Price of Anarchy in Algorithmic Matching of Romantic Partners

2

ANDRÉS ABELIUK, Department of Computer Science, University of Chile, and National Center for Artificial Intelligence (CENIA), Chile
KHALED ELBASSIONI, Khalifa University of Science & Technology, UAE
TALAL RAHWAN, Computer Science, New York University, UAE
MANUEL CEBRIAN, Department of Statistics, Universidad Carlos III de Madrid, Spain
IYAD RAHWAN, Center for Humans & Machines, Max-Planck Institute for Human Development, Germany

Algorithmic matching is a pervasive mechanism in our social lives and is becoming a major medium through which people find romantic partners and potential spouses. However, romantic matching markets pose a principal-agent problem with the potential for moral hazard. The agent's (or system's) interest is to maximize the use of the matching website, while the principal's (or user's) interest is to find the best possible match. This creates a conflict of interest: the optimal matching of users may not be aligned with the platform's goal of maximizing engagement, as it could lead to long-term relationships and fewer users using the site over time. Here, we borrow the notion of price of anarchy from game theory to quantify the decrease in social efficiency of online algorithmic matching sites where engagement is in tension with user utility. We derive theoretical bounds on the price of anarchy and show that it can be bounded by a constant that does not depend on the number of users in the system. This suggests that as online matching sites grow, their potential benefits scale up without sacrificing social efficiency. Further, we conducted experiments with human subjects in a matching market and compared the social welfare achieved by an optimal matching service against a self-interested matching algorithm. We show that introducing competition among matching sites aligns the self-interested behavior of platform designers with their users and increases social efficiency.

CCS Concepts: • **Theory of computation** → **Quality of equilibria**; **Online algorithms;**

Additional Key Words and Phrases: Online matching markets, online dating, price of anarchy

ACM Reference format:
Andrés Abeliuk, Khaled Elbassioni, Talal Rahwan, Manuel Cebrian, and Iyad Rahwan. 2024. Price of Anarchy in Algorithmic Matching of Romantic Partners. *ACM Trans. Econ. Comput.* 12, 1, Article 2 (March 2024), 25 pages.
https://doi.org/10.1145/3627985

A. Abeliuk and K. Elbassioni are joint first authors who contributed equally to this work.
This work was partially supported by National Center for Artificial Intelligence CENIA FB210017, Basal ANID.
Authors' addresses: A. Abeliuk, Department of Computer Science, University of Chile, and National Center for Artificial Intelligence (CENIA), Santiago, Chile; e-mail: aabeliuk@dcc.uchile.cl; K. Elbassioni, Khalifa University of Science & Technology, Abu Dhabi, UAE; e-mail: khaled.elbassioni@ku.ac.ae; T. Rahwan, Computer Science, New York University, Abu Dhabi, UAE; e-mail: talal.rahwan@nyu.edu; M. Cebrian, Department of Statistics, Universidad Carlos III de Madrid, Madrid, Spain; e-mail: manuel.cebrian@uc3m.es; I. Rahwan (Corresponding author), Center for Humans & Machines, Max-Planck Institute for Human Development, Berlin, Germany; e-mail: rahwan@mpib-berlin.mpg.de.

2167-8375/2024/03-ART2
https://doi.org/10.1145/3627985

1 INTRODUCTION

With the rise of online dating over the past decade, the pool of potential partners has grown from a few thousand people you might meet in your daily life to millions of people. Space and time constraints to meet a partner through your existing social circles are no longer a limitation with online dating sites; you can potentially be matched with any user in any part of the world. Online dating has enormous potential to ameliorate what for many people is a time-consuming and often frustrating activity. "Online dating is pervasive, and has fundamentally altered both the romantic acquaintance process and the process of compatibility matching" [9]. Indeed, 11% of American adults and 38% of those currently single and looking for a partner have used online dating sites or mobile dating apps [30]. In short, online dating is changing the way we find romantic partners [27].

The vast amount of people using dating services poses a challenging problem for service providers. There are millions of potential matches—too many for a single user to search exhaustively—so the system must decide on which potential matches to display to the user in order to present a manageable subset. This means that online dating sites must use an algorithm to curate what is presented to users, based on the large amount of data the system collects from them. For example, *OkCupid* claims[1] that their algorithm matches people using "match percentages," which basically quantifies how much users have in common. However, none of these algorithms are fully disclosed.

In the early 1970s, Herbert Simon presaged that wealth of information would create attention scarcity and thus become a new kind of currency [28]. Indeed, we see technology companies competing for people's attention, clicks, and loyalty. This notion, coined the "attention economy" [8, 11], is more relevant than ever and is considered one of the most important determinants of business success. In e-commerce, for example, more emphasis is often placed on users' future visits and purchases than on maximizing their long-term satisfaction [12, 21, 29]. The economic theory of "planned obsolescence" predicts that oligopolists will profit from producing goods with shorter durability to induce consumers to purchase more frequently [4]. There is evidence of planned obsolescence in the textbook market [17], where publishers introduce new editions to decrease the value of old ones. In online dating sites, an example of planned obsolescence could be that matches with shorter expected "durability" are deliberately suggested, and better matches are subsequently introduced to reduce the "value" of the old matches, all in the hope of increasing the rate at which users return to the site. This raises the following question:

How would the social efficiency of online matching sites be affected if system designers maximized the likelihood of users returning to the site?

Put differently, since the business model of most online dating services is heavily dependent on the number of users they have [26], this creates a fundamental conflict of interest: *optimally matching users may lead to long-term couples and fewer singles using the site, which is detrimental to the business.* Thus, a self-interested dating service provider maximizing its revenue will have as one of its goals maximizing user *engagement*, which may not always align with the goal of maximizing user *utility* [19]. As such, while dating sites offer users access to an unprecedented number of potential mates, they also pose a principal-agent problem with a potential moral hazard (e.g., [1]) because the agent (the system designer) has access to more information than the principal (the user) and different goals.

Driven by these observations, this article aims to quantify the social welfare loss that users might experience from a dating site that maximizes engagement instead of user utility. In doing so, we take a first step toward quantifying the maximum social welfare loss that could result from

[1]https://theblog.okcupid.com

online dating sites acting selfishly. In particular, we model the centralized matching as a classical weighted matching on bipartite graphs [22]—with the novel difference that self-interested behavior is represented by a different objective function than the classical objective of maximizing the sum of weights. Namely, the goal of dating sites is to maximize the number of users remaining in the system according to a Markov decision process. To our knowledge, this is the first attempt to model the detrimental effects of selfish (centralized) matching.

2 RESULTS

2.1 Algorithmic Matching Model

Matching in dating sites is an inherently online problem since there is not complete information about the arriving users in the system in advance, but it is obtained incrementally. In contrast, we first consider an offline version of the matching problem, where all users are known in advance and are in the system from the beginning. The online matching problem is analyzed in Section 2.4.

From the perspective of a social designer, the problem of finding partners for everyone to maximize social welfare is modeled as a weighted matching on bipartite graphs, where weights represent how well two users are matched. We assume that users' preferences are known.

Formally, let $\mathcal{M} := [m]$ and $\mathcal{W} := [n]$ be two finite sets of m "men" and n "women," respectively. Each user $i \in \mathcal{M} \cup \mathcal{W}$ is characterized by a vector of d dimensions: $\chi(i) \in \mathbb{R}_+^d$ if $i \in \mathcal{M}$ and $\kappa(i) \in \mathbb{R}_+^d$ if $i \in \mathcal{W}$. This vector representation is very general and characterizes each user through a set of d relevant features, such as age, personality traits, film preferences, and so forth. Let $w_{ij}^{\mathcal{M}} := w^{\mathcal{M}}(\chi(i), \kappa(j)) \in [0,1]$ be a measure of how well user i matches user j, and let $w_{ij}^{\mathcal{W}} := w^{\mathcal{W}}(\chi(i), \kappa(j)) \in [0,1]$ be a measure of how well user j matches user i. Note that users' preferences are not symmetric.

For the offline model, assume that all elements in $\mathcal{W}$ and $\mathcal{M}$ are known to the system. The system decides the (partial) distribution $\mathbf{x} := (x_{ij} : i \in \mathcal{M}, j \in \mathcal{W})$, where x_{ij} is the probability that an element $i \in \mathcal{M}$ is assigned to $j \in \mathcal{W}$. Define the *satisfaction* or *utility* of $i \in \mathcal{M}$ as $u_i := \sum_{j \in \mathcal{W}} w_{ij}^{\mathcal{M}} x_{ij} \in [0,1]$ and of $j \in \mathcal{W}$ as $u_j := \sum_{i \in \mathcal{M}} w_{ij}^{\mathcal{W}} x_{ij} \in [0,1]$.

A *social welfare maximizing* system (from the users' point of view) would assign a distribution $\mathbf{x}$ to maximize the utility of all users: $\sum_{i \in \mathcal{M}} u_i + \sum_{j \in \mathcal{W}} u_j$. That is, it attempts to solve the following optimization problem:

$$
\begin{aligned}
z_{sw}^*(\mathbf{w}) := \max_{\mathbf{x}} \ & \textstyle\sum_{i \in \mathcal{M}} u_i + \sum_{j \in \mathcal{W}} u_j && (O_{sw}(\mathbf{w}))\\
\text{s.t.} \quad & u_i := \textstyle\sum_j w_{ij}^{\mathcal{M}} x_{ij}, \ \text{ for all } i \in \mathcal{M} && (1)\\
& u_j := \textstyle\sum_i w_{ij}^{\mathcal{W}} x_{ij}, \ \text{ for all } j \in \mathcal{W} && (2)\\
& \textstyle\sum_{j \in \mathcal{W}} x_{ij} \le 1, \ \text{ for all } i \in \mathcal{M}, && (3)\\
& \textstyle\sum_{i \in \mathcal{M}} x_{ij} \le 1, \ \text{ for all } j \in \mathcal{W}, && (4)\\
& x_{ij} \ge 0, \ \text{ for all } i \in \mathcal{M} \text{ and } j \in \mathcal{W}. && (5)
\end{aligned}
$$

Since we are working on bipartite graphs, it is well known that this linear program will yield integer solutions when the constraint matrix is totally unimodular [15]. In other words, there always exists an optimal integer solution to the social welfare maximizing matching problem where every element $i \in \mathcal{W}$ is matched with only one element $j \in \mathcal{M}$ or is not matched at all.

Modeling Self-interested Behavior. Next, we introduce the concept of "user engagement" into the matching problem. The objective function of a self-interested designer of a matching service is to maximize this value. User engagement refers to the extent to which a user actively uses and participates in the matching system.

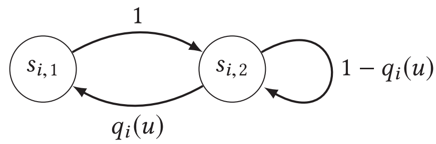

Fig. 1. Markov process with two states $s_{i,1}$ and $s_{i,2}$, representing the states of being in and out of the system, for user i having received utility u.

We use a *Markov decision process* [24] to model the user engagement with the system (i.e., matching service). The model has two states: $s_{i,1}$ and $s_{i,2}$, representing the states of being in and out of the system for user i, respectively. The actions of the system are to determine what utility to give to the user. The Markov chain is depicted in Figure 1. At state $s_{i,1}$ there is an infinite set of actions, each corresponding to a utility $u \in [0, 1]$. The user moves to state $s_{i,2}$ with probability 1, where he or she stays there with probability $1 - q_i(u)$ (either for a successful relationship or for a complete dissatisfaction with the matching service) or returns to the system (state $s_{i,1}$) with probability $q_i(u)$ after receiving utility u.

Let $q_i(u)$ be the probability that user $i \in \mathcal{M} \cup \mathcal{W}$ will return to the system after receiving utility u. We will make the natural assumptions[2] that $q_i(u)$ satisfies:

(A1) $q_i(0) = q_i(1) = 0$;
(A2) $q_i(u)$ is twice differentiable in $u \in [0, 1]$;
(A3) $q_i(u)$ is strictly concave in $u \in [0, 1]$.

The intuition behind these assumptions is that users who are dissatisfied with the service (i.e., receive a low utility) are less likely to return, and in the extreme case of receiving zero utility, the user will never return to the system. On the other hand, users who are paired with good matches, such as a long-term relationship in the dating context, have less incentive to search for a new partner. In the extreme case of receiving the maximum utility of one (i.e., a perfect match), the user will never return to the system. Figure 2 (left) shows a family of functions satisfying these assumptions.

The limiting probability of a Markov model represents the long-term behavior of the system. In the context of the online dating model, the limiting probability of the Markov model represents the long-term probability of a user being in or out of the matching service. For $k \in \{1, 2\}$, let $\pi_{i,k}(u)$ be the limiting probability that the user will be in state $s_{i,k}$, given that the decision made by the system yields the user utility u. Then we have

$$\pi_{i,1}(u) = \pi_{i,2}(u) q_i(u), \quad \pi_{i,1}(u) + \pi_{i,2}(u) = 1,$$

which gives

$$\pi_{i,1}(u) = \frac{q_i(u)}{1 + q_i(u)} \text{ and } \pi_{i,2}(u) = \frac{1}{1 + q_i(u)}. \tag{6}$$

Observation 1. *$\pi_{i,1}(u)$ satisfies (A1), (A2), and (A3) (with q_i replaced by $\pi_{i,1}$).*

Proof. (A1) and (A2) are immediate from Equation (6). To see that $\pi_{i,1}(u)$ is concave in $u \in [0, 1]$, note that its second derivative is given by

$$\pi''_{i,1}(u) = \frac{(1 + q_i(u)) q''_i(u) - 2(q'_i(u))^2}{(1 + q_i(u))^3},$$

which is strictly negative for $u \in [0, 1]$ due to the strict concavity assumption on $q_i(\cdot)$. □

[2] Some of these assumptions can be relaxed, but we will not do so here to keep the presentation simple.

The system would *selfishly* allocate users to maximize the *expected number of users returning to the system* $s(\mathbf{u};\mathbf{q}) := \sum_i \pi_{i,1}(u_i)$. That is, the selfish central designer attempts to solve the following optimization problem:

$$
\begin{aligned}
z_s^*(\mathbf{w};\mathbf{q}) := \max_{\mathbf{x}} \sum_{i\in\mathcal{M}} \pi_{i,1}(u_i) + \sum_{j\in\mathcal{W}} \pi_{j,1}(u_j) \qquad\qquad (O_s(\mathbf{w};\mathbf{q}))\\
\text{s.t.} \qquad & u_i := \textstyle\sum_j w_{ij}^{\mathcal{M}} x_{ij}, \quad \text{for all } i \in \mathcal{M}\\
& u_j := \textstyle\sum_i w_{ij}^{\mathcal{W}} x_{ij}, \quad \text{for all } j \in \mathcal{W}\\
& \textstyle\sum_{j\in\mathcal{W}} x_{ij} \le 1, \quad \text{for all } i \in \mathcal{M},\\
& \textstyle\sum_{i\in\mathcal{M}} x_{ij} \le 1, \quad \text{for all } j \in \mathcal{W},\\
& x_{ij} \ge 0, \quad \text{for all } i \in \mathcal{M} \text{ and } j \in \mathcal{W}.
\end{aligned}
$$

In the *offline* case, both problems ($O_{sw}(\mathbf{w})$) and ($O_s(\mathbf{w};\mathbf{q})$) are subject to the same set of constraints (Equations (1)–(5)) and can be solved optimally in polynomial time via linear programming and convex programming, respectively. Next, we quantify the loss of efficiency in the system caused by the selfish behavior of the central designer with the notion of **price of anarchy (PoA)** [20].

Definition 1 (Price of Anarchy). Let $\mathbf{u}_{sw}^*$ and $\mathbf{u}_s^*(\mathbf{q})$ be the optimal solutions of the problem for the social welfare maximizing system ($O_{sw}(\mathbf{w})$) and the selfish system ($O_s(\mathbf{w};\mathbf{q})$), respectively.[3]

Define the *price of anarchy* $\text{PoA} = \text{PoA}(\mathbf{q}) \in [0,1]$ as the ratio of the worst possible selfish matching to the social optimum. Formally,

$$
\text{PoA}(\mathbf{q}) := \min_{\mathbf{w}\ge\mathbf{0}} \frac{f(\mathbf{u}_s^*(\mathbf{q}))}{f(\mathbf{u}_{sw}^*)},
$$

where $f(\mathbf{u}) := \sum_i u_i$.

2.2 Bounds on the Price of Anarchy

In this section, we present our main theoretical result, which states that the price of anarchy can be bounded by a constant that depends only on the functions q_i. This result has two implications: (1) as online dating sites grow, their potential benefits scale without sacrificing social efficiency; (2) the loss of social utility is primarily caused by users' actions, and thus can be improved by modifying users' behavior. In other words, the system's ability to exploit users is limited by their own actions. This idea is further explored in the next section by modeling user behavior in a competitive market instead of in a monopoly. It should be noted that no platform can ever have a monopoly in the dating space, as there will always be alternative ways to meet in person.

Formally, we show that for any vector of probability functions $\mathbf{q} = (q_1, \dots, q_m)$ satisfying (A1) through (A3), the price of anarchy ($\text{PoA}(\mathbf{q})$) can be bounded by a constant that depends only on the functions q_i but not on the number of users m.

Theorem 1. *Let* $H(\mathbf{q}) := \max_{i\in\mathcal{M}\cup\mathcal{W}}\{q_i'(0)\} > 0$*, and* $c > 0$ *be the unique solution of the equation*

$$
c = \frac{H(\mathbf{q})}{2} \cdot L(\mathbf{q}, c), \tag{7}
$$

where $L(\mathbf{q}, c) = \min_{i\in\mathcal{M}\cup\mathcal{W}} \overline{u}_i(c)$*, and where for* $i \in \mathcal{M}\cup\mathcal{W}$*,* $\overline{u}_i(c) \in [0,1]$ *is the unique positive solution of the equation*

$$
c = \pi_{i,1}'(u). \tag{8}
$$

[3]Note that u_s^* is unique by the strict concavity of $s(\mathbf{u};\mathbf{q})$.

Then,

$$\text{PoA}(\mathbf{q}) \geq L(\mathbf{q}, c)/2.$$

Proof. First, note that the solutions in Equation (7) and (8) are well defined, since the strict concavity of $q_i(u)$ in $[0,1]$ implies that the function $\pi'_{i,1}(u) = \frac{q'_i(u)}{(1+q_i(u))^2}$ is strictly decreasing in $[0,1]$, which in turn implies that the function $L(\mathbf{q}, c)$ is also strictly decreasing in $c \in [0, H(\mathbf{q})]$. It follows that the right-hand side of Equation (7) is strictly decreasing in $c \in [0, H(\mathbf{q})]$ and takes a positive value when $c = 0$ and a value of 0 when $c = H(\mathbf{q})$, implying that the root c in Equation (7) exists.

Fix any optimal solutions $(\mathbf{u}^*, \mathbf{x}^*)$ and $(\widehat{\mathbf{u}}, \widehat{\mathbf{x}})$ for problems ($O_{sw}(\mathbf{w})$) and ($O_s(\mathbf{w}; \mathbf{q})$), respectively. Then our to goal is to lower-bound

$$\frac{\sum_{i \in \mathcal{M}} \widehat{u}_i + \sum_{j \in \mathcal{W}} \widehat{u}_j}{\sum_{i \in \mathcal{M}} u_i^* + \sum_{j \in \mathcal{W}} u_j^*}. \tag{9}$$

Fix $c > 0$ as given by Equation (7), and let $\mathcal{S}_{\mathcal{M}}(c) := \{i \in \mathcal{M} : \pi'_{i,1}(\widehat{u}_i) > c\}$ and $\mathcal{S}_{\mathcal{W}}(c) := \{j \in \mathcal{W} : \pi'_{j,1}(\widehat{u}_j) > c\}$. If one tries to bound the ratio directly by using the simple (local) lower bound $\min_i \frac{\widehat{u}_i}{u_i^*}$, one runs into the difficulty that some of the $\widehat{u}_i$s could be very close to zero (though not exactly zero, since the function $s(\mathbf{u}; \mathbf{q})$ is strictly concave). To overcome this, we use the local bound only for users not in $\mathcal{S}(c) := \mathcal{S}_{\mathcal{M}}(c) \cup \mathcal{S}_{\mathcal{W}}(c)$, while for users in $\mathcal{S}(c)$ we use a global bound derived from the **Karush-Kuhn-Tucker (KKT)** optimality conditions [33]. This is formalized in the following two claims.

Claim 1. *Suppose that $i \in \mathcal{M} \setminus \mathcal{S}_{\mathcal{M}}(c) \cup \mathcal{W} \setminus \mathcal{S}_{\mathcal{W}}(c)$; then $\widehat{u}_i \geq \frac{L(\mathbf{q},c)}{2}(u_i^* + \widehat{u}_i)$.*

Proof. Immediately from the definitions, since $\pi'_{i,1}(\widehat{u}_i) \leq c$, this implies, due to the concavity of q_i, that

$$\widehat{u}_i \geq \overline{u}_i(c) \geq L(\mathbf{q}, c) \geq L(\mathbf{q}, c)(u_i^* + \widehat{u}_i)/2.$$

The last inequality follows from the definition of the utilities, which are bounded by 1, i.e., $u_i \in [0,1]$. □

Claim 2.

$$\sum_{i \in \mathcal{M}} u_i^* + \sum_{j \in \mathcal{W}} u_j^* \leq \frac{1}{c}\left(\sum_{\substack{i \in \mathcal{S}_{\mathcal{M}}(c) \\ j \in \mathcal{W}}} \pi'_{i,1}(\widehat{u}_i) w_{ij} \widehat{x}_{ij} + \sum_{\substack{j \in \mathcal{S}_{\mathcal{W}}(c) \\ i \in \mathcal{M}}} \pi'_{j,1}(\widehat{u}_i) w_{ij} \widehat{x}_{ij} \right) + \sum_{i \in \mathcal{M} \setminus \mathcal{S}_{\mathcal{M}}(c)} (\widehat{u}_i + u_i^*) + \sum_{j \in \mathcal{W} \setminus \mathcal{S}_{\mathcal{W}}(c)} \left(\widehat{u}_j + u_j^*\right) \tag{10}$$

Proof. Consider problem ($O_s(\mathbf{w}; \mathbf{q})$). By the KKT (necessary) conditions for optimality (which are also sufficient since the functions $\pi_{i,1}(\cdot)$ are concave by Observation (1)), there exists $\widehat{\beta}_i \geq 0$, $\widehat{\sigma}_j \geq 0$, and $\widehat{\mu}_{ij} \geq 0$, for $i \in \mathcal{M}$ and $j \in \mathcal{W}$, such that Equations (1) through (5) hold for $(\mathbf{u}, \mathbf{x}) =$

$(\widehat{\mathbf{u}}, \widehat{\mathbf{x}})$, and

$$-\pi'_{i,1}(\widehat{u}_i)w_{ij}^{\mathcal{M}} - \pi'_{j,1}(\widehat{u}_j)w_{ij}^{\mathcal{W}} + \widehat{\beta}_i + \widehat{\sigma}_j = \widehat{\mu}_{ij}, \quad \text{for all } i \in \mathcal{M} \text{ and } j \in \mathcal{W}, \tag{11}$$

$$\widehat{\mu}_{ij}\widehat{x}_{ij} = 0, \quad \text{for all } i \in \mathcal{M} \text{ and } j \in \mathcal{W}, \tag{12}$$

$$\widehat{\beta}_i \left(\sum_{j \in \mathcal{W}} \widehat{x}_{ij} - 1 \right) = 0, \quad \text{for all } i \in \mathcal{M}, \tag{13}$$

$$\widehat{\sigma}_j \left(\sum_{i \in \mathcal{M}} \widehat{x}_{ij} - 1 \right) = 0, \quad \text{for all } j \in \mathcal{W}. \tag{14}$$

Note that Equations (11) through (14) imply that

$$\sum_{i \in \mathcal{M},\ j \in \mathcal{W}} \pi'_{i,1}(\widehat{u}_i)w_{ij}^{\mathcal{M}}\widehat{x}_{ij} + \sum_{j \in \mathcal{W},\ i \in \mathcal{M}} \pi'_{j,1}(\widehat{u}_j)w_{ij}^{\mathcal{W}}\widehat{x}_{ij} = \sum_{i \in \mathcal{M}} \widehat{\beta}_i + \sum_{j \in \mathcal{W}} \widehat{\sigma}_j, \tag{15}$$

$$\widehat{\beta}_i + \widehat{\sigma}_j \geq \pi'_{i,1}(\widehat{u}_i)w_{ij}^{\mathcal{M}} + \pi'_{j,1}(\widehat{u}_j)w_{ij}^{\mathcal{W}}, \quad \text{for all } i \in \mathcal{M} \text{ and } j \in \mathcal{W}. \tag{16}$$

Let us next define the following optimization problem:

$$z^*_{\widetilde{sw}}(\mathbf{w}) := \max_{\mathbf{x}} \sum_{i \in \mathcal{S}_{\mathcal{M}}(c)} \sum_{j \in \mathcal{S}_{\mathcal{W}}(c)} (w_{ij}^{\mathcal{M}} + w_{ij}^{\mathcal{W}})x_{ij} \qquad (O_{\widetilde{sw}}(\mathbf{w}))$$
$$\text{s.t.} \qquad \text{Constraints (1)–(5).}$$

Let $(\widetilde{\mathbf{u}}, \widetilde{\mathbf{x}})$ be an optimal solution to problem $O_{\widetilde{sw}}(\mathbf{w})$. Clearly, by the feasibility of x^*,

$$\sum_{i \in \mathcal{S}_{\mathcal{M}}(c)} \sum_{j \in \mathcal{S}_{\mathcal{W}}(c)} (w_{ij}^{\mathcal{M}} + w_{ij}^{\mathcal{W}})x^*_{ij} \leq \sum_{i \in \mathcal{S}_{\mathcal{M}}(c)} \sum_{j \in \mathcal{S}_{\mathcal{W}}(c)} (w_{ij}^{\mathcal{M}} + w_{ij}^{\mathcal{W}})\widetilde{x}_{ij}.$$

Thus,

$$\begin{aligned} \sum_{i \in \mathcal{M}} u_i^* + \sum_{j \in \mathcal{W}} u_j^* &= \sum_{i \in \mathcal{M}} \sum_{j \in \mathcal{W}} (w_{ij}^{\mathcal{M}} + w_{ij}^{\mathcal{W}})x^*_{ij} \\ &\leq \sum_{i \in \mathcal{S}_{\mathcal{M}}(c)} \sum_{j \in \mathcal{S}_{\mathcal{W}}(c)} (w_{ij}^{\mathcal{M}} + w_{ij}^{\mathcal{W}})\widetilde{x}_{ij} + \sum_{i \in \mathcal{M} \setminus \mathcal{S}_{\mathcal{M}}(c)} u_i^* + \sum_{j \in \mathcal{W} \setminus \mathcal{S}_{\mathcal{W}}(c)} u_j^*. \end{aligned} \tag{17}$$

Let us write the dual for LP $O_{\widetilde{sw}}(\mathbf{w})$ (in the variables $\beta := (\beta_i : i \in \mathcal{S}_{\mathcal{M}}(c))$ and $\sigma := (\sigma : j \in \mathcal{S}_{\mathcal{W}}(c))$):

$$z^*_{\widetilde{sw}}(\mathbf{w}) := \min_{\beta, \sigma} \sum_{i \in \mathcal{S}_{\mathcal{M}}(c)} \beta_i + \sum_{j \in \mathcal{S}_{\mathcal{W}}(c)} \sigma_j \qquad (D_{\widetilde{sw}}(\mathbf{w}))$$

$$\text{s.t.} \qquad \beta_i + \sigma_j \geq w_{ij}^{\mathcal{M}} + w_{ij}^{\mathcal{W}}, \quad \text{for all } i \in \mathcal{S}_{\mathcal{M}}(c) \text{ and } j \in \mathcal{S}_{\mathcal{W}}(c), \tag{18}$$

$$\beta_i \geq 0,\ \sigma_j \geq 0, \quad \text{for all } i \in \mathcal{S}_{\mathcal{M}}(c) \text{ and } j \in \mathcal{S}_{\mathcal{W}}(c). \tag{19}$$

From Equation (16), we have for $i \in \mathcal{S}_{\mathcal{M}}(c)$ and $j \in \mathcal{S}_{\mathcal{W}}(c)$: $\widehat{\beta}_i + \widehat{\sigma}_j \geq c \cdot (w_{ij}^{\mathcal{M}} + w_{ij}^{\mathcal{W}})$. Thus, the pair of vectors $(\beta', \sigma') := ((\frac{\widehat{\beta}_i}{c} : i \in \mathcal{S}_{\mathcal{M}}(c)), (\frac{\widehat{\sigma}_j}{c} : j \in \mathcal{S}_{\mathcal{W}}(c)))$ are feasible for the dual LP $D_{\widetilde{sw}}(\mathbf{w})$. It follows that

$$\begin{aligned}
\sum_{\substack{i\in\mathcal{S}_{\mathcal{M}}(c)\\ j\in\mathcal{S}_{\mathcal{W}}(c)}} (w_{ij}^{\mathcal{M}} + w_{ij}^{\mathcal{W}})\widetilde{x}_{ij} &= z^*_{\widetilde{sw}}(\mathbf{w}) \text{ (by LP duality)}\\
&\leq \frac{1}{c}\left(\sum_{i\in\mathcal{S}_{\mathcal{M}}(c)} \widehat{\beta}_i + \sum_{j\in\mathcal{S}_{\mathcal{W}}(c)} \widehat{\sigma}_j\right) \text{ (by the feasibility of } (\beta',\sigma') \text{ for } (D_{\widetilde{sw}}(\mathbf{w})))\\
&\leq \frac{1}{c}\left(\sum_{\substack{i\in\mathcal{M}\\ j\in\mathcal{W}}} \pi'_{i,1}(\widehat{u}_i) w_{ij}^{\mathcal{M}}\widehat{x}_{ij} + \sum_{\substack{j\in\mathcal{W}\\ i\in\mathcal{M}}} \pi'_{j,1}(\widehat{u}_j) w_{ij}^{\mathcal{W}}\widehat{x}_{ij}\right) \text{ (by Equation (15))}\\
&\leq \frac{1}{c}\left(\sum_{\substack{i\in\mathcal{S}_{\mathcal{M}}(c)\\ j\in\mathcal{W}}} \pi'_{i,1}(\widehat{u}_i) w_{ij}^{\mathcal{M}}\widehat{x}_{ij} + \sum_{\substack{j\in\mathcal{S}_{\mathcal{W}}(c)\\ i\in\mathcal{M}}} \pi'_{j,1}(\widehat{u}_i) w_{ij}^{\mathcal{W}}\widehat{x}_{ij} + \right.\\
&\qquad \left. c\sum_{\substack{i\in\mathcal{M}\setminus\mathcal{S}_{\mathcal{M}}(c)\\ j\in\mathcal{W}}} w_{ij}^{\mathcal{M}}\widehat{x}_{ij} + c\sum_{\substack{j\in\mathcal{W}\setminus\mathcal{S}_{\mathcal{W}}(c)\\ i\in\mathcal{M}}} w_{ij}^{\mathcal{M}}\widehat{x}_{ij}\right) \text{ (by definition of } \mathcal{S}_{\mathcal{W}}(c), \mathcal{S}_{\mathcal{M}}(c))\\
&= \frac{1}{c}\left(\sum_{\substack{i\in\mathcal{S}_{\mathcal{M}}(c)\\ j\in\mathcal{W}}} \pi'_{i,1}(\widehat{u}_i) w_{ij}^{\mathcal{M}}\widehat{x}_{ij} + \sum_{\substack{j\in\mathcal{S}_{\mathcal{W}}(c)\\ i\in\mathcal{M}}} \pi'_{j,1}(\widehat{u}_i) w_{ij}^{\mathcal{W}}\widehat{x}_{ij} + \right.\\
&\qquad \left. c\sum_{i\in\mathcal{M}\setminus\mathcal{S}_{\mathcal{M}}(c)} \widehat{u}_i + c\sum_{j\in\mathcal{W}\setminus\mathcal{S}_{\mathcal{W}}(c)} \widehat{u}_j\right) \text{ (by feasibility of } \widehat{\mathbf{x}} \text{ for } (O_{sw}(\mathbf{w}))).
\end{aligned} \tag{20}$$

Using Equation (20) in Equation (17), we arrive at the claim. □

Using the above two claims, we get

$$\frac{\sum_{i\in\mathcal{M}}\widehat{u}_i + \sum_{j\in\mathcal{W}}\widehat{u}_j}{\sum_{i\in\mathcal{M}}u_i^* + \sum_{j\in\mathcal{W}}u_j^*} \geq$$

$$\geq \frac{\sum_{i\in\mathcal{S}_{\mathcal{M}}(c)_{j\in\mathcal{W}}} w_{ij}^{\mathcal{M}}\widehat{x}_{ij} + \sum_{j\in\mathcal{S}_{\mathcal{W}}(c)_{i\in\mathcal{M}}} w_{ij}^{\mathcal{W}}\widehat{x}_{ij} + \sum_{i\in\mathcal{M}\setminus\mathcal{S}_{\mathcal{M}}(c)}\widehat{u}_i + \sum_{j\in\mathcal{W}\setminus\mathcal{S}_{\mathcal{W}}(c)}\widehat{u}_j}{\frac{1}{c}\sum_{i\in\mathcal{S}_{\mathcal{M}}(c)_{j\in\mathcal{W}}} \pi'_{i,1}(\widehat{u}_i)w_{ij}^{\mathcal{M}}\widehat{x}_{ij} + \frac{1}{c}\sum_{j\in\mathcal{S}_{\mathcal{W}}(c)_{i\in\mathcal{M}}} \pi'_{j,1}(\widehat{u}_i)w_{ij}^{\mathcal{W}}\widehat{x}_{ij} + \sum_{i\in\mathcal{M}\setminus\mathcal{S}_{\mathcal{M}}(c)}(\widehat{u}_i + u_i^*) + \sum_{j\in\mathcal{W}\setminus\mathcal{S}_{\mathcal{W}}(c)}(\widehat{u}_j + u_j^*)}$$

$$\geq \min\left\{c\cdot\min_{i\in\mathcal{S}_{\mathcal{M}}(c)}\frac{1}{\pi'_{i,1}(\widehat{u}_i)},\ c\cdot\min_{j\in\mathcal{S}_{\mathcal{W}}(c)}\frac{1}{\pi'_{j,1}(\widehat{u}_j)},\ \min_{i\in\mathcal{M}\setminus\mathcal{S}_{\mathcal{M}}(c)}\frac{\widehat{u}_i}{\widehat{u}_i + u_i^*},\ \min_{j\in\mathcal{W}\setminus\mathcal{S}_{\mathcal{W}}(c)}\frac{\widehat{u}_j}{\widehat{u}_j + u_j^*}\right\}$$

$$\geq \min\left\{\frac{c}{\max_{i\in\mathcal{S}_{\mathcal{M}}(c)}\pi'_{i,1}(0)},\ \frac{c}{\max_{j\in\mathcal{S}_{\mathcal{W}}(c)}\pi'_{j,1}(0)},\ \frac{L(\mathbf{q},c)}{2},\ \frac{L(\mathbf{q},c)}{2}\right\} \tag{21}$$

$$\geq \min\left\{\frac{c}{H(\mathbf{q})},\ \frac{c}{H(\mathbf{q})},\ \frac{L(\mathbf{q},c)}{2},\ \frac{L(\mathbf{q},c)}{2}\right\} = \frac{L(\mathbf{q},c)}{2}, \tag{22}$$

where Equation (21) follows from the concavity of q_i and Claim 1; Equation (22) follows from Equation (7). The theorem follows. □

The following example illustrates an application of the theorem.

Example 1. Suppose that $q_i(u) := u(1-u)^{(1-\alpha)}$ for all $i \in \mathcal{M} \cup \mathcal{W}$, which satisfies assumptions (A1) through (A3) for $0 \leq \alpha < 1$. Figure 2 (left) shows the probabilities for different αs. As α

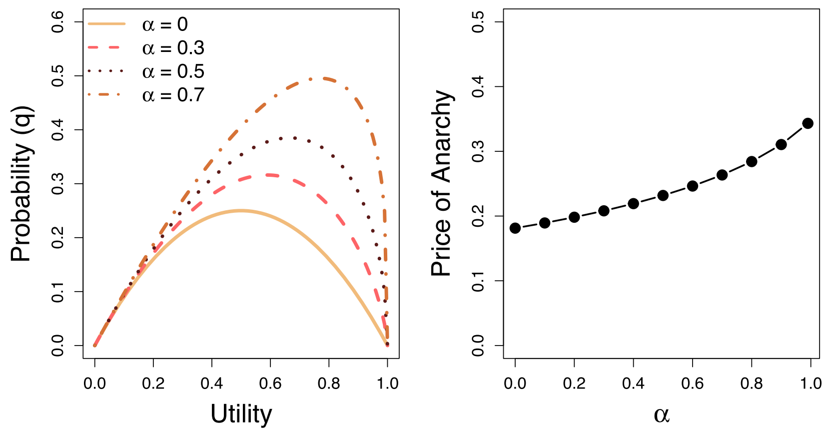

Fig. 2. **Left:** The figure shows probability functions $q_i(u) := u(1-u)^{(1-\alpha)}$ for different αs. **Right:** The price of anarchy as a function of α.

increases, the user's probability of returning to the system is maximized with a higher utility. Next, we find the price of anarchy. We have that

$$\pi'_{i,1}(u) = \frac{q'_i(u)}{(1+q_i(u))^2} = \frac{(1-u)^{(-\alpha)}(\alpha u - 2u + 1)}{(1+u(1-u)^{(1-\alpha)})^2}.$$

Let $\overline{u}$ be the unique solution of Equation (8): $c = \pi'_{i,1}(\overline{u})$. We also have that $H(\mathbf{q}) = q'_i(0) = 1$, and c is chosen to satisfy Equation (7), which means that $c = \overline{u}/2$. Thus, we have to find the unique solution in $(0,1)$ for the equation: $\overline{u}/2 = \pi'_{i,1}(\overline{u})$. Finally, the PoA is given by $L(\mathbf{q}, c)/2$, where $L(\mathbf{q}, c) = \min_{i \in \mathcal{M} \cup \mathcal{W}} \overline{u}_i(c) = \overline{u}$.

For $\alpha = 0$, the above equation becomes, $\overline{u}/2 = (1-2\overline{u})/((1-\overline{u})\overline{u}+1)^2$, whose unique solution is $\overline{u} = 0.363$, which implies that $\text{PoA}(q) \geq 0.1815$. Figure 2 (right) shows the PoA results for all αs in $(0,1)$. Intuitively, as α increases, so does the PoA, because the peak of $q_(u)$ is closer to 1.

2.3 The Invisible-hand Effect

So far, we have considered the monopolistic case where users have only one platform in the market to get matched. One might expect that a competitive market, in which users are free to choose among competing dating sites, would benefit users' utility. Next, we examine the effect of competition among dating sites. To do so, we assume that a user can only be on one site at a time and that the probability of the user remaining on a dating site is proportional to the utility that the user has previously received from that site. That is, users who have received better service from the system in the past are more likely to return in the future.

Formally, the conditional probability that a user will return to the same dating site having previously received a utility u, given that he or she wants to date again, is $\Pr(\text{return}|\text{dates again}) = u$. The probability of dating again is defined as $\Pr(\text{dates again}) = q(u)$. Thus,

$$\Pr(\text{return}) = \Pr(\text{return}|\text{dates again}) \Pr(\text{dates again}) = u \cdot q(u).$$

Similarly, if a user wants to date again, he or she will consider a competing dating site with a probability of $1-u$.

We extend the Markov decision process defined earlier to model the competitive situation for user i as a three-state Markov process. The first state represents matchmaking with the dating

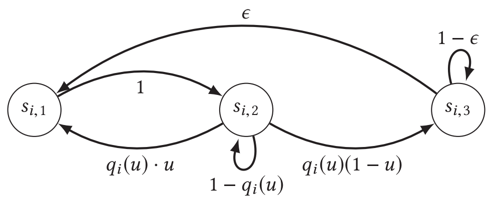

Fig. 3. Markov process with three states $s_{i,1}$, $s_{i,2}$, and $s_{i,3}$, representing the states of actively dating in the system, not seeking a partner, and being out of the system (or in the competition), respectively, for user i having received utility u.

system; the second represents being on the system but not matchmaking; the third represents leaving the system to join a competing dating site. We assume that once a user leaves, there is a small probability ϵ that he or she will return. $\epsilon > 0$ ensures that the Markov chain is ergodic (i.e., any state can be reached from any other state) and thus has a stationary distribution. The stationary distribution of a Markov chain is a probability distribution representing the proportion of time spent in each state in the long run. The Markov chain is shown in Figure 3.

From state $s_{i,1}$, the user moves with probability 1 to state $s_{i,2}$, where he or she stays with probability $1 - q_i(u)$ (for either a successful relationship or complete dissatisfaction with the system) or returns to the system (state $s_{i,1}$) with probability $q_i(u) \cdot u$ (note that without competition this probability was $q_i(u)$). The user can also move from state $s_{i,2}$ to state $s_{i,3}$ (for a user switching to the competition) with probability $q_i(u)(1-u)$, where he or she stays there with probability $1-\epsilon$ (ϵ being smaller as the competition is higher) or returns to the system with probability ϵ. Then, the limiting probability in each state is given by

$$\begin{aligned}
\pi_{i,1}(u) =& \pi_{i,2}(u)q_i(u) \cdot u + \pi_{i,3}(u)\epsilon, \\
\pi_{i,2}(u) =& \pi_{i,1}(u) + \pi_{i,2}(u)(1-q_i(u)), \\
\pi_{i,3}(u) =& \pi_{i,2}(u)q_i(u) \cdot (1-u) + \pi_{i,3}(u)(1-\epsilon), \\
1 =& \pi_{i,1}(u) + \pi_{i,2}(u) + \pi_{i,3}(u),
\end{aligned}$$

which yields

$$\pi_{i,1}(u,\epsilon) = \frac{q(u)}{1 + q_i(u) + \frac{q_i(u)}{\epsilon}(1-u)}. \tag{23}$$

Note that $\pi_{i,1}(u)$ is not necessarily a concave function, and thus, the result from Theorem 1 cannot be applied directly. Importantly, however, we show in the following lemma that as competition increases, the self-interested behavior of system designers converges to maximize the social utility of users.

LEMMA 1. *Let $\pi_{i,1}(u,\epsilon)$, defined in Equation (23), be the stationary distribution of state* 1, *representing being in the system under the Markov chain with competition. Then, for all* $i \in \mathcal{M}$, $\operatorname{argmax}_u \pi_{i,1}(u,\epsilon) = 1$, *as* $\epsilon \to 0$.

PROOF. To derive the above result, we show that the global optimum u^* of $\pi_{i,1}(u,\epsilon)$ is the unique solution satisfying

$$\pi'_{i,1}(u^*,\epsilon) = \frac{q'_i(u^*) + \frac{q_i^2(u^*)}{\epsilon}}{\left(1 + q_i(u^*) + \frac{q_i(u^*)}{\epsilon}(1-u^*)\right)^2} = 0,$$

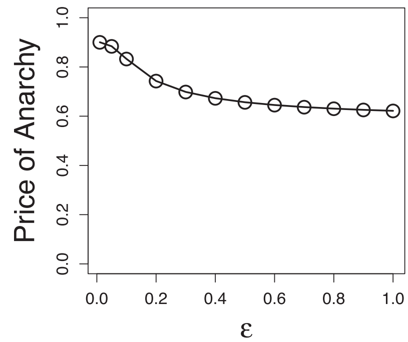

Fig. 4. The figure shows the empirical price of anarchy for the probability function $q_i(u) := u(1-u)$ for different ϵs. The empirical price of anarchy is obtained as the minimum utility ratio between the selfish and the social welfare optimal matching, over several instances, randomizing the preference weights.

or equivalently,

$$\epsilon = -\frac{q_i^2(u^*)}{q_i'(u^*)}. \tag{24}$$

Since $\epsilon > 0$, Equation (24) holds only if $q_i'(u^*) < 0$. This in turn means that $u^* \in [u', 1]$, where u' is the solution to $q_i'(u') = 0$. It follows from the concavity of q that $-\frac{q_i^2(u)}{q_i'(u)}$ is decreasing in the interval $u \in [u', 1]$. Therefore, u^* must be unique. Furthermore, since $\epsilon \to 0$, we have that $q_i^2(u^*) \to 0$, or equivalently, $u^* \to 1$. □

Example 2. We can extend Example 1, with $q_i(u) := u(1-u)$, to the competitive case. The direct application of Theorem 1 is not possible, since $\pi_{i,1}(u)$ is not a concave function for all values of *epsilon* $\in [0, 1]$. Figure 4 shows the empirical PoA for the probability function $q_i(u) := u(1-u)$. The empirical PoA is obtained as the minimum utility ratio between the selfish and social welfare maximizing matching, over several instances, randomizing the preference weights. That is, for each instance, we generate a set of random weights $\mathbf{w}$ sampled from a Beta distribution (with parameters $(2, 2)$). We then solve the optimal matching $O_{sw}(\mathbf{w})$ and the selfish matching $O_s(\mathbf{w}; \mathbf{q})$ and obtain the ratio of their utilities.

2.4 Bounds on PoA in the Online Model

In this section, we will focus on an online extension of the matching problem. In the online model, elements in $\mathcal{M} \cup \mathcal{W}$ are not now known in advance but arrive sequentially and must be matched upon arrival. Inspired by the notion of competitive ratio, we are interested in quantifying the efficiency of assignments for the online selfish matching problem ($O_s(\mathbf{w}; \mathbf{q})$) with respect to assignments for the offline social welfare maximizing matching problem ($O_{sw}(\mathbf{w})$).

Our model is a special case of the online bipartite matching problem with edge weights. Most work focuses on one-sided online problems, where one set of nodes (out of the two) is known in advance and the other set is revealed online. In the unweighted version, the greedy algorithm is known to have a competitive ratio of 0.5. When edge weights are considered and nodes arrive in an adversarial order, any algorithm can be arbitrarily bad without any additional assumptions [7]. However, with the addition of the "free disposal" assumption, where multiple arriving nodes can be assigned to the same node but only the maximum weight is considered, the greedy algorithm has a competitive ratio of 0.5 [7]. For the two-sided online matching problem, where both sets of nodes arrived online, when fractional matches are allowed, the greedy algorithm has a competitive ratio of 0.5 [34], and no online algorithm can improve a competitive ratio of approximately 0.63 [16].

The question of whether it is possible to design a weighted online bipartite matching algorithm that outperforms the greedy competitive ratio of 1/2 has remained an intriguing open problem. For that reason, we will focus on the simple and practical online greedy algorithm. The greedy algorithm is a simple and efficient algorithm that can be applied to a wide range of optimization problems with minimal effort in implementation and interpretation. In practice, the greedy algorithm is often used as a heuristic or as an initial solution that can be further refined by other methods [2].

Definition 2 (Online Greedy Selfish Matching). Upon arrival of an element $k \in \mathcal{M} \cup \mathcal{W}$, the algorithm decides a partial match for k based on all previously revealed elements. Let $\mathcal{M}' := [m'] \subseteq \mathcal{M}$ and $\mathcal{W}' := [n'] \subseteq \mathcal{W}$ be the revealed elements at a given step. When an element $k \in \mathcal{M} \cup \mathcal{W}$ arrives, if $k \in \mathcal{M}$, the system decides the (partial) distribution $\mathbf{x}_k := (x_{kj} : j \in \mathcal{W}', \sum_{j \in \mathcal{W}'} x_{kj} \leq 1)$. Otherwise, if $k \in \mathcal{W}$, the (partial) distribution is $\mathbf{x}_k := (x_{ik} : i \in \mathcal{M}', \sum_{i \in \mathcal{M}'} x_{ik} \leq 1)$. Upon arrival of element $k \in \mathcal{M}$, the algorithm maximizes the retention rate of k: $\pi_{k,1}(u_k)$ (defined in Equation (6)), and the retention of all elements potentially matched to k: $\sum_{j \in \mathcal{W}'} \pi_{j,1}(u_j)$. That is, the greedy algorithm solves

$$\max_{\mathbf{x}_i} \ \pi_{k,1}(u_k) + \sum_{j \in \mathcal{W}'} \pi_{j,1}(u_j) \tag{25}$$

$$\text{s.t.} \qquad u_k := \sum_{j \in \mathcal{W}'} w_{kj}^{\mathcal{M}} x_{kj}, \tag{26}$$

$$u_j := \sum_{i \in \mathcal{M}'} w_{ij}^{\mathcal{W}} x_{ij}, \quad \text{for all } j \in \mathcal{W}', \tag{27}$$

$$\sum_{j \in \mathcal{W}'} x_{kj} \leq 1, \tag{28}$$

$$\sum_{i \in \mathcal{M}'} x_{ij} \leq 1, \quad \text{for all } j \in \mathcal{W}', \tag{29}$$

$$x_{kj} \geq 0, \quad \text{for all } j \in \mathcal{W}'. \tag{30}$$

For the arrival of element $k \in \mathcal{W}$, the greedy algorithm is analogous.

Note that the difference of this problem with respect to the offline model comes from restricting the set of constraints to specify the availability of elements $\mathcal{M}' \cup \mathcal{W}'$ and to maximize the engagement of element k and the revealed potential matches of k. Also, note that by allowing partial matches, an element can be partially matched to newly arriving elements.

Next, we analyze the PoA of the greedy policy for the selfish (probabilistic) matching problem. We extend the definition of the price of anarchy for the online model as follows.

Definition 3 (Price of Anarchy for the Online Model). Let $\mathcal{A}$ be an online matching algorithm for solving the selfish online matching problem ($O_s(\mathbf{w};\mathbf{q})$) and $(\mathbf{u}_s^{\mathcal{A}};\mathbf{q})$ be the objective value of the online strategy. Let $f(\mathbf{u}) := \sum_i u_i$ and $u_{sw}^* \in \operatorname{argmax}\{f(\mathbf{u}) : \mathbf{u}$ satisfies Equations (1)–(5)$\}$. That is, u_{sw}^* is the social optimum for the offline model. Define the *price of anarchy* $\text{PoA}^{\mathcal{A}} = \text{PoA}^{\mathcal{A}}(\mathbf{q}) \in [0,1]$ (from the users' point of view), with respect to algorithm $\mathcal{A}$, as

$$\text{PoA}^{\mathcal{A}}(\mathbf{q}) := \min_{\mathbf{w} \geq \mathbf{0}} \frac{f(\mathbf{u}_s^{\mathcal{A}})}{f(\mathbf{u}_{sw}^*)}.$$

Theorem 2. *Let $\mathcal{A}$ be the greedy policy for the selfish (probabilistic) matching problem. Then*

$$\mathrm{PoA}^{\mathcal{A}}(\mathbf{q}) \geq L(\mathbf{q}, c)/2,$$

where $L(\mathbf{q}, c)$ is defined as in Theorem 1.

Proof. By the KKT (necessary) conditions for optimality, there exists $\widehat{\beta}_{i,j} \geq 0$, $\widehat{\sigma}_{i,j} \geq 0$, and $\widehat{\mu}_{ij} \geq 0$, for $i \in \mathcal{M}$ and $j \in \mathcal{W}$, such that Equations (26) through (30) hold for $(\mathbf{u}, \mathbf{x}) = (\widehat{\mathbf{u}}, \widehat{\mathbf{x}})$, and

$$-\pi'_{i,1}(\widehat{u}_i) w_{ij}^{\mathcal{M}} - \pi'_{j,1}(\widehat{u}_j) w_{ij}^{\mathcal{W}} + \widehat{\beta}_i + \widehat{\sigma}_{i,j} = \widehat{\mu}_{ij}, \quad \text{for all } i \in \mathcal{M} \text{ and } j \in \mathcal{W}, \tag{31}$$

$$\widehat{\mu}_{ij} \widehat{x}_{ij} = 0, \quad \text{for all } i \in \mathcal{M} \text{ and } j \in \mathcal{W}, \tag{32}$$

$$\widehat{\beta}_{i,j} \left(\sum_{k \leq j} \widehat{x}_{ij} - 1 \right) = 0, \quad \text{for all } i \in \mathcal{M} \text{ and } j \in \mathcal{W}, \tag{33}$$

$$\widehat{\sigma}_{i,j} \left(\sum_{k \leq i} \widehat{x}_{kj} - 1 \right) = 0, \quad \text{for all } i \in \mathcal{M} \text{ and } j \in \mathcal{W}. \tag{34}$$

For any $j \in \mathcal{W}$, taking the sum of the constraints in Equation (34) over $i \in \mathcal{M}$ gives

$$\sum_{i \in \mathcal{M}} \widehat{\sigma}_{i,j} \widehat{x}_{ij} = \sum_{i \in \mathcal{M}} \widehat{\sigma}_{i,j} \left(1 - \sum_{k < i} \widehat{x}_{kj} \right) \geq \sigma_{1,j},$$

which implies that

$$\sum_{i \in \mathcal{M},\ j \in \mathcal{W}} \widehat{\sigma}_{i,j} \widehat{x}_{ij} \geq \sum_{j \in \mathcal{W}} \sigma_{1,j}.$$

In the same way, taking the sum of the constraints in Equation (33) over $j \in \mathcal{W}$ gives

$$\sum_{i \in \mathcal{M},\ j \in \mathcal{W}} \widehat{\beta}_{i,j} \widehat{x}_{ij} \geq \sum_{i \in \mathcal{M}} \beta_{i,1}.$$

From these two formulas, together with Equations (28), (29), and (31) through (34), it follows that

$$\sum_{i \in \mathcal{M},\ j \in \mathcal{W}} \pi'_{i,1}(\widehat{u}_i) w_{ij}^{\mathcal{M}} \widehat{x}_{ij} + \sum_{j \in \mathcal{W},\ i \in \mathcal{M},} \pi'_{j,1}(\widehat{u}_j) w_{ij}^{\mathcal{W}} \widehat{x}_{ij} \geq \sum_{i \in \mathcal{M}} \widehat{\beta}_{i,1} + \sum_{j \in \mathcal{W}} \widehat{\sigma}_{1,j}, \tag{35}$$

$$\widehat{\beta}_{i,1} + \widehat{\sigma}_{1,j} \geq \pi'_{i,1}(\widehat{u}_i) w_{ij}^{\mathcal{M}} + \pi'_{j,1}(\widehat{u}_j) w_{ij}^{\mathcal{W}}, \quad \text{for all } i \in \mathcal{M} \text{ and } j \in \mathcal{W}. \tag{36}$$

From Equation (36), we have for $i \in \mathcal{S}_{\mathcal{M}}(c)$ and $j \in \mathcal{S}_{\mathcal{W}}(c)$: $\widehat{\beta}_{i,1} + \widehat{\sigma}_{1,j} \geq c \cdot (w_{ij}^{\mathcal{M}} + w_{ij}^{\mathcal{W}})$. Thus, the pair of vectors $(\beta', \sigma') := ((\frac{\widehat{\beta}_{i,1}}{c} : i \in \mathcal{S}_{\mathcal{M}}(c)), (\frac{\widehat{\sigma}_{1,j}}{c} : j \in \mathcal{S}_{\mathcal{W}}(c))$ are feasible for the dual LP $O_{sw}(\widetilde{\mathbf{w}})$. As such, the result follows from Equations (20), (17), and (22). □

Analogous to the offline matching model, Theorem 2 states that in the online matching model, the price of anarchy is bounded by a function of q representing the behavior of users. The next section tests the results with an experimental study on human subjects.

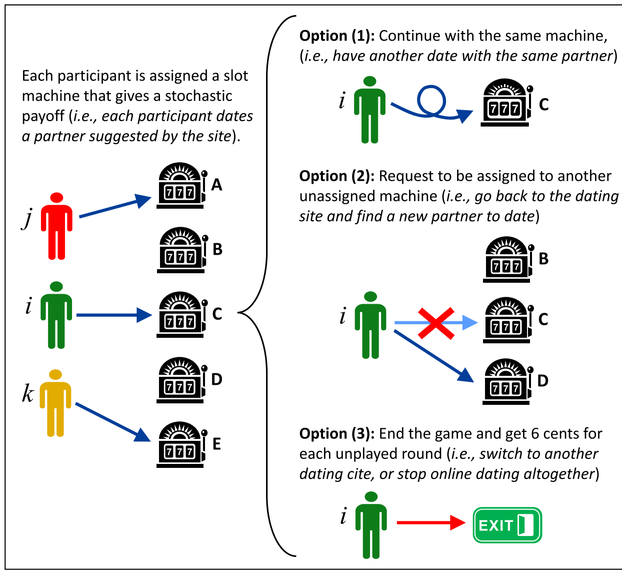

Fig. 5. An illustration of the experimental study, showing participants' initial assignments and the three options available in each of the 10 rounds of the experiment.

3 EXPERIMENTAL STUDY

Our theoretical results provide bounds on the price of anarchy, which is defined as the worst-case ratio of selfish matching to the social optimum. We showed that even in a monopolistic setting, the price of anarchy can be bounded by a constant that does not depend on the number of users, only on the implicit expectations of users about the system. However, our theoretical results are constrained by the functional form $q(u)$ that dictates user behavior, which in practice can vary from user to user and even change over time (although we expect the overall behavior to follow assumptions (A1)–(A3) on average). To complement our analytical results, we designed a human subject experiment to test our results without making any modeling assumptions about user behavior. In our experiment, $q(u)$ is learned empirically from the participants' past choices.

Our experiment is designed to evaluate two hypotheses derived from the theoretical analysis of the model. First, we want to test that the PoA depends only on the user's decisions about how to interact with the matching system, and not directly on the distribution of matching weights. Specifically, we want to examine the empirical relationship between the probability that the user returns to the system to request a new match after being matched and receiving a utility u (i.e., $q(u)$) and the PoA. Second, we want to validate whether $q(u)$ satisfies assumptions (A1) to (A3).

3.1 Experimental Design

In a series of experiments, we compared the social welfare achieved by the social welfare maximizing matching algorithm against the selfish matching algorithm. We operationalized the algorithmic matching market as a multi-armed bandit problem, where the system selects matches and then participants decide whether to accept the match or request a new match. Figure 5 provides a high-level illustration of the experiment. More specifically, participants arrive at the market and are assigned to a slot machine that gives them a stochastic payoff. Participants play for 10 rounds, and at each round, they choose between three options: (1) continue with the assigned slot machine, (2) request to be assigned to a new slot machine, or (3) take the risk-free outside option and

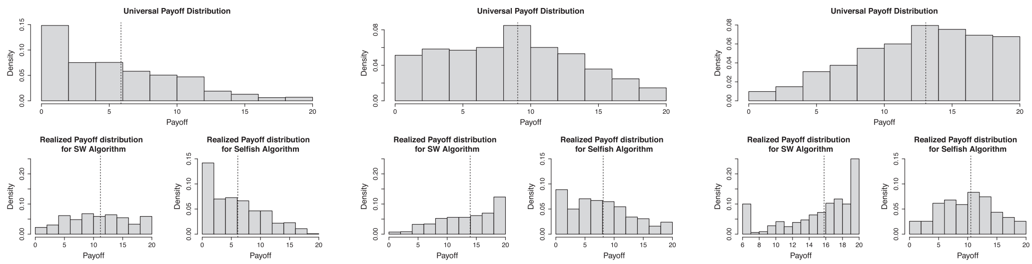

(a) Study A (b) Study B (c) Study C

Fig. 6. Distribution of payoffs. The top distributions correspond to the payoff distribution of all user-slot pairs for each study. The bottom distributions show the realized payoffs experienced by participants. That is, it corresponds to the payoff distribution of user-slot pairs that were matched. The vertical dashed lines represent the mean of the distribution. The universal distributions correspond to the realized payoffs of a matching algorithm that randomly assigns users to slots. Note that for the selfish algorithm, there is a shift to the left compared to the universal payoffs (slightly decreasing the mean), whereas for the social welfare maximizing algorithm, the shift is to the right with a significant increase in the mean. Study A features a scarce market with a right-skewed Beta distribution of payoffs $\beta(1, 2)$; Study B has a symmetric Beta distribution $\beta(2, 2)$; and Study C represents a non-scarce market with a left-skewed Beta distribution $\beta(2, 1)$.

exit the game. To make this analogy more concrete, the payoffs capture going out on a date with someone and observing the value of their match; option (1) would represent having multiple dates with the same match and engaging in a long-term relationship; option (2) represents going back to the dating site to find a new match; and finally, option (3) would be users either moving to another dating site or stopping using online dating altogether. We contextualized the experiment as an "interplanetary mining game," where slot machines are represented as planets that participants mine for a monetary reward. The advantage of this design is that we have prior knowledge of the users' preferences, which allows us to compute the optimal match. Figure 9 in the appendix shows screenshots of the experiment.

The algorithm used to match players takes two forms. The social welfare maximizing matching algorithm assigns users to slot machines to maximize the sum of the social utilities of all users. The selfish matching algorithm assigns users to slot machines to maximize the number of users requesting to be re-matched to the system. At each round, participants can choose to stop being matched by the algorithm and accept an outside payment of 6 cents for each round remaining to be played, or the user can choose to continue playing in the hope of making a larger profit.

One of the advantages of this experimental design is that the central designer has perfect information about the expected payoffs, which allows us to quantify the impact of the matching algorithms without having to deal with the challenges of inferring and predicting user preferences, which is beyond the scope of this article. The interested reader is referred to Hitsch et al. [14] for a paper on estimating mate preferences from an online dating service.

We conducted three studies with groups of participants playing simultaneously, recruited using Amazon Mechanical Turk and filtered for English speakers and good past performance on Mturk. Each study differed only in how the payoffs w_{ij} were distributed. Study A represents a scarce market, where few slot machines have high payoffs, following a right-skewed Beta distribution $\beta(1, 2)$; Study B has payoffs following a symmetric Beta distribution $\beta(2, 2)$; and Study C represents a non-scarce market where most slots have high payoffs, following a left-skewed Beta distribution $\beta(2, 1)$. Figure 6 shows the distribution for each study. The distribution of payoffs also corresponds to what the payoffs would look like if the matching algorithm randomly assigned users to slots. Table 1 shows a summary of the users and parameters in each of the three studies.

Table 1. Summary of Participants for Each Experiment

	Participants	Instances	Players per group	Outside payoff	Distribution of payoffs
Study A	113	20	3	6 cents	$\beta(1, 2)$
Study B	87	17	3	6 cents	$\beta(2, 2)$
Study C	87	14	3	6 cents	$\beta(2, 1)$

Each instance of a game consists of participants that are evenly randomized between two experimental conditions: the social welfare maximizing matching algorithm and the selfish matching algorithm. Initially, each player $i \in [n]$ is assigned a mean payoff $w_{ij} = p_j + \epsilon_i$ for each of the $j \in [m]$ slot machines, where p_i is a realization of a Beta distributed variable, and ϵ_i is a realization of a normally distributed variable with a mean of zero. Thus, the rewards from the machines are correlated across participants, implying that there is competition for high-yielding slot machines. To compare across the experimental conditions, the mean payoffs of players and slot machines are the same across the two matching conditions; e.g., the first participant assigned to the social welfare maximizing condition has the same mean payoff per slot machine $j \in [m]$ as the first participant assigned to the selfish condition. The number of participants per condition is $n = 3$ and they play for 10 rounds. To ensure that a participant is not assigned to the same slot more than once, we choose $m = 13$ slot machines.

The game is played in real time (or synchronously) with other players, so each slot machine can only be assigned to one player at a time in one of the experimental conditions. This means that if a participant decides to remain in that slot in the next round, that slot will not be available to other participants. If the participant decides to switch from a slot machine, then he or she will never be assigned to the same slot again. To continue our analogy with dating services, the above represents the case where users only date one partner at a time, the partner does not date anyone else, and some users are more universally appealing than others.

The rewards of the slots are stochastic, and their mean payoffs are known in advance to the central designer, but not to the participants. Formally, when player i is assigned to slot j, he or she receives a payoff following a normal distribution with mean $w_{ij} = p_j + \epsilon_i$ (and variance set to 3). Hence, users need repeated interactions to estimate the underlying reward of a slot machine.

The social welfare maximizing matching algorithm is given by solving problem $O_{sw}(\mathbf{w})$, where $\mathbf{w}$ are the mean payoffs. However, to solve the selfish matching problem $O_s(\mathbf{w}; \mathbf{q})$, we need to learn the probability $q_i^k(u)$ that a user i who has received utility u will return to the system in round k. In the context of our experiment, this is the probability of a user requesting to be re-matched by the system. Given that individual data is sparse, we assume that users are homogeneous for the experiment, i.e., $\forall i, j, k,\ q_i^k(u) = q_j^k(u)$. Thus, for ease of notation, we can drop the user index i in $q(u)$ and start round 1 with a prior probability $q^1(u) = u(1-u)$ satisfying assumptions (A1) to (A3). We then update the probability based on the past choices of the participants as follows:

$$q^{k+1}(u) = \alpha q^k(u) + (1-\alpha) f_k(u), \tag{37}$$

where $f_k(u)$ is the fraction of participants who, having received a payoff u at round k, requested to be reassigned. In other words, $f_k(u)$ is the empirical probability of switching, defined as the fraction of participants who requested a rematch from the system after receiving a payoff of u in round k. The parameter $\alpha = 0.15$ controls the learning rate of the update rule. This update rule is based on a common strategy, the Delta rule [31], for estimating the probability in a dynamic environment.

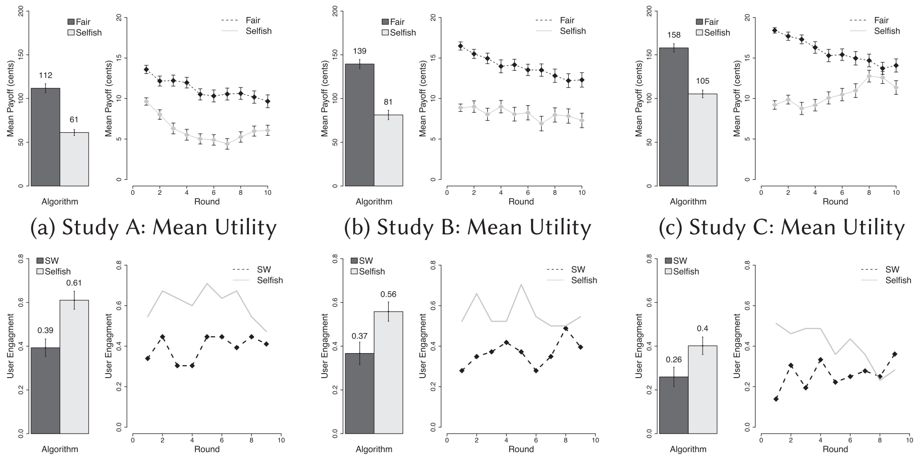

(a) Study A: Mean Utility (b) Study B: Mean Utility (c) Study C: Mean Utility

(d) Study A: Engagement Rate (e) Study B: Engagement Rate (f) Study C: Engagement Rate

Fig. 7. Figures a–c: In each sub-figure, the left plot shows the mean overall payment received by participants; the right plot shows the mean payoffs received per round (standard error bars are calculated over the mean per group). Figures d–f: In each sub-figure, the left plot shows the mean overall engagement rate of participants; the right plot shows the engagement rate per round, where engagement rate is defined as the percentage of users who request to be re-matched.

3.2 Experimental Results

Figure 7 shows the mean payments and engagement rates of the participants in each study. The engagement rate is defined as the percentage of users who request to be re-matched. Regarding the payments, the social welfare maximizing matching algorithm achieves an overall higher utility than the selfish algorithm. This difference is at most twice the utility achieved by the selfish algorithm, and it decreases across studies as the market becomes less scarce. Naturally, in the social welfare maximizing algorithm, payoffs decrease over time because, by definition, each assignment maximizes social welfare. Thus, any subsequent changes in assignments can only decrease the overall utility, but not necessarily decrease the payment of an individual participant. The lower histograms in Figure 6 show the distribution or realized payoffs that users received in each experimental condition. Compared to the underlying distribution, the social welfare maximizing algorithm can transform the distribution of payoffs, increasing the mean payoff and shifting the distribution to the right in all three studies. On the other hand, the selfish algorithm performs worse than the underlying distribution and is therefore inferior to random matching.

Focusing on engagement rates, defined as the percentage of users requesting to be re-matched, we find that the selfish algorithm (whose objective function is to maximize engagement) yields at least 50% higher rates than the social welfare maximizing matching (see Figure 7). At the same time, the selfish algorithm does not induce any significant change in the rate of users taking the outside option (i.e., drop rates) compared to the social welfare maximizing algorithm; see Figure 13 in the appendix. Drop rates are slightly higher for the social welfare maximizing algorithm; however, the difference is statistically insignificant. Overall, the selfish algorithm increases engagement rates without increasing the drop rates.

Figure 8(b) shows the average ratio of total utility over pairs of instances with the same initial conditions. The error bars depict the maximum and minimum values. Thus, the PoA is represented by the minimum value of each study, which is the worst empirical ratio. Note that for some rounds

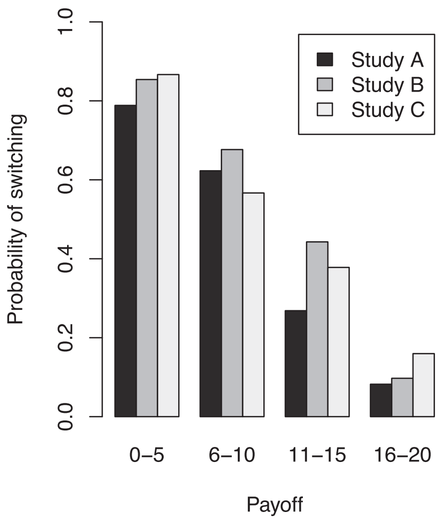

(a) Probability of switching

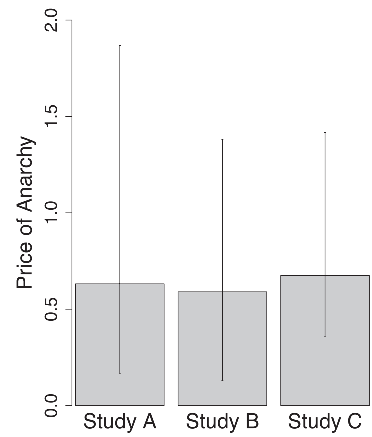

(b) Price of Anarchy

Fig. 8. **Left:** Probability of switching for each study. Distribution of the fraction of participants who received a payoff u in round i that requested to be reassigned. **Right:** Average price of anarchy for the three different studies. The plot shows the mean ratio of total utility between the social welfare maximizing algorithm and the selfish algorithm, controlling for initial conditions. The error bars depict the maximum and minimum levels. The minimum level represents the PoA, which is the worst ratio of each study.

the maximum was greater than 1, meaning that in some cases the selfish algorithm outperformed the social welfare maximizing algorithm. Figure 8(a) shows the empirical distribution of switching $q(u)$ for the different studies. We can observe that the concavity assumption is not satisfied empirically. Contrasting these empirical results with the analytical results shown in Figure 2, the theory suggests that the increase in the PoA in Study C is driven by the higher probability of switching for larger payoffs ($11 - 20$ cents).

Figures 10, 11, and 12 in the appendix show the empirical distribution of switching $f_k(u)$ for the different rounds and algorithms in Studies A, B, and C, respectively. As can be seen, the switching rates of participants who received payoffs below 6 cents are surprisingly high, given that they could choose the risk-free outside option of 6 cents. In addition, we observe a decreasing trend in switching rates for high payments (16–20 cents) as the number of rounds advances. Taken together, these results suggest that participants are risk takers, but their propensity to take risks decreases in the final rounds of the game.

Overall, our main findings are as follows: As the markets become less scarce, the social welfare loss generated by the selfish matching algorithm decreases, while the gain in engagement rates across studies remains constant. At the same time, the selfish algorithm does not induce any significant change in the rate of users taking the outside option compared to the social welfare maximizing algorithm (see Figure 13 in the appendix). In the scarcer case (Study A), the mean utility of users under the selfish algorithm was precisely what they could have received as an outside payment had they decided not to play. This is a surprising outcome, given that the matching algorithm is agnostic about the outside payment and its value. Rather, the algorithm in Equation (37) learns to maximize the probability that a user will click again based on the implicit feedback provided by the users' actions. Through this feedback loop, the algorithm self-regulates the system to keep payments close to, but above, the baseline to keep engagement high at a minimal cost to the system. This suggests that users' implicit expectations of their potential utility, both in and out of the system, drive the selfish matching algorithm to perform in different ways.

4 DISCUSSION

Millions of people use dating websites worldwide. While these algorithmic matching sites offer users access to an unprecedented number of potential partners, these sites are not altruistic, but rather self-interested, driven by economic incentives. A fundamental conflict of interest arises as the optimal matching of users may result in long-term couples and a decrease in the number of singles using the site. Consequently, to maximize the likelihood of users returning to the site, the service providers are incentivized to use algorithms that suggest matches that lead to short-lived relationships. Indeed, a recent study suggests that online couples are less likely to marry than offline couples [23]; however, conditional on being married, online married couples are slightly less likely to have their marriage result in a break-up [5]. These incentives go beyond online dating and also apply to repeated oligopolistic seller-buyer interactions, from laptops and phones to subscription apps and software providers. In these cases, the online platform can match users in a mediocre way so that they are likely to return to the service when their match is not optimal.

Our results shed new light on how the self-interested behavior of online dating sites impacts the social welfare of their users. We introduce and study a model based on the classical matching problem on graphs to represent the self-interested behavior of dating sites. We quantify the social efficiency of dating sites using the notion of price of anarchy to capture how much the system degrades users' utility due to the self-interested behavior of the system. We establish theoretical bounds on the price of anarchy that are independent of the number of users, regardless of whether the system knows all users in advance or whether they arrive sequentially (see Section 2.4). Furthermore, we experimentally test our model and create a matching market with human subjects to compare the social welfare achieved by an optimal matching service with a selfish matching algorithm.

Closely related to our question is the concern about inefficiencies that may arise due to user behavior on online dating platforms. Motivated by group fairness considerations in online bipartite matching, Esmaeili et al. [6] propose a matching algorithm with fair guarantees on both sides of the market. Another study proposes a two-sided matching market in which strategic users on both sides can screen potential partners. The authors characterize the equilibria as a function of the screening costs that users on both sides of the market incur when searching for matches and show that simple interventions such as blocking one of the sides from screening, leading to a one-sided searching, can improve the social welfare [18]. There is empirical evidence of such asymmetries in screening for partners; for example, women are much more selective than men on *Tinder* [32]. Empirical studies in online dating sites have found that both women and men have a strong preference for similarity along many attributes, yet they tend to pursue more desirable partners [3, 14]. Using the estimated mate preferences based on a particular online dating site, one study shows that the online dating site achieves near-optimal matching compared to the optimal matching predicted by the Gale-Shapley algorithm in a two-sided market [13]. Freedman et al. [10] propose a methodology for estimating the weights of participant profiles in a kidney exchange matching problem based on human-elicited value judgments. In the context of recommendation systems that match users to content providers, recent work has focused on modeling and optimizing users' long-term utility, as opposed to the standard myopic matching that maximizes user engagement [21, 35].

We recognize that our findings are subject to several limitations, and future studies may extend these results in useful ways. First, while we study offline and online centralized matching, we leave for future work extensions to a dynamic environment where preferences may change over time. Second, one can extend the complete information setting from the perspective of the central designer to account for uncertainty about users' preferences. Third, future work can explore more realistic or complex user interactions with the platform, as well as models of user behavior, by relaxing some of the assumptions required for analytically tractable purposes. Fourth, our goal was to present a stylized and simplified model to capture online dating in a general sense, not to model any specific platform. Lastly, the experiments conducted with human subjects are limited in scope and may not fully capture the complexity and diversity of real-world two-sided matching markets. Replicating the study in more diverse contexts could provide further insights into the mechanisms of user behavior in algorithmic matching systems.

This study is a step toward understanding the impact of algorithms on society and opens several problems and challenges. Further studies on how machines shape human behavior and vice versa are crucial for the design of AI for society [25]. Our daily lives are increasingly being affected by black-box systems, and we advocate the need for more regulation of these systems. It is our responsibility as researchers to understand, quantify, and inform policymakers about the possible—intentional or otherwise—side effects of algorithms.

Ethical Compliance

This study was approved by the Institute Review Board (IRB) at the Massachusetts Institute of Technology (MIT). The authors complied with all relevant ethical considerations.

Data Availability

The data and code that support the findings of this study are available at https://github.com/aabeliuk/Algorithmic-Matching-Exp.

APPENDIX

A EXPERIMENTAL STUDY

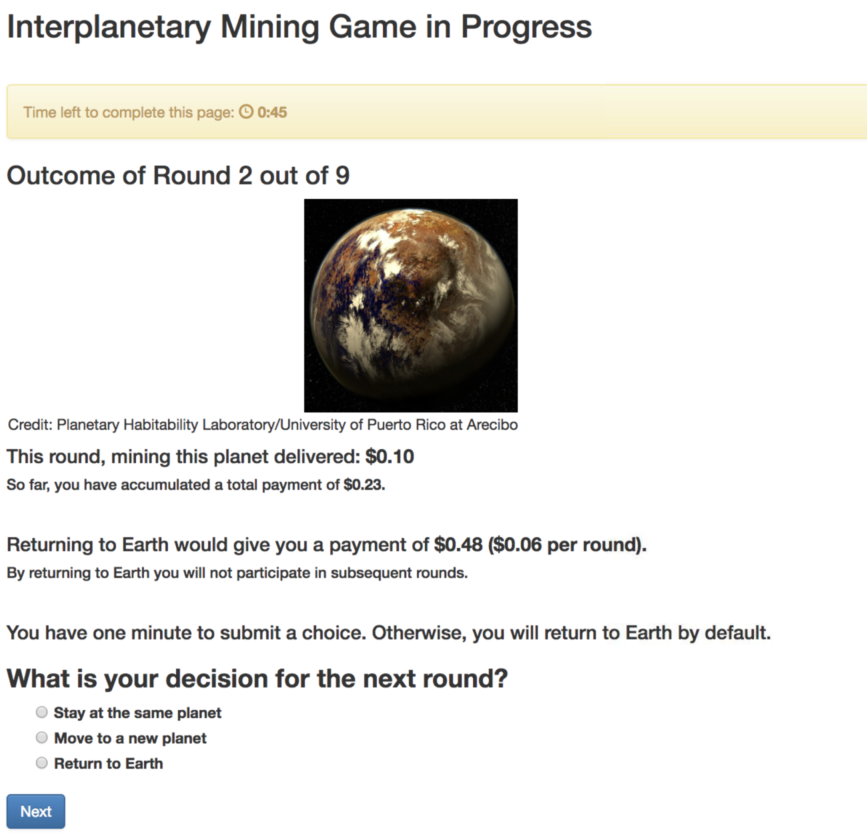

(a) Screenshot of the interface shown to participants at each round. Users are shown their last payoff, their accumulated payment so far, and their three options for the next round.

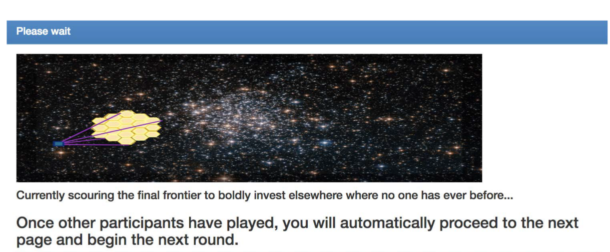

(b) Screenshot of the waiting page after participants have made a decision for each round. This is displayed until all 6 users in the game have made a decision for the current round.

Fig. 9. Screenshots of the experiments shown to participants.

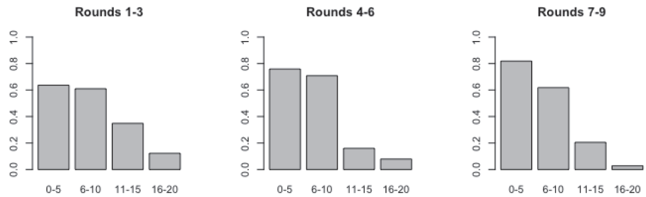

(a) SW Algorithm

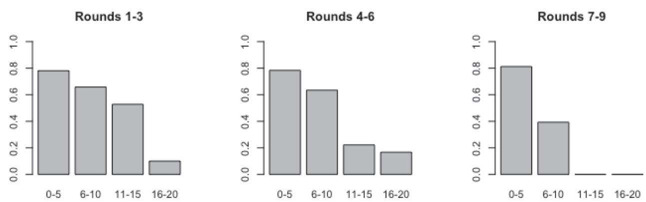

(b) Selfish Algorithm

Fig. 10. Probability of switching for Study A. Distribution of the fraction of participants who received a payoff u in round i that requested to be reassigned.

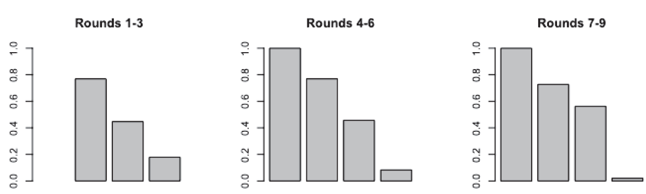

(a) SW Algorithm

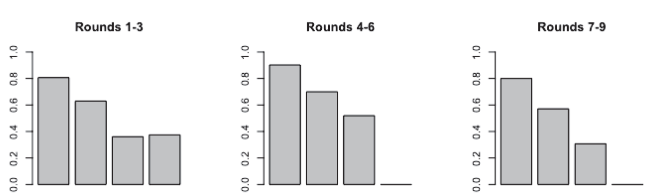

(b) Selfish Algorithm

Fig. 11. Probability of switching for Study B. Distribution of the fraction of participants who received a payoff u in round i that requested to be reassigned. Blank columns correspond to no user receiving a payoff in that corresponding range.

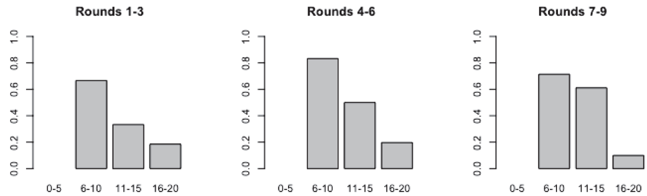

(a) SW Algorithm

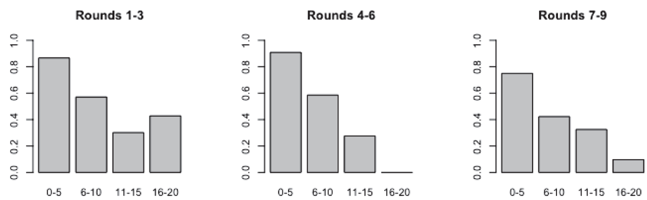

(b) Selfish Algorithm

Fig. 12. Probability of switching for Study C. Distribution of the fraction of participants who received a payoff u in round i that requested to be reassigned. Blank columns correspond to no user receiving a payoff in that corresponding range.

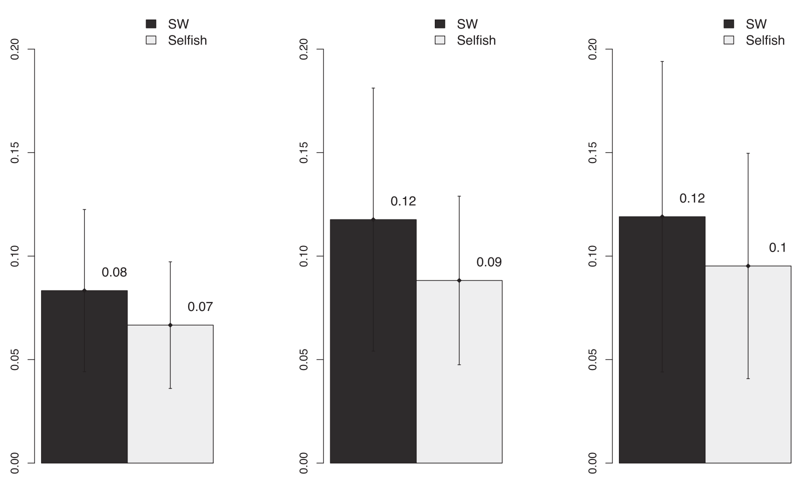

(a) Study A (b) Study B (c) Study C

Fig. 13. Drop rate. The drop rate is defined as the percentage of users who take the outside payment and thus leave the system. Error bars depict the standard error of the mean.

ACKNOWLEDGMENTS

We wish to thank Ciara Mulcahy for implementing the experiments.

REFERENCES

[1] Felipe Balmaceda, Santiago R. Balseiro, Jose R. Correa, and Nicolas E. Stier-Moses. 2010. The cost of moral hazard and limited liability in the principal-agent problem. In *International Workshop on Internet and Network Economics*. Springer, 63–74.

[2] Korte Bernhard and Jens Vygen. 2008. *Combinatorial Optimization: Theory and Algorithms*. Springer.

[3] Elizabeth E. Bruch and M. E. J. Newman. 2018. Aspirational pursuit of mates in online dating markets. *Science Advances* 4, 8 (2018), eaap9815.

[4] Jeremy Bulow. 1986. An economic theory of planned obsolescence. *Quarterly Journal of Economics* 101, 4 (1986), 729–749.

[5] John T. Cacioppo, Stephanie Cacioppo, Gian C. Gonzaga, Elizabeth L. Ogburn, and Tyler J. VanderWeele. 2013. Marital satisfaction and break-ups differ across on-line and off-line meeting venues. *Proceedings of the National Academy of Sciences* 110, 25 (2013), 10135–10140.

[6] Seyed A. Esmaeili, Sharmila Duppala, Vedant Nanda, Aravind Srinivasan, and John P. Dickerson. 2022. Rawlsian fairness in online bipartite matching: Two-sided, group, and individual. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems*. 1583–1585.

[7] Matthew Fahrbach, Zhiyi Huang, Runzhou Tao, and Morteza Zadimoghaddam. 2020. Edge-weighted online bipartite matching. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS '20)*. IEEE, 412–423.

[8] Josef Falkinger. 2007. Attention economies. *Journal of Economic Theory* 133, 1 (2007), 266–294.

[9] Eli J. Finkel, Paul W. Eastwick, Benjamin R. Karney, Harry T. Reis, and Susan Sprecher. 2012. Online dating: A critical analysis from the perspective of psychological science. *Psychological Science in the Public Interest* 13, 1 (2012), 3–66.

[10] Rachel Freedman, Jana Schaich Borg, Walter Sinnott-Armstrong, John P. Dickerson, and Vincent Conitzer. 2020. Adapting a kidney exchange algorithm to align with human values. *Artificial Intelligence* 283 (2020), 103261.

[11] Michael H. Goldhaber. 1997. The attention economy and the net. *First Monday* 2, 4 (1997).

[12] Rebecca W. Hamilton, Roland T. Rust, and Chekitan S. Dev. 2017. Which features increase customer retention. *MIT Sloan Management Review* 58, 2 (2017), 79–84.

[13] Gunter J. Hitsch, Ali Hortaçsu, and Dan Ariely. 2010. Matching and sorting in online dating. *American Economic Review* 100, 1 (2010), 130–63.

[14] Günter J. Hitsch, Ali Hortaçsu, and Dan Ariely. 2010. What makes you click? Mate preferences in online dating. *Quantitative Marketing and Economics* 8, 4 (2010), 393–427.

[15] Alan J. Hoffman and Joseph B. Kruskal. 2010. Integral boundary points of convex polyhedra. In *50 Years of Integer Programming 1958-2008*. Springer, 49–76.

[16] Zhiyi Huang, Ning Kang, Zhihao Gavin Tang, Xiaowei Wu, Yuhao Zhang, and Xue Zhu. 2018. How to match when all vertices arrive online. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*. 17–29.

[17] Toshiaki Iizuka. 2007. An empirical analysis of planned obsolescence. *Journal of Economics & Management Strategy* 16, 1 (2007), 191–226.

[18] Yash Kanoria and Daniela Saban. 2017. *Facilitating the Search for Partners on Matching Platforms: Restricting Agent Actions*. Technical Report.

[19] Jon Kleinberg, Sendhil Mullainathan, and Manish Raghavan. 2022. The challenge of understanding what users want: Inconsistent preferences and engagement optimization. In *Proceedings of the 23rd ACM Conference on Economics and Computation (EC '22)*. Association for Computing Machinery, New York, NY, 29. https://doi.org/10.1145/3490486.3538365

[20] Elias Koutsoupias and Christos Papadimitriou. 1999. Worst-case equilibria. In *Annual Symposium on Theoretical Aspects of Computer Science*. Springer, 404–413.

[21] Martin Mladenov, Elliot Creager, Omer Ben-Porat, Kevin Swersky, Richard Zemel, and Craig Boutilier. 2020. Optimizing long-term social welfare in recommender systems: A constrained matching approach. In *Proceedings of the 37th International Conference on Machine Learning*. 6987–6998.

[22] Christos H. Papadimitriou and Kenneth Steiglitz. 1982. *Combinatorial Optimization: Algorithms and Complexity*. Courier Corporation.

[23] Aditi Paul. 2014. Is online better than offline for meeting partners? Depends: Are you looking to marry or to date? *Cyberpsychology, Behavior, and Social Networking* 17, 10 (2014), 664–667.

[24] Martin L. Puterman. 1994. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley. https://doi.org/10.1002/9780470316887

[25] Iyad Rahwan, Manuel Cebrian, Nick Obradovich, Josh Bongard, Jean-François Bonnefon, Cynthia Breazeal, Jacob W. Crandall, Nicholas A. Christakis, Iain D. Couzin, Matthew O. Jackson, Nicholas R. Jennings, Ece Kamar, Isabel M. Kloumann, Hugo Larochelle, David Lazer, Richard McElreath, Alan Mislove, David C. Parkes, Alex 'Sandy' Pentland, Margaret E. Roberts, Azim Shariff, Joshua B. Tenenbaum, and Michael Wellman. 2019. Machine behaviour. *Nature* 568, 7753 (2019), 477–486.

[26] Jean-Charles Rochet and Jean Tirole. 2006. Two-sided markets: A progress report. *RAND Journal of Economics* 37, 3 (2006), 645–667.

[27] Michael J. Rosenfeld and Reuben J. Thomas. 2012. Searching for a mate: The rise of the internet as a social intermediary. *American Sociological Review* 77, 4 (2012), 523–547.

[28] Herbert A. Simon. 1996. Designing organizations for an information-rich world. *International Library of Critical Writings in Economics* 70 (1996), 187–202.

[29] Alan D. Smith. 2002. Loyalty and e-marketing issues: Customer retention on the web. *Quarterly Journal of Electronic Commerce* 3 (2002), 149–162.

[30] Aaron Whitman Smith and Maeve Duggan. 2013. *Online Dating & Relationship*. Pew Research Center, Washington, DC.

[31] Richard S. Sutton and Andrew G. Barto. 2018. *Reinforcement Learning: An Introduction.* MIT Press.

[32] Gareth Tyson, Vasile C. Perta, Hamed Haddadi, and Michael C. Seto. 2016. A first look at user activity on tinder. In *2016 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM '16).* IEEE, 461–466.

[33] Lieven Vandenberghe and Stephen Boyd. 2004. *Convex Optimization*. Vol. 1. Cambridge University Press, Cambridge.

[34] Yajun Wang and Sam Chiu-wai Wong. 2015. Two-sided online bipartite matching and vertex cover: Beating the greedy algorithm. In *Automata, Languages, and Programming: 42nd International Colloquium (ICALP '15), Proceedings, Part I 42*. Springer, 1070–1081.

[35] Ruohan Zhan, Konstantina Christakopoulou, Ya Le, Jayden Ooi, Martin Mladenov, Alex Beutel, Craig Boutilier, Ed Chi, and Minmin Chen. 2021. Towards content provider aware recommender systems: A simulation study on the interplay between user and provider utilities. In *Proceedings of the Web Conference 2021.* 3872–3883.

Received 19 August 2020; accepted 13 October 2023