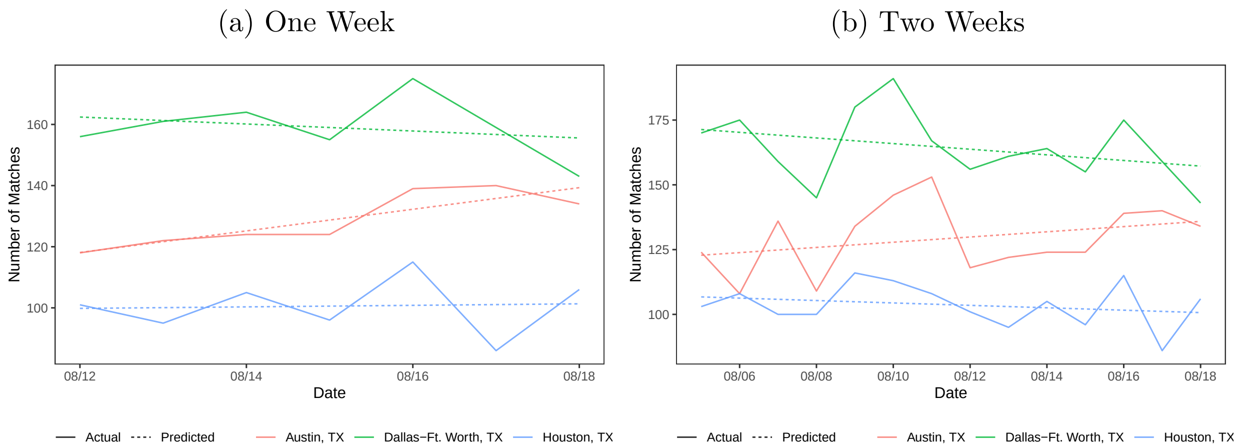

Table 15 First Field-Experiment: Placebo Test

	Dependent Variable: Number of Matches	
	(1)	(2)
Post	7.857	-
	(7.126)	-
Post×Treated	−17.143	−17.143
	(10.078)	(9.726)
Austin	128.714***	121.857***
	(5.039)	(9.726)
Dallas	159.000***	152.143***
	(5.039)	(9.726)
Observations	28	28

Note: Column (1) includes a dummy for the periods after the start of the experiment, while column (2) considers date fixed effects. Significance reported: *p<0.1; **p<0.05; ***p<0.01.

D.3.2. Robustness to Control Markets. As an additional robustness check, we estimate our difference-in-differences model by excluding one of the control markets at a time. The results are reported in Table 16. As expected, we observe that the estimates are relatively similar to those reported in Table 2, supporting our finding that the parallel-trends assumption holds.

Table 16 First Field-Experiment: Robustness Excluding Markets

	Dependent Variable: Number of Matches			
	Excluding Dallas		Excluding Austin	
	(1)	(2)	(3)	(4)
Post	-9.286^{*}	-	7.857	-
	(4.945)	-	(7.262)	-
Post×Treated	35.857^{***}	35.857^{***}	18.714^{*}	18.714^{**}
	(6.993)	(6.673)	(10.269)	(8.076)
Austin	128.714^{***}	123.571^{***}	-	-
	(3.496)	(6.673)	-	-
Dallas	-	-	159.000^{***}	157.714^{***}
	-	-	(5.135)	(8.076)
Houston	100.571^{***}	95.429^{***}	100.571^{***}	99.286^{***}
	(3.496)	(6.673)	(5.135)	(8.076)
Observations	28	28	28	28

Note: Columns (1) and (2) exclude data on Dallas, while columns (3) and (4) exclude data on Austin. Columns (1) and (3) include a dummy for the periods after the start of the experiment, while columns (2) and (4) consider date fixed effects. Significance reported: $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$.

D.3.3. Robustness to Days Considered. In Table 17 we report the results of estimating the difference-in-differences model for different subsets of data. More specifically, columns (1) and (2) report the results excluding data from the last day (to avoid end-of-horizon effects), while columns (3) and (4) report the results excluding the first two days (to avoid exploitation of existing backlog). These results are consistent with those reported in our main analysis.

Table 17 First Field-Experiment: Robustness Excluding Days

	Dependent Variable: Number of Matches			
	Excluding Last Day		Excluding First Two Days	
	(1)	(2)	(3)	(4)
Post	-3.024	-	1.143	-
	(6.398)	-	(5.402)	-
Post×Treated	29.786^{***}	29.786^{***}	25.886^{***}	25.886^{***}
	(6.398)	(7.677)	(5.402)	(6.471)
Austin	125.396^{***}	120.967^{***}	126.190^{***}	121.762^{***}
	(4.060)	(7.013)	(4.256)	(7.006)
Dallas	162.319^{***}	157.890^{***}	161.524^{***}	157.095^{***}
	(4.060)	(7.013)	(4.256)	(7.006)
Houston	100.571^{***}	96.143^{***}	100.571^{***}	96.143^{***}
	(4.629)	(7.300)	(4.783)	(7.262)
Observations	39	39	36	36

Note: Columns (1) and (2) exclude data from the last day of the experiment, while columns (3) and (4) exclude data from the first two days of the experiment. Columns (1) and (3) include a dummy for the periods after the start of the experiment, while columns (2) and (4) consider date fixed effects. Significance reported: $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$.

D.4. Robustness Relative to Other Markets

In Table 18 we report summary statistics for the weekly number of matches generated in all markets with more than 400 matches per week on average. The table also includes the number of matches generated in weeks 33 and 34, i.e., *before* and *after* the first experiment. Finally, we also include the percentage of variation in the week of the intervention relative to the week before, and we sort regions in decreasing order. We observe that the improvement in Houston is 26.4%, the highest among all the markets considered. Moreover, we observe that the second-largest variation is in Seattle, but the change is only 10.2%, significantly less than in Houston. These results support the fact that the improvement observed in Houston in the week of the intervention is not driven by noise but rather is a consequence of our algorithm.

Table 18 Difference in matches by region

	Overall		Intervention		
Region	Mean	Std.	Before	After	Diff [%]
Houston	776.143	62.039	704	890	0.264
Seattle	637.640	95.232	666	734	0.102
Toronto	418.480	74.479	381	414	0.087
Philadelphia	1343.760	213.345	1357	1455	0.072
Dallas	1189.000	38.518	1113	1168	0.049
San Diego	803.680	134.531	891	933	0.047
Los Angeles	6507.240	816.911	6739	7023	0.042
London	2471.720	338.357	2483	2540	0.023
Washington	3017.040	366.897	3162	3229	0.021
New York	16381.240	1809.329	17229	17420	0.011
Denver	1312.640	203.630	1418	1433	0.011
Chicago	3111.280	430.827	3413	3397	-0.005
Boston	2795.440	369.321	2908	2884	-0.008
San Francisco	4130.600	502.646	4194	4154	-0.010
Miami	1671.480	258.330	1764	1747	-0.010
Atlanta	980.640	168.892	1124	1105	-0.017
Phoenix	413.200	71.169	425	399	-0.061
Austin	907.429	52.526	901	836	-0.072

Note: Overall summary statistics of weekly matches considering data from week 23 to week 47. Before represents the weekly number of matches obtained in week 33 (before intervention). After represents the weekly number of matches obtained in week 34 (during intervention). Diff is computed as (After-Before)/Before.

D.5. Second Field Experiment

We ran a second field experiment in collaboration with our industry partner. In this case, Austin was the treated market (i.e., we used our algorithm DH to select the assortments) and Houston and Dallas were the control markets. The experiment ran for seven days in the week of September 9 to September 15, 2020. To estimate log-in and like probabilities we used the data from our empirical analysis. This data source is the same one that we used in the first field experiment, which yielded the same estimated coefficients. We also used the same parameters for the DH heuristic.

In Table 19 we report the overall results of the second experiment. The label *Before* represents the week prior to the intervention (in this case, the week from September 2 to September 8, 2020), while the label *After* represents the week of the intervention (i.e., from September 9 to September 15, 2020).

Table 19 Second Field-Experiment: Overall Results

Market	Period	Queries	Likes	Matches
Austin	After	26837	9993	1256
Austin	Before	27394	11200	975
Dallas	After	38676	15349	1079
Dallas	Before	37494	15310	1216
Houston	After	27719	10940	701
Houston	Before	26979	10881	742

We observe that the number of likes obtained in the week of the intervention in Austin is smaller than in the week before. This is similar to what we observe in the first experiment, and the reason is that DH prioritizes maximizing the match rate by taking into account the like probabilities of both users, whereas the default algorithm places more weight on maximizing likes. In addition, we observe a significant increase in the number of matches generated in Austin, whereas the number of matches decreased in the other two markets. This change is illustrated in Figure 16, where we plot the daily number of matches generated in each market (left) and the number of users that were active on each day (right).

To assess whether the increase in the number of matches is significant and due to the implementation of DH, in Table 20 we report the results of estimating our difference-in-differences model (Panel 1) and also of performing the placebo test (Panel 2). As previously discussed, this placebo test involves removing the data

Figure 16 Second Field-Experiment: Number of Matches and Active Users

Note: Number of matches and number of users who logged in between September 2 and September 15, 2020. The vertical line marks the day when the experiment started in the treated market (Austin).

of the treated market (in this case Austin), and labeling one of the control markets as treated (in this case, Houston).

Table 20 Second Field-Experiment: Difference-in-Differences Results

Panel 1: Main Results

	Dependent Variable: Number of Matches	
	(1)	(2)
Post	-12.714^{**}	-
	(6.089)	-
Post×Treated	52.857^{***}	52.857^{***}
	(6.089)	(7.318)
Austin	139.286^{***}	131.952^{***}
	(5.821)	(9.835)
Dallas	170.286^{***}	162.952^{***}
	(5.042)	(9.416)
Houston	109.429^{***}	102.095^{***}
	(5.042)	(9.416)
Observations	42	42
R^2	0.436	0.637

Panel 2: Placebo Test

	Dependent Variable: Number of Matches	
	(1)	(2)
Post	-19.571^{**}	-
	(8.754)	-
Post×Treated	13.714	13.714
	(12.380)	(10.300)
Dallas	173.714^{***}	160.857^{***}
	(6.190)	(10.300)
Houston	106.000^{***}	93.143^{***}
	(6.190)	(10.300)
Observations	28	28
R^2	0.185	0.718

Note: Panel 1 reports the estimation results, while Panel 2 reports the results of the placebo test. In both panels, column (1) includes a dummy for the periods after the start of the experiment, while column (2) considers date fixed effects. Significance reported: *p<0.1; **p<0.05; ***p<0.01.

From Panel 1 we observe that the variable of interest is positive and significant, consistent with the first field experiment results. Moreover, we observe that the magnitude of the effect relative to the market fixed effect is higher than in the first experiment (37.95% vs. 27.13% in the first experiment). This is because, on average, 49.57% of the profiles shown during the time window of the experiment were chosen freely by DH (after the business constraints were satisfied), compared to 38.30% in the first experiment. Also, from Panel 2, we observe that the variable of interest is not significant in the placebo test, suggesting that the increase in the number of matches is due to the implementation of DH in Austin.

As for the first field-experiment, we perform a series of robustness checks to assess the strength of our results. First, we confirm that the parallel-trends assumption also holds in the second experiment by doing the same analysis as in the first experiment. Specifically, we obtain that the p-values for the Autor (2003) tests are 0.1229 and 0.7515 for one and two weeks before the experiment, respectively. Moreover, from Table 21 we observe that none of the pre-trends are significantly different from the pre-trends in Austin, supporting the validity of the parallel-trends assumption. A graphic representation of this is provided in Figure 17, where we plot the number of matches and also the estimated trends for each market.

Table 21 Second Field-Experiment: Pre-Trends Analysis

	Dependent Variable: Number of Matches	
	(1)	(2)
Period	1.036	0.712
	(3.188)	(1.173)
Period × Dallas	0.857	0.090
	(4.508)	(1.658)
Period × Houston	−0.143	2.064
	(4.508)	(1.658)
Austin	127.893***	133.945***
	(35.641)	(9.985)
Dallas	152.893***	166.912***
	(35.641)	(9.985)
Houston	96.179***	78.967***
	(35.641)	(9.985)
Observations	21	42
R^2	0.035	0.152

Note: Columns (1) and (2) report the results for one week and two weeks of the pre-treatment period, respectively. Significance reported: *p<0.1; **p<0.05; ***p<0.01.

Figure 17 Second Field-Experiment: Pre-Trends

(a) One Week (b) Two Weeks

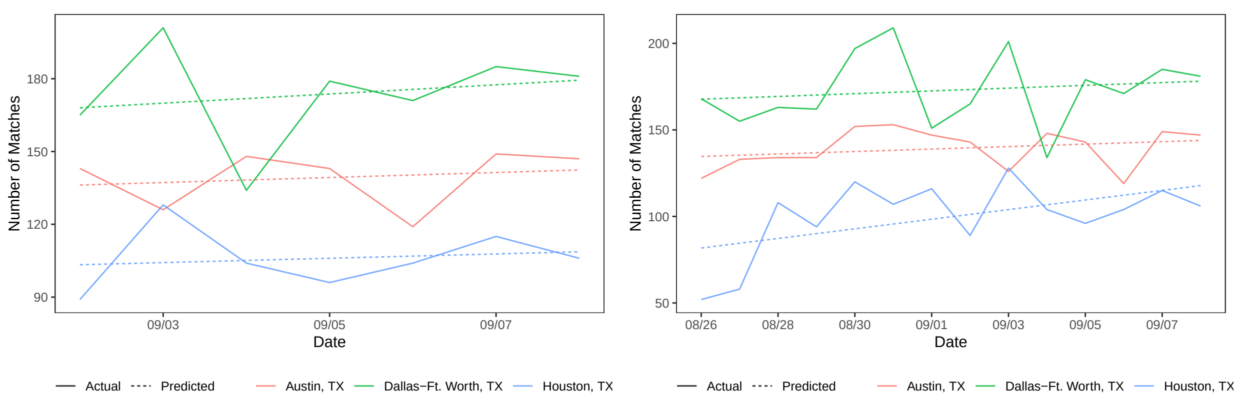

As an extra robustness check, we also estimated our results by excluding one control market at a time. As in the first experiment, Table 22 shows that our results are robust to the control markets considered. Similarly, from Table 23 we observe that the results are robust to excluding some days of the experiment (specifically, the last day or the first two).

Finally, Table 24 confirms that the extra matches generated by DH do not come from exploiting the backlog that existed before the start of the experiment. Indeed, 34.15% of the matches generated in the After period result from backlog queries generated in the same period, compared to 29.6% in the Before period. Hence, these results confirm that the improvement in DH is due to exploiting the previously generated backlog.

Table 22 Second Field-Experiment: Robustness Excluding Markets

	Dependent Variable: Number of Matches			
	Excluding Dallas		Excluding Houston	
	(1)	(2)	(3)	(4)
Post	−4.714	-	−19.571**	-
	(8.111)	-	(8.170)	-
Post×Treated	43.857***	43.857***	59.714***	59.714***
	(11.471)	(11.523)	(11.555)	(12.237)
Austin	139.286***	132.643***	139.286***	136.786***
	(5.374)	(11.121)	(5.777)	(12.237)
Dallas	-	-	173.714***	171.214***
	-	-	(5.777)	(12.237)
Houston	104.857***	94.000***	-	-
	(5.374)	(11.121)	-	-
Observations	28	28	28	28
R^2	0.496	0.746	0.555	0.750

Note: Columns (1) and (2) exclude data from Dallas, while columns (3) and (4) exclude data from Houston. Columns (1) and (3) include a dummy for the periods after the start of the experiment, while columns (2) and (4) consider date fixed effects. Significance reported: *p<0.1; **p<0.05; ***p<0.01.

Table 23 Second Field-Experiment: Robustness Excluding Days

	Dependent Variable: Number of Matches			
	Excluding Last Day		Excluding First Two Days	
	(1)	(2)	(3)	(4)
Post	−11.774	-	−10.957**	-
	(7.807)	-	(4.567)	-
Post×Treated	49.488***	49.488***	48.071***	48.071***
	(7.807)	(9.368)	(4.567)	(6.413)
Austin	139.286***	131.952***	139.286***	131.952***
	(5.911)	(9.828)	(6.044)	(9.940)
Dallas	169.665***	162.332***	171.149***	163.815***
	(5.184)	(9.443)	(5.378)	(9.589)
Houston	110.049***	102.716***	108.565***	101.232***
	(5.184)	(9.443)	(5.378)	(9.589)
Observations	39	39	36	36
R^2	0.398	0.625	0.373	0.617

Note: Columns (1) and (2) exclude data from the last day of the experiment, while columns (3) and (4) exclude data from the first two days of the experiment. Columns (1) and (3) include a dummy for the periods after the start of the experiment, while columns (2) and (4) consider date fixed effects. Significance reported: *p<0.1; **p<0.05; ***p<0.01.

Table 24 Second Field-Experiment: Backlog and Non-Backlog Queries

		Backlog			Non-Backlog					
Market	Period	Queries	Likes	Matches	Queries	Likes	LR	NB	Shown	Matches
Austin	After	5561	1179	1179	21276	8813	0.414	8813	1875	429
Austin	Before	3972	921	921	23422	10279	0.439	10279	901	289
Dallas	After	5192	1026	1026	33484	14323	0.428	14323	1761	443
Dallas	Before	5075	1158	1158	32419	14152	0.437	14152	1561	449
Houston	After	3119	661	661	24600	10279	0.418	10279	1176	321
Houston	Before	3131	694	694	23848	10187	0.427	10187	1002	282

Note: Number of backlog and non-backlog queries, together with the resulting number of likes and matches. Shown refers to non-backlog queries that were liked and ultimately shown to the other user as backlog queries within the same period (e.g., September 2-8 for Before). LR stands for like rate; NB stands for the new backlog generated within the corresponding period.

D.6. Additional Results

Table 25 Regression Results: Including and Excluding Houston

Panel 1: First Stage

	Dependent Variable: Num. of Matches			
	Including Houston		Excluding Houston	
	OLS	Neg. Bin.	OLS	Neg. Bin.
Treated	0.228***	0.555***	0.228***	0.553***
	(0.011)	(0.027)	(0.011)	(0.027)
Constant	−0.312	−2.687***	−0.308	−2.715***
	(0.424)	(0.940)	(0.410)	(0.888)
Observations	14,810	14,810	14,659	14,659
Pseudo-R^2	0.189	0.209	0.190	0.210

Panel 2: Second Stage

	Dependent Variable: Likes			
	Including Houston		Excluding Houston	
	OLS	Neg. Bin.	OLS	Neg. Bin.
$\hat{M}_i^{t_2}$	-0.423***	-0.222**	-0.408***	-0.244**
	(0.147)	(0.112)	(0.148)	(0.114)
Constant	-3.020***	-2.961***	-2.828**	-2.781**
	(1.148)	(1.141)	(1.125)	(1.120)
Observations	51,561	50,533	51,057	51,057
Pseudo-R^2	0.323	0.322	0.322	0.322

Note: The first stage is estimated using OLS and negative binomial regressions. The second stage for negative binomial regression addresses the forbidden regression problem. Standard errors are clustered at the user level. Significance reported: *p<0.1; **p<0.05; ***p<0.01.