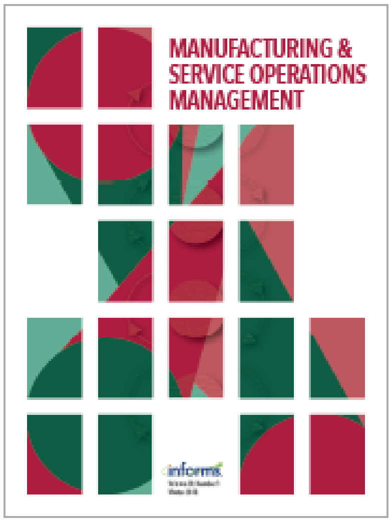

Manufacturing & Service Operations Management

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Improving Match Rates in Dating Markets Through Assortment Optimization

Ignacio Rios, Daniela Saban, Fanyin Zheng

To cite this article:
Ignacio Rios, Daniela Saban, Fanyin Zheng (2022) Improving Match Rates in Dating Markets Through Assortment Optimization. Manufacturing & Service Operations Management

Published online in Articles in Advance 07 Apr 2022

. https://doi.org/10.1287/msom.2022.1107

MANUFACTURING & SERVICE OPERATIONS MANAGEMENT
http://pubsonline.informs.org/journal/msom
Articles in Advance, pp. 1–20
ISSN 1523-4614 (print), ISSN 1526-5498 (online)

Honorable Mention–2021 *M&SOM* Practice-Based Research Competition

Improving Match Rates in Dating Markets Through Assortment Optimization

Ignacio Rios,[a,*] Daniela Saban,[b] Fanyin Zheng[c]

[a] Naveen Jindal School of Management, University of Texas at Dallas, Richardson, Texas 75080; [b] Graduate School of Business, Stanford University, Stanford, California 94305; [c] Columbia Business School, Columbia University, New York, New York 10027
*Corresponding author

Contact: ignacio.riosuribe@utdallas.edu, https://orcid.org/0000-0002-8526-9353 (IR); dsaban@stanford.edu, https://orcid.org/0000-0002-5217-4112 (DS); fz2225@gsb.columbia.edu (FZ)

Received: September 23, 2020
Revised: April 29, 2021; October 20, 2021; December 13, 2021
Accepted: February 13, 2022
Published Online in Articles in Advance: April 7, 2022

https://doi.org/10.1287/msom.2022.1107

Abstract. ***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 toward this end. Our proposed algorithm was piloted by our collaborator, a major online dating company in the United States. ***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.

History: This paper is an honorable mention in the 2021 *Manufacturing & Service Operations Management* Practice-Based Research Competition.
Supplemental Material: The online appendix is available at https://doi.org/10.1287/msom.2022.1107.

Keywords: display optimization • assortments • online platforms • matching • dating • behavioral operations

1. Introduction

In the past two decades, hundreds of dating services have emerged, 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 who met in the United States 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).

The size and relevance of the dating market highlight the need to make these platforms more efficient. To contribute toward this goal, in September 2018 we partnered with a major dating app to help them optimize the assortments to be shown to their users.[1] Our partner's platform offers a limited set of profiles (ranging from three to nine) 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.).

1.1. 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 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.

1.1.1. 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 matching markets.

1.1.2. 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 quasiexperiment 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 quasiexperiments in matching markets without interference, which is an exciting new area of research.

1.1.3. Proposed Algorithms. Based on our previous findings—namely, the estimated like probabilities, the fact that the platform's problem is computationally hard, and the fact 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%–45% relative to our partner's current algorithm. Roughly, 80% of the improvement comes from finding better matches (via (i) and (ii)), and the remaining 20% of the improvement is because of accounting for the history effect.

1.1.4. 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.

1.1.5. 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 for optimization of 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 1.2 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.

1.2. 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), and online selection of personalized assortments (Golrezaei et al. 2014, Berbeglia and Joret 2020). 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, although most of the assortment optimization literature focuses on settings where consumers are short lived and limited to one choice, users in our setting have repeated interactions with the platform and can like as many alternatives as they want from their daily assortment. This introduces several complications (e.g., the set of feasible assortments must be updated dynamically depending on users' past decisions). Moreover, because of 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 display optimization in dating platforms. In this context, Kanoria and Saban (2021) study how hiding quality information can considerably improve the platform's outcomes, and Halaburda et al. (2018) show that 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 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. 2010b), and that there exist strong assortative patterns (Hitsch et al. 2010a). Other papers empirically show the impact of design decisions and information on matching outcomes. Lee and Niederle (2014) show that the number of matches generated can increase by allowing users to signal their preferences, whereas 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 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 affects the way that choices are made. We contribute to these literatures by empirically showing that the context—through the history—can shape users' behavior. Moreover, our paper contributes to the nascent literature that analyzes how behavioral aspects can affect optimal assortment decisions (Ovchinnikov 2019).[2] 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.

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, Singh et al. (2019) and Cohen et al. (2021) in taxi or ride-sharing, and Martinez et al. (2021) 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, the user decides whether to like or not like each profile in the assortment.[3] A match between two users occurs if both users like each other. When a match is formed, both users are notified.[4]

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 who 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 consist of two parts.

1. User characteristics. For each user, we observe their profile information, including their age, height, location, education, religion, and 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 these data, we compute a set of usage metrics for the recent past, including the number of days active, the number of matches obtained, and 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, and the last month, among others. Finally, we also use these data to determine the *backlog* of each user on each day.

A unique feature of these data is that they allow 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),[5] $\mathcal{T} = \{1, \ldots, T\}$. 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, 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 who 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 v_i^t. We assume that these variables are independent across users and periods and that the parameters v_i^t can be estimated accurately 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 $\boldsymbol{S}_i^t \subseteq \mathscr{P}_i^t$ be the assortment selected by the platform to be offered to user i in period t. If user i logs in during period t, user i observes the assortment previously chosen by the platform, $\boldsymbol{S}_i^t$, and decides whether to like/not like each profile $j \in \boldsymbol{S}_i^t$. 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 the backlog by adding those users who 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 the like/not like decision for each user j in the assortment based on the (random) utility that user i gets 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 the last session before period t; in general, M_i^t can be defined using any measure of previous activities.[6] 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 the user may have less bandwidth to pursue new conversations or 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

Figure 1. Time Line of the Within-Period Dynamics of the Model

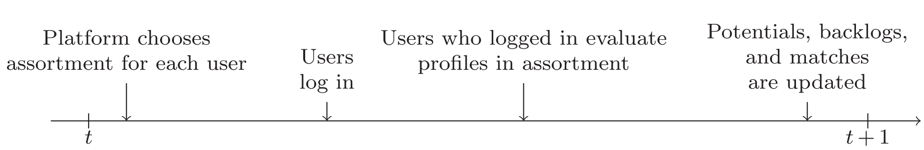

number of matches since the last session, and a random error. Hence, we write it as[7] $U_{ijt} = U(X_{ij}, M_i^t)$. Then, user i decides to like j in period t if and only if $U(X_{ij}, M_i^t) \geq u_{i0}$, where u_{i0} is user i's outside option. In Section 4, we provide an expression for $U(X_{ij}, M_i^t)$, which will be validated using our partner's data.

Let $\vec{\mathbf{\Phi}}_i^t = \{\mathbf{\Phi}_{ij}^t : j \in \mathcal{I}\}$ denote the vector of random variables representing whether user i liked each profile in period t (i.e., $\mathbf{\Phi}_{ij}^t = 1$ 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; that is, $j \in S_i^t$). Thus, we assume that

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

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 the last session: that is,

$$\phi_{ij}(M) = \mathrm{P}\Big(U_{ijt} \geq u_{i0} \,\Big|\, X_{ij}, j \in \mathbf{S}_i^t, M_i^t = M, \Upsilon_i^t = 1\Big). \quad (1)$$

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 $\mathbf{\Phi}_{ij}^t$ and $\mathbf{\Phi}_{ji}^{t'}$ are independent conditional on the vector of time-invariant characteristics X_{ij}, the assortments $\mathbf{S}_i^t, \mathbf{S}_j^{t'}$, 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 $\mathbf{\Phi}_{ji}^0 = 1$ 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:

$$\{\mathbf{\Phi}_{ij}^t = 1 \text{ and } \mathbf{\Phi}_{ji}^t = 1\} \text{ or } \cup_{\tau<t}\{\mathbf{\Phi}_{ij}^t = 1 \text{ and } \mathbf{\Phi}_{ji}^\tau = 1\} \text{ or } \cup_{\tau<t}\{\mathbf{\Phi}_{ij}^\tau = 1 \text{ and } \mathbf{\Phi}_{ji}^t = 1\}.$$

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 the last session can be expressed as $M_i^{t+1} = \sum_{j\in\mathcal{I}} \mu_{ij}^t + (1 - \Upsilon_i^t) \cdot M_i^t$.

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 $\left\{\vec{S}^{t,\pi}\right\}_{t=1}^T = \{\{S_i^{t,\pi}\}_{i\in\mathcal{I}}\}_{t=1}^T$ 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,

$$\sup_{\pi\in\Pi} \mathbb{E}\left[\sum_{t\in\mathcal{T}} \sum_{i\in\mathcal{I}} \sum_{j\in\mathcal{I}: j<i} \mu_{ij}^{t,\pi}\right], \quad (2)$$

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 this problem could be modeled as a Markov decision problem, where the state space is given by the set of potentials, the backlogs, and the number of matches since the last session for each user.[8] However, as described in Section 5, our focus is on designing 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 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 the 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.

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 quasiexperimental 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 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).[9]

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, it must be estimated from the data. In particular, we model U_{ijt} as

$$U_{ijt} = X'_{ij}\beta + M_i^t\gamma + \xi_{it} + \epsilon_{ijt}, \tag{3}$$

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 the last session, and X_{ij} encodes a set of time-invariant observable characteristics of users i and j that includes three groups of covariates). First, we include time-invariant characteristics of users i and j, including their height, race, etc. Second, we include measures of the absolute and relative attractiveness of users i and j, namely their attractiveness score (ratio of likes to evaluations received during their time on the platform) and their quintile of attractiveness compared with all users of the same gender. Lastly, we include the interactions between the characteristics of users i and j (e.g., user i might be looking for partners from the same religion or age group). Following Hitsch et al. (2010a), for each numerical variable x_k included in X_{ij}, we include the squared positive and negative differences, $|x_{jk} - x_{ik}|_+^2$ and $|x_{jk} - x_{ik}|_-^2$, 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. The term ξ_{it} are unobserved user time-specific characteristics, which can include how many people the user is currently dating, whether there is someone to whom the user feels especially connected, and 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 the online appendix. 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 the 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, user i might have become pickier when evaluating profiles on the platform in recent periods, which may have led to fewer matches since the 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 in period t. As a result, directly estimating Equation (3) 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 with 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. Thus, to generate matches for a user, one would need to manipulate the like decisions of other users, which is unfeasible to implement in practice. Moreover, 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 quasiexperiment 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 quasiexperiment to estimate the causal effect of an extra match on users' subsequent like decisions. Our analysis also provides an important example that shows that, although randomized experiments and quasiexperiments 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. Quasiexperiment. 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 backlog profiles in terms of the attractiveness score and uses this metric to decide which backlog profiles to show (in decreasing order of score). 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 because of 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 quasiexperiment 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 quasiexperimental design.

Next, we describe the quasiexperiment in detail. The time line of the quasiexperiment 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 pretreatment 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 posttreatment period where the outcome is measured. To avoid potential interference in the quasiexperiment, 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

Figure 2. Time Line of the Quasiexperiment

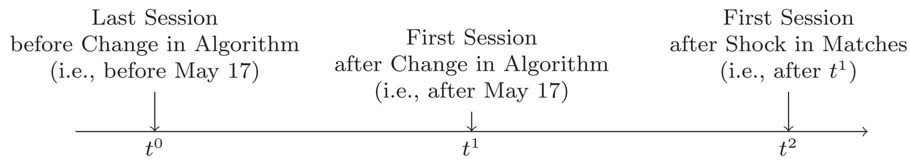

other users responding to the change in the focal users' like decisions in the quasiexperiment.[10] As t_i^1 is either May 17 or May 18 for most users, we add the constraints that t_i^0 is no sooner than May 14 and that t_i^2 is no later than May 21. As the variation in t_i^τ, $\tau \in \{0,1,2\}$ across users is small, we drop the subscript i for brevity, and we exclude from the analysis all users who 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 the 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 the 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.

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 because of the exogenous reasons listed. Thus, 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 (PSs) 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 with the control users, the treated users have a higher probability of obtaining matches in period t^1. Moreover, because 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 the online appendix, 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 pretreatment period t^0, we estimate the propensity that users are treated in the quasiexperiment. 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.[11] Next, we describe the details of the three-step estimation procedure.

4.1.3.1. Step 1. 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 pretreatment 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 a detailed description of the full set of variables in the online appendix. $M_i^{t^1}$ represents the number of matches obtained by user i since the last session in the pretreatment 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 the online appendix).

4.1.3.2. Step 2: First-Stage Regression of 2SLS. Following Hirano et al. (2003), we compute weights ω_i using the estimated propensity scores $\hat{e}_i$: that is,

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

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 the last session before t^2 for each user i (i.e., $\hat{M}_i^{t^2}$).

4.1.3.3. Step 3: Second-Stage Regression of 2SLS. Using the estimated weights in step 1, we estimate the following model: $\phi_{ij}^{t^2} = \mathrm{P}(\gamma\hat{M}_i^{t^2} + X_{ij}\beta_2 + Z_i^{t^1}\delta_2 + \varepsilon_{ij})$, 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 quasiexperiment consists of 8,398 control and 6,412 treated users. In the online appendix, 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 control groups after propensity score weighting. We find that there is no statistically significant difference between the two groups across all these statistics.

In Table 1, we report the first-stage results of 2SLS estimated using an ordinary least squares 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

Table 1. Quasiexperiment Estimation Results

	OLS	Negative binomial
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

Notes. The first stage is estimated using ordinary least squares (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 is reported with asterisks.
**$p < 0.05$; ***$p < 0.01$.

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 Table 1, we also 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).

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%–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 the online appendix, 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 in (2), 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., $|S_i^t| \le K_i^t$ 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 these constraints, we use the notation introduced in Section 3. We can write the last two constraints simply as $S_i^t \subseteq \mathcal{P}_i^t$, where we assume that the set of potentials $\mathcal{P}_i^t$ is updated every period by removing both the profiles that were shown to the user in the last period (if any) and also, all users who 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, can be used to obtain an upper bound for the platform's problem in (2) and that also plays a fundamental role in constructing our heuristics in

Section 5.3:

$$
\begin{aligned}
\tilde{\pi} := \max_{x,y,z} \quad & \sum_{t\in\mathcal{T}} \sum_{i\in\mathcal{I}} \sum_{j\in\mathscr{P}_i^1} v_i^t \cdot \phi_{ij}^0 \cdot y_{ij}^t \\
& + \frac{1}{2} \cdot v_i^t \cdot v_j^t \cdot \phi_{ij}^0 \cdot \phi_{ji}^0 \cdot z_{ij}^t \\
\text{s.t.} \quad & y_{ij}^t \le \mathbb{1}_{\{j\in\mathscr{B}_i^1\}} - \sum_{\tau=1}^{t-1} y_{ij}^\tau \cdot v_i^\tau \\
& + \sum_{\tau=1}^{t-1} \left(x_{ji}^\tau - v_i^\tau \cdot z_{ji}^\tau\right) \cdot v_j^\tau \cdot \phi_{ji}^0, \\
& \qquad \forall i \in \mathcal{I},\, j \in \mathscr{P}_i^1,\, t \in \mathcal{T} \\
& \sum_{t=1}^{T} \left(x_{ij}^t + y_{ij}^t\right) \cdot v_i^t \le 1, \\
& \qquad \forall i \in \mathcal{I},\, j \in \mathscr{P}_i^1,\, t \in \mathcal{T} \\
& \sum_{j\in\mathscr{P}_i^1} x_{ij}^t + y_{ij}^t \le K_i^t, \\
& \qquad \forall i \in \mathcal{I},\, j \in \mathscr{P}_i^1,\, t \in \mathcal{T} \\
& z_{ij}^t \le x_{ij}^t, \quad z_{ij}^t \le x_{ji}^t, \quad z_{ij}^t = z_{ji}^t, \\
& x_{ij}^t, y_{ij}^t, z_{ij}^t \in [0,1], \\
& \qquad \forall i \in \mathcal{I},\, j \in \mathscr{P}_i^1,\, t \in \mathcal{T}. \qquad (4)
\end{aligned}
$$

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 nonbacklog 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 $\sum_{\tau=1}^{t-1} y_{ij}^\tau \cdot v_i^\tau$ denotes the probability that i saw j as part of a backlog query in the past, and the term $\sum_{\tau=1}^{t-1} (x_{ji}^\tau - v_i^\tau z_{ji}^\tau) \cdot v_j^\tau \cdot \phi_{ji}^0$ 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 (4) 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 $\{v_i^t\}_{i\in\mathcal{I},t\in[\mathcal{T}]}$, 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 (2), and let $\tilde{\pi}$ be an optimal solution to (4). Then, $\pi^* \le \tilde{\pi}$.*

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, because of 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 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.

5.3.1. Optimization. The first step is to solve an optimization problem similar to that in (4), 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 $\{v_i^t\}_{i\in\mathcal{I}}$, and the like probabilities $\{\phi_{ij}^t\}_{i\in\mathcal{I},j\in\mathscr{P}_i^t}$, where $\phi_{ij}^t = \phi_{ij}(M_i^t)$. 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. Although 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 predefined 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 (5) in Algorithm 1.

5.3.2. Rounding. After solving the optimization problem described, 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. After 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.[12]

Algorithm 1 (DH)

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

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

Step 1. Optimization. Solve

$$\begin{aligned} x^*, y^*, z^* &= \arg\max_{x,y,z} \sum_{\tau=t}^{t+1} \sum_{i\in\mathcal{I}} \sum_{j\in\mathscr{P}_i^t} v_i^\tau \cdot \phi_{ij}^t \cdot y_{ij}^\tau + \frac{1}{2} \cdot v_i^\tau \cdot v_j^\tau \cdot \phi_{ij}^t \cdot \phi_{ji}^t \cdot z_{ij}^\tau \\ &\quad + \bar{\xi} \cdot \left(\Psi^t\left(\vec{x}^{\,t}, \vec{y}^{\,t}, \vec{z}^{\,t}, \vec{M}^{\,t}\right) \right. \\ &\quad \left. + \Psi^{t+1}\left(\vec{x}^{\,t}, \vec{y}^{\,t}, \vec{z}^{\,t}, \vec{x}^{\,t+1}, \vec{y}^{\,t+1}, \vec{z}^{\,t+1}, \vec{M}^{\,t}\right)\right) \\ \text{s.t.} &\quad \sum_{j\in\mathscr{P}_i^t} x_{ij}^\tau \cdot a_{ijl}^{\tau,x} + y_{ij}^\tau \cdot a_{ijl}^{\tau,y} \geq b_{il}^\tau, \\ &\quad \forall i \in \mathcal{I},\ \tau \in \{t, t+1\},\ l \in \mathcal{L}, \\ &\quad (x, y, z) \in \mathcal{S}^\tau. \end{aligned} \tag{5}$$

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

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.

5.3.3. 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 $M_i^{\tau+1}$,

$$\phi_{ij}(M_i^{\tau+1}) - \phi_{ij}(M_i^\tau) \approx (M_i^{\tau+1} - M_i^\tau) \cdot \gamma_{ij}(M_i^\tau), \tag{6}$$

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

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

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

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

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. 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

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

Using a similar reasoning, we define

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

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.

6.1. Data and Simulation Setting

We use a data set similar to that described in Section 4, which includes all heterosexual users in Houston, Texas who observed at least 100 profiles from their potentials between September 1, 2019 and April 1, 2020 and who logged in at least once between March 1 and April 1, 2020. For each of these users, we assume that their initial set of potentials is composed 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; 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 and 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. (We cannot disclose what these business constraints are because of the terms in our Non-Disclosure Agreement (NDA).)

Each simulation can be summarized as follows. In each period, we start by choosing the assortments that will be shown to each user who logs in. Then, the subset of users who are active in that period is realized, and each of these users makes like/not like decisions about the profiles shown in their assortment. Based on these decisions, we compute the number of matches generated, and we also update the sets of potentials (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.

6.2. Benchmarks

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

1. Partner. Implementation of our partner's current algorithm.

2. Naïve. 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 nor a penalty.

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 five). These values are the ones leading to the maximum number of matches for these policies (see the online appendix for further details).

6.3. Results

Our main simulation results are summarized in Figure 3. Figure 3(a) reports the average number of matches generated over 100 simulations by each policy, whereas Figure 3(b) reports box plots 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.[13]

To identify the sources of improvement, we aim to quantify how much is because of (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 3(c), we plot the average number of matches obtained by Greedy and DH from 100 simulations (the 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 $i \in \mathcal{I}$

Figure 3. (Color online) Simulation Results

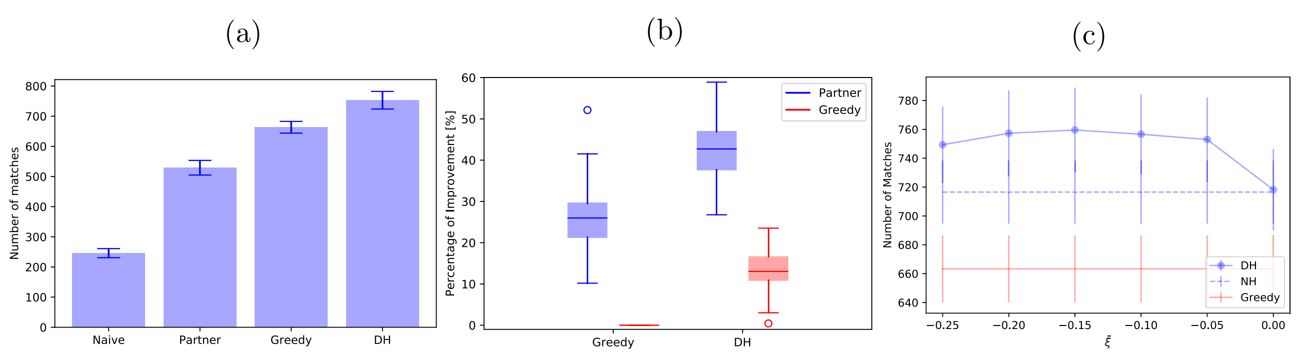

Notes. (a) Number of matches. (b) Relative improvement. (c) Sources of improvement.

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 because of better matches from the improvement because of including the history effect.

First, we observe that our 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).[14] 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 that the remaining 20% of the improvement is because of accounting for the history effect.

Finally, an additional set of simulations reported in the online appendix shows that the improvement obtained by DH is also meaningful and stable over a longer time horizon. We 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.

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.[15] 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, whereas the other two markets were assigned to the control group. We chose these three markets because of their geographic proximity and their similarity in the distribution of the main variables of interest. In the online appendix, we show that there are no significant differences in the main demographics across the three markets. In the interest of space, in the next sections we focus on the first field experiment, and we report the results of the second field experiment in the online appendix.

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, whereas 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 the online appendix.[16] As the estimates are very similar across models, we believe that this choice does not affect the results of the experiment. 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 lines mark 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. (Color online) First Field Experiment: Number of Matches and Active Users

Notes. The number of matches and the number of users who logged in between August 12 and August 25, 2020 are shown. The vertical lines mark the day when the experiment started in the treated market (Houston). (a) Number of matches. (b) Number of users.

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.

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: (1) $M_{mt} = \alpha_m + Z_t \cdot \gamma + W_m \cdot Z_t \cdot \delta + \epsilon_{mt}$ and (2) $M_{mt} = \alpha_m + \lambda_t + W_m \cdot Z_t \cdot \delta + \epsilon_{mt}$, 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.[17]

In Table 3, we report the estimation results. The first column provides the results of the first model, whereas 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

Table 2. First Field Experiment: Summary of Results

Market	Period	Queries	Likes	Matches
Austin	After	27,633	11,243	836
Austin	Before	27,580	11,440	901
Dallas	After	38,234	14,796	1,168
Dallas	Before	38,931	15,067	1,113
Houston	After	27,337	10,091	890
Houston	Before	27,796	11,083	704

Note. The table 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.

Table 3. First Field Experiment: Difference-in-Differences Results

	(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

Notes. The table reports the estimation results. Column (1) includes a dummy for the periods after the start of the experiment, whereas column (2) considers date fixed effects. Both columns include market fixed effects. Significance is reported with asterisks.
$^{***}p < 0.01$.

Table 4. First Field Experiment: Backlog and Nonbacklog Queries

		Backlog			Nonbacklog					
Market	Period	Queries	Likes	Matches	Queries	Likes	LR	NB	Shown	Matches
Austin	After	3,787	810	810	23,846	10,433	0.438	10,433	922	269
Austin	Before	3,712	869	869	23,868	10,571	0.443	10,571	923	297
Dallas	After	4,828	1,115	1,115	33,406	13,681	0.410	13,681	1,477	452
Dallas	Before	4,936	1,063	1,063	33,995	14,004	0.412	14,004	1,646	439
Houston	After	3,942	839	839	23,395	9,252	0.395	9,252	1,946	436
Houston	Before	2,929	682	682	24,867	10,401	0.418	10,401	954	267

Notes. The table shows the number of backlog and nonbacklog 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.

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 the online appendix 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 the online appendix). 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 the online appendix).

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 4, 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 4 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 "depleting" the existing backlogs 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 4 we provide a picture of the nonbacklog queries. We focus our discussion on Houston. The total numbers of nonbacklog 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 shown 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 toward maximizing likes. In total, of the 23,395 nonbacklog queries during the After period, 9,252 were liked, resulting in 9,252 new additions to the backlog compared with 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 shows that our algorithm is obtaining matches not by depleting the backlog previously generated but rather, by exploiting the backlog generated by itself.

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 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. We propose a family of algorithms to optimize the assortments offered by the platform that leverage this finding. We show through simulations that our algorithm considerably outperform relevant benchmarks; the results of two field experiments confirm that the improvements translate into our partner's platform.

Overall, our results showcase the importance of accounting for the effect that algorithmic decisions have on user input and behavior 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.

Acknowledgments
The authors are especially grateful to their industry partner and its chief of data science for providing constant support and guidance throughout this project. Also, the authors thank the area editor, the associate editor, and the referees for their constructive feedback.

Endnotes

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

[2] This is in contrast to other classic problems in operations management, including auction design (Elmaghraby and Katok 2019), procurement (Engelbrecht-Wiggans and Katok 2006, Beer et al. 2021), and pricing (Katok et al. 2014, Özer and Zheng 2016). See part III in Donohue et al. (2019) for an overview of applications of behavioral operations.

[3] 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.

[4] 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.

[5] 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 log-in 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.

[6] 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.

[7] One may also conjecture that the utility user i gets from being matched with user j can depend on the other users who 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.

[8] 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 (2)? Using a reduction from the exact cover problem, we establish in the online appendix that the aforementioned problem is NP complete.

[9] 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.

[10] 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.

[11] 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 because of 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.

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

[13] In the online appendix, we show that these differences remain for different values of the history effect and that similar results are obtained in other markets.

[14] 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 $\phi_{ij}(0)$.

[15] 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.

[16] The results in the online appendix 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.

[17] To validate the DID approach, in the online appendix we show that there are no significant differences in the trends before the intervention, proving that the parallel-trends assumption holds in our setting.

References

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

Beer R, Rios I, Saban D (2021) Increased transparency in procurement: The role of peer effects. *Management Sci.* 67(12): 7291–7950.

Berbeglia G, Joret G (2020) Assortment optimisation under a general discrete choice model: A tight analysis of revenue-ordered assortments. *Algorithmica* 82(4):681–720.

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

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

Cohen MC, Fiszer MD, Kim BJ (2021) Frustration-based promotions: Field experiments in ride-sharing. *Management Sci.*, ePub ahead of print August 31, https://doi.org/10.1287/mnsc.2021.3991.

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

Davis JM, Gallego G, Topaloglu H (2014) Assortment optimization under variants of the nested logit model. *Oper. Res.* 62(2):250–273.

Donohue K, Katok E, Leider S, eds. (2019) *The Handbook of Behavioral Operations* (Wiley, Hoboken, NJ).

Elmaghraby WJ, Katok E (2019) Behavioral research in competitive bidding and auction design. Donohue K, Katok E, Leider S, eds. *The Handbook of Behavioral Operations* (Wiley, Hoboken, NJ), 525–553.

Engelbrecht-Wiggans R, Katok E (2006) E-sourcing in procurement: Theory and behavior in reverse auctions with noncompetitive contracts. *Management Sci.* 52(4):581–596.

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

Fisman R, Iyengar SS, Kamenica E, Simonson I (2008) Racial preferences in dating. *Rev. Econom. Stud.* 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 Oper. Management* 20(4):767–787.

Golrezaei N, Nazerzadeh H, Rusmevichientong P (2014) Real-time optimization of personalized assortments. *Management Sci.* 60(6): 1532–1551.

Halaburda H, Piskorski MJ, Yildirim P (2018) Competing by restricting choice: The case of search platforms. *Management Sci.* 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 (2010a) Matching and sorting in online dating. *Amer. Econom. Rev.* 100(1):130–163.

Hitsch GJ, Hortaçsu A, Ariely D (2010b) What makes you click? Mate preferences in online dating. *Quant. Marketing Econom.* 8(4):393–427.

Kanoria Y, Saban D (2021) Facilitating the search for partners on matching platforms: Restricting agents' actions. *Management Sci.* 67(10):5969–6627.

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

Kök AG, Fisher ML, Vaidyanathan R (2015) Assortment planning: Review of literature and industry practice. Agrawal N, Smith SA, eds. *Retail Supply Chain Management: Quantitative Models and Empirical Studies* (Springer, Boston), 175–236.

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

Lin M (2018) Online dating industry: The business of love. Accessed November 11, 2019, https://www.toptal.com/finance/business-model-consultants/online-dating-industry.

Martinez M, Neilson C, Larroucau T, Rios I (2021) Application mistakes and information frictions in college admissions. Working paper, Department of Economics, University of Pennsylvania, Philadelphia.

McAlister L (1982) A dynamic attribute satiation model of variety-seeking behavior. *J. Consumer Res.* 9(2):141–150.

Ovchinnikov A (2019) Incorporating customer behavior into operational decisions. Donohue K, Katok E, Leider S, eds. *The Handbook of Behavioral Operations* (Wiley, Hoboken, NJ), 587–617.

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

Rosenfeld MJ, Thomas RJ, Hausen S (2019) Disintermediating your friends: How online dating in the United States displaces other ways of meeting. *Proc. Natl. Acad. Sci. USA* 116(36): 17753–17758.

Rusmevichientong P, Shen ZJM, Shmoys DB (2010) Dynamic assortment optimization with a multinomial logit choice model and capacity constraint. *Oper. Res.* 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 Oper. Management* 23(11):2023–2039.

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

Talluri K, van Ryzin G (2004) Revenue management under a general discrete choice model of consumer behavior. *Management Sci.* 50(1):15–33.

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

Wang R (2018) When prospect theory meets consumer choice models: Assortment and pricing management with reference prices. *Manufacturing Service Oper. Management* 20(3):583–600.

Wang R, Sahin O (2018) The impact of consumer search cost on assortment planning and pricing. *Management Sci.* 64(8):3649–3666.

Yu J (2018) Search, selectivity, and market thickness in two-sided markets. Working paper, Stanford Graduate School of Business, Stanford, CA.