When Online Dating Meets Nash Social Welfare: Achieving Efficiency and Fairness

Yongzheng Jia
Tsinghua University
jiayz13@mails.tsinghua.edu.cn

Xue Liu
McGill University
xueliu@cs.mcgill.ca

Wei Xu
Tsinghua University
wei.xu.0@gmail.com

ABSTRACT

Mobile dating applications such as Coffee Meets Bagel, Tantan, and Tinder, have become significant for young adults to meet new friends and discover romantic relationships. From a system designer's perspective, in order to achieve better user experience in these applications, we should take both the *efficiency* and *fairness* of a dating market into consideration, so as to increase the overall satisfaction for all users. Towards this goal, we investigate the nature of diminishing marginal returns for online dating markets (i.e., captured by the submodularity), and trade-off between the efficiency and fairness of the market with Nash social welfare. We further design effective online algorithms to the apps. We verify our models and algorithms through sound theoretical analysis and empirical studies by using real data and show that our algorithms can significantly improve the ecosystems of the online dating applications.

ACM Reference Format:
Yongzheng Jia, Xue Liu, and Wei Xu. 2018. When Online Dating Meets Nash Social Welfare: Achieving Efficiency and Fairness. In *WWW 2018: The 2018 Web Conference, April 23–27, 2018, Lyon, France.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3178876.3186109

1 INTRODUCTION

Online dating applications are more and more ubiquitous and becoming an integral part of young adults' everyday lives. These applications [1], such as Coffee Meets Bagel [1], Tantan [2], and Tinder [3], provide platforms for people to make new friends with various purposes including meeting new friends and developing personal or romantic relationships. The report in [4] shows that the percentage of online dating users in the USA triples (i.e., from 5% to 15%) from 2013 to 2016. Tinder consistently ranks as one of the top 10 grossing apps in Apple's online store, with more than 50 million active users [5]. It is also reported that Coffee Meets Bagel has created 997 million matches and more than 50,000 happy couples in long-standing relationships [6]. Tantan claims 60 million registered users, of which 6 million are active on a daily basis [7].

Behind the great success of the online dating apps is the *double opt-in* design, which provides the users with appealing online dating experiences: For instance, all of the above three apps present potential matches with profile cards that a user can swipe through within the app. Each profile card includes a deck of photos and an optional text-based biography. Each user has two main activities: pass (i.e., a "swipe-left") and like (i.e., a "swipe-right") on the profile cards. These dating apps only notify the users when both sides like each other (i.e., referred to as a *match*). Only matched users can start conversations. The double opt-in design turns the complex, anxious, and a sometimes awkward act of introducing oneself to another person into a simple yet playful experience [8]. This design motivates users to discover more profiles and get more matches.

To achieve better user experience in these emerging dating apps, a system designer should consider both *efficiency* and *fairness* of the dating markets. Efficiency means that the apps should make as many matches as possible in the market. This is similar to maximizing the social welfare.

In the meantime, we should also consider fairness of the dating market, which is often ignored by the designers of the apps. Most dating apps follow the *freemium strategy*, in which the basic features are free to all, while the paying users get premium services. Examples of premium services include Tinder *Boost*, Tinder *Plus* and Coffee Meets Bagel *Woo*. The dating apps give more preferences to the paying users. However, these preferences may introduce unfairness, causing the non-paying users more difficult to get matches. For example, those who use Boost will have much more opportunities to be shown to others and hence get more matches than those who do not. A more "fair" situation is that the app should help both active paying users and non-paying users get a number of matches.

Keeping the non-paying users a number of matches is important to the system designer as it leads to higher retention: We analyze the correlation between the users' retention rates and number of matches based on the data from online dating apps. We find that users with few matches (usually, these are "less attractive users") are often frustrated, hence they tend to become inactive with low retention rates (We show the detailed analysis in Section 3.2).

Each online dating application is a sophisticated ecosystem resulting from the interactions of many factors. To better understand the fairness problems discussed above, we classify the factors into two categories: the *uncontrollable* factors (i.e., mainly determined by the user's attractiveness), and the *controllable* factors (i.e., caused by the policies and algorithms used in the app). For the uncontrollable factors, there exists a natural and intrinsic unfairness as a consequence of the attractiveness for people's appearance. Moreover, it is reported in [9] that users tend to pursue attractive users regardless of their own appearance in online dating. This tendency causes that a small group of attractive users can get matches much more easily than others. For the controllable factors, Tinder and some other online dating apps can control each recommendation by showing a user's profile card to another or not. Furthermore, the dating apps

[1]As these dating applications are all mobile-phone based, we will use "apps" as the abbreviation of "applications" throughout the rest of the paper

ACM ISBN 978-1-4503-5639-8/18/04.
https://doi.org/10.1145/3178876.3186109

can also control the privileges to the paying users, so as to trade-off the revenue and the fairness of the dating market. In this work, we only focus on the discussion of the controllable factors: to improve both efficiency and fairness for the online dating ecosystems based on economic models and online algorithm design. We show that our algorithms can significantly improve the efficiency and fairness of the online dating market, and the online dating apps can use them to relieve the effect of the uncontrollable factors to provide better user experiences.

Existing studies such as [10] [11] and [12] observe the imbalanced distributions of matches on Tinder, implying the importance of fairness. However, they do not propose sound solutions to the problem.

It is challenging to design and implement an online dating market to be both efficient and fair. There are three key challenges to overcome. *First*, the objectives of efficiency and fairness do not often align with each other. It is difficult to present appropriate performance metric to trade off these objectives within one systematic framework. *Second*, the algorithms deployed for the apps should run fast enough and scale to enormous user activities. For example, Tinder processes billions of events per day, generating terabytes of data [8], hence a slow algorithm degrades the user's experience significantly. *Last but not least,* the algorithms should be *online* to deal with unpredictable user activities. The online requirement is important because it is hard to predict when the users begin/stop swiping; how many profile cards they will swipe on. Furthermore, their preferences for matches may also vary over time.

To the best of our knowledge, this is the first work to present a generalized model to achieve both efficient and fair online dating markets based on the data-driven studies, with the goal of designing fast online algorithms:

First, we present a systematic and generalized model for the dating markets to trade off the objectives between efficiency and fairness. We find the match goal based on the correlation between users' retention rates and number of matches from data of online dating apps, and discover the property of *diminishing marginal returns* for the online dating markets. We further set up match goals for different user groups, and define the utility function to measure the satisfaction for each user in the dating app. Then we present the objective function to maximize the overall satisfaction (i.e., welfare) in the market, which indicates both efficiency and fairness.

Second, by discovering the diminishing marginal returns when a user gets more and more matches, we reduce our problem to the online submodular welfare maximization problem. Then we present a $\frac{1}{2}$-competitive online greedy algorithm to solve the problem. We further show that the online greedy algorithm is effective both in theory and practice.

Third, we adapt the *Nash social welfare* to the online dating markets, which provides a natural balance between efficiency and fairness. We further reduce the problem of maximizing the Nash social welfare to a special case of the submodular welfare optimization, and adjust our online algorithm to solve the problem.

Last but not least, we present data-driven empirical studies to evaluate the performance of our model by using the data from an online dating app. To this goal, we define generalized performance metrics, as well as discuss the selection of appropriate utility functions and parameters. More interestingly, we discover an *equilibrium* when we are evaluating the performance of Nash social welfare. The equilibrium indicates a market configuration where both non-paying users and paying users are satisfied. Finally, we evaluate the improvement of applying the Nash social welfare by comparing the performance with the distributions of our dataset. The results show significant improvement for both efficiency and fairness by using Nash social welfare.

The rest of the paper is organized as follows: First, we discuss related work in Section 2 and present the problem model in Section 3. Next, we reduce the problem to the online submodular welfare maximization problem and present an efficient online algorithm to solve it in Section 4. Then we discuss how to leverage the Nash social welfare to the context of online dating markets in Section 5. Finally, we present the effectiveness of our approaches based on data-driven studies in Section 6. We conclude the paper and present the open questions for future work in Section 7.

2 RELATED WORK

The online dating market has attracted broad and interdisciplinary research interest in social networking, communications, economics and even psychology, sociology and anthropology.

Online dating applications. In recent years, the studies for online dating apps emerge both in academia and social media [6][10]. Researchers study the user motivations [13][14], social impacts [15][16], and privacy issues [17] for online dating. They also investigate the gender differences between males and females, including different selection strategies [18][12], as well as conversation behaviors [19].

Some research uses economic models to analyze user behavior for the dating markets, [20] investigates an economic matching model to explain the matching patterns and evaluate the efficiency of the matches. The authors in [9] analyze and predict the user preferences in online dating based on data-driven studies.

Some articles demonstrate the *imbalanced* distributions of the matches in online dating, and shows that it is hard for some less attractive males to get a match [10] [11]. The authors of [12] further present a hypothesis of the "feedback loop" in the online dating market: The males are forced to be *less* selective in the hope of getting a match, while females are becoming *more* selective, as they know that any males they like will result in a match with high probabilities. All these findings imply that fairness is a crucial factor to consider for the online dating apps.

Two-sided Markets: models and algorithms. The online dating market is typical two-sided (matching) market [21]. To better understand the models and challenges of the online dating market, we compare it to other two-sided markets. One is the well-studied online ride-sharing market (e.g., Uber and Lyft) [22][23][24]. Much simpler than the online dating market, the online ride-sharing market is based on a *centralized matching design*, in which the market maker (i.e., the platform) *decides* all matches. However, in the online dating market, the platform only *recommends* potential matches by showing profile cards, and all the (swiping) activities from the users are uncontrollable by the platform.

Another similar two-sided market is the online advertising market, such as Google's Adwords [25]. The authors in [26] summarizes various models and fundamental online algorithms for the online advertising market. [27] presents theoretical studies to design online ad allocation algorithms to achieve both efficiency and fairness.

Methodologies. Submodular welfare maximization is a framework for resource allocation with decreasing marginal utilities. Existing studies investigate the complexity [28] [29], offline algorithms [30] [31], and online algorithms [26] [32] [33] [34].

Nash social welfare is a sound criterion to trade off between efficiency and fairness, which is first proposed by John Nash [35]. Recently, researchers figure out the its properties in [36] [37] [38]. They also design offline algorithms to maximize the Nash social welfare with different settings in [39] [40] [41] [42]. We will discuss more details of the methodologies in Section 4 and 5.

3 MARKET, OBSERVATIONS, AND MODELS

In this section, we present the model of each user's utility (i.e., degree of satisfaction) in online dating apps based on our results from data-driven studies. We first introduce the preliminaries for the market configurations of online dating. Then we analyze the correlations between users' retention rate and number of matches by using the data from a popular online dating app, and discover the match goal (i.e., expected number of matches within a period) from our results. We further define each user's utility by his match goal and actual matches. Finally, we formulate our objective function, which maximizes the total degrees of satisfaction for all users.

3.1 Market Configurations for Online Dating

The online dating market consists of two groups, and we call them *male* and *female*. Both groups behave the same way in the market. In this paper, we only consider the matches happening between the two groups. We consider a dating market with M male and F female users. We adopt the convention $[X] = \{1, 2, \ldots, X\}$ to denote the set of X elements throughout the paper, e.g., $[M] = \{1, 2, \cdots, M\}$ is the set of all males and $[F] = \{1, 2, \cdots, F\}$ is that of all females.

To simplify the discussion, we divide time into fixed-length time slots that we call *rounds*. We consider a user *active* in round t if he/she swipes at least once in round t.

Note that as both groups are in symmetry in most of our discussions, without loss of generality, we focus on the algorithms for recommending males to females in this paper, and the other side works the same way.

We also define the number of swipes (or the number of profiles reviewed) in round t as the *(swiping) capacity* of the user in that round. We denote the capacity of a male m or a female f in round t as $c_m^{(t)}$ and $\bar{c}_f^{(t)}$, respectively.

Estimate achieved matches. To estimate the number of achieved matches for each user, we take a closer look at the match making process under the *double opt-in* model.The app has a mechanism to estimate the probability of whether female f will swipe right (i.e., like) on a male m in round t. We denote the estimation as $\bar{p}_{f,m}^{(t)}$ (and the other way around as $p_{m,f}^{(t)}$), $\forall m \in [M]$ and $\forall f \in [F]$. In practice, collaborative filtering is a good algorithm to predict the probabilities. The detail of the prediction algorithms is beyond the scope of the paper, and we only take the predicted probabilities as our input.

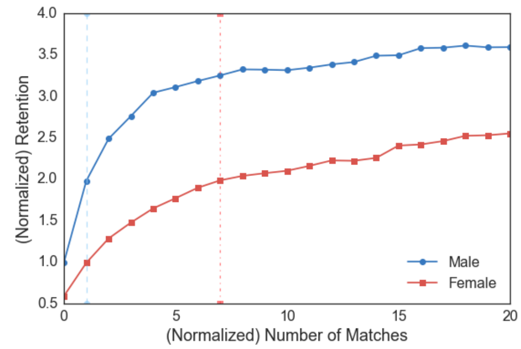

Figure 1: (Normalized) retention vs. number of matches

We further define the *match score* at round t as the product of the predicted (like) probabilities in both directions, i.e. $w_{m,f}^{(t)} = p_{m,f}^{(t)} \cdot \bar{p}_{f,m}^{(t)}$. Intuitively, the match score captures the degree of the mutual-likes between a pair of users. When we recommend male m to female f at round t, the probability of a match achieved between m and f is $w_{m,f}^{(t)}$. The *decision variable* is whether the online dating app should recommend m to f at round t, denoted as a binary variable $x_{m,f}^{(t)} \in \{0, 1\}, \forall m \in [M], \forall f \in [F]$. To sum up, we estimate the expectation of achieved match of male m at round t as follows:

$$a_m^{(t)} = \sum_{f \in [F]} w_{m,f}^{(t)} \cdot x_{m,f}^{(t)} \quad \forall m \in [M]. \tag{1}$$

3.2 Finding the Match Goal

Retentions vs. matches. To discover the insights of the online dating market, and explore how the number of matches influences users' satisfaction in online dating. We analyze the correlations between the retention rate and the number of matches per week. In practice, the retention rate is a sound and widely-used indicator, providing us with quantitative measurements of users' satisfaction, as a user tends to continue using the app if he is satisfied.

We collect two weeks' activity data for the users from a popular online dating app. For each user, we count his matches during the first week, and count his retention as a binary of whether he opens the app during the second week, that is, with 1 indicating a "yes", and 0 for "no". Then the question is: What is the probability of a user to use the app (i.e., retention rate) in the second week with a given number of matches during the first week?

Figure 1 shows the correlation between the (normalized) retention rate and the (normalized) number of matches for both male and female users[2]. The x-axis is the (normalized) number of matches received by a user in the first week, and the y-axis is the (normalized) average retention rate[3] (i.e., whether they opens the app in

[2]In this paper, some of the observations are based on data from online dating apps. However, all the data presented is normalized in order not to reveal any of the business-related details.

[3]We will use "matches" as the abbreviation of "(normalized) matches" and "retention rate" as the abbreviation of "(normalized) retention rate" in the rest of Section 3.2.

the second week) of the users with the corresponding number of matches. Due to the confidentiality of the data, we normalize the retention rate by dividing each average retention rate by males' average retention rate with no match. The two solid curves in Figure 1 illustrate the correlation between retention rates and number of matches. We limit the range of x-axis (i.e., the number of matches) to $\{0, 1, 2, \cdots, 20\}$, as the curves for both males and the females stay flat when the number of weekly matches exceeds 20. Moreover, the two vertical dotted lines show the median number of weekly matches for both males and females.

From the data, we have the following key observations based on the relationships between the retention rate and the number of weekly matches for both males and females:

- A male's retention rate increases fast when he gets a new match with less than 7 matches during the week. When a male user gets more than 7 matches. The retention rate stays stable.
- The average number of matches for a male user is about 5, and the median is 1 (i.e., half of the males have less than 1 match in the first week). If we improve a male's number of matches from 0 to 7, the retention rate will triple. If we improve it from 1 to 7, the retention rate will also improve as high as 65%, which indicates that helping each male user get one match per day, will increase the retention rate by 65% for half of the males. Meanwhile, improving the a male's number of matches from the median level to the average level will also increase his retention rate by 49%.
- However, the effect of improving the number of matches for females is not significant. The average number of matches for a female user is about 15, and the median is 7. The slope of the curve for females after the median (i.e., 7) is much more moderate. Moreover, even if we improve a female's number of matches from 7 to 20, the retention rate will only increase 27%.

Based on the above observations, we have the following conclusions: 1) For both males and females, more matches lead to a higher retention rate. 2) We find that the males' retention rate is much more sensitive than that of females in terms of number of matches, as the number of matches for females is much more optimistic. 3) Improving the weekly number of matches for each male to about 7 (i.e., we call this the *magic number* for males' matches) will promote the males' retention rate significantly. In the meantime, if a male gets more matches than the magic number, then the improvement of his retention rate is rather modest. 4) The retention curves for both males and females are *concave*, indicating the *diminishing marginal returns* when a user gradually gets more matches, and we will discuss the details in Section 4. 5) Our observations also illustrate why we care much more on males' numbers of matches, as the improvement of retention rate is moderate even though we promote females' number of matches a lot from the current level.

To give a formal definition of the observed magic number for following theoretical studies, we introduce the concept of the *match goal* for the males as follows:

Match goal. For a male $m \in [M]$ using the dating app, we define his match goal as $g_m^{(t)}$ in round t, which is his expectation for the number of matches in that round. In practice, we can set the length for each round as one day, or set the length for each round as a week if we consider the seasonality factor for a user to get a match. This is because the distribution of matches for the online dating apps varies over the days of the week, as there are both more male and female users online on weekends than weekdays. It is easier for a user to get a match on Saturday night than Monday morning. For a common non-paying user on the app, the system can set his match goal as one match per day (or 7 to 8 matches per week), that is, a non-paying user will be happy to get one match per day. For a paying user (e.g., a Tinder Boost user) or a new user, the system can set a higher match goal for him (e.g., 2 to 4 matches per day), since a paying user deserves more matches and a new user is more motivated when getting more matches when he starts to use the app. In practice, we can also dynamically adjust the match goals in different regions, as the market configurations may vary.

A user's match goal is an estimation for his expectation of matches set by the system. To capture the differences between the match goal and the actual matches of a (male) user m, we also define the *achieved matches*, $a_m^{(t)}$, for the actual number of matches for m during round t. We further define the *(match) achievement rate* $r_m^{(t)}$ as the ratio of the achieved match (i.e., $a_m^{(t)}$) to match goal (i.e., $g_m^{(t)}$) for each male m, such that $r_m^{(t)} = \frac{a_m^{(t)}}{g_m^{(t)}}, \forall m \in [M]$.

3.3 Problem Formulation

Satisfaction and utility function. We define each user's degree of satisfaction (i.e., $s_m^{(t)}$) in (2) with a utility function $u_m^{(t)}$ times a weight factor $\alpha_m^{(t)}$. The utility function is defined on $[0, +\infty)$ with input $r_m^{(t)}$. The weight factor $\alpha_m^{(t)}$ is used to distinguish the priorities for different users.

$$s_m^{(t)} = \alpha_m^{(t)} \cdot u_m^{(t)}(r_m^{(t)}) = \alpha_m^{(t)} \cdot u_m^{(t)}\Big(\frac{\sum_{f \in [F]} w_{m,f}^{(t)} x_{m,f}^{(t)}}{g_m^{(t)}}\Big). \tag{2}$$

Note that the value of each $w_{m,f}^{(t)}$ and $g_m^{(t)}$ is maintained by the system and is fixed, and therefore $s_m^{(t)}$ is only determined by the decision variables $x_{m,f}^{(t)}$. To simplify the model, we can further assume that the utility functions are *symmetric* over different males in all rounds, such that $u_m^{(t)}(\cdot) \triangleq u(\cdot), \forall m \in [M], t \in \mathbb{N}^+$.

Objective. Our objective function is to maximize the overall weighted degrees of satisfaction for males (i.e., (3)), subjected to the (swiping) capacity constraint for each female (i.e., (4a)), and the (swiping) capacity constraint for each male (i.e., (4b)):

$$max: \sum_{m \in [M]} s_m^{(t)} \tag{3}$$

s.t.,

$$\sum_{m \in [M]} x_{m,f}^{(t)} \leq \bar{c}_f^{(t)}, \quad \forall f \in [F]; \tag{4a}$$

$$\sum_{f \in [F]} x_{m,f}^{(t)} \leq c_m^{(t)}, \quad \forall m \in [M]; \tag{4b}$$

$$x_{m,f}^{(t)} \in \{0, 1\}, \quad \forall m \in [M], \forall f \in [F]. \tag{4c}$$

In practice, when we design the algorithm of recommending males to females, we can neglect the constraint for each male's capacity (i.e., (4b)). This is because it is difficult to estimate each male's capacity $c_m^{(t)}$ by only looking at the females' swiping activities in an online setting, and a male tends to spontaneously increase his (swiping) capacity if he has not get enough matches. However, even though we add constraint (4b), our online algorithm in Section 4 still works by applying the methodologies in [32].

As we discussed before, we want to mention that the mirroring model and algorithm of recommending females to males throughout this paper is theoretically correct, though practically it impacts the retention much less significantly. In the following section, for concise and practical considerations, we only discuss the algorithms for the recommendations from males to females.

4 ONLINE SUBMODULAR WELFARE MAXIMIZATION

In the previous section, we define the problem model to maximize the overall degrees of satisfaction of the market (i.e., the welfare). Based on the results in Section 3.2, the marginal returns (i.e., degree of satisfaction) for each user decrease when he gradually achieves more matches during a certain round. Moreover, when the number of his matches reaches a threshold (e.g., he can hardly keep up with all the conversations with his matches), his utility stops increasing.

4.1 Define the Submodularity

We have discovered the property of diminishing marginal returns for online dating markets, to capture this property, we introduce the *submodularity* in this section. We first define the *impression set* $I_m^{(t)}$ for each $m \in [M]$, denoting the set of females that the app recommends m to:

$$I_m^{(t)} = \{f | x_{m,f}^{(t)} = 1\} \quad \forall m \in [M]. \tag{5}$$

It is easy to find that $\forall I_m^{(t)}, I_m^{(t)} \subseteq [F]$ and $I_m^{(t)} \in 2^{[F]}$. Furthermore, we present each user's *weighted utility function on impression set* (i.e., equivalent to (2)), which is a function defined on $2^{[F]}$:

$$s_m^{(t)} = \mu_m^{(t)}(I_m^{(t)}) = \alpha \cdot u_m^{(t)}\Big(\frac{\sum_{f \in I_m^{(t)}} w_{m,f}^{(t)}}{g_m^{(t)}}\Big) \quad \forall m \in [M]. \tag{6}$$

We then present the property of *submodularity and monotonicity* for μ_m as follows:

Definition 1. *(monotone submodular) μ is* **submodular** *iff for each male $m \in [M]$ with any impression sets $\tilde{I}_m^{(t)} \subseteq I_m^{(t)}$, and a female $f \notin I_m^{(t)}$, such that:*

$$\mu_m(I_m^{(t)} \cup \{f\}) - \mu_m(I_m^{(t)}) \leq \mu_m(\tilde{I}_m^{(t)} \cup \{f\}) - \mu_m(\tilde{I}_m^{(t)}). \tag{7}$$

Furthermore, we say that μ is **monotone submodular** *if additionally, $\mu_m(I_m^{(t)}) \geq \mu_m(\tilde{I}_m^{(t)})$.*

To recap, $\mu_m(I_m^{(t)} \cup \{f\}) - \mu_m(I_m^{(t)})$ is the marginal utility of m when the app recommends him to one more female f after showing his profile to the females in his impression set $I_m^{(t)}$, and (7) presents the diminishing marginal returns for m. Additionally, m's utility is non-decreasing when his impression set grows, so μ_m is *monotone submodular.*

As our utility function is monotone submodular, our problem of maximizing the objective function (3) becomes a *submodular welfare maximization* problem. We summarize existing studies in both offline and online settings. We further present Algorithm 1 and show that it is effective and practical for our app, especially in an online setting.

Algorithm 1: Gready Algorithm for Online Submodular Welfare Mazimization - *GA*

```
Initialization: Set each I_m^(t) = ∅, ∀m ∈ [M].
When a female f ∈ [F] logs into the application at round t,
while f keeps swiping do
    (a) Select the male m* ∈ [M], such that
        m* = argmax_{m∈[M]} ( μ_m(I_m^(t) ∪ {f}) − μ_m(I_m^(t)) ).
    (b) Recommend male m* to f, I_{m*}^(t) = I_{m*}^(t) ∪ {f}
end
```

4.2 Solutions: Offline and Online

We first depict the offline and online settings for the online dating apps when we recommend males to females. In an offline setting, we know all the values of inputs such as $w_{m,f}^{(t)}$, $\bar{c}_f^{(t)}$, and μ in advance. Whereas in an online setting, the app presents real-time recommendations to each female f when she arrives with a subset of input values. Some values (e.g., $\bar{c}_f^{(t)}$, $w_{m,f}^{(t)}$) are not revealed until female f logs into the app and starts swiping, without knowing when a female becomes active/inactive, or her preference over time.

To solve the submodular welfare maximization problem, there is a greedy algorithm that is intuitive but efficient, in which each time we recommend a male m to f with maximum marginal returns, and then update his impression set. The algorithm can solve both offline and online cases. We present the online greedy algorithm in Algorithm 1.

We show that Algorithm 1 is promising both in theory and in practice, by comparing it with existing results for both offline and online settings.

Offline setting. In the offline setting, the submodular welfare maximization problem is NP-hard to approximate better than $1 - \frac{1}{e}$ [28], and [30] provides an algorithm in the value oracle model to achieve this bound. [29] proves that beating the $1 - \frac{1}{e}$ bound needs exponential communication. [31] proposes a randomized local search algorithm, which is simpler than [30], also achieving the bound. However, these algorithms have high complexity, and thus are impractical in real-life apps.

Online setting. In the online setting, [32] shows that Algorithm 1 achieves a ratio of $\frac{1}{2}$ for the online submodular welfare maximization problem even under adversarial input orders. Furthermore, in [33], the authors proved that there is no (randomized) algorithm achieving a competitive ratio of better than $\frac{1}{2}$ for the online submodular welfare maximization problem, unless $NP = RP$.

In comparison, Algorithm 1 is a solid online algorithm in theory, which guarantees a tight $\frac{1}{2}$-competitive bound for any submodular function μ_m.

Time complexity. We then analyze the time complexity for Algorithm 1. For each recommendation of female f, the algorithm goes through M iterations to select the best candidate, so it takes $M \cdot \bar{c}_f^{(t)}$ iterations for all the recommendations of female f. If we take all the females $f \in [F]$ into consideration, then the overall time complexity for Algorithm 1 is $O(M\bar{C}_F^{(t)})$, where $\bar{C}_F^{(t)}$ is the total capacities for all females in $[F]$, such that $\bar{C}_F^{(t)} = \sum_{f\in[F]} \bar{c}_f^{(t)}$. Therefore it is fast enough and practical to apply Algorithm 1 in the app.

5 NASH SOCIAL WELFARE

To further trade-off the efficiency and fairness, we adapt the Nash social welfare to the online dating markets, which provides a natural balance between the efficiency and fairness.

5.1 Definition and Properties

Eq. (8) presents the formulation of Nash social welfare for online dating markets: allocate the recommendations to maximize the geometric mean of the users' utilities (i.e., $s_m^{(t)}$).

$$\mathsf{NSW}([M]) = \left(\Pi_{m\in[M]} s_m^{(t)}\right)^{\frac{1}{M}}. \tag{8}$$

The idea dates back to John Nash's famous solution to the bargaining problem in [35]. Recently, the Nash social welfare is captured by a family of the *generalized (power) mean function* [38] with an exponent τ:

$$A_\tau([M]) = \left(\frac{1}{M}\cdot\sum_{m\in[M]}(s_m^{(t)})^\tau\right)^{\frac{1}{\tau}}. \tag{9}$$

In particular, $\mathsf{NSW}([M])$ corresponds to $A_0([M])$, which is the limit of $A_\tau([M])$ as τ goes to zero.

In order to study the trade-offs between efficiency and fairness for Nash social welfare, we extend the treatment for its properties. Two other well-studied functions captured by $A_\tau([M])$ include: (i) the *egalitarian* (i.e., max-min) objective when $\tau \to -\infty$ and (ii) *utilitarian* (i.e., average) objective when $\tau = 1$ [43]. They correspond to extreme fairness and extreme efficiency, respectively. However, $\tau \to -\infty$ may cause large inefficiency, such that the total matches in the dating market may be poor. While $\tau = 1$ neglects how unhappy some males might be, causing unfairness such that they can hardly get a match. The Nash social welfare lies between the two extremities and strikes a natural balance. This is because maximizing the geometric mean leads to more balanced recommendations (i.e., fairness), and also takes efficiency into consideration.

Futhermore, [36] presents game-theoretic properties for Nash social welfare, and proves that each allocation with the maximum Nash social welfare is both *Pareto optimal* (indicating efficiency) and *EF1* (i.e., *Envy-Freeness up to One Good*, indicating fairness).

5.2 Maximizing the Nash Social Welfare

Reduce maximizing $\mathsf{NSW}([M])$ to (3). Revisit the definition of $s_m^{(t)}$ in (2), to simplify, we first set a uniform weight parameter for all users, such that $\alpha_m^{(t)} = 1$ and $s_m^{(t)} = u_m^{(t)}(r_m^{(t)}) = r_m^{(t)}, \forall m \in [M]$. We then introduce a logarithm utility function $\tilde{s}_m^{(t)} = \tilde{u}_m^{(t)} = log(r_m^{(t)})$. Therefore maximizing the Nash social welfare (i.e., $(\Pi_{m\in[M]} s_m^{(t)})^{\frac{1}{M}}$) is equivalent to the objective function (3) by using $\tilde{s}_m^{(t)}$ (i.e., maximizing $\sum_{m\in[M]} \tilde{s}_m^{(t)}$). Since $\tilde{s}_m^{(t)} = log(s_m^{(t)})$, $r_m^{(t)} \geq 0$ and M is a constant, the reduction holds.

To consider different weight parameters, we set $s_m^{(t)} = u_m^{(t)}(r_m^{(t)}) = (r_m^{(t)})^{\alpha_m^{(t)}}$, and introduce $\tilde{s}_m^{(t)} = \alpha_m^{(t)} \cdot log(r_m^{(t)})$ and $\tilde{u}_m^{(t)}(r_m^{(t)}) = log(r_m^{(t)})$. Then we can also reduce the problem to maximize the Nash social welfare into (3) using $\tilde{s}_m^{(t)}$ (i.e., $\sum_{m\in[M]} \tilde{s}_m^{(t)}$) as $\tilde{s}_m^{(t)} = log(s_m^{(t)})$. Since each $\tilde{s}_m^{(t)}$ is a monotone submodular utility function, we can reduce the problem of maximizing Nash social welfare as a special case of online submodular welfare maximization problem if we set $s_m^{(t)} = (r_m^{(t)})^{\alpha_m^{(t)}}$ in (8) for each user.

Note that this setup is reasonable for real-life dating apps. For instance, we set $\alpha_m^{(t)} > 1$ for the paying users and $\alpha_m^{(t)} = 1$ for the non-paying users, indicating that the app takes on a higher priority for each paying user. Therefore when we maximize the Nash social welfare in (8), we are prone to first make $r_m^{(t)}$ as large as possible for the paying users, since the value of $s_m^{(t)}$ grows faster when $\alpha_m^{(t)} > 1$. We can further adjust $\alpha_m^{(t)}$ to set different priorities for paying users.

Utility caps. However, if there exist a male m with $r_m^{(t)} > 1$ and $\alpha_m^{(t)} > 1$, it will reduce fairness of the market to further increase $r_m^{(t)}$ to maximize Nash social welfare. To tackle this problem and further improve the fairness, it is practical to introduce utility caps by setting $r_m^{(t)} = 1$ when $\frac{a_m^{(t)}}{g_m^{(t)}} > 1$, which makes $s_m^{(t)} = (r_m^{(t)})^{\alpha_m^{(t)}} \leq 1$. Thus for a male reaching his match goal, making more match for him will not improve the Nash social welfare any more.

Offline solution. Recent years, designing approximation algorithms to maximize the Nash social welfare attracts lots of research interest. [37] shows that maximizing Nash social welfare for indivisible items with additive utilities is both NP-hard and APX-hard. [39] presents offline approximation algorithms to maximize the Nash social welfare with additive utilities, and [40] further handle the cases with utility caps. However, note that these algorithms apply the Eisenberg-Gale program [44] to achieve a fractional optimal solution, followed by a carefully designed rounding algorithm. Thus they all need the values of all inputs (e.g., $w_{m,f}^{(t)}$, $\bar{c}_f^{(t)}$) in advance and can not be directly applied in the online settings for dating apps.

Discussions on the online solution. It is worth noting that existing work mostly studies the offline solutions to maximize the Nash social welfare, which is not applicable to the dating markets as we discussed in Section 1 for the requirement of online algorithms.

In this work, our contribution is to reduce the Nash social welfare maximization to (3) based on the reductions above, and use Algorithm 1 to get an online solution. In practice, we set $\tilde{u}_m^{(t)}(r_m^{(t)}) = log(\epsilon + r_m^{(t)})$, where $\epsilon = 10^{-4}$ or other small values. This is because $log(0) = -\infty$ and if there exists a male m with $a_m^{(t)} = 0$, the output is always $-\infty$, hence Algorithm 1 makes no improvements in initial steps. Adding a small ϵ will tackle this dilemma, letting Algorithm 1 make improvements in each iteration.

6 DATA-DRIVEN EMPIRICAL STUDIES

In this section, we conduct evaluations based on the combined models and techniques discussed above.

6.1 Evaluation Setups

We first describe the dataset configuration for our evaluations. We use the distribution based on the user's swiping data from an online dating app during one week in a small region. We also normalize the number of matches in our evaluation like what we do in Section 3.2. There are about 3,800 (i.e., $M = 3,800$) male active users[4] and 1,700 (i.e., $F = 1,700$) female active users using the app. We partition the male users in $[M]$ into two user groups, with $[M_{\text{npu}}]$ denoting the M_{npu} non-paying users, and $[M_{\text{pay}}]$ denoting the M_{pay} paying users. We define the paying rate $\gamma = \frac{M_{\text{pay}}}{M}$, and we calculate $\gamma \approx 0.26$ based on the dataset we use. We setup uniform match goal for the male users in each group, using g_{pay} to denote the match goal for each paying user all the time (i.e., $g_m^{(t)} = g_{\text{pay}}, \forall m \in [M_{\text{pay}}]$) and g_{npu} for the match goal of each non-paying user. We set $g_{\text{npu}} = 7$, and use $\eta = \frac{g_{\text{pay}}}{g_{\text{npu}}}$ to capture the *goal gap* between the paying and non-paying users. We can change η to adjust the added value for premium services, and increasing η will decrease the fairness of the market.

The preference prediction scores (i.e., $p_{m,f}^{(t)}, \bar{p}_{f,m}^{(t)}$) come from the backend recommendation system of the app. Then we calculate each match score (i.e., $w_{m,f}^{(t)}$) and estimate the mean of a match score (i.e., $\mathbb{E}_w$). Deriving from the system, we get $\mathbb{E}_w = 0.05$. To measure the ratio between the supply (i.e., expected matches from females' swipes) and demand (i.e., sum of match goals for all males), we define $\Psi_{[M]}^{(t)}$ as the *expected overall happiness* for the males in the market at round t. We calculate the value of $\Psi_{[M]}^{(t)}$ by dividing the total expected matches achieved by the total match goals for all males:

$$\Psi_{[M]}^{(t)} = \frac{\mathbb{E}_w \cdot \sum_{f \in [F]} \bar{c}_f^{(t)}}{M \cdot \left(\gamma \cdot g_{\text{pay}} + (1-\gamma) \cdot g_{\text{npu}}\right)}. \tag{10}$$

We have $\Psi_{[M]}^{(t)} \in (0,1)$, as the limited capacities for females make it difficult to reach each male's match goal. The value of $\Psi_{[M]}^{(t)}$ also varies over time. In our traces, the value of $\Psi_{[M]}^{(t)}$ is mostly in $[0.3, 0.7]$ for various market configurations, and 0.5 is a good estimation for common cases.

6.2 Performance Metrics

To design better online algorithm using Algorithm 1, we may try different monotone submodular utility functions and parameters. Therefore we need a set of performance metrics (i.e., indicators), each of which is independent of the utility functions and parameters, to evaluate both the efficiency and fairness of the market.

Efficiency indicators. We use the *uniform happiness indicators* to measure the efficiency of the market. For each male $m \in [M]$, his uniform happiness (i.e., denoted as $h_m^{(t)}$) is $r_m^{(t)}$ if $r_m^{(t)} \leq 1$, and $h_m^{(t)} = 1$ if $r_m^{(t)} > 1$. The happiness indicator for all users is $H_{[M]}^{(t)} = \frac{1}{M} \cdot \sum_{m \in [M]} h_m^{(t)}$, which implies the ratio of satisfied users in the dating market. In the same way, we also define the happiness indicator for paying users (i.e., $H_{[M_{\text{pay}}]}^{(t)}$) and for non-paying users (i.e., $H_{[M_{\text{npu}}]}^{(t)}$).

[4]If a user logs in and swipes at least once, then we count him/her as an active user.

Fairness Indicators. To evaluate the overall fairness for the male users, we define the match fairness (i.e., $J_{[M]}^{(t)}$) and impression fairness (i.e., $\tilde{J}_{[M]}^{(t)}$) based on *Jain's fairness index* [45]. The metric indicates the fairness for the number of matches for each male and the cardinality of each male's impression set. Jain's index provides a promising system-level overview of the fairness, which is both easy to compute (i.e., has an explicit form) and normalized (i.e., ranges in $[0, 1]$). Therefore it is more appropriate to our model than other fairness indicators such as Gini coefficient [46], max-min fairness [45], or α-fairness [45]. Based on the definition of Jain's index, we calculate $J_{[M]}^{(t)}$ and $\tilde{J}_{[M]}^{(t)}$:

$$J_{[M]}^{(t)} = \frac{\left(\sum_{m \in [M]} a_m^{(t)}\right)^2}{M \cdot \left(\sum_{m \in [M]} (a_m^{(t)})^2\right)} \tag{11}$$

$$\tilde{J}_{[M]}^{(t)} = \frac{\left(\sum_{m \in [M]} |I_m^{(t)}|\right)^2}{M \cdot \left(\sum_{m \in [M]} |I_m^{(t)}|^2\right)} \tag{12}$$

6.3 Choosing the Utility Function

Different utility functions imply various trade-offs between efficiency and fairness. However, since the match goal in each group varies, if we recommend male $m \in [M_{\text{npu}}]$ and $m' \in [M_{\text{pay}}]$ to female f with $w_{m,f}^{(t)} = w_{m',f}^{(t)}$, the increment in $r_m^{(t)}$ and $r_{m'}^{(t)}$ varies, and therefore the marginal utility may be biased. To eliminate the bias, we introduce the *indifferent coefficient* (i.e., denoted as β), which depends on the utility function u and the goal gap λ, but is independent of achieved matches $a_m^{(t)}$.

Definition 2. **(Indifferent Coefficient)** *For a differentiable and monotone submodular utility function u and goal gap η, if there exists a constant β, such that:*

$$u'\left(\frac{a}{g_{\text{npu}}}\right) = \beta \cdot u'\left(\frac{a}{\eta \cdot g_{\text{npu}}}\right), \quad \forall a \geq 0, \tag{13}$$

and we define β as the indifferent coefficient of u and η.

While β indicates the same accelerated speed for marginal utility gained for different user groups, note that not every monotone submodular function has an indifferent coefficient. For instance, the *tanh* function (i.e., $u(r) = tanh(r)$) does not have an indifferent coefficient. It is easy to verify that of power functions (i.e., $u(r) = r^\tau, \forall r \in [0, +\infty)$) with exponent $\tau \in (0, 1]$ are all monotone submodular with $\beta = \eta^\tau$. Additionally, the logarithmic utility function (i.e., $u(r) = log(r)$, reduced from the Nash social welfare) has a consistent indifferent coefficient $\beta = 1$.

Priority parameter. Let us revisit the definition for the weight parameters $\alpha_m^{(t)}$ in (2). If we set $\alpha_m^{(t)} = 1, \forall m \in [M_{\text{npu}}]$, and $\alpha_m^{(t)} = \beta, \forall m \in [M_{\text{pay}}]$, then the non-paying users and paying users will have the same accelerated speed to get a match. In practice, to make the premium services more valuable, the app needs to help the paying users get a match more quickly, otherwise, it is hard for a paying user to reach his match goal with the same speed. Therefore we introduce the *priority parameter* for the paying users (i.e., λ), indicating the level of accelerated speed for paying users, comparing to the non-paying users. Then we set $\alpha_m^{(t)} = \beta \cdot \lambda, \forall m \in [M_{\text{pay}}]$. In this way, we decompose the value for each weight coefficient $\alpha_m^{(t)}$ into two independent components: one component is the indifferent coefficient β that only depends on the utility function (i.e., a constant

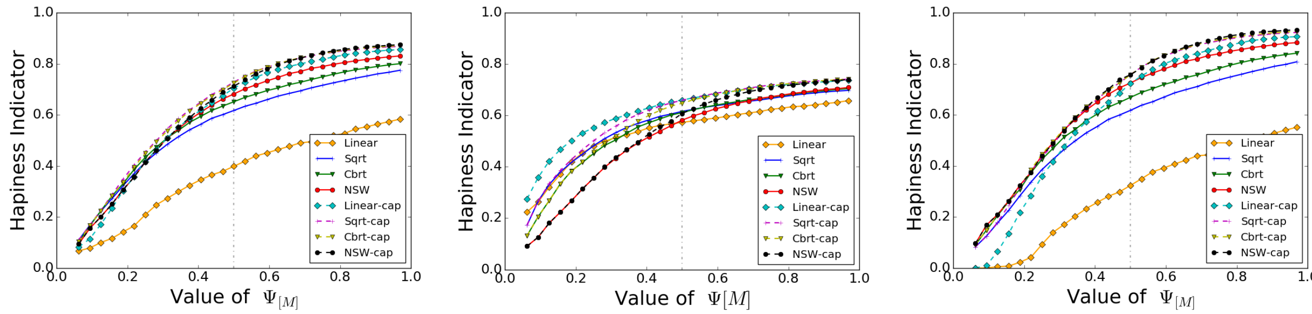

(a) All users (b) Paying users (c) Non-paying users

Figure 2: Happiness indicators for different user groups.

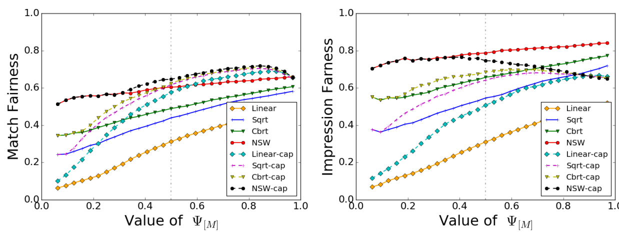

(a) Match fairness (b) Impression fairness

Figure 3: Impression fairness for all users.

part), and the other component is the priority, measuring the level of privileges to the paying users (i.e., the tunable parameter). In practice, we can dynamically adjust the value of λ to strike a good balance between the happiness for non-paying and paying users, (i.e., like tuning a *hyperparameter* in machine learning algorithms). We will show the details about how to tune it in Section 6.4.

Evaluate the functions. Then we evaluate a set of monotone submodular utility functions, where each of them has an indifferent coefficient β. We choose `Linear` (i.e., $u(r) = r$), `Sqrt` (i.e., $u(r) = r^{\frac{1}{2}}$), `Cbrt` (i.e., $u(r) = r^{\frac{1}{3}}$), and `NSW` (i.e., $u(r) = log(r)$), as well as these functions incorporating with utility caps, (i.e., denoted as `Linear-cap`, `Sqrt-cap`, `Cbrt-cap`, and `NSW-cap`).

To show the performance of different utility functions, we fix $\eta = 3$, $\lambda = 3$, and vary the value of $\Psi^{(t)}_{[M]} \in (0, 1)$, to evaluate the results with various supply-demand conditions.

Figure 2(a) to Figure 2(c) show the efficiency indicators for the functions. For overall happiness and non-paying users' happiness, both `NSW` and `NSW-cap` perform well and `NSW` beats all the other functions without utility caps. As `NSW` and `NSW-cap` care more about the overall happiness, especially the non-paying users, sometimes they do not perform well on the happiness for paying users. This is also because $\lambda = 3$ is not an appropriate parameter, and we show how to choose the best λ under different η and $\Psi^{(t)}_{[M]}$ in Section 6.4.

Figure 3(a) and Figure 3(b) demonstrate the performance for fairness. We see that `NSW` beats all the other functions on impression fairness $\tilde{J}^{(t)}_{[M]}$. For match fairness $J^{(t)}_{[M]}$, `NSW-cap` beats all the other functions, and `NSW` beats the other functions without utility caps.

6.4 Parameters and Equilibriums

From the results above, we select `NSW` and `NSW-cap` to further evaluate how to select appropriate parameters.

Tuning the priority parameter λ. We first fix $\Psi^{(t)}_{[M]} = 0.5$ and $\eta = 3$. To capture how λ impacts the priority for the paying users in reality, comparing to the expected goal gap η, we introduce the deviation factor δ, denoting the real gap (i.e., for $a^{(t)}_m$) between paying and non-paying users divided by η.

For `NSW` (i.e., Figure 5(a)), easy to find that the fairness indicators both go down when λ increases. A good value for λ is around 5, where $\delta \approx 1$, and the values of $H^{(t)}_{[M]}$, $H^{(t)}_{[M_{\text{npu}}]}$, $H^{(t)}_{[M_{\text{pay}}]}$ are almost equal (i.e., around 0.64). This interesting phenomenon shows an *equilibrium* (i.e., when $\lambda = \lambda^*$) between non-paying users and paying users, where the three happiness indicators reach around an *equilibrium happiness* (i.e., denoted as H^*). The equilibrium provides a promising selection of λ (i.e., λ^*), as when $\lambda < \lambda^*$, the paying users are not satisfied. Additionally, when $\lambda > \lambda^*$, the happiness of non-paying users drops faster than the happiness of paying users can increase.

For `NSW-cap` (i.e., Figure 5(b)), we also find the equilibrium when $\lambda \approx 6$ with $\delta \approx 1$ and $H^* \approx 0.66$, indicating that 66% of both the non-paying and paying users are happy, which is better than the equilibrium happiness of `NSW`.

Finding equilibriums. Then we find the equilibriums under different configurations (i.e., $\Psi^{(t)}_{[M]}$ and η). We show the results in Table 1 (`NSW`) and Table 2 (`NSW-cap`).

Table 1: Finding equilibriums for `NSW`.

$\Psi^{(t)}_{[M]}$	η	λ^*	H^*	$J^{(t)}_{[M]}$	$\tilde{J}^{(t)}_{[M]}$
0.3	2.0	2.5	0.56	0.62	0.83
0.3	3.0	4.0	0.47	0.51	0.67
0.3	4.0	5.5	0.40	0.45	0.60
0.5	2.0	3.0	0.72	0.62	0.80
0.5	3.0	5.0	0.64	0.50	0.65
0.5	4.0	7.0	0.56	0.44	0.57
0.7	2.0	3.5	0.80	0.61	0.77
0.7	3.0	5.5	0.73	0.51	0.62
0.7	4.0	8.5	0.71	0.44	0.55

When we fix the value of η, increasing $\Psi^{(t)}_{[M]}$ will result in higher equilibriums for both λ^* and H^*. This is because improving the supply (i.e., total capacities of females) will make it easy to get a match for each user, so as to improve the happiness indicators.

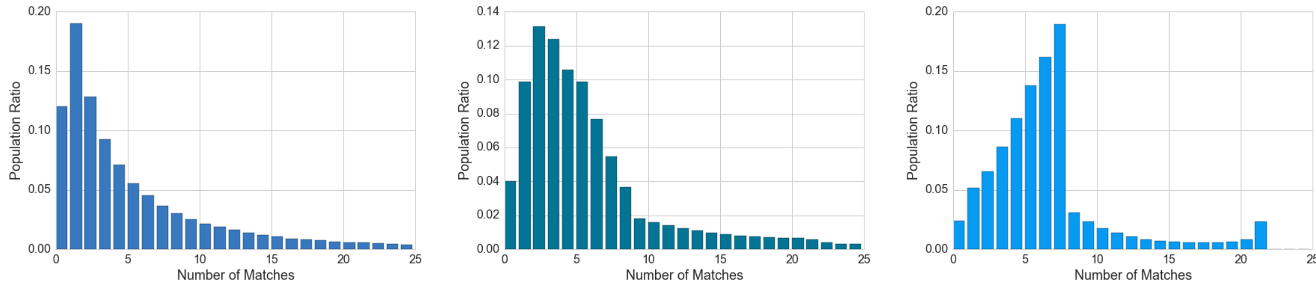

(a) Dataset (b) NSW (c) `NSW-cap`

Figure 4: Match distribution for the males.

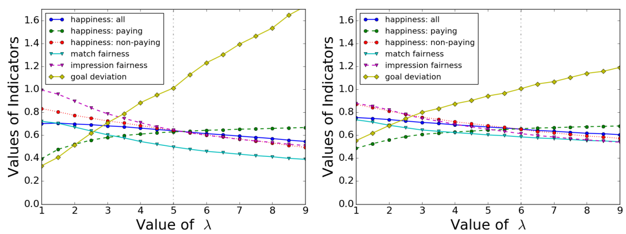

(a) NSW (b) `NSW-cap`

Figure 5: Indicators for `NSW` and `NSW-cap` with different λ.

Table 2: Finding equilibriums for `NSW-cap`.

$\Psi^{(t)}_{[M]}$	η	λ^*	H^*	$J^{(t)}_{[M]}$	$\tilde{J}^{(t)}_{[M]}$
0.3	2.0	3.0	0.57	0.66	0.79
0.3	3.0	4.5	0.48	0.52	0.66
0.3	4.0	5.5	0.40	0.45	0.60
0.5	2.0	4.5	0.75	0.75	0.69
0.5	3.0	6.0	0.66	0.59	0.62
0.5	4.0	7.5	0.57	0.48	0.55
0.7	2.0	10.0	0.82	0.80	0.58
0.7	3.0	12.0	0.74	0.63	0.54
0.7	4.0	14.0	0.71	0.51	0.50

Furthermore, since non-paying user can get a match more easily, we need to raise λ^* to make $\delta \approx 1$.

Additionally, when we fix the value of $\Psi^{(t)}_{[M]}$, we see that a larger η will increase λ^*. While increasing η leads to a more imbalanced market, the fairness indicators and H^* both go down, as non-paying users will get less match and therefore unsatisfied.

6.5 Evaluation of the Performance Gains

In this section, we evaluate the performance gains of using Nash social welfare maximization (i.e., `NSW` and `NSW-cap`). From the dataset and the differences between paying and non-paying in the region, we calculate that $\Psi^{(t)}_{[M]} = 0.5$, $\eta = 3$, and $g_{\text{pay}} = \eta \cdot g_{\text{npu}} = 21$. By the definitions in (11) and (12), we calculate that the match fairness for the males in the dataset is 0.297, and the impression fairness is 0.342 by using Jain's fairness indicators.

We then plot the match distribution for our dataset (Figure 4(a)), as well as for the result of `NSW` (Figure 4(b)) and `NSW-cap` (Figure 4(c)). The x-axis is the number of matches, and the y-axis is the population ratio given a fixed number of matches. We also zoom in the x-axis to $\{0, 1, 2, \cdots, 24\}$, as $g_{\text{npu}} = 7$ and $g_{\text{pay}} = 21$, and the population ratio is very small when the matches exceed g_{pay}.

For `NSW-cap`, the match fairness can improve by 98.7% to 0.59 comparing to the matches of the dataset, and the impression fairness increases 81.3%. Furthermore, we estimate that the average number of matches of the males will increase 32.0% for `NSW-cap` by using H^* divided by $\Psi^{(t)}_{[M]}$. From Figure 4(c), we can observe two peaks of population ratios for 7 (i.e., g_{npu}) and 21 matches (i.e., g_{pay}), and also find that the population ratio becomes 0 when the number of matches exceeds 21. These are all caused by the *utility caps*.

For `NSW` (i.e., without introducing the utility caps), it increases 68.4% of the match fairness and 28.0% of the average number of matches, which underperforms the `NSW-cap`. Whereas it increases the impression fairness by 90.1%, which outperforms the `NSW-cap`. Furthermore, as `NSW` does not introduce the utility caps, we do not see clear peaks at g_{npu} and g_{pay} from Figure 4(b).

7 CONCLUSION AND FUTURE WORK

Modern mobile dating applications provide appealing double opt-in design to make the match making simple and enjoyable. We strongly believe that a better trade-off between the efficiency and fairness of the dating markets can further improve the users' satisfactions, removing the frustration when a user can not obtain a match or his desired number of matches. This paper focuses on the discussion on designing and implementing efficient and fair online dating markets. We discuss the submodularity and leverage the Nash social welfare to trade off between the two objectives through both theoretical analysis and empirical studies.

This work leads to many open questions to future directions to build a more efficient and fair online dating market: 1) We have not discussed the *uncontrollable* factors affecting the fairness in this work, and we can classify the users based on their attractiveness levels, and design better recommendation algorithms for each level of users. 2) We mentioned that the females' retention rate is not much sensitive to the number of matches, and we want to design better algorithms or policies to improve the females' retention rate.

ACKNOWLEDGMENT

This research is supported in part by the Tsinghua Initiative Research Program Grant 20151080475. We thank Maria Renhui Zhang for the insightful discussions.

REFERENCES

[1] *Coffee Meets Bagel*. https://coffeemeetsbagel.com/.
[2] *TanTan*. https://tantanapp.com/en/.
[3] *Tinder*. https://www.gotinder.com/.
[4] 15% of American Adults Have Used Online Dating Sites or Mobile Dating Apps. *Retrived from http://www.pewinternet.org/2016/02/11/15-percent-of-american-adults-have-used-online-dating-sites-or-mobile-dating-apps/*, 2016.
[5] Marie Claire. Tinder: The Online Dating App Everyone's Talking About . *Retrived from http://www.marieclaire.co.uk/life/sex-and-relationships/tinder-the-online-dating-app-that-everyone-s-talking-about-112522*, 2017.
[6] Cosmo's Dating App School. *Cosmopolitan*, 2017.
[7] Tantan, China's top dating app, raises $70M for monetization push and overseas expansion. *Retrived from https://techcrunch.com/2017/06/21/chinas-tantan-raises-70-million/*, 2017.
[8] Maria Zhang. Retrived from http://tech.gotinder.com/building-tinder/. In *Tinder Tech Blog: Building Tinder*, 2016.
[9] Günter J. Hitsch, Ali Hortaçsu, and Dan Ariely. What makes you click? mate preferences in online dating. *Quantitative marketing and Economics*, 2010.
[10] Medium. Retrived from https://medium.com/@worstonlinedater/tinder-experiments-ii-guys-unless-you-are-really-hot-you-are-probably-better-off-not-wasting-your-2ddf370a6e9a. In *Tinder Experiments II: Guys, unless you are really hot you are probably better off not wasting your time on Tinder – a quantitative socio-economic study*, 2015.
[11] Jessica Strubel and Trent A Petrie. Love me tinder: Body image and psychosocial functioning among men and women. *Body Image*, 2017.
[12] Gareth Tyson, Vasile C. Perta, Hamed Haddadi, and Michael C. Seto. A first look at user activity on tinder. In *Proceedings of IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, 2016.
[13] S. R. Sumter, L. Vandenbosch, and L. Ligtenberg. Love me Tinder: Untangling emerging adults motivations for using the dating application Tinder. *Telematics and Informatics*, 2017.
[14] Christopher J. Carpenter and Bree McEwan. The players of micro-dating: Individual and gender differences in goal orientations toward micro-dating apps. *First Monday*, 2016.
[15] J. Ward. What are you doing on Tinder? impression management on a matchmaking mobile app. *Information, Communication & Society*, 2016.
[16] Gábor Orosz, István Tóth-Király, Beáta Bőthe, and Dóra Melher. Too many swipes for today: The development of the problematic tinder use scale. *Journal of Behavioral Addictions*, 2016.
[17] Camille Cobb and Tadayoshi Kohno. How Public Is My Private Life? Privacy in Online Dating. In *Proceedings of the 26th International Conference on World Wide Web (WWW)*, 2017.
[18] Olga Abramova, Annika Baumann, and Hanna Krasnova. Gender differences in online dating: What do we know so far? . In *Hawaii International Conference on System Sciences (HICSS)*, 2016.
[19] Jennie Zhang and Taha Yasseri. What happens after you both swipe right: A statistical description of mobile dating communications. *arXiv preprint arXiv:1607.03320*, 2016.
[20] Günter J. Hitsch, Ali Hortaçsu, and Dan Ariely. Matching and Sorting in Online Dating. *The American Economic Review*, 2010.
[21] Jean-Charles Rochet and Jean Tirole. Two-Sided Markets: An Overview. *mimeo: MIT*, 2004.
[22] M. Keith Chen. Dynamic Pricing in a Labor Market: Surge Pricing and Flexible Work on the Uber Platform. In *Proceedings of the ACM Conference on Economics and Computation*, 2016.
[23] Zhixuan Fang, Longbo Huang, and Adam Wierman. Prices and Subsidies in the Sharing Economy. In *Proceedings of the 26th International Conference on World Wide Web (WWW)*, 2017.
[24] Yongzheng Jia, Wei Xu, and Xue Liu. An Optimization Framework For Online Ride-sharing Markets. In *IEEE International Conference on Distributed Computing Systems (ICDCS)*, 2017.
[25] *Google Adwords*. https://adwords.google.com.
[26] Aranyak Mehta. Online matching and ad allocation. *Foundations and Trends in Theoretical Computer Science*, 2013.
[27] Jon Feldman, Monika Henzinger, Nitish Korula, Vahab S. Mirrokni, and Cliff Stein. Online Stochastic Packing Applied to Display Ad Allocation. In *European Symposium on Algorithms (ESA)*, 2010.
[28] Subhash Khot, Richard J. Lipton, Evangelos Markakis, and Aranyak Mehta. Inapproximability results for combinatorial auctions with submodular utility functions. In *Proceedings of the international conference on Internet and Network Economics (WINE)*, 2005.
[29] Vahab Mirrokni, Michael Schapira, and Jan Vondrak. Tight information-theoretic lower bounds for welfare maximization in combinatorial auctions. In *Proceedings of the conference on Electronic commerce (EC)*, 2008.
[30] Jan Vondrak. Optimal approximation for the submodular welfare problem in the value oracle model. In *Proceedings of the annual ACM symposium on Theory of computing (STOC)*, 2008.
[31] Yuval Filmus and Justin Ward. A tight combinatorial algorithm for submodular maximization subject to a matroid constraint. In *Proceedings of the IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, 2012.
[32] Benny Lehmann, Daniel Lehmann, and Noam Nisan. Combinatorial auctions with decreasing marginal utilities. In *Proceedings of the ACM conference on Electronic Commerce (EC)*, 2001.
[33] Michael Kapralov, Ian Post, and Jan Vondrák. Online submodular welfare maximization: Greedy is optimal. In *Proceedings of the annual ACM-SIAM symposium on discrete algorithm (SODA)*, 2013.
[34] Nitish Korula, Vahab Mirrokni, and Morteza. Online submodular welfare maximization: Greedy beats 1/2 in random order. In *Proceedings of the annual ACM symposium on Theory of computing (STOC)*, 2015.
[35] John Nash. The Bargaining Problem. *Econometrica 18, 2*, 1950.
[36] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. The unreasonable fairness of maximum nash welfare. In *Proceedings of the ACM Conference on Economics and Computation (EC)*, 2016.
[37] Euiwoong Lee. APX-Hardness of Maximizing Nash Social Welfare with Indivisible Items. *Information Processing Letters 122, 17-20*, 2017.
[38] Simina Branzei, Vasilis Gkatzelis, and Ruta Mehta. Nash social welfare approximation for strategic agents. In *Proceedings of the ACM Conference on Economics and Computation (EC)*, 2017.
[39] Richard Cole and Vasilis Gkatzelis. Approximating the nash social welfare with indivisible items. In *Proceedings of the annual ACM symposium on Theory of computing*, 2015.
[40] Jugal Garg, Martin Hoefer, and Kurt Mehlhorn. Approximating the nash social welfare with budget-additive valuations. In *Arxiv: 1707.04428v1*, 2017.
[41] Siddharth Barman, Sanath Kumar Krishnamurthy, and Rohit Vaish. Finding fair and efficient allocations. In *Arxiv: 1707.04731*, 2017.
[42] Sara Ramezani and Ulle Endriss. Nash Social Welfare in Multiagent Resource Allocation. *Lecture Notes in Business Information Processing*, 2010.
[43] Hervé Moulin. Fair Division and Collective Welfare. *Public Choice*, 2004.
[44] Edmund Eisenberg and David Gale. Consensus of Subjective Probabilities: The Pari-Mutuel Method. *The Annals of Mathematical Statistics 30, 1*, 1959.
[45] Tian Lan, David Kao, Mung Chiang, and Ashutosh Sabharwal. An axiomatic theory of fairness in network resource allocation. In *Proceedings of the IEEE International Conference on Computer Communications (INFOCOM)*, 2010.
[46] Vinod Thomas, Yan Wang, and Xibo Fan. Measuring education inequality : Gini coefficients of education. *Policy Research Working Papers*, 2001.