Search, Selectivity, and Market Thickness in Two-Sided Markets

Jessica Yu*

For the most up-to-date version, please email the author at jhyu1@stanford.edu

March 6, 2019

Abstract

This paper investigates search and matching in online marketplaces, emphasizing how user behavior responds to the presence of others on the platform, which I call "market thickness". Unlike standard settings in which firms typically benefit from increasing their customer base, in two-sided markets, changes in market thickness can induce complex effects in matching due to the endogenous adjustment of search and selectivity. Motivated by the observed correlation between individual's selectivity and the number of potential matches and competitors in this market, this paper causally measures the impact of market thickness on behavior and explores its implications for platform design in the context of online dating. I design and implement a field experiment that varies information sent to platform participants about the number of potential matches (market size) and number of competitors. Consistent with intuition and observational patterns, the experiment shows that individuals become more selective when they believe they have more potential matches, and less selective when they believe they have more competition. The effect of changes in selectivity on matching is however an equilibrium outcome. I therefore use the exogenous variation generated by the experiment to identify the parameters of a microfounded model, which then allows me to estimate the equilibrium and evaluate platform-design-linked counterfactuals. I find that when accounting for changes in selectivity in response to market thickness, matching exhibits decreasing returns to scale, and an increase in market size does not necessarily increase match quality. I then show how changing selectivity through a platform design feature can mitigate negative effects of changes in market and competition size.

1 Introduction

This paper investigates search and matching on decentralized platforms with a goal of improving platform design. Decentralized platforms are two-sided marketplaces in which individuals search to find matches, which are formed only upon the explicit agreement of both parties in response to

*Doctoral student in Marketing, Stanford Graduate School of Business, Email: jhyu1@stanford.edu. I am grateful to my advisors, Harikesh Nair and Wes Hartmann, and my committee, Navdeep Sahni and Bart Bronnenberg, for their guidance and support. Thanks also to Pedro Gardete, James Lattin, Sridhar Narayanan, James Sams, and the participants at the Stanford Marketing seminar. Financial support was provided by the firm. The paper presents the views of the author, not of the firm. All remaining errors are my own.

proposals. In contrast, centralized markets assign matches to agents in a centralized mechanism, obviating both the need for search, and the uncertainty from agents' perspective of whether match proposals will be accepted. Decentralized platforms are ubiquitous. An example is job search. Workers and firms search to fill vacant jobs, and it is costly for workers to apply to jobs and for firms to conduct interviews. Another example is dating. In order to go on a date, the man or woman needs to propose a date, and in order for the date to occur, both the man and woman need to agree. Since matches are determined through an agent's search behavior, understanding the nature of search is critical to platform design. I design and implement a field experiment on a popular dating app, which I combine with a microfounded, econometric model of search in order to evaluate the impact of platform size and design on matching.

Understanding agent behavior in two-sided markets is challenging in that an agent's actions are based off beliefs about other agents' behavior. Because search and match proposals are costly, agents are what I call *strategically selective*. An agent's decision to propose a match with another agent depends on (1) his beliefs about whether the other agent will accept the match, and (2) the likelihood that he will find another match if the current potential match is not realized. The likelihood of finding another match depends on the availability of other agents on the platform, which has been referred to in the literature as *market thickness*. Specifically, market thickness is comprised of two constructs: the number of potential matches, which I call "market size", and the number of competitors, which I call "competition size".[1] For example, a seller's market size is the number of buyers, and competitor size is the number of other sellers on the platform; thus, the seller's market thickness is the number of buyers *and* the number of sellers. I refer to an increase in market thickness as an increase in both elements of the tuple (i.e. an equal increase in both buyers and sellers). The endogenous adjustment of search, selectivity, and matching to contexts with varying market and competition size forms the core essence of behavior on a two-sided platform. I explore this dependence using data from a large, worldwide dating application.

Using historical observational data from this dating app, I find a strong correlation between market thickness and selectivity.[2] Individuals who have a larger market size seem to be more

[1] More specifically, market thickness is a tuple: (market size, competition size).

[2] I define selectivity to be the threshold at which an individual uses to decide whether to propose a match. Individuals who are more selective have a higher threshold; they are more likely to propose a match with higher quality users and less likely to propose a match with lower quality users, compared to a less selective individual.

selective, and users who have more competition seem to be less selective. These raw, stylized patterns in the data stand in contrast from the conventional wisdom in the theoretical literature on market thickness and matching. Broadly speaking, the theory suggests that markets with more potential matches generate more and better matches. Intuitively, a market with more potential matches has a larger potential choice set, so it is more likely to have an agent who is willing to match. Conversely, when an agent has more competition, he is less likely to find a match and conditional on matching, the matches are lower quality. However, this intuition does not hold when individuals can endogenously adjust their search and selectivity in response to market and competition size. In this paper, I show that such endogenous adjustment is occurring; and that allowing for the adjustment in a microfounded model of search and matching can rationalize these patterns observed in the historical data.

In contrast, the extant literature has either focused on centralized matching contexts that abstract away from search and selectivity (Gale and Shapley 1962; Abdulkadiroğlu et al. 2005; Roth and Peranson 1999); do not model search as endogenous to market thickness (Halaburda, Jan Piskorski, and Yildirim 2017; Kanoria and Saban 2017); abstracted away from search (Rochet and Tirole 2003; Hagiu 2006); or studied matching in the aggregate, abstracting away from these issues (Petrongolo and Pissarides 2001). In Section 2, I review these this literature in more detail.

There are challenges to be confronted in order to formally investigate the effect of market thickness on matching. First, one needs high frequency, individual-level data. In many contexts, only market-level outcomes are observed, or if data at the individual level is available, typically only the match is observed. For example, in job search data, the researcher can typically observe only the job that is accepted, not the jobs that were rejected, the applications the worker sent, or the interviews that were conducted. One of the points of this paper is to present a new dataset, which contains microlevel data on search. I can observe the sequence in which profiles were viewed, to whom match proposals were sent, and the matches that were realized. Second, quantifying the causal impact of market thickness on search behavior and matching is difficult because matching is confounded with the endogeneity of market thickness. The worry is that heterogeneous agents may self-select into locations of varying market thickness depending on their propensity to search or value for finding a match. For example, people with a high propensity to search may self select into locations with many people. It is difficult to estimate the causal effect from observational data alone due to the challenge of finding exogenous variation in market and competition size. Third, measuring the effect of market thickness requires a broad support of market and competition size in the data. That is, to measure the heterogeneous effect of market thickness across small and large markets, we need to see exogenous variation in market thickness across a variety of different markets.

I address these problems using a three-fold strategy: (1) develop and implement a randomized controlled trial that generates exogenous variation in information about market thickness in 2 neighboring countries served by the app, (2) track each individual's searches and outcomes, and (3) estimate a structural model of search and selectivity to analyze counterfactuals involving market thickness. The dating platform is designed such that an agent views profiles sequentially, and he must either `like` or not `like` the agent before viewing the next profile. In this setting, a `like` is the action that an agent takes to propose a match. The sequential nature of the search process ensures that I am able to observe each search instance. For each agent, I observe not only the experimental variation but also each profile he views, whether he `liked` the profile, and whether they match.

The experiment shows that beliefs in market thickness generate changes in behavior. While individuals, on average, are not significantly changing their search intensity, When users believe their market size increases by 50%, they become 3% less likely to like low quality users and 2.8% more likely to like high quality users. In contrast, when they believe they have 50% more competition, they become 2.3% more likely to like low quality users, and 4.5% less likely to like high quality users. Due to this change in selectivity, individuals get fewer matches when they think they have a larger market size, and more matches when they think they have more competitors. However, the effect of market thickness on matching must be evaluated under an equilibrium, as matches are equilibrium outcomes.

To find equilibrium implications when the actual, rather than beliefs about, market thickness changes, I develop a dynamic model of sequential search that is expanded to handle two-sided markets through incorporating beliefs about other agents' actions on the platform. Assuming the treatment changes beliefs about market thickness, by giving the model an explicit role for beliefs, I can leverage the exogenous variation in beliefs about market thickness from the experiment to estimate the parameters in the model. In this sense, this setup creates a bridge between experimental and theory-driven components of my investigation.

While beliefs are a key part of this model, beliefs also introduce complications in estimation. The observed agents' search behavior and matches are the outcomes of a Bayesian Nash Equilibrium of a dating game between men and women. In this equilibrium, agents have beliefs about other agents' actions, the likelihood another agent of a given quality likes them; and beliefs about quality of profiles they will see. Given these beliefs, they search and decide who to propose matches with optimally. However, when there are many agents in the game, solving the equilibrium becomes computationally challenging. I obviate some of the challenges in inference by borrowing from the literature on two-step estimation methods (Hotz and Miller 1993; Bajari, Benkard, and Levin 2007; Aguirregabiria and Mira 2007). The detailed data on agent states, along with assumptions on agent rationality, allow me to estimate equilibrium beliefs about other agents' actions in the first stage. This enables inference of a complex dynamic game at lower computational costs. Rich variation in data from prior to the experiment facilitates making the first-stage estimation flexible to the extent possible.

The estimates from the model show that the effects of market thickness are generally consistent with the experimental results, and are heterogeneous across market types and gender. Women respond more to competition, and men respond more to market size. In addition, competition has bigger effect in large markets.

Given the estimates from the model, I simulate how changes in market thickness impact an agent's number of matches and the average quality of the the agents they match with. The counterfactuals that I simulate revolve around increasing users on both sides of the market, or just one side, which can be done through targeted advertising or decreasing fees to join. I find that matching exhibits decreasing returns to scale. When 25% more men and 25% more women join the platform, the number of matches formed total increases by 10% in small markets, and 2% in large markets; agents get fewer and lower quality matches as well.

In addition, when only the number of women increase, men do not necessarily get higher quality matches as theory would have predicted. In large markets, where competition has a bigger effect, the average quality of women that a man matches with decreases by 10%. This effect is due to changes in selectivity in response to competition. Low quality types increase their like rate the most when competition increases. As a result, men have a higher likelihood of matching with lower quality types, as they are the ones more likely to reciprocate the match proposal.

These counterfactuals imply that limiting platform membership can increase matching outcomes for some users. However, the platform monetizes through premium users, so limiting membership may mean a loss in potential revenue. In the last counterfactual, I explore how changing the limit on the number of proposals an agent can send can mitigate negative effects of platform growth. I show when there are 25% more women on the platform, doubling the like limit for women can mitigate the negative impact on match quality for men.

In summary, this paper makes four contributions. First, I bring new data, combined with a field experiment and a model that simulates counterfactual outcomes, to analyze search and selectivity on a two-sided matching platform. The existing empirical literature on this topic has not treated the endogenous adjustment of search in such platforms in a satisfactory way. Second, the paper's results demonstrates that search and selectivity are endogenous to beliefs about market thickness, and presents a microfounded model that links search and selectivity to market thickness. This contributes to the literature on matching that has emphasized the importance of linking the aggregate number of agents on each side to the total matches formed, such as in the matching function[3], to microfoundations. Third, using the individual-level effect of beliefs about market thickness on search behavior and through the simulation of counterfactuals, this paper re-examines the conventional wisdom in the matching markets literature that individuals are better off when they have more potential matches. While the magnitude of the effect is specific to this empirical setting, my analysis suggests that it is possible for agents to find a better match when they have more competition. When designing the platform, I show that is important for firms to consider how agents react to market thickness in their setting rather than assume competition (market) size strictly reduces (increases) match quantities and qualities. Fourth, the paper integrates variation in beliefs generated from a field experiment, into a microfounded model, which then incorporates those exogenous beliefs to estimate the parameters. The variation in beliefs not only extends transparently to the model, but also is straightforward for the firm to implement, as experimentally changing market thickness is difficult to do. I show that changing beliefs about market thickness is sufficient to influence behavior and to estimate the model. This strategy of inducing belief variation in an experiment and incorporating it into a theory-driven model for identification may be useful

[3] See Petrongolo and Pissarides (2001) for a comprehensive literature review.

in other contexts as well.

The remainder of the paper is organized as follows. First, I review relevant literature and its relation to this paper. In Section 3, I describe the empirical setting and the motivating data patterns in more detail before presenting the experimental design in Section 4. Section 5 summarizes the data and presents the reduced-form results from the experiment. In Section 6, I propose the model of search and selectivity that is based on sequential search foundations. Estimation details are discussed in Section 7. The results of this model are shown in Section 8. Section 9 illustrates the counterfactuals produced from the estimated parameters. Section 10 concludes by summarizing the key findings and directions for future research.

2 Relevant Literature

This paper contributes to three separate but related literatures: job search, matching markets, and consumer search.

2.1 Job Search

Market thickness primarily has been studied in the labor literature as the matching function. This function relates the aggregate number of unemployed workers U and vacant job openings V to the number of jobs filled by $M = m(U, V)$. The idea of this function is to create an aggregate function that captures frictions, such as heterogeneity, congestion effects, etc, rather than modeling out all these factors separately. Many papers have either made functional-form assumptions about $m(U, V)$, (Burdett, Shi, and Wright 2001; Diamond 1982; Acemoglu and Shimer 1999; Buchholz 2016; Howitt and Mcafee 1987; Pissarides 1984), or empirically estimate the returns to scale of the matching function (Coles and Smith 1998; Blanchard and Diamond 1994; Gregg and Petrongolo 1997; Bleakley and Fuhrer 1997; Burda and Profit 1996).[4]

My paper is related to a subset of papers in the matching function literature that go beyond aggregate unemployed worker and job openings, by focusing on the microfoundations that may contribute to the shape of the matching function. While the simplest matching function only includes aggregate unemployed workers and vacant jobs, Petrongolo and Pissarides (2001) state that

[4]See Petrongolo and Pissarides (2001) for a comprehensive survey of matching function literature.

there has been literature focused on finding other variables that affect the match rate, such as search intensity. For example, Stevens (2007) theoretically develops a new matching function where search intensity directly affects the match rate. Pissarides (1984) develops a two-sided search model and theoretically shows that when both sides of the market search, it creates too much employment due to low search intensity in equilibrium.Gautier, Moraga-González, and Wolthoff (2007) empirically estimate a structural model of search using wage data from the Dutch labor market to derive the socially-optimal distribution search intensity. They find simultaneously creating more vacant jobs and increasing search intensity can create large welfare gains. While search intensity is clearly an important factor in determining match rates, to the best of my knowledge, no paper in this literature has been able to observe individual-level search and relate it to market thickness. Therefore, my contribution to this stream of literature is to shed light on the matching function's microfoundations. This paper demonstrates how the number of choices (analogous to aggregate job openings) and the number of competitors (analogous to aggregate unemployed workers) affect and selectivity, which may guide future work in understanding the "black box".

2.2 Matching Markets

There is a rich theory literature on matching markets. Much of the theory literature focuses on competition, pricing, and critical mass using simplifying assumptions of matching and consumer search. This area of research focuses on platform participation rather than matching. To simplify the matching process, the number of matches on a platform is commonly assumed to be a fraction of the product of the number of agents on each side (Rochet and Tirole 2003). There are also theory papers that focus on search and matching. Halaburda, Jan Piskorski, and Yildirim (2017) build a theoretical, microfounded model in a setting where agents propose a match to the agent with the highest match value. The authors show that increasing choice for all agents on the platform has two separate effects. It increases the chances of matching because the agent is more likely to find an attractive match (choice effect). On the other hand, because everyone has more choice, the other agent is less likely to accept the match, because she also has more choice (competition effect). Kanoria and Saban (2017) theoretically evaluates platform design in relation to which side of the market should send match proposals. In contrast to this literature, my paper empirically examines the role of more agents on the platform in settings where agents engage in costly search to evaluate a subset of other agents and are strategically selective about who they propose to match with. Shimer and Smith (2001) build a model of search on a decentralized platform where agents have heterogeneous productivity. The authors show that the equilibrium in decentralized market is inefficient and how the thick market externality and congestion have differing effects for agents of different productivity levels.

There are limited empirical papers on search and matching on platforms. One stream of literature focuses on platform growth. Tucker and Zhang (2010) implement a field experiment where a two-sided network advertises information about sellers and/or buyers and examine the impact the treatment has on entry. They find that if sellers do not know how many buyers are in the market, sellers are more likely to enter the market when there are more sellers as they believe more sellers mean more buyers. Chu and Manchanda (2016) study cross and direct network effects on an e-commerce platform and find that growth in buyers is primarily driven by growth in sellers. Cullen and Farronato (2015) study how a two-sided platform can create matches, even when supply and demand on the platform are not stable. They find that on TaskRabbit, supply is much more elastic than demand, so the firm should focus on increasing demand rather than supply. My paper builds on these findings by looking at not only how the two sides affect how many matches are made, but also match quality, which is not focused on in their setting. In addition, my paper builds on this area of research through (1) independently measuring the effects of more agents on the platform, as the treatments for market size and competition size are independent, and (2) studying outcomes for agents once on the platform, rather than the decision to join.

Another stream of empirical matching markets literature focuses on matching outcomes Gan and Li (2004) focus on the economics job market and find evidence for a positive effect of market thickness on the probability of matching at the aggregate level, but are not able to observe the individual search process. Fradkin (2015) creates a model of search on Airbnb and focuses on how ranking algorithms can increase matches in the marketplace. My paper is different from these in that I highlight the role of strategic selectivity; the decision to propose a match depends on the belief about how likely the other agent is to accept, which has not been emphasized.

There is a subset of empirical papers that focus on online dating in particular. Hitsch, Hortaçsu, and Ariely (2010a) use data from an online dating website to measure the stability of matches. They find that the predicted stable matches are similar to matches made on the website. There has also been research on preferences in dating. Fisman et al. (2006) measure differences in preferences between men and women through a speed dating experiment, and Lee and Niederle (2014) show that preference signaling increases an individual's match rate through a field experiment on a Korean dating website. Hitsch, Hortaçsu, and Ariely (2010b) estimate mate preferences by modeling a person's probability of contacting another person with a threshold-crossing rule. However, these aforementioned papers have not studied the role of market thickness in influencing matching outcomes. In addition, Hitsch, Hortaçsu, and Ariely (2010b) find no evidence of strategic behavior in their setting, but this paper shows that information about market thickness is able to change the type of people that a user proposes a match to, such suggests strategic selectivity in this setting.

Lastly, the stable matching literature often ignores the search process as it is focused on centralized markets, where search is not present. There is, however, a subset of the matching literature that focuses on developing theoretical models of search and matching. In these papers, agents follow a sequential search model where the agent's reservation value depends on search costs, arrival rates, and the distribution of quality of agents on the other side of the market (see Burdett and Coles (1999), Adachi (2003), and Mortensen (1982) for examples). To the best of my knowledge, there does not exist any theoretical or empirical literature that studies matching through the lens of strategic selectivity and the availability of other agents on the platform.

2.3 Consumer Search

Many economics papers in consumer search do not involve two-sided matching, as purchasing a product is not a two-sided decision. Behavioral economics and psychology have explored the role of the number of choices on search and purchase behavior. Reutskaja et al. (2011) study how the computation processes consumers use during search change when the number of options change. They conduct a lab experiment where hungry subjects have three seconds to choose a snack out of sets of 4, 9, or 16 snacks. They have subjects rank all snacks prior to the experiment. Using eye-tracking technology, they find that as the set size increased, subjects searched across more snacks, but because of the time constraint and only being able to see a subset of the choices, they made less optimal choices. Iyengar and Lepper (2000) show that subjects often postpone choice, or do not make a choice at all, when the choice sets are large. Diehl and Poynor (2010) find that individuals are less satisfied with their purchases when the product is selected from a smaller selection, indicating they have higher expectations about match values they can attain in a larger set. This paper contributes the literature on consumer search by building on sequential search models and extending them to allow for (1) two-sided matching decisions and (2) the number of choices.

3 Empirical Setting

The setting of this paper is a mobile online dating application that has millions of active users worldwide. The types of people that use this app are mainly heterosexual adults between ages 20 to 30. Based on a user's gender, age, and distance criteria, the app serves other agents' profiles one at a time. Each profile consists of a name, pictures, and a short bio. Information such as age, occupation, and education may also be displayed.

A user may view each profile only once. Once shown the profile, he must decide whether to anonymously `like` or not `like` that profile. In this context, a `like` is defined as an action that the user takes to propose a match. A match occurs when two users mutually `like` each other. If he chooses not to `like`, then another agent's profile is displayed. An important feature of this platform is that when a user is looking at another user's profile, he does not know if the other user has liked him, which introduces uncertainty from the focal user's perspective. Ex-ante, he does not know whether he will match with that other user. For the remainder of this paper, I refer to "`like`" and "not `like`" as these actions, unless specified otherwise. Once an agent chooses to not `like` a profile, he will never be shown that profile again. Messages can only be exchanged between matched users.

3.1 Design

Figure 1 illustrates the layout of the app. When the agent opens the app, she will be shown a profile. In this example, she sees "Adam"'s profile. She then decides whether to `like` Adam. If she either does not `like` Adam or does `like` Adam but Adam has not `liked` her back, then the profile will immediately serve her the next profile, "Bob". However, if she does `like` Adam and Adam has already liked her, then she will get a notification saying she and Adam have matched, after

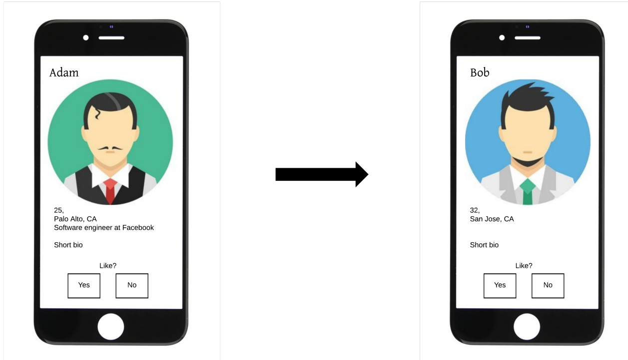

Figure 1: Example layout of the app. When the user opens the app, she is first presented with Adam's profile. The profile includes pictures, age, location, and a short bio. If the user does not `like` Adam, then she is shown the next profile, Bob.

which she and Adam can exchange messages. Note that matches may not occur instantaneously, i.e. Adam does not see and `like` her profile until the next day.

3.1.1 `like` limit

The platform monetizes through a "freemium" payment plan. Users can use the app for free, but paid users get access to more features. A difference between paid premium version and the free version is the `like` limit. Free users have a limit on how many agents' profiles they can `like` within a 12 hour period, while premium users do not. I refer to $\bar{L}$ as the cap on the number of `likes`; the exact number is confidential, as requested by the firm. The analysis in this paper will be conducted on the proportion of `likes` a user has left.

Once the agent runs out of `likes`, he cannot search for the next 12 hours. Clearly, agents value not having a limit on `likes`, so they are willing to pay for the premium version. Since this limit is important for the platform's monetization strategy and is an institutional feature of the data, I take this policy into account in the model and counterfactuals.

3.2 Historical Data

The observational, historical data displays strong correlations between selectivity and market thickness. Selectivity, in this section, is measured by the individual's `like` rate, which is the proportion of profiles that he `likes` out of all the profiles he views. I use data from heterosexual app users in the same two countries the experiment is implemented in. Specifically, the sample contains over

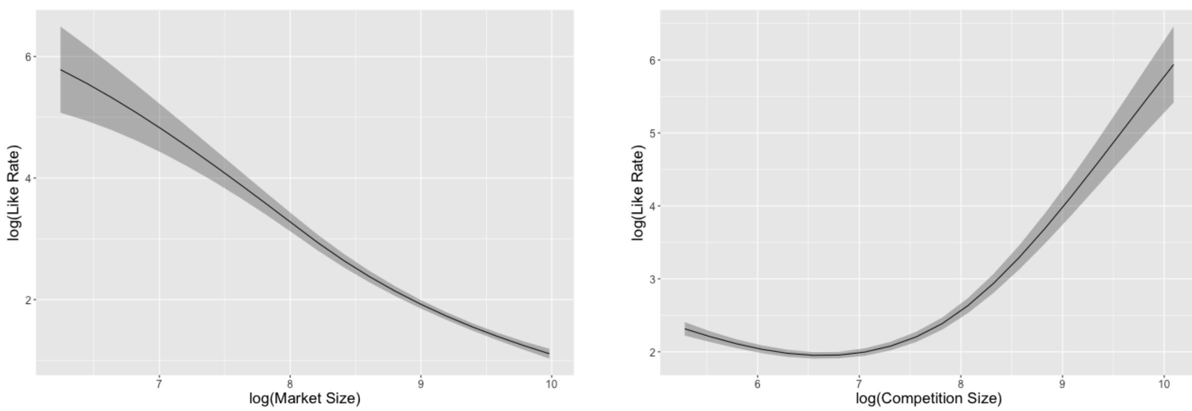

Figure 2: This figure shows the correlation between market size and competition size with selectivity, for markets with a market/competition size of at least 200. The lines are the results of a locally linear regression of an individual's `like` rate, which is the proportion of profiles that an individual `likes`, on his market size, competition size, age, and gender. The plot on the left shows the correlation between market size and the `like` rate, while holding competition constant at the average value, and the plot on the right shows the correlation between competition size and the `like` rate, while holding the market size constant at the average value. The shaded regions are 95% confidence intervals.

50,000 women and 100,000 men who viewed at least one profile in a 34-day time period in 2015. To explore this correlation, I estimate the following locally linear regression.

$$LikeRate_i = f\Big(\log(ms_i), \log(cs_i), age_i, gender_i\Big) + \epsilon_i \tag{1}$$

ms_i is i's average market size. Substantively, I determine a user's market size at t, the time of each profile view, which is the number of users of the opposite gender within 100 miles at t. ms_i is then the average of i's market size over time. Similarly, cs_i is i's average competition size, which is determined in the same method as market size, but is instead the number of users of the *same* gender within 100 miles.

Figure 2 displays the results from Equation 1. The `like` rate is log-transformed by an unreported base to preserve data confidentiality. The figure shows a strong correlation between both market size and competition size, with selectivity. An individual with more potential matches seems to be less likely to `like` a profile than another individual with fewer potential matches. On the other hand, an individual with more competition, on average, seems more likely to `like` another profile than an individual with less competition.

Due to the two-sidedness of matching, this correlation may result in matching outcomes that differ from the common wisdom from theory on matching and market thickness. Individuals who are more selective may get fewer matches; they may be more likely to `like` "higher quality" users who are less likely to `like` them back.

3.2.1 Endogeneity

While Figure 2 shows a strong correlation between market thickness and selectivity, these results cannot be interpreted causally. The challenge in measuring the causal effect of market thickness is that individuals self-select into markets. For example, individuals who have high propensities to search may be more likely to travel to a market with more potential matches. It could also be the case that individuals who in large cities are more selective on the dating app because they have more outside options. Thus, to measure the causal effect of market thickness on behavior, I need to leverage exogenous changes in market thickness. However, finding natural sources of exogenous variation in market and competition size in this setting is difficult, as there are network effects between the two sides of the market, where if the firm randomly increased marketing towards women only, more men may join the platform as well, making it difficult to disentangle the separate effects of market and competition size. Due to this lack of natural variation in market thickness, I design and implement a randomized control trial to generate the exogenous variation needed to quantify these relationships.

4 Experiment Design

An ideal experiment exogenously changes an agent's market thickness. However, it is not feasible to experimentally add or remove agents from the platform, so this experiment varies *information* about market thickness. Upon opening the app, a user in the treatment group sees a pop-up notification with the message "There are at least m men and w women nearby!". Users in the treatment group see varying values of m and w. I refer to the pop-up message as the *treatment*, and m and w as the *treatment values*. Figure 3 provides an illustration of the design of the pop-up message. The following section describes this experiment in more detail.

4.1 Treatment Values

In the ideal experiment, m and w would be randomly drawn. However, the agents must see a reasonable number of his or her market and competition sizes to ensure that experimental effects

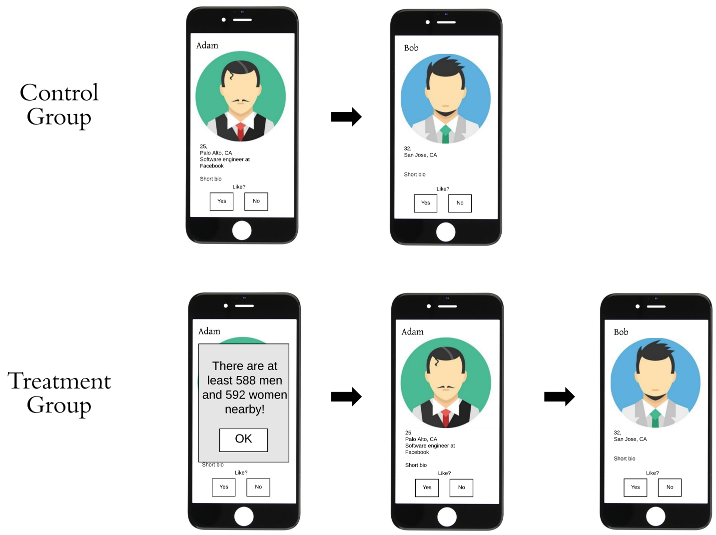

Figure 3: Example of how the treatment condition is different from the control condition. The treatment condition receives the pop-up message while control users see no change in the app.

are not influenced by the Hawthorne effect (McCarney et al. 2007). A way that the Hawthorne effect may arise in this study is through unrealistic treatment assignments: agents may not believe the treatment values they are shown. For example, individuals who live in rural areas would not believe there are 10,000 men or women around them, and similarly, users in large cities would not believe there are 10 men or women around them. To obtain realistic values, m and w are drawn in the following way.

The experiment is implemented in two neighboring countries. Each country is divided into approximately 625 mi^2 grids.[5] For each grid g, I determine the population size M_g, which is the number of men who have opened the app from within that grid in the two weeks prior to the start of the experiment. Not all grids are selected to be in the experiment. Details on grid selection are in Section 4.2. When a treated user opens the app from a selected grid g, the treatment values m and w are drawn from a distribution centered around M_g. Specifically, both m and w deviate from M_g by factors, F^{ms} and F^{cs}. F^{ms} is the market size factor, and F^{cs} is the competition size factor. The treatment values that a user sees in the pop-up message are the treatment factor times the population size; the number of men and women are determined by the corresponding market and competition size factors based on the user's gender. For a male user, the treatment value for

[5]Grids are formed by flooring each latitude and longitude to the nearest 0.75 (roughly a distance of 25 miles).

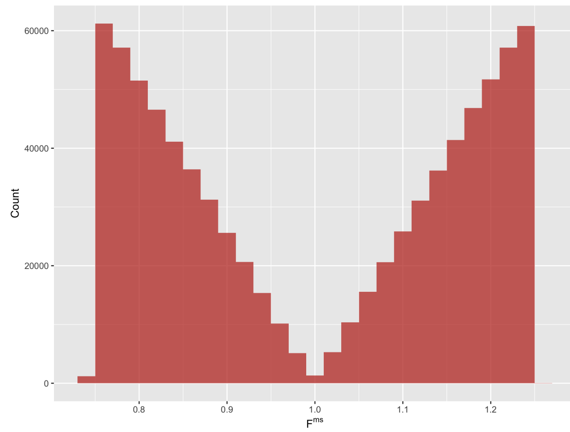

Figure 4: Histogram of F^{ms} across all sessions. Each observation is at the user-session level.

the number of women is $F^{ms} \times M_g$ and the number of men is $F^{cs} \times M_g$, while for a female user, the number of women is $F^{cs} \times M_g$, and number of men is $F^{ms} \times M_g$. F^{ms} and F^{cs} are drawn from a V-shaped distribution between 0.75 and 1.25.[6] The V-shaped distribution, as opposed to the uniform distribution, increases statistical power by increasing the variance of the treatment values. Figure 4 displays the histogram of the draws of F^{ms}, where F^{ms} is the market size factor. Similarly, F^{cs} is the competition size factor.[7] In summary, the values displayed to the agent in the treatment message are correlated with the true number of agents in that grid, but the factors (F^{ms} and F^{cs})by which they deviate are randomized.

4.2 Sample Selection

4.2.1 Locations

The experiment is run in 2 neighboring countries where this app is popular. The identities of these countries are kept anonymous. Not all grids in these countries are selected for the experiment. The criteria by which grids are selected are that the grid must have a population size of at least 200, and must not include a large city.[8] A total of 292 grids fit this criteria.

Figure 5 displays a histogram of the population sizes of the 292 selected grids. Most grids are relatively small, with a population size of less than 1,000. One grid contains a relatively large city, with a population size of 37,000. The advantage of implementing this experiment with

[6]The variation in the treatment values is similar, if not smaller, than the natural variation in market thickness in this setting. Figure 14 in the Appendix displays how market thickness varies over time using historical data.

[7]The market size for men is w, and competition size is w. Vice versa for women.

[8]Whether a grid contains a large city is determined by the firm.

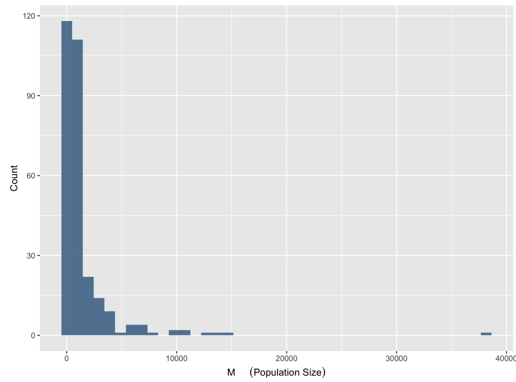

Figure 5: Histogram of the population size of each of the selected grids. This figure validates that the treatment values are indeed being drawn from a V-shaped distribution.

large variation in population sizes is that I can better measure the heterogeneous effects of market thickness across different types of markets.

4.2.2 Users

The sample is selected from 225,680 active users who are seeking matches with members of the opposite sex. These are men and women who have opened the app within a month prior to the start of the experiment from one of the selected grids. I randomly select 90% into the treatment group, and 10% into the control group, resulting in 203,170 and 22,510 users, respectively. I select 90/10 instead of 50/50 to maximize the number of people who are exposed to the treatment. The treatment effects of interest are measured by comparisons between users within the treatment group, rather than comparisons between users in the treatment vs control groups. In other words, the treatment effect of market size is measured by comparing someone who saw a high value of their market size compared to someone else in the treated group who saw a low value of their market size.

If the user does not open the app during the span of the experiment or does not open the app from one of the selected grids, he is not exposed to the treatment. After running the experiment for 3 weeks, 84,589 users were exposed to the treatment. Figure 6 summarizes the user sample selection process.

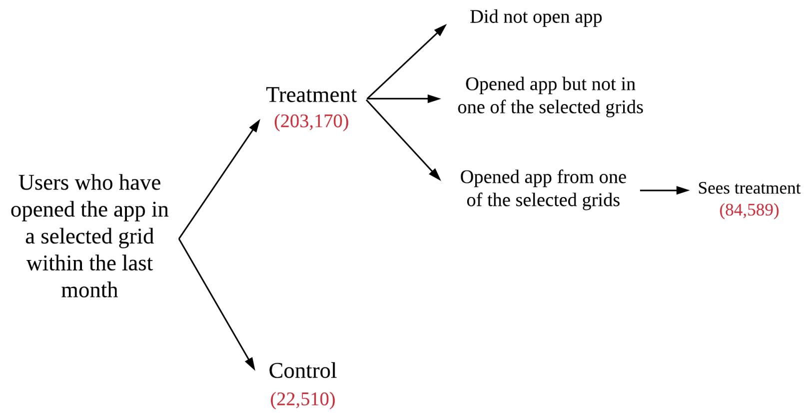

Figure 6: A summary of how users are selected into the experiment. Out of all the users who opened the app from one of the selected grids, 90% were selected into the treatment, and 10% into the control. Of the users in the treatment group, a user is exposed to the treatment if he opens the app from within one of the selected grids.

5 Data

5.1 Data Summary

A treated user is exposed to the treatment every 6 hours, where the treatment value are re-randomized. I conduct the remainder of the data description and analysis on data from the agent's first session only to avoid potential selection problems, similar to Sahni and Nair (2016).[9] Selection bias may arise in the following way. Individuals who see a lower draw of the market size may become less selective and be more likely to find a date. However, the low quality types, who are the least likely to find a date, will come back to the app and start a second search session. Thus, the sample of users who are exposed to the treatment in the second session are no longer random; the treatment in the first session affects which users are exposed to the treatment in the second session. This would violate the stable unit treatment value assumption (Rubin 1974), where in this case, each unit is a user-session. To avoid these potential issues, data from the second session onwards is excluded from the analysis.

After removing outliers and premium users from the data, there are 26,092 women and 40,647 men remaining.[10] Table 1 reports summary statistics for the number of profiles viewed, the proportion of profiles that were `liked` (`like` rate), and the number of matches made during

[9]More details on the session-level randomization is included in the Appendix.

[10]Premium users are removed from the data because they do not face the same dynamic forces as freemium users. A source of dynamics is the `like` limit; premium users do not have a `like` limit.

	Mean	SD	Median
Men (N = 40,647)			
Views	34.46	54.23	13.00
Like Rate	3.88	2.41	4.02
Matches	0.35	0.38	0.17
Women (N = 26,092)			
Views	44.25	78.74	11.00
Like Rate	0.69	1.14	0.26
Matches	0.49	0.54	0.17

Table 1: Descriptive statistics of agent actions in the first session. The `like` rate and number of matches are log-transformed by an unreported base to maintain confidentiality. On average, women search more, are more selective, and get more matches than men.

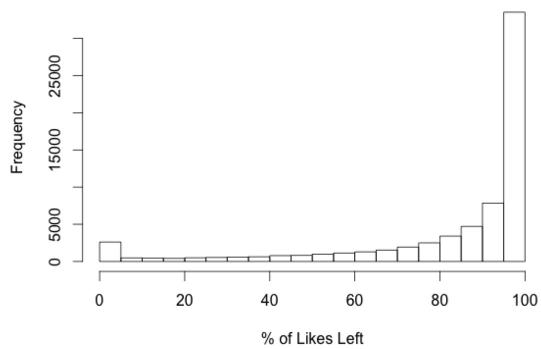

Figure 7: Histogram of the proportion of the `like` limit the user has remaining at the end of the session. Having 100% left indicates the user did not use any `likes` during the session.

the session. The `like` rate and number of matches are log transformed by an unreported base to maintain data confidentiality. The summary statistics demonstrate that there is considerable heterogeneity in search behavior both within and across genders. For example, on average, men view 35 profiles per session, but the median man views only 13 profiles. Women are also much more selective and get more matches than men.

The behavior also changes according to the "`like` limit". Figure 7 plots a histogram of the proportion of the `like` limit the user has remaining at the end of his session. For example, if the user has 100% of the `like` limit left at the end of the session, that means he did not `like` any other users during his session. The user hits the `like` limit when he has 0% left. The histogram shows that the `like` rate does bind for a subset of users. Moreover, while agents are not aware of the exact number of `likes` they have left, the data suggests that agents are cognizant of how

	Dependent variable:	
	$\mathbb{1}\{Like\}$	
% of Likes Remaining	0.072***	0.00000
	(0.005)	(0.00003)
Profiles Already Viewed	−0.00002***	−0.0001***
	(0.00000)	(0.00002)
User Type	Free	Paid
User FE	Y	Y
Observations	2,571,467	598,544
R^2	0.478	0.481
Note:		*p<0.1; **p<0.05; ***p<0.01

Table 2: OLS of whether the user `likes` the other profile on the proportion of the `like` quota he has left, and how many profiles he has viewed during the session. The first column is for free users, and the second column is premium users only. Standard errors are clustered at the user level.

many `likes` they have left. Table 2 presents the results of a linear regression of whether a user i `likes` another user j on the percentage of `likes` he has left, and how many profiles he has already viewed. The first column is freemium users, and the second column is for premium users. The coefficient on the percentage of `likes` left for freemium users is significant and positive, meaning the user is more likely to `like` another profile when he has more `likes` left. However, premium users do not exhibit this same behavior; there is no correlation between how many people they have already liked, and their propensity to like the next profile.

5.2 First Cut of Experimental Data

Before introducing the model, I first assess whether the treatment values shift behavior. The notation for the remainder of this section is the following. An agent i's market size is the number of agents of the opposite gender, and i's competition size is the number of agents of the same gender. M_i denotes the population size of the i's location when he was exposed to the treatment, and is measured in 10,000's. ms_i and cs_i are the market and competition size treatments, respectively, and F_i^{ms} is the market size factor, and F_i^{cs} the competition size factors. For example, if a male agent in a location with $M = 1000$ sees the treatment "There are at least 900 men and 1200 women nearby", his market size treatment is 1200, competition size treatment is 900, market size factor F^{ms} is 1.2, and competition size factor F^{cs} is 0.9. To reiterate, F^{ms} and F^{cs} are between 0.75 and 1.25 and are truly random, while ms and cs are *random conditional on location.*

5.2.1 Selectivity

I first measure how selectivity changes with market thickness. How do individuals change their overall propensity to `like` another person, and does that change vary by the other individual's quality type? In other words, how does market size and competition size change the likelihood that the individual likes a high quality user compared to a low quality user? To measure the effect on selectivity, I estimate the following logistic regression.

$$Like_{ij} = \alpha_\ell + \gamma_1 F_i^{ms} + \gamma_2(F_i^{ms} \times q_j) + \gamma_3 F_i^{cs} + \gamma_4(F_i^{cs} \times q_j) + \gamma_5 q_j + \gamma_6 X_i + \eta_{ij} \qquad (2)$$

Each observation in this regression is one profile view, where i sees j's profile. $Like_{ij}$ is an indicator variable for whether i likes j. α_ℓ is a location (grid) random effect, and X_i is a matrix of i's gender and q_i. The effects of interest in Equation 2 are the interaction terms of market size and competition size with q_j, where q_j is user j's quality type.

For the remainder of this paper, I measure an individual's "quality" using his or her observed "desirability". i's quality score is the likelihood that another user likes i's profile, conditional on seeing i's profile, which is calculated using data prior to the experiment. For example, if $q_i = 1$, then every user who has seen i's profile has liked i. One concern of this metric is that there are systematic differences in behavior across different sized markets. For example, if users in cities are more selective, then those users will have lower quality scores, relative to users in rural areas where the like rates are higher. Thus, to remove systematic differences in quality scores based on market conditions, q is expressed in percentiles, relative to markets with similar population sizes. This allows a user's quality score to be compared with another user in a different location.[11]

To measure heterogeneous treatment effects across markets of different population sizes, I estimate this regression separately for agents in small and large markets, where small markets are defined to be markets with below average population size and large to be above average.[12] The average population size for all users exposed to the treatment is 11,157.

[11] The specifics of how q is calculated is in the Appendix.

[12] An alternative specification to measure heterogeneous effects across population size is to include an interaction between F^{ms}, q_j and M. I choose to estimate the specification in Equation 2 because 3-way interaction effects are more difficult to interpret.

5.2.2 Search Intensity and Matches

To measure the average treatment effects, I estimate the following regression.

$$Views_i = \alpha_\ell + \beta_1 F_i^{ms} + \beta_2 F_i^{cs} + \beta_3 (F_i^{ms} \times M_i) + \beta_4 (F_i^{cs} \times M_i) + \beta_5 X_i + \epsilon_i \tag{3}$$

i is the individual who is exposed to to the treatment. $Views_i$ is the number of profiles that i views during the session. α_ℓ denotes location fixed effects. X_{it} are controls for individual characteristics. Specifically, I include controls for gender and i's quality type, in percentiles. An individual's quality type is measured by the ratio of users who have `liked` i's profile over the number of users who have seen i's profile. For instance, if $q_i = 1$, every individual who has seen i's profile has `liked` his profile. To obtain the percentile, i's quality type is compared to the quality type of all other users of the same gender. All quality measures are obtained from historical data, prior to the experiment. I also estimate the same regression to measure the impact of the treatment values on $Matches_i$, the number of matches made during the session.

5.3 Results

5.3.1 Selectivity

Table 3 presents the effect of market thickness on selectivity. Column (1) presents the results for users in markets with a below-average population size, (2) for users in markets with above-average population size.[13]

The effect of the market size factor is negative and significant in both columns, meaning that users who received a higher market size factor are less likely to like another user. However, the interaction of the market size factor and q_j is positive in both columns, and significant in the first column. When users see information about a larger market size, they are less likely to like another user, but *more* likely to like a user of higher quality. To provide a sense of the magnitude of these effects, in large markets, when the market size factor increases from 75% to 125%, the average user becomes 3.25% less likely to like another user in the 10^{th} percentile of quality and 0.1% more likely to like a user with the 90^{th} percentile quality.

[13] Coefficients for gender effects are included in the estimation, but are removed from the results table by request of the company

	Dependent variable:	
	$\mathbb{1}\{\text{Like}\}$	
	(1)	(2)
F^{ms}	-0.166^{***}	-0.075^{*}
	(0.029)	(0.039)
$F^{ms}:q_j$	0.144^{***}	0.082
	(0.046)	(0.061)
F^{cs}	-0.189^{***}	0.091^{**}
	(0.028)	(0.039)
$F^{cs}:q_j$	0.053	-0.186^{***}
	(0.046)	(0.061)
q_j	2.914^{***}	3.219^{***}
	(0.065)	(0.087)
q_i	-1.613^{***}	-1.602^{***}
	(0.008)	(0.010)
Location RE	Y	Y
M	Lower	Upper
Observations	1,606,307	949,003
Log Likelihood	−666,991.000	−362,586.100
Note:	*p<0.1; **p<0.05; ***p<0.01	

Table 3: OLS regression of Equation 2

In small markets, the overall effect of competition size is negative and significant, and the interaction between the competition size factor and q_j is positive but not significant. A potential explanation for the direction of this effect is that individuals are less likely to like anyone because the probability that they match is not large enough to outweigh the cost of sending a like (due to the like limit). Another explanation is that the relationship between F^cs and q_j is not able to be captured by a linear specification. In large markets, the effect of competition is significant and positive, and the interaction term is significantly negative. This means that those individuals are more likely to like lower quality users (3.1% more likely to like a 10^{th} percentile quality user), and less likely to like higher quality users (1.25% less likely to like a 90^{th} percentile quality user). These results show that the effect of market thickness may be different for small and large markets, so this heterogeneity should be included in the model, which is presented in the following section.

	Dependent variable:	
	Views'	Matches'
	(1)	(2)
F^{ms}	-0.007	0.010
	(0.023)	(0.024)
$F^{ms} : M$	-0.006	-0.015^{*}
	(0.007)	(0.008)
F^{cs}	-0.018	-0.039
	(0.026)	(0.028)
$F^{cs} : M$	0.012^{*}	0.023^{***}
	(0.007)	(0.009)
q_i	0.136^{***}	0.358^{***}
	(0.017)	(0.016)
Location FE	Y	Y
Observations	66,739	66,739
R^2	0.037	0.041
Note:	*p<0.1; **p<0.05; ***p<0.01	

Table 4: OLS estimates from Equation 3. Standard errors clustered by location.

5.3.2 Search Intensity and Matches

Table 4 displays the OLS estimates of Equation 3. The dependent variable in column (1) is the standardized number of profile views. (2) is the standardized number of matches.

There is no significant effect of market thickness on search intensity. The effect of $F^m s$ on search intensity is negative but not statistically significant. On average, a 50% increase in competition size increases the profile views by 1% (0.006 standard deviations).

As shown in Column (2), for the average market ($M = 11,157$), increasing an individual's beliefs about market size by 50% decreases matches by 2% (0.008 standard deviations), and increasing beliefs about competition size by 50% results increases matches by 3% (0.012 standard deviations). These individual effects are significant but small in magnitude, as matches are relatively rare in the data as compared to profile views and `likes`.

5.4 Discussion

The takeaway from this experiment is that individuals respond to information about market thickness by adjusting their search behavior. In general, when an individual is told she has larger market size, she becomes more selective, and when she believes she has more competition, she becomes less selective. The absolute magnitudes of these effects are small, but I believe these effects are relatively large and surprising, given the strong role that inherent preferences play in this setting. The magnitude of these effects of market thickness may change once these preferences are controlled for in a model. Another potential factor is that q, the quality measure, may be noisily measured with respect to the "true" quality, resulting in attenuation bias.

While the experiment is able to identify a causal effect on search behavior, we are not able to draw conclusions on matching outcomes from the experiment alone because the treated users' behavior is out of equilibrium. Because matches are two-sided, *both* users' behaviors must be taken into account. The treated user i has information on her market thickness, unlike the vast majority of other users on the platform. In the counterfactual, if j also received the same information about his market thickness, he might have changed his actions, which would affect the matching outcomes. Thus, in order to evaluate matching outcomes for counterfactual analysis, equilibrium behavior for users on both sides of the market must be simulated. The following section describes in more detail the role of the model in counterfactual analysis.

6 Model

6.1 Motivation

The goal of this model is to guide platform design through the analysis of counterfactuals, and to organize the various empirical results through a set of parsimonious parameters that can explain the behavior. While the platform can answer some policy questions through experimentation, many relevant counterfactuals may be too costly, or even impossible to test. The experimental results are not enough to estimate matching outcomes under different counterfactuals for two reasons. First, interpreting individual effects of market thickness on matching as a whole is not straightforward. In particular, if market size increases for one side of the market, by definition, competition increases for the other side. Thus, the direction of the overall impact of changing market thickness is not clear due to the different individual effects of market and competition sizes. Second, matching is the result of a Bayesian Nash equilibrium. The idea that each agent's behavior depends on his beliefs about other agents' behavior is inherent to matching markets. Matching is a two-sided decision, and an agent does not know whether the other agent is willing to match. Therefore, he forms beliefs about j's actions and behaves optimally according to these beliefs, which is the equilibrium of the game. Thus, the effects of a policy implementation must be interpreted as the result of an equilibrium of a dating game of imperfect information between men and women.

6.2 Model Details

This paper builds a model of search and selectivity in order to simulate matching outcomes. In other words, matching outcomes are treated as a statistical process determined by the agent's joint decision for how much to search and which agents to `like`. Modeling matches in this way allows me to pool all the actions an agent can take in a cohesive way. The structural model also allows me to incorporate a current policy that is a driver of agent behavior and a source of dynamics: the `like` limit. The agent's actions in the current time period affects how many `likes` he has left in the next time period, and Table 2 provides evidence that agents behave accordingly. Not only does this policy influence agent behavior, but the platform also uses this policy as a source of monetization. Agents can pay to subscribe to the premium version of this app, which does not have a `like` limit. Motivated by the fact that the `like` limit adds an additional source of dynamics that drives selectivity and is an important policy for monetization, the `like` limit is an integral aspect of this model for counterfactual analysis.

The foundation of this model is built on sequential search and is extended to allow for two-sided matching. The exogenous variation in market thickness in the experiment is introduced at the search session level, so I model an agent's search for a single session as well. That is, I model one episode of search for each agent, and I assume that agents are not forward-looking across search episodes. This mirrors the analysis strategy in the descriptive section. In each time period t in the session, the agent can view one profile. In a session with T time periods, the agent can view up to T profiles. The agent's objective is to find another agent to go on a date with and to maximize the utility from that date (date utility), subject to costs incurred from search.

In the following paragraphs, I first an overview of the model by describing how the experiment changes beliefs, and the timing of the agents' actions. I then discuss the agent's states, actions, and utilities in more technical terms. Lastly, I explain how market thickness enters the model and discuss the equilibrium in more detail.

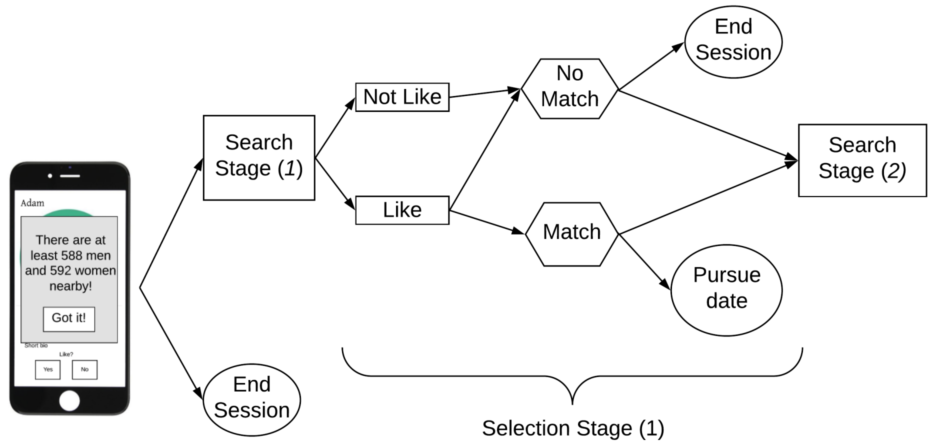

Figure 8: This figure shows the actions the agent can take after seeing the treatment. The rectangles represent the actions in the model. The circles depict endpoints in the search session. The hexagons represent uncertainty from i's perspective. After being exposed to the treatment, he can either enter the search stage or end the session. If he enters the search stage, he evaluates the profile and then enters the selection stage. The selection stage can result in the following possible actions: (1) the agent matches and pursues a date with the other agent, (2) the agent matches and decides to continue searching in the next time period, (3) no match is formed, and the agent continues to the search stage in the second time period, or (4) no match is formed, and the agent ends the session.

6.3 Timing of Agent Actions

6.3.1 Belief Updation

The experiment varies information about market thickness but to draw implications of market thickness, I make the following assumption about user beliefs.

Assumption 1. *Agents update their beliefs about their market size and competition size to the minimum number of men and women, respectively, provided by the experiment.*

In other words, I assume that users completely believe the treatment values and update their beliefs accordingly. Prior to the experiment, an agent has beliefs about her market thickness, denoted by $\tilde{m}_t$ and $\tilde{c}_t$, which are beliefs about her market and competition size, respectively, at time t. For the remainder of the paper, tilde's denote beliefs. At time τ, she is exposed to the treatment message of "There are at least m^* men and c^* women nearby". She then updates her beliefs to m^*_τ and c^*_τ.

6.3.2 Actions during the Session

Figure 8 summarizes the timing of the steps in the model starting from when the agent is exposed to the treatment. Each user starts her session with a "blank slate".

Assumption 2. *At time t, the agent does not consider pursuing a date with matches formed at $t' < t$.*

This is a strong assumption that is made to simplify the model. In reality, users make take a long time to respond to their matches, so they may start another session while waiting to hear back from previous matches. By making this assumption, it may make the model less predictive of actual behavior, but should have no impact on the estimates of the market thickness parameters, as those values are randomized.

Given beliefs about market and competition size, i decides whether to search (view a profile) or end the search session. If she ends the search session, she receives the utility from the outside option, which I normalize to 0.

There are two stages within each time period: the search stage and then the selection stage. In the search stage, the agent i views an agent j's profile. In doing so, she incurs a search cost of viewing the profile c^v and then continues to the selection stage. In the selection stage, she chooses to `like` or not `like` j. If she `likes` j, and j `likes` i a match occurs. i has some belief $\tilde{\pi^m}_{ij}$ that j likes i. If i and j match, i screens j's profile at a cost c^m. This may involve looking more closely at j's profile or exchanging messages. After screening, i then decides whether to pursue a date with j. i receives the date utility from j if and only if i and j mutually pursue a date with each other. At the time i decides to pursue a date with j, she does not know whether j pursues a date with her. Thus, if she chooses to pursue a date with j, she receives the expected date utility ρ_{ij}. By structuring the model in this way, I make the following assumption.

Assumption 3. *An agent gets positive utility from a match only if they go on a date.*

This is a strong assumption, as it is plausible that agents go on the platform just to browse or chat online. If she does not pursue a date with j, then she enters the search stage again in time $t + 1$.

However, if either i does not like j or j does not `like` i, no match is formed. The agent can then decide to end the search session without pursuing a date, in which she chooses the outside option, or continue to the search stage in $t + 1$.

6.4 Market Thickness

In a centralized market, an agent's belief about how likely he is to find a date depends on the ratio of men to women. For instance, a male agent may believe that he has a better chance of being assigned a date if the ratio of men to women decreases. Allowing agents to search for a date complicates this process, as an agent's search and selectivity decisions will be affected by this belief. Consider a simple search model where an agent `likes` a profile if the expected date utility is greater than the expected value of continuing search. Intuitively, if the market to competition size ratio is high, the agent believes he has a good chance of finding a date. As a result, he may be less likely to `like` the current profile because (1) `likes` are costly, and (2) he believes he should be able to end up with a better date if he continues searching.

To link this intuition to the model, market thickness enters through its effect on an agent's value of continuing search, and the probability that another agent sees i's profile. Also, due to the nature of the game, market thickness has indirect and direct effects. I refer to indirect effects as effects that change i's behavior through changing i's beliefs about other agents' actions. For instance, if i has more potential matches, that means j has more competitors. Thus, i may believe that j may be more likely to `like` i's profile. Conversely, direct effects change i's utility, even if beliefs about how other agents' behave remained constant. In the following section, I describe the direct effects of market and competition size.

6.4.1 Market Size

Market size can enter the model through two ways: quality and quantity. The most obvious way is through quantity. When there are more agents in the market, the agent has more search opportunities. While this aspect of market size may be relevant in smaller markets, in the markets in this experiment, the quantity of agents is not binding. That is, the market size is so much greater than the agent's search intensity during a session that running out of search opportunities is not a concern. Therefore, I model market size through its effect on the agent's belief about the type of agent he will see in the next search opportunity.

There are a multitude of reasons why market size would affect beliefs on the quality type of agents in the next time period. First, when the market size increases, new agents are joining the platform, and those new agents may be of a different type than existing agents. It is plausible that the incremental users may be different from the ones who have already on the platform. Second, the platform's matching algorithm may augment the effect of market size. When there are more agents on the platform, there are also more high quality agents. If the agent believes that the matching algorithm sorts the order in which profile are shown by quality, when market size increases, the quality of the next agent shown should be on average higher than if the market size had decreased. Third, psychology literature shows evidence that when individuals have more choice, they have greater expectations about the ability to find a match that better aligns with their preferences (Diehl and Poynor 2010). That is, when there are more alternatives, individuals believe they will find a better match.

The data cannot distinguish between these mechanisms, but the overall effects would manifest in similar ways: market size affects beliefs about the quality of agents shown in the next time period. I cannot measure how much of the overall effect comes from each mechanism. Therefore, I measure the aggregate effect of market size and am agnostic about how much each mechanism contributes to the overall effect. I model the effect with the following specification.

$$q_{it} \sim TN(\tilde{\mu}_i, \tilde{\sigma}_i^2)$$
$$\tilde{\mu}_i = \hat{\mu}_i + \delta_1 F_i^{ms} + \delta_2 F_i^{ms} \times M_i$$

TN denotes a truncated normal distribution, bounded between 0 and 1, with mean $\tilde{\mu}_i$ and variance σ_i^2. $\hat{\mu}_i$ is the agent's prior belief about the mean quality of the distribution of q, which can be estimated by the true quality types of profiles that i sees prior to the treatment. Post treatment, the mean is shifted by F_i^{ms}. The interaction term δ_2 captures heterogeneity in the effect of market size across markets of different population sizes.

6.4.2 Competition Size

When there is more competition, the chance that another agent has seen a given agent's profile decreases. For example, in the labor market, when an unemployed worker has more competition, it becomes less likely the firm sees his resume. In this model, for each profile that he sees, when the agent believes that when there is more competition, it is less likely that the other agent will match with him because the other agent is less likely to see his profile.

As described previously, if i `likes` j, i and j match only if j sees i's profile and j `likes` i's profile. Thus, i's belief about the probability that they match is specified as the following.

$$\begin{aligned}\tilde{\pi}^{m}_{ij} =& \tilde{\pi}^{like}_{ji}\tilde{\pi}^{see}_{ji}(cs)\\ =& \tilde{\pi}^{like}_{ji}\frac{\bar{s}_i}{cs_i}\end{aligned} \tag{4}$$

$\bar{s}_i$ measures the extent that competition cs affects i's belief that an agent in his market will see his profile. It can be interpreted as an approximation of i's belief about the combination of likelihood that the platform will serve i's profile to j, which depends both on j's search intensity and the platform's matching algorithm. If i believes that j views many profiles, competition has a small effect on whether j sees i's profile. If $\bar{s}_i$ is much greater than cs, competition does not have a direct effect i's belief on matching because it does not alter whether j sees i's profile.

6.5 Model Assumptions

Before going into more technical description of the model, I first lay out the modeling assumptions.

6.6 Actions and States

There are four states in this model, $x_{it} = \{L_{it}, q_{it}, ms_i, cs_i\}$

1. $L_{it} \in \mathbb{Z}_0^{\bar{L}}$: the `like` limit, which is an integer between 0 and $\bar{L}$
2. $q_{it} \in [0, 1]$: quality type of the agent whose profile is served to agent i
3. ms_i: the agent's market size at the start of the session
4. cs_i: the agent's competition size at the start of the session

At the beginning of time t, agent i knows his state L_{it}, which is the number of `likes` he has left. An agent is forward looking in how many `likes` he has left. If he has very few `likes` left, he may become more selective in who he `likes` in the current period to ensure that he can continue searching. Thus, L_{it} is the only "endogenous" state, where the actions at t affect $L_{i,t+1}$.

The set of actions an agent can take depends on the state L. If $L_{it} = 1$, and the agent `likes` the profile at t, i is not able to search in $t+1$. Conditional on entering the search stage, the agent's states are $\{L_{it}, q_{it}\}$, and the set of actions in the selection stage is $a_{it}^{selection} = \{l, nl\}$, where l is `like` and nl is not `like`.

The state transitions are the following. Every time the agent `likes` another agent, L decreases by 1. If $L_{it} = 1$ and $a_{it} = nl$, then the agent's search session ends at t.

$$L_{i,t+1} = \begin{cases} L_{it} & a_{it} = nl \\ L_{it} - 1 & a_{it} = l \end{cases}$$

I make the following assumption for the state transition for q.

Assumption 4. *Conditional on i, and within a session, i believes the draws of q are independent across t.*

This assumption is plausible because of the large variance in the quality of profiles that are shown. Because users are more likely to see users with similar quality, this implies that there is an ordering of quality as the user searches through the entire market. However, the rate that quality changes across searches depends on how precisely the algorithm sorts. If the algorithm presents profiles in order of strictly decreasing quality, then this assumption would not be valid. I validate in the data that profiles are not presented in strictly decreasing quality. In addition, conditional on searching, the median number of searches per session is 29, so most agents do not search enough within a session to see a decline in average quality. Table 7 in the Appendix shows that there is no statistically significant correlation between q_t and t within a search session.
Thus, the state transition for q_{it} is the following.

$$q_{i,t+1} \sim N(\mu_i, \sigma_i^2) \tag{5}$$

μ_i and σ_i^2 describe the distribution q of profiles that i believes the platform will serve to him.

6.7 Current Period Utilities

Before the agent decides whether to end the session or continue searching, he observes idiosyncratic, choice-specific shocks ϵ_{it}^{ss} and ϵ_{it}^{s}. If he chooses to stop searching, he receives utility $u^{ss} = \epsilon_{it}^{ss}$.

Search Stage

In the search stage t, the agent views a profile j and observes and forms beliefs about the following.

1. q_j: the quality type of j
2. $\rho_i(q_j)$: the expected utility that i receives from pursuing a date with j
3. $\tilde{\pi}_{ij}^{m}$: i's belief that i and j will match, conditional on i `liking` j
4. π_{ij}^{d}: the probability that i will pursue a date with j, after matching and learning the true expected date utility

$\rho_i(q_j)$, the expected date utility, captures i's belief on whether the date will be realized. Formally,

$$\rho_i(q_j) = \pi_{ji}^{d}\tilde{\rho}_{ij} \tag{6}$$

where $\tilde{\rho_{ij}}$ is the utility that i receives from a realized date with j, and π_{ji}^{d} is the probability that j also pursues a date with i. However, I cannot observe π_{ji}^{d} the data, so I am unable to separate π_{ji}^{d} from $\tilde{\rho}_{ij}$. For the remainder of the paper, I denote $\rho_i(q_j)$ by ρ_{ij} for ease of exposition.

After observing viewing the profile, i incurs a search cost c^v. Thus, the current period utility for searching is the following.

$$u^{s} = -c^{v} + \epsilon_{it}^{s} \tag{7}$$

Selection Stage

In the selection stage, the agent decides whether to `like` j. If she `likes` j, she gets the following utility.

$$u_{it}^{l}(q_j) = \alpha_i^l + \tilde{\pi}_{ij}^{m}(-c^m + \pi_{ij}^{d} E[\rho_{ij} + \epsilon_{ij}^{d}]) + \epsilon_{ijt}^{l} \tag{8}$$

α_i^l captures heterogeneity in an individual's propensity to `like` other agents. If i `likes` j, i believes there is a probability $\tilde{\pi}_{ij}^{m}$ that they match.

The second part of the right-hand side of Equation 8 is the utility that i receives if they match. If i and j match, i screens j's profile. This may involve having a conversation with j or looking at j's profile more closely. Doing so incurs a screening cost c^m. During screening, i forms a belief about the likelihood that j pursues a date with i. He then updates his expected date utility to $\rho_{ij} + \epsilon_{ij}^{d}$, where ϵ_{ij}^{d} is a shock that is observed by i but unobserved by the econometrician. i then decides whether to pursue a date with j by comparing the expected date utility and utility shock to the value of continuing search in the next time period. Specifically, i pursues a date with j if

$$\Pr\left(\rho_{ij} + \epsilon_{ij}^{d} > V_{i,t+1}^{s} + \epsilon_{i,t+1}^{s}\right) \tag{9}$$

At the time that i decides to `like` j, he does not observe ϵ_{ij}^{d}. Hence the expected value of the expected date utility over ϵ_{ij}^{d} in Equation 8.

ϵ_{ijt}^{l} shock represents a shock that is observed by i before the screening stage. Examples of this shock are the contents of j's profile, such as pictures or the bio, that influence the propensity of i to `like` j. The implicit assumption is that these shocks are independent of i's beliefs about the likelihood of matching and j pursuing a date.

If i does not `like` j, he receives an idiosyncratic shock.

$$u_{it}^{nl} = \epsilon_{ijt}^{nl} \tag{10}$$

Figure 9 provides a detailed summary of the steps of the model.[14] Specifically, it shows during each step, what actions the agent takes and the utilities received from those actions.

[14]Figure 9 is located at the end of the document.

6.8 Value Functions

Each agent i maximizes the following optimization problem.

$$\max_{\vec{a}=(a_0,a_1,a_2...)} E\big[\sum_{t=0}^{T} u_i^a(x_{it})\big]$$

i chooses an action a during each time period to maximize his overall expected utility. The value functions for stopping search, search, and not `like` are below. i subscripts are dropped for ease of exposition.

$$V_t^s(L) = -c^v + \int_q \max\big(V_t^l(L,q) + \epsilon_{j,t}^l, V_t^{nl}(L)) + \epsilon_{j,t}^{nl}\big)dF(q;ms) \tag{11}$$

$$V_t^{nl}(L) = E[\max(V_{t+1}^s(L) + \epsilon_{t+1}^s, \epsilon_{t+1}^{ss})] \tag{12}$$

Equation 11 is the value function for searching. If i searches, he incurs a search cost c^v. The continuation value is the expected maximum utility from the selection stage, where i either `likes` or does not `like` the agent shown at t. However, before deciding to search, he does not know what the quality type of profile that will be shown, so he forms his expectation based on his belief about the distribution of q, which is a function of his market size ms.

Equation 12 depicts the value function for not `liking`. If he does not like, his utility is the expected maximum utility from searching in the next time period, or ending the search session.

$$\begin{aligned} V_{it}^l(q,L) =& \alpha_i^l + \tilde{\pi}_j^{like}\tilde{\pi}_j^{see}(cs)(\pi_i^d\rho - c^m) + \\ & (1-\tilde{\pi}_j^{like}\tilde{\pi}_j^{see}(cs)\pi_i^d)E[\max(V_{i,t+1}^s(L-1) + \epsilon_{i,t+1}^s, \epsilon_{i,t+1}^{ss})] \end{aligned} \tag{13}$$

Equation 13 is the value function for `liking` the profile of type q. As described in the previous section, the first line in this equation is the current period utility that an agent receives. $(1 - \pi_j^{like}\pi_j^{see}(cs)\pi_i^d)$ is the probability that i does not pursue a date with j. If i does not pursue date, his continuation value is the expected maximum utility of stopping search, or entering the search stage in $t+1$.

I parameterize ρ as the following, where λ_i represents the extent that the quality q_j affects the expected date utility.

$$\rho_{ij} = \lambda_{1i} q_j$$

6.9 Beliefs about Market Thickness in Equilibrium

This section demonstrates the indirect effects of market thickness. Clearly, an agent's behavior depends on his beliefs about other agents' behavior. In particular, an agent's decision to `like` another agent depends on his beliefs about whether the other agent will `like` him back. I illustrate the equilibrium concept with the following system of equations. Each agent on this platform has imperfect information about how other agents will behave. Therefore, they form beliefs about other agents' behavior. In particular, i's decision to `like` agent j, conditional on seeing j's profile, is a function of i's belief about j will `like` i, ($\tilde{Like}_{ji}$), and i's expected utility from continuing to search. Similarly, j's decision to `like` i, conditional on seeing i's profile, depends on $\tilde{Like}_{ij}$, and her value of continuing search. If i and j are symmetric, their actions can be written as the following.

$$Like_{ij} = g\big(\tilde{Like}_{ji}, V^s_{i,t+1}\big) \tag{14}$$

$$Like_{ji} = g\big(\tilde{Like}_{ij}, V^s_{j,t+1}\big) \tag{15}$$

As described in the previous section, market and competition have direct effects on the value of continuing search, and the likelihood that another agent sees i's profile. Thus, $V^s_{i,t+1}$ is a function of the market and competition size. Equations 14 and 15 can be rewritten as the following.

$$Like_{ij} = g\big(\tilde{Like}_{ji}(ms_j, cs_j), V^s_{i,t+1}(ms_i, cs_i)\big) \tag{16}$$

$$Like_{ji} = g\big(\tilde{Like}_{ij}(ms_i, cs_i), V^s_{j,t+1}(ms_j, cs_j)\big) \tag{17}$$

This system of equations clearly demonstrates the equilibrium concept of decentralized matching markets. Because ms_i and cs_i affect i's decision to `like` j, j's decision to `like` i is not only a direct function of ms_j and cs_j, but also ms_i and cs_i. Since i does not know ρ_{ji}, j's date utility from going on a date with i, and i does not know j's actions, this platform can be modeled as a Bayesian game with incomplete and imperfect information. In this game, each agent has beliefs over the actions of other agents, given their quality. The equilibrium consists of a strategy profile where an agent with quality q_i has an action that maximizes his expected payoff for each quality profile that he sees. Section 11.4.1 in the Appendix goes into more detail on the characterization of this game.

As the econometrician, each agent's actions that I observe are the optimal response given their beliefs about the actions of other agents. Since market thickness can change beliefs about the actions of other agents', the estimation method must take these indirect effects in the equilibrium into account.

7 Estimation

7.1 Overview

Estimating the equilibrium of a dynamic game is not trivial, especially when there are many agents. Two-step methods, such as Hotz and Miller (1993) and Bajari, Benkard, and Levin (2007), drastically reduce the computational burden of making inference on parameters by avoiding equilibrium estimation. The main idea of these estimators to estimate the equilibrium beliefs in the first stage from the observed data.

Before going into details on the estimation procedure, I provide intuition on how I apply the two-step estimator to this paper. Given an agent's beliefs about his market and competition size, he also has beliefs about how other agents in his market behave, which are conditional on his own market thickness. That is, when a man thinks there are more women in his market, he also knows that the women in his market have more competition. Thus, he updates his beliefs about the likelihood that a woman will `like` him, and optimizes his behavior given these beliefs. But what are these beliefs? Given the assumption about rational beliefs, the econometrician can approximate the beliefs with the observed data. The man's beliefs about how a woman behaves in a market with m men and w women are the same as a woman's actual actions in that market.

A challenge arises from the fact that the treated users beliefs are different from the non-treated users' beliefs. To elaborate, the experiment changes a treated agent's beliefs about his market thickness to m^* and w^*, so he believes other agents are also in a market with m^* men and w^* women and behave accordingly. However, in reality, the other agents in his market behave based on the true market thickness, which is not equal to m^* and w^*. Thus, a treated agent's beliefs about others' actions cannot be approximated with the other agent's actions.

To overcome this estimation challenge, I estimate beliefs using historical data where I can observe agents in markets where the true market thickness is equal to the experimental market thickness. That is, for an agent with market thickness beliefs m^* and w^*, I approximate his belief about other agents' actions with observed actions from agents in markets with m^* men and w^* women. An added advantage of using historical data is the volume of data; this large amount of data allows me to estimate the first stage as flexibly as possible. However, to leverage this data set, I must make the assumption that the prior to the experiment, agent's beliefs about his market thickness are equal to his true market thickness.

Assumption 5. *Prior to the experiment, the agent has rational beliefs about how many men and women are in his market.*

In line with two-step estimation literature, I also make the assumptions that agents have rational beliefs about other agents' actions, and the data, both historical and post-treatment, is generated from the same equilibrium profile (Bajari, Benkard, and Levin 2007).

Another caveat of the pre-experimental data is that the variation in market thickness is not exogenous. However, I do not think this is a major concern due how individuals make inference about other's actions. It seems reasonable that an individual does not update his beliefs from $\tilde{\pi}^{like}(q|\tilde{m},\tilde{w})$ to $\tilde{\pi}^{like}(q|m^*,w^*)$ based on the *causal* effect of market thickness. They may update their beliefs about what they know from past experience or from other markets.[15]

7.2 Identification

There exists significant hetereogeneity in this data, as evidenced by the large standard deviations in Table 1, so I estimate parameters at the individual level. For each agent i, the parameters to be estimated are $\theta_i = [\delta_{1i}, \delta_{2i}, c_i^v, c_i^m, \bar{s}_i, \lambda_i, \alpha_i^l]$. δ_{1i} and δ_{2_i} are identified by the randomization introduced by the experiment. I leverage the exogenous variation generated by the experiment to estimate the parameters. For example, without the experimental variation, I would not be able to

[15] Figure 16 in the Appendix plots the difference between observed actions in pre-experimental data vs post-experimental data.

identify whether an individual's propensity to `like` other agents is due to an inherent characteristic (α_i^l), or a causal effect of his market size.

δ_{i1} captures how changes in market size correlates with the agent's belief about μ_q. c^v is identified by how much an agent searches. To illustrate, if two agents see the same treatment values, the difference in how much they search would inform the search cost parameter. c^m is identified by the frequency at which an agent `likes` another agent and by the quality types of the agents that are liked. When the screening cost increases, agents would become more likely to `like` agents with higher expected date utilities. Intuitively, λ is identified by the extent an agent likes higher quality agents over lower quality agents. If high quality agents are much more likely to be `liked` than low quality agents, I would expected λ to be positive and large. Lastly, α_i^l is identified by individual-level differences in the probability of `liking` that exist after accounting for the other parameters. To further ensure each parameter can be identified, I estimated the model on simulated data and was able to recover the parameters.

7.3 Two-Step Estimation

7.3.1 First Step

The goal of the first step is to estimate an agent's beliefs about other agent's behavior, given his treatment values. I do so by using pre-experimental data from individuals in both the treatment and control group. The relevant beliefs in the second stage is the likelihood that an agent of quality q `likes` i, given values of market thickness. While the number of likes a person has left is an endogenous state variable and affects a person's decision to like, I assume that the i's beliefs about j's behavior is not explicitly dependent on their L state. Because there are so many agents in the market, and L is private information, i optimizes his actions based on the beliefs about the average long-run behavior of other agents over the distribution of other agents' L. This concept has been referred to as a *stationary equilibrium* (Hopenhayn 1992).

Let $\hat{\pi}_m^{like}(q_i, q_j, m_j, c_j)$ denote the belief that a man, with quality q_i, has about the likelihood a woman with quality q_j will `like` him, given her market size ms_j and competition size cs_j. Likewise, $\hat{\pi}_w^{like}(q_i, q_j, m_j, c_j)$ is a woman's belief that a man of type q_j will `like` her. I estimate $\hat{\pi}_m^{like}(q_i, q_j, m_j, c_j)$ in the following way.

1. Divide quality types for men, women and number of men and women into bins.

   (a) $q^w = \{q_1^w, ..q_N^w\}$, $q^m = \{q_1^m, ...q_N^m\}, m = \{m_1, ...m_{N'}\}, w = \{w_1, ..., w_{N'}\}$

2. For each grid point n in q^m, estimate the following logit regression for all men i such that $q_{n-1}^m \leq q_i < q_n^m$. This gives the probability that a woman of type q_j in a location l with m_l men and w_l women `likes` i.

$$Like_{jil} = \beta_1 + \beta_2 q_j + \beta_3 q_j^2 + \beta_4 m_l + \beta_5 m_l^2 + \beta_6 w_l + \beta_7 w_l^2 + \eta_{jil}$$

   (a) Predict for all grid points: $\hat{\pi}_m^{like}(q_n^m, q^w, m, w) = L\hat{ike}(q^w, m, w)$

This results in a 4-dimensional matrix $\hat{\pi}_m^{like}(q^m, q^w, m, w)$, which is the probability that a woman of type q^w `likes` a man of type q^m in a market with m men and w women. The process is similar for estimating women's beliefs.

7.3.2 Second Step

I estimate the parameters in the second step for men and women separately. The effect of the market size factor F^{ms} is specified as follows. Prior to the treatment, the agent's belief about q_{t+1} is that is drawn *iid* from a truncated normal distribution between 0 and 1, with mean $\hat{\mu}_i$ and standard deviation $\hat{\sigma}_i^2$. After being exposed to the market size m_i, his belief about q_{t+1} shifts.

$$\begin{aligned} q_{t+1} &\sim TN(\mu_i, \sigma^2) \\ \mu_i &= \frac{\exp(\tilde{\mu}_i + \delta_1 F_i^{ms} + \delta_2 F_i^{ms} M_i)}{1 + \exp(\tilde{\mu}_i + \delta_1 F_i^{ms} + \delta_2 F_i^{ms} M_i)} \\ \tilde{\mu}_i &= \log(\frac{\hat{\mu}_i}{1 + \hat{\mu}_i}) \end{aligned}$$

μ_i is transformed such that it is between 0 and 1. $\tilde{\mu}_i$ is specified such that if $\delta_1 = 0$ and $\delta_2 = 0$, then $\mu_i = \hat{\mu}_i$. In other words, δ_1 and δ_2 describe the extent the change in market size beliefs shifts μ_i from $\hat{\mu}_i$. σ_i^2 is estimated as the observed standard deviation of quality types of agents of the opposite gender. That is, if i is a man, then $\sigma_i^2 = \hat{\sigma}_w^2$, which is the standard deviation of q for all women in the treatment and control groups. The specifications of the remaining parameters are straightforward.

Likelihood Function

Given the estimated agent beliefs from the first step, I estimate the parameters via maximum likelihood. The value functions are estimated as a finite horizon problem with $T = 500$ with no time discounting across periods within a session. I select $T = 500$ because it is much greater than the observed maximum searches per session. Because the errors are assumed to be iid EV Type 1, the probability of taking each action given states at t can be simplified to the following.

$$\Pr(s_{it}, \ell_{it}) = \frac{\exp(V^s(L_{it}))}{1+\exp(V^s(L_{it}))} \times \frac{\exp(V^\ell(L_{it}, q_{it}))}{\exp(V^\ell(L_{it}, q_{it})) + \exp(V^{n\ell}(L_{it}))} \tag{18}$$

$$\Pr(s_{it}, n\ell_{it}) = \frac{\exp(V^s(L_{it}))}{1+\exp(V^s(L_{it}))} \times \frac{\exp(V^{n\ell}(L_{it}))}{\exp(V^\ell(L_{it}, q_{it})) + \exp(V^{n\ell}(L_{it}))} \tag{19}$$

$$\Pr(d_{it}) = \frac{\exp(\rho_i(q_{it}))}{\exp(V^s(L_{it}-1)) + \exp(\rho_i(q_{it}))} \tag{20}$$

Equation 18 is the probability that i searches and `likes` the agent at t. The first term on the right side is the probability that the agent searches, and second is the probability that the agent `likes`, conditional on searching. The following equation for the probability of searching and not `liking` is similar. $\Pr(d_{it})$ is the probability that i decides to pursue going on a date with the agent shown at t, conditional on matching.

Let y^ℓ_{it} be an indicator for whether i `likes` the profile shown at time t, y^s_{it} be an indicator for whether i searches, y^m_{it} be whether i matches, and y^d_{it} be the indicator for whether i pursues a date. Note that if $y^\ell_{it} = 1$ or $y^m_{it} = 1$, then $y^s_{it} = 1$. Given the observed actions, parameters θ_i, and the probability of each action, the likelihood for an individual is the following.

$$L_i(\theta_i) = \begin{cases} \prod_{t=1}^{T_i} \left[\left(\Pr(s_{it}, \ell_{it}) \Pr(d_{it})^{y^m_{it} y^d_{it}} (1 - \Pr(d_{it})^{y^m_{it}(1-y^d_{it})} \right)^{y^s_{it} y^\ell_{it}} \left(\Pr(s_{it}, n\ell_{it}) \right)^{y^s_{it}(1-y^\ell_{it})} \right] & TotalViews_i > 0 \\ 1 - (\Pr(s_{it}, \ell_{it}) + \Pr(s_{it}, n\ell_{it})) & TotalViews_i = 0 \end{cases}$$

$TotalViews_i$ is the number of profiles views per session. If $TotalViews_i = 0$, that means i did not search at all during the session.

Ackerberg's Importance Sampling Method

Each individual has a unique draw of ms_i and cs_i, so in order to calculate the likelihood of the data, the value functions need to be calculated for each individual. This poses a computational challenge;

with 66,000 agents in the data, if the value functions take 0.1 seconds to converge for each agent, each iteration in the likelihood optimization function 1.8 hours. For an optimization that takes 1000 iterations to converge, the entire estimation procedure would take 75 days. To reduce the computational burden, I estimate each parameter as a random effect using Ackerberg's importance sampling method (Ackerberg 2005), along with parallelization of value function iteration. Details of this estimation are included in the Appendix. While random effects can capture heterogeneity in the data, it does not allow me to estimate how the parameters vary with population size. For instance, agents in larger markets may inherently be more selective than those in smaller markets. Ideally, I would estimate all parameters as a flexible function of M, but the data is not rich enough for such a flexible model. Therefore, in addition to estimating the parameters separately for men and women, I also estimate parameters separately for agents in small ($M < 11,157$) vs large markets ($M \geq 11,157$).[16]

8 Results

8.1 Parameter Estimates

Table 5 presents the estimates of the parameters from the structural model. Standard errors are estimated from the information matrix. Agents classified as in the small market on average had a population size of 4,000, and agents in the large market have an average population size of 30,000. The key parameters are δ_1, δ_2, and $\bar{s}$. For all agents in both types of markets, the market size factor F^{ms} is positive and significant. This implies that agents who believe they have a larger population size behave as if they have higher expectations for the quality of the profile shown in the next time period. The estimate of $\bar{s}$ informs the effect of competition size. If $\bar{s} > cs$, then competition does not have an effect on an agent's beliefs about whether an agent in his market sees his profile. In small markets, the mean of the estimated distribution of $\bar{s}$ is greater than the average competition size, indicating that that competition does not have a large direct effect on behavior for the average agent. However, in large markets, especially for women, the estimate of the mean of $\bar{s}$ is greater than the average competition size, which reflects that competition does have an effect on beliefs about visibility.

[16] The average population size for all the agents exposed to the treatment is 11,157.

	Men		Women	
Parameter	Mean	SD	Mean	SD
Small Market ($\bar{M}$ = 4,000)				
Market size coefficient (δ_1)	0.336***	5.080***	1.254***	4.964***
	(0.022)	(0.008)	(0.050)	(0.008)
Market size coefficient - Heterogeneity with M (δ_2)	-0.098***	5.126***	-0.016***	5.079***
	(0.042)	(0.029)	(0.050)	(0.030)
Search cost (c^v)	2.478***	1.131***	1.689***	1.264***
	(0.017)	(0.014)	(0.015)	(0.014)
Screening cost (c^m)	4.096***	3.150***	4.234***	3.081***
	(0.035)	(0.014)	(0.044)	(0.029)
Belief about j's search in 10,000's ($\bar{s}$)	1.191***	1.280***	0.406***	0.501***
	(0.017)	(0.022)	(0.003)	(0.004)
Date utility coefficient (λ_i)	13.395***	35.778***	9.304***	22.455***
	(0.139)	(10.524)	(0.218)	(7.155)
Individual RE for liking (α_i^l)	1.541***	2.170***	0.093***	2.551***
	(0.025)	(0.058)	(0.024)	(0.078)
Log-likelihood	-474,268		-284,427	
Large Market ($\bar{M}$ = 30,000)				
Market size coefficient (δ_1)	0.460***	5.150***	0.997***	4.807***
	(0.064)	(0.011)	(0.074)	(0.011)
Market size coefficient - Heterogeneity with M (δ_2)	0.013***	5.189***	0.173***	4.923***
	(0.064)	(0.039)	(0.074)	(0.045)
Search cost (c^v)	2.146***	1.273***	1.301***	1.526***
	(0.016)	(0.014)	(0.039)	(0.030)
Screening cost (c^m)	4.150***	3.163***	4.279***	3.098***
	(0.059)	(0.039)	(0.068)	(0.045)
Belief about j's search in 10,000's ($\bar{s}$)	1.186***	1.308***	0.251***	0.853***
	(0.007)	(0.019)	(0.013)	(0.013)
Date utility coefficient (λ_i)	5.048***	47.958***	11.388***	34.496***
	(0.588)	(35.601)	(0.508)	(24.063)
Individual RE for liking (α_i^l)	1.628***	2.219***	-0.152***	2.531***
	(0.025)	(0.086)	(0.036)	(0.116)
Log-likelihood	-249,465		-121,214	

Note: *p<0.1; **p<0.05; ***p<0.01

Table 5: Model estimation results. The top half of the table present the parameter estimates for men and women (separately) for agents in markets of below average population size, which I call "small markets". The average market size of the small markets is 4,000. The bottom half of the parameters are for men and women in large markets, where the average population size is 30,000.

8.2 Predicted Actions

The magnitudes of the parameter estimates themselves are not straightforward to interpret, so I illustrate the effects of market thickness on selectivity with the predicted `like` rates, given the parameter estimates, in Figures 10 and 11. These figures illustrate the effect of market thickness for the average agent with the median quality type.

In these figures, the x-axis is the quality of another agent in percentiles, and the y-axis is the relative probability that the agent `likes` the agent of that quality percentile. The lines represent either changes in market size, while holding competition size constant, or competition size, while holding market size constant. To maintain the confidentiality of the data, the baseline `like` probability is obscured. All reported `like` probabilities are relative to the `like` probability of the 10^{th} percentile quality, for the smallest market or competition size depicted in the plot. To illustrate, the plot in the top left corner of Figure 10 depicts how the `like` rate changes for a male agent in a small market when he sees profiles of different qualities. Each line represents a different value of market size, while holding competition size constant at $cs = 4,000$. The plotted `like` rate is the `like` rate relative to the `like` rate for the 10^{th} percentile q_j when the market size is 3,000. Thus, for all plots, the `like` rate for the 10^{th} percentile q_j will always be 0. In a market size of 5,000, the probability that the agent `likes` a 10^{th} percentile q_j is 5% less than when the market size is 3,000. Figure 17 in the Appendix shows that when beliefs about market size increases, women are more likely to search.

In large markets, both the average man and women become more selective when market size increases, and less selective when competition sizes increases. However, the magnitudes of the effects are different for men and women. A 50% increase in market size decreases the absolute `like` rate for the median quality user by about 8% for men, and 2% for women. However, women generally have a lower `like` rate than men, so proportionally, the 2% decrease in the `like` rate is relatively a larger effect for women than the 8% is for men. Conversely, a 50% increase in competition increases the `like` rate for the median quality user by 1% for men, and 5% for women.

These patterns estimated by the model are consistent with the linear regression estimates. Effects on selectivity are larger in magnitude in large markets. In addition, the estimated `like` rates are similar to what is observed in data for both men and women, indicating that the model is a reasonable approximation to the true data generating process. The plots of the predicted `like` rates from the model compared to the actual `like` rates observed in the data are presented in Figure 19 in the Appendix.

A key takeaways from these plots is behavior in small markets is different from behavior in large markets. In small markets, there is the effect of market thickness on selectivity for women with the median quality is much smaller than the effect for women in large markets. In addition, men and women behave very different in responses to changes in market thickness. Men react more strongly to changes in market size, while women are more responsive to changes in competition size.

9 Implications for Platform Design

The previous section shows that actions change when *beliefs* about market thickness change. In this section, I simulate how actual changes in market thickness affect outcomes. The goal of these counterfactuals is to inform platform design by first evaluating how changes in market thickness affect matching outcomes in equilibrium, and second, how changing the like limit can mitigate the negative effects of selectivity when more people join the platform.

In the first counterfactual, I simulate a 25% increase in market thickness: there are 25% more men and 25% more women on the platform. In these counterfactuals, the new users have the same quality distribution as the exisiting users on the platform.

The second counterfactual involves increasing agents on one side of the market. Many marketplaces, especially dating markets, have imbalances in the number of agents on each side of the market. To try to balance the ratio, platforms may target marketing towards women or disincentivize men from joining. I refer to selective targeting towards one side of the market as "gender gating". I simulate how increasing the number of women on the platform by 25%, while holding the number of men constant.

However, as more women join the platform, that introduces more competition to women, which may make matching more difficult. How can the platform mitigate negative effects of market thickness? One way that the platform can influence selectivity, other than through market thickness, is through the like limit. Therefore, in the last counterfactual, I evaluate how increasing the like limit for women when women have more competition can affect matching outcomes for all platform members. More specifically, there are 25% more women on the platform, and the like limit doubles for women only.

9.1 Counterfactual Estimation

I compare all counterfactuals to a baseline, which is a market of population size M with $0.75M$ women and M men. Rather than selecting a baseline of a market with the same number of men and women, I choose to simulate markets with more men than women, since that is more realistic to actual online dating conditions. I estimate all counterfactuals for both a large market ($M = 30,000$) and a small market ($M = 4,000$).

To calculate the counterfactuals, I first solve for the equilibrium for a market with population size M with m men and w women by iterating the value functions until the actions of women equal men's beliefs about women's actions, and vice versa. I then simulate the actions for the m men and w women for $T = 500$, meaning they have a maximum of 500 search opportunities. Details on estimation of the counterfactuals are provided in the Appendix. Note that the outcomes of the simulations should be interpreted with the possibility of many equilibriums, and the estimated equilibrium in this paper may just be one of them.

The baseline market is 3,000 women and 4,000 men for a small market, and 22,500 women and37,500 men for a large market. The outcomes variables I consider are the percentage change in the number of matches a man or woman gets compared to the baseline, and the change in average match quality, conditional on getting a match; The goal of the counterfactuals presented in this paper is not to prescribe the optimal solution for the firm, but to show that market thickness and policies can have different effects on match quality and quality, and firms can design their policies to fit their needs.

9.2 Counterfactual Results

Figures 12 and 13 display the changes in match quantity and match quality, respectively, under the different counterfactuals.[17] The error bars represent the 2.5 and 97.5 percentiles of the 50 simulations for each counterfactual. Note that these are not traditional standard errors, but they illustrate the range of outcomes due to the differences in simulations. For the remainder of this section, I refer to "significant" as the 95% confidence interval for the 50 simulations do not overlap with 0.

In both small and large markets, when 25% more men and women join the platform, both match quantity and quality decrease for men, while there is essentially no change in matching outcomes for women. In terms of overall matches, I find that matching exhibits decreasing returns to scale as market thickness increases; in small markets, the total number of matches formed on the platform increase by 11%, and by 2.5% in large markets.

[17] Tables 8 and 9 in the Appendix report the numbers.

When only more women join the platform, men get more matches, as expected. On the other hand, women get fewer matches, but the difference is significant in small markets only. However, contrary to what theory predicts, the match quality for men in large markets significantly decreases. In small markets, match quality for men increases, but is not significantly different from 0. This difference in matching outcomes between small and large markets is because women respond much more strongly to competition in large markets. Not only do women increase their like rate, there is heterogeneity in how much they increase their like rate based on their own quality type. Lower quality types increase their like rate the most in response to competition, so men become more likely to match with lower quality types.

Lastly, third bar in Figures 12 and 13 show the effects on matches when there are more women, but also the like limit doubles for women only. Intuitively, an increase the like limit decreases the implicit cost of each like. Thus, women can be more "risky" and like a high quality user who is less likely to like her back, compared to a lower quality user. In both markets, men get more matches, and get matches with higher quality women. The increase between match quantity and quality is significantly different compared to the counterfactual where there is an increase in women but the like limit stays the same only for large markets. There is not a significant difference in matching outcomes for women.

9.3 Discussion

In summary, given the causal effects of market thickness estimated from the model, the counterfactuals show that the same policy implemented in small versus large markets may lead to different outcomes, indicating that platforms should customize policies based on market thickness. The main difference between small and large markets is their sensitivity to the direct effect of competition: increasing competition has a smaller direct effect for agents in small markets than agents in larger markets.

There are four main takeaways from the counterfactuals. (1) Increasing the number of men and women on the platform may reduce match quantity and quality. (2) Matching has decreasing returns to scale as there are more agents on both sides of market. (3) Increasing market size does not necessary imply an increase in match quality, as agents on both sides to the market react to this change in the market. (4) Lastly, all these effects are mainly driven by changes in selectivity in response to market and competition size, and by platform design features. By inducing agents to become less selective through changing their like limit, the firm may be able to grow one or both sides of the market without significantly worsening match quantity and quality.

10 Conclusion

Understanding search in large, decentralized matching markets is important to the success of the platform, as agents search to find a match. Search in matching markets is complicated by the fact that matching is two-sided. An individual's decision in how much to search and to propose a match to depends on his beliefs about the availability of other agents on the platform. This papers studies how the role of beliefs about market thickness, defined as the number of agents on the same and opposite sides of the market, affects search and selectivity, and guides how firms should design their platforms in light of this effect.

Through a field experiment that varies information about market thickness, I document a causal effect between an individual's belief about how many potential matches and competitors and their behaviors and outcomes on the platform. When agents believe they have more choices, hence a larger market size, they become more selective in who they want to match with. On the other hand, when they believe they have more competitors, they become less selective.

To measure equilibrium matching outcomes under changes in market thickness and different platform policies, I build a structural model of search and selectivity which incorporates beliefs about market thickness. When taking changes in selectivity in response to market thickness into account, this paper finds decreasing returns to scale for matching when market thickness increases. In addition, matching outcomes are not always as predicted by theory. For instance, an increase in market size does not necessarily increase match quality because users on both sides of the market are changing their selectivity in response to the change in their market and competition sizes. Lastly, since changes in selectivity drives these results, policies that affect selectivity can be used to mitigate negative effects from platform growth. One way the platform can influence selectivity is by placing a limit on how many match proposals a person can send.

A limitation of this paper is that a user's quality type is measured vertically. I assume that users prefer to match with higher quality users. While the model allows for some individual-level heterogeneity in the evaluation of another user's quality, I cannot fully capture horizontal preferences. Second, a limitation of the data is that I do not observe a user's characteristics, besides age and gender. If I could observe more detailed characteristics, the quality score could be measured using those characteristics, which is a better way to capture preferences. Lastly, I cannot observe when dates happen, which is a common problem in online dating sites. Firms cannot observe when a match on the platform results in a date or marriage, as people often communicate with their matches through offline channels. Due to the lack of data, I make strong assumptions about the likelihood of a match resulting in a date.

One avenue of future research is to study the long term effects of changes in market thickness, and how people learn about market conditions and adapt their search behavior over time. In addition, this paper does not allow for changes in the types of users who join or leave the market. Hopefully further research can expand on how market thickness not only changes the behavior of agents on the platform, but also how it affects the composition of agents who join and leave.

References

Abdulkadiroğlu, Atila et al. (2005). "The Boston public school match". In: *American Economic Review* 95.2, pp. 368–371.

Acemoglu, Daron and Robert Shimer (1999). "Efficient Unemployment Insurance". In: *Journal of Political Economy* 107.5, pp. 893–928.

Adachi, Hiroyuki (2003). "A search model of two-sided matching under nontransferable utility". In: *Journal of Economic Theory* 113.2, pp. 182–198.

Aguirregabiria, Victor and Pedro Mira (2007). "Sequential estimation of dynamic discrete games". In: *Econometrica* 75.1, pp. 1–53.

Bajari, Patrick, C Lanier Benkard, and Jonathan Levin (2007). "Estimating dynamic models of imperfect competition". In: *Econometrica* 75.5, pp. 1331–1370.

Blanchard, O.J. and Peter Diamond (1994). "Ranking, unemployment duration, and wages". In: *The Review of Economi* 61.3, pp. 417–434.

Bleakley, Hoyt and Jeffrey C Fuhrer (1997). "Shifts in the Beveridge Curve, Job Matching, and Labor Market Dynamics". In: *New England Economics Review* September, pp. 3–19.

Buchholz, Nicholas (2016). "Spatial Equilibrium, Search Frictions and Efficient Regulation in the Taxi Industry". In: *Working paper*, pp. 1–64.

Burda, Michael C. and Stefan Profit (1996). "Matching across space: Evidence on mobility in the Czech Republic". In: *Labour Economics* 3.3, pp. 255–278.

Burdett, Kenneth and Melvyn G Coles (1999). "Long-Term Partnership Formation: Marriage and Employment". In: *The Economic Journal* 109.456, F307–F334.

Burdett, Kenneth, Shouyong Shi, and Randall Wright (2001). "Pricing and Matching with Frictions". In: *Journal of Political Economy* 109.5, pp. 1060–1085.

Chu, Junhong and Puneet Manchanda (2016). "Quantifying cross and direct network effects in online consumer-to-consumer platforms". In: *Marketing Science* 35.6, pp. 870–893.

Coles, Melvyn G and Eric Smith (1998). "Marketplaces and Matching". In: *Source: International Economic Review INTERNATIONAL ECONOMIC REVIEW* 39.1, pp. 239–254.

Cullen, Zoe and Chiara Farronato (2015). “Outsourcing Tasks Online : Matching Supply and Demand on Peer-to-Peer Internet Platforms”. In: *Working Paper* October.

Diamond, Peter (1982). “Aggregate Demand Management in Search Equilibrium”. In: *Journal of Political Economy* 90.5, pp. 881–894.

Diehl, Kristin and Cait Poynor (2010). “Great expectations?! Assortment size, expectations, and satisfaction”. In: *Journal of Marketing Research* 47.2, pp. 312–322.

Fisman, Raymond et al. (2006). “Gender differences in mate selection: Evidence from a speed dating experiment”. In: *The Quarterly Journal of Economics* 121.2, pp. 673–697.

Fradkin, Andrey (2015). “Search Frictions and the Design of Online Marketplaces”. In: *Working Paper*.

Gale, David and Lloyd S Shapley (1962). “College admissions and the stability of marriage”. In: *The American Mathematical Monthly* 69.1, pp. 9–15.

Gan, Li and Qi Li (2004). “Efficiency of Thin and Thick Markets”. In:

Gautier, Pieter A., Jose L. Moraga-González, and Ronald Wolthoff (2007). “Estimation of Search Intensity: Do Non-Employed Workers Search Enough?” In: *SSRN eLibrary* November, pp. 1 –61.

Gregg, Paul and Barbara Petrongolo (1997). *Random or Non-random Matching? Implications for the Use of the UV Curve as a Measure of Mathing Effectiveness*. Tech. rep. 348.

Hagiu, Andrei (2006). “Pricing and commitment by two-sided platforms”. In: *The RAND Journal of Economics* 37.3, pp. 720–737.

Halaburda, Hanna, Mikołaj Jan Piskorski, and Pinar Yildirim (2017). “Competing by Restricting Choice: The Case of Matching Platforms”. In: *Management Science*.

Hitsch, Günter, Ali Hortaçsu, and Dan Ariely (2010a). “Matching and sorting in online dating”. In: *American Economic Review* 100.1, pp. 130–163.

— (2010b). “What makes you click?-Mate preferences in online dating”. In: *Quantitative Marketing and Economics*, pp. 1–35.

Hopenhayn, Hugo A (1992). “Entry, exit, and firm dynamics in long run equilibrium”. In: *Econometrica: Journal of the Econometric Society*, pp. 1127–1150.

Hotz, V Joseph and Robert A Miller (1993). “Conditional choice probabilities and the estimation of dynamic models”. In: *The Review of Economic Studies* 60.3, pp. 497–529.

Howitt, Peter and R Preston Mcafee (1987). “Costly Search and Recruiting”. In: *International Economic Review* 28.1, pp. 89–107.

Iyengar, Sheena S and Mark R Lepper (2000). “When Choice is Demotivationg: Can One Desire Too Much of a Good Thing?” In: *Journal of Personality and Social Psychology* 79.6, pp. 995–1006.

Kanoria, Yash and Daniela Saban (2017). “Facilitating the search for partners on matching platforms : Restricting agent actions”. In: pp. 1–71.

Lee, Soohyung and Muriel Niederle (2014). “Propose with a rose? Signaling in internet dating markets”. In: *Experimental Economics* 18.4, pp. 731–755. arXiv: `arXiv:1011.1669v3`.

McCarney, Rob et al. (2007). “The Hawthorne Effect: a randomised, controlled trial”. In: *BMC medical research methodology* 7.1, p. 30.

Mortensen, Dale T. (1982). “The Matching Process as a Noncooperative Bargaining Game”. In: *NBER Chapters*, pp. 233–258. arXiv: `arXiv:1011.1669v3`.

Petrongolo, Barbara and Christopher Pissarides (2001). “Looking into the black box: a survey of the matching function”. In: *Journal of Economic Literature* XXXIX.June, pp. 390–431.

Pissarides, Christopher A (1984). “Search intensity, job advertising, and efficiency”. In: *Journal of labor Economics* 2.1, pp. 128–143.

Reutskaja, Elena et al. (2011). “Search Dynamics in Consumer Choice under Time Pressure: An Eye- Tracking Study”. In: *The American Economic Review* 101.2, pp. 900–926. arXiv: `00368075`.

Rochet, Jean-Charles and Jean Tirole (2003). "Platform competition in two-sided markets". In: *Journal of the european economic association* 1.4, pp. 990–1029.
Roth, Alvin E and Elliott Peranson (1999). "The redesign of the matching market for American physicians: Some engineering aspects of economic design". In: *American economic review* 89.4, pp. 748–780.
Rubin, Donald B (1974). "Estimating causal effects of treatments in randomized and nonrandomized studies." In: *Journal of educational Psychology* 66.5, p. 688.
Sahni, Navdeep S and Harikesh Nair (2016). "Native advertising, sponsorship disclosure and consumer deception: Evidence from mobile search-ad experiments". In: *Sponsorship Disclosure and Consumer Deception: Evidence from Mobile Search-Ad Experiments (February 23, 2016).*
Shimer, Robert and Lones Smith (2001). "Matching, Search, and Heterogeneity". In: *Advances in Macroeconomics* 1.1.
Stevens, Margaret (2007). "New microfoundations for the aggregate matching function". In: *International Economic Review* 48.3, pp. 847–868.
Tucker, Catherine and Juanjuan Zhang (2010). "Growing two-sided networks by advertising the user base: A field experiment". In: *Marketing Science* 29.5, pp. 805–814.

11 Appendix

11.1 Natural Variation in Market Thickness

I use a 2-week sample of the historical data to provide more information on how much the market thickness naturally varies in this setting. The sample is all users who have opened the app from one of the countries the experiment is implemented in in the two-week window. For each location[18], I determined the number of distinct users who opened the app, for each day-hour. I then calculated the average and standard deviation of number of distinct users over each location - day. Figure 14 plots the coefficient of variation ($\frac{SD}{Mean}$) for each hour in a selected location that seems representative of all other locations. That is, each observation in that plot is the standard deviation of the number of distinct users who open that app from that location-hour divided by the average number of users who open the app from that location-hour. There appears to be significant variation in the number of users at any given hour. This gives an idea of much the market thickness varies for a given hour. To compare this to the experiment, the maximum variation in the experiment is 25%. Thus much of the natural variation at the hour level is greater than the variation induced by the experiment.

[18] 25x25 mile

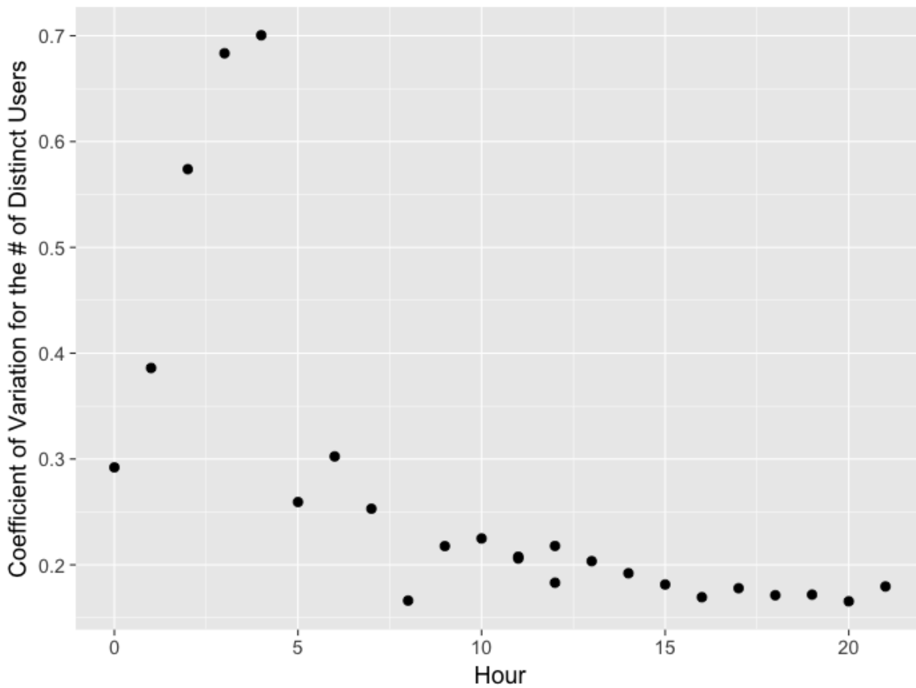

Figure 14: The coefficient of variation in the number of distinct users who open the app for each hour of the day, averaged over a two-week period. The sample in this figure is all users who open the app from a location in one of the countries that the experiment is implemented in during the two-week time period. The selected location to be representative of the other locations. The time zone has been shifted to preserve confidentiality.

11.1.1 Session-level Randomization

There are two levels of randomization. The first is across users. Users in the treatment group can be exposed to the treatment, while users in the control group are not exposed. The second level of randomization is within the treated users at the session level. That is, treatment values m and w vary across users in the treatment group, and also within users at the session level.

Treated users are exposed to the pop-up message if (1) they open the app from within one of the selected grids, and (2) it has been at least 6 hours since they have seen the pop-up. A new session starts if it has been at least 6 hours since the user was last exposed to the treatment. Each time the treatment is shown, the treatment values are re-randomized. Figure 15 shows how the treatment values differ across sessions. For a treated user, at $t_{treatment1}$, he sees the treatment for the first time, with $m = 588$ and $w = 592$. After viewing some profiles, the user opens the app again at $t_{treatment2}$. If $t_{treatment2}$ is more than 6 hours after $t_{treatment1}$, the user is then shown the pop-up message again with new values of m and w. Conditional on the user's location, draws of m and w are independent across users and across sessions.

11.2 Randomization Check

I ensure the treatment within the treatment group is sufficiently randomized by comparing the correlations between the treatment variables and pre-treatment user characteristics.

	ρ	P-val
F^{cs}	0.001	0.552
Age	0.002	0.079
Gender	0.000	0.940
Account Age	0.000	0.867
Profile Views	0.000	0.839
Like Rate	-0.002	0.168
Matches	-0.000	0.803

Table 6: The correlation between the market size factor F^{ms} and other treatment variables, and pre-experimental characteristics.

The correlations are shown in Table 6. Table 6 confirms that there is no correlation between F^{ms}, F^{cs}, user characteristics, and pre-treatment behavior, such as the number of profiles the users have viewed, their `like` rates, and the number of matches they've made.

	Dependent variable:	
	q_j	
	(1)	(2)
Index	-0.00000^{***}	
	(0.00000)	
q_i		0.176^{***}
		(0.021)
Agent FE	Y	N
Location FE	N	Y
Filter FE	N	Y
Observations	19,877,176	20,932
R^2	0.780	0.240
Note:	*p<0.1; **p<0.05; ***p<0.01	

Table 7: Linear regressions of q_j, the quality of the profile shown to i, on the order that the profiles are shown (Index) and on i's quality (q_i). The coefficient of Index in Column 1 is -1.519e-06. Column 1 includes agent fixed effects (standard errors are clustered by agent), and Column 2 uses only the first profile view for each agent.

11.3 Quality Score Calculation

Quality scores are calculated based on the observed pre-experimental data for how likely another user is to like the focal user.

$$q_i = \frac{\text{\# users who like } i\text{'s profile}}{\text{\# users who see } i\text{'s profile}}$$

The objective for the quality score is to get a measure on how "desirable" a user is, and to be able to compare quality scores across users. As mentioned in the paper, one concern is that since markets may be different in their search behavior and selectivity, the quality scores for users across markets may not be comparable. For example, if users in large cities are more selective, then the users in those cities will have smaller proportion of people who have liked their profile. Compared to a user in a small town, where people are less selective, users in large cities will have a lower raw quality score, which can lead to the potentially-incorrect conclusion that people in large cities are of lower quality.

To address this concern, quality scores are represented by percentiles, where a user's raw quality score is compared to all other users' who are in grids of a similar population size. More specifically, I separate the grids in to three groups (small, medium, and large), based on the tercile of their population size. For each user's raw quality score, I compare it to the quality score of all other users who are in the same tercile of population size grids.

11.4 Structural Model

11.4.1 Bayesian Game

I describe this setting more formally as a Bayesian game of incomplete and imperfect information. Each agent i has a type λ_i, where $\rho_i(q) = \lambda_i q$. In other words, each agent values going on a date with another player of type q differently. Each agent knows his own λ_i but not those of other agents. q can be though of as a noisy signal for λ. i also does not know the history of any other agents.

When agent i, with quality q_i meets another agent j with quality q_j, i knows ρ_{ij}, i's date utility from going on a date with j, but does not know ρ_{ji}, j's date utility. Thus, he has incomplete information about ρ_{ji} so he does not know how j will behave. But given q_j, i has beliefs about λ_j, and thus, also has beliefs about whether j `likes` i.

This Bayesian game consist of

- A set of agents $\mathcal{I}$
- A set of actions for each agent i: $a_i = \{l, nl\}$
- A set of states: $x = \{L, ms, cs\}$
- A set of types for each i: $\lambda_i \in \mathbb{R}$
- A probability distribution over q: $f(q)$
- Each i has beliefs over the likelihood that another agent of type q_j will `like` i
- A payoff function for each i and strategy, given states and beliefs about another agent of type q_j: $u_i^l(q_j, L, ms, cs), u_i^{nl}(q_j, L, ms, cs)$

In the BNE, each agent with type q has a strategy that maximizes his expected utility for each quality of profile that he sees. Since this is a finite game (there are a finite number of agents, and actions), a BNE is guaranteed to exist. However, because the type set (λ) is not compact, then a pure strategy BNE is not guaranteed to exist.

11.4.2 Ackerberg Importance Sampling Method

I estimate the random effects using Ackerberg's importance sampling method. In some applications, the set of draws from the importance density is not individual specific. In words, $\theta_{ir} = \theta_r$ for all i. To promote more mixing, in this paper, I made separate draws of θ_{ir} for each individual. Below are the steps for each individual i with parameters θ_i.

1. Make R draws of $\theta_{ir} \sim f(\bar{\theta}_h, \Sigma_h)$ where h represents the importance density. $\Omega_h = \{\bar{\theta}_h, \Sigma_h\}$ denotes the hyperparameters of the importance density.

    - The search and screening costs, and $\bar{s}$ are non-negative, so those are drawn from truncated normal importance densities, while $\alpha^l, \delta_1, \delta_2, \lambda$ have unbounded support. The latter parameters are drawn from a multivariate normal distribution. To simplify computation, the covariance between $c^v, c^m, \bar{s}$ are fixed to be 0, while the covariance matrix for $\alpha^l, \delta_1, \delta_2, \lambda$ is estimated. Σ_h denotes the covariance matrix for all parameters, while the covariance between $c^v, c^m, \bar{s}$ are fixed to be 0.

2. Compute the likelihood $\mathscr{L}_i(\theta_{ir})$ for all R draws.

3. Let $\Omega_g = \{\bar{\theta}_g, \Sigma_g\}$ be the hyperparameters of θ. Compute $\mathscr{L}_i(\Omega_g, \theta_i) = \frac{1}{R}\sum_{r=1}^{R} \mathscr{L}_i(\theta_{ir}) \times \frac{h(\theta_{ir};\Omega_h)}{g(\theta_{ir};\Omega_g)}$

4. Maximize the following log likelihood.

$$\max_{\Omega_g} \sum_i \log \mathscr{L}_i(\Omega_g)$$

In this paper, some of the parameters are restricted to be positive, such as $\bar{s}$, the screening cost, and the search cost, while others have infinite support. Since I cannot draw from a multivariate distribution that is truncated for some parameters, I set the diagonals for the covariance matrix to 0.

11.4.3 Estimating Beliefs from Pre vs. Post Experiment Data

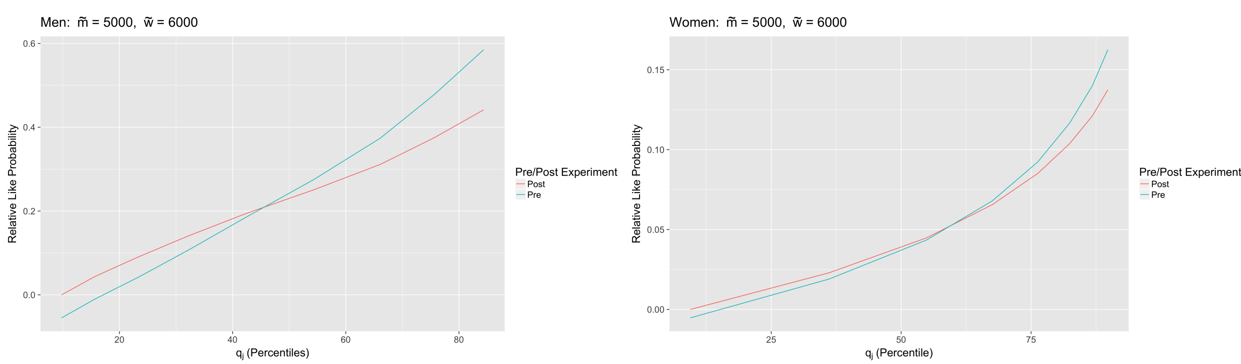

Figure 16: A comparison of the observed `like` rates using pre-experimental and post-experimental data for men (left) and women (right). The plotted `like` probability is relative to the `like` probability for the 10^{th} percentile quality using in the post-experimental data.

11.4.4 Value Functions for Search

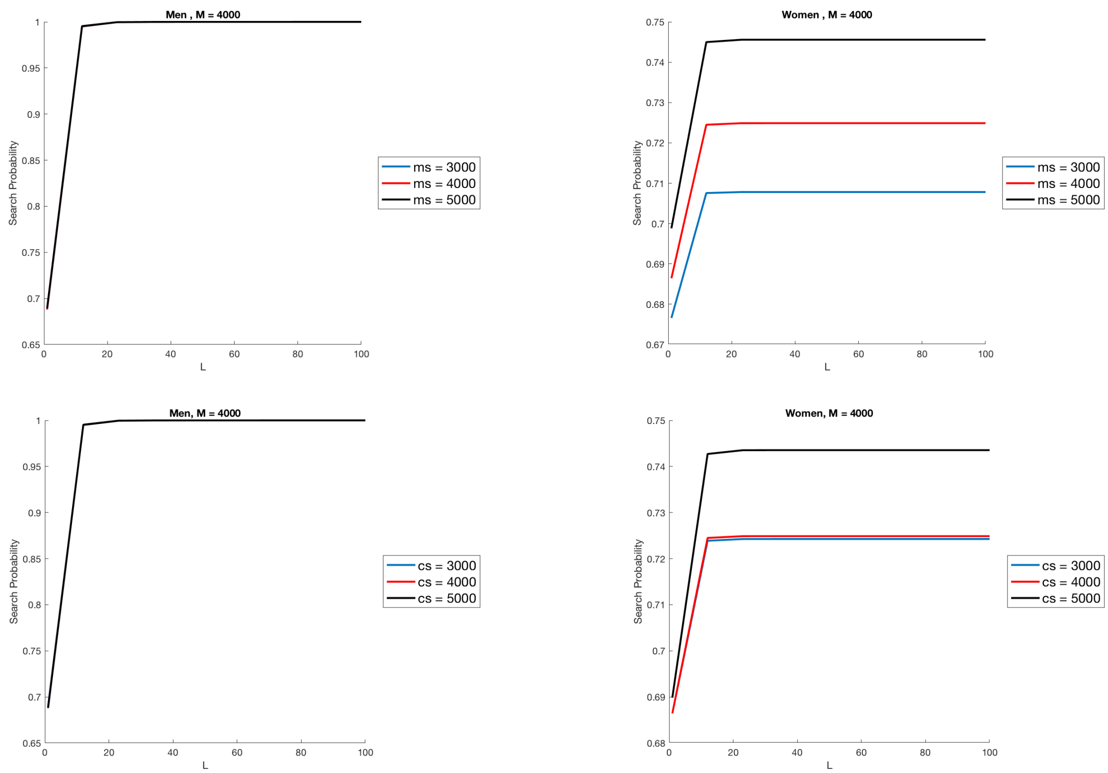

Figure 17: Search and market size in a small market.

11.4.5 Comparing Model Estimates to Observed Data

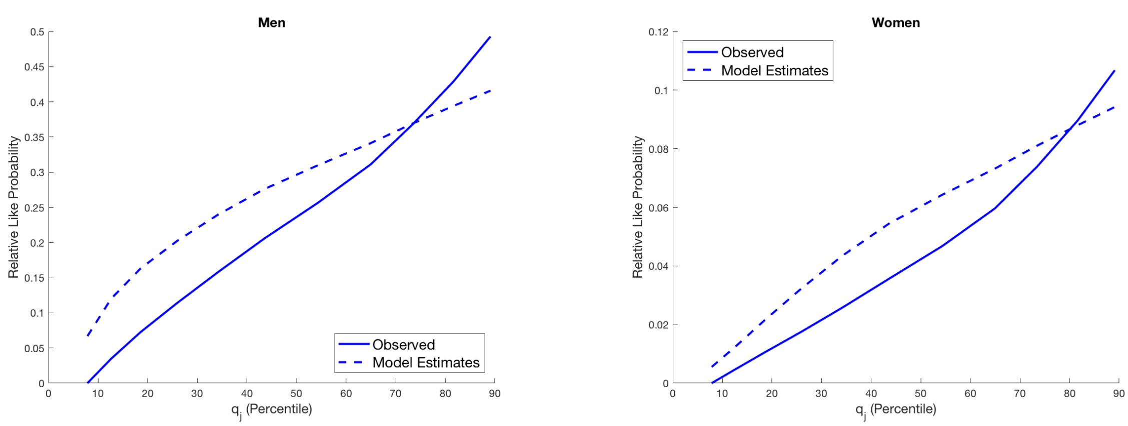

Figure 19: Model estimates versus observed data.

Figure 19 shows the predicted `like` rates based on model estimates versus the observed `like` rates. The goal of these plots is to evaluate whether the model and estimates are able to recover observed patterns in the data. I do by plotting the probability that an average man or woman `likes` another agent of type q_j for the observed data, and as predicted by the model estimates. The observed data is plotted based on the following logistic regression.

$$Like_{ij} = \beta_0 + \beta_1 q_i + \beta_2 q_j + \beta_3 q_j^2 + \beta_4 ms_i + \beta_5 cs_i + \epsilon_{ij} \tag{21}$$

The graph on the left in Figure 19 plots the `like` probability for men, and for the women on the right. The x-axis is the quality of an agent that an agent sees, and the y-axis is the relative probability that the agent `likes` that profile for the an agent with an average quality and with an average market and competition size. The `like` probability (or `like` rate) is relative to the `like` probability for a 20^{th} percentile quality profile. The solid line is the observed data, which is the predicted `like` probability from Equation 21, and the dotted line is the predicted `like` probability from the model.

Figure 19 shows that the model does a decent job at simulating the data generating process to produce the observed data. Especially for women, the model estimates are very similar to the observed data.

Counterfactual Estimation

11.4.6 Equilibrium Estimation

While the equilibrium was not calculated in order to estimate the model parameters, the equilibrium must be calculated for counterfactual analysis. The general idea for equilibrium is that the an agent's beliefs about the behavior of an agent on the other side of the market equals the actions of that agent, and vice versa. When agent i changes his actions in response to a change on the platform, the other agent j's picks the optimal action in response to i's actions. In response, because j changed his actions in response to i, i then changes his actions in response to his beliefs about j's behavior. This sequence iterates until actions and beliefs converge. The following provides details on this estimation procedure.

For ease of exposition, I illustrate the procedure with two agents: man i and woman j. Let policy p denote the original policy, and p' denote the new policy. $\pi_i^{like}(q_j|\theta, p, b_{ij}^{like}, b_{ij}^{search})$ is the probability that i `likes` j, given parameters θ, policy p, i's beliefs about whether j `likes` i, conditional on searching b_{ij}^{like}, and i's beliefs about j's search intensity b_{ij}^{search}. $\pi_i^{search}(q_j|\theta, p, b_{ij}^{like}, b_{ij}^{search})$ is the probability that i searches in each time period.

1. Set i's initial beliefs about j's actions: $b_{ij}^{like(0)} = \pi_j^{like(0)}$, $b_{ij}^{search(0)} = \pi_j^{search(0)}$
2. Calculate $\pi_i^{like(1)}(q_j|\theta, p, _{ij}, \pi_j^{like(0)}, \pi_j^{search(0)})$ and $\pi_i^{search(1)}(q_j|\theta, p, \pi_j^{like(0)}, \pi_j^{search(0)})$.
3. Given $\pi_i^{like(0)}$ and $\pi_i^{search(0)}$, calculate $\pi_j^{like(1)}(q_i|\theta, p, \pi_i^{like(1)}, \pi_i^{search(1)})$ and $\pi_j^{search(1)}(\theta, p, \pi_i^{like(1)}, \pi_i^{search(1)})$
4. $\pi_j^{like(0)} = \pi_j^{like(1)}$, $\pi_j^{search(0)} = \pi_j^{search(1)}$, $\pi_i^{like(0)} = \pi_i^{like(1)}$, $\pi_i^{search(0)} = \pi_i^{search(1)}$
5. Repeat steps 2-4 until π_i^{like},π_i^{search}, π_j^{like}, π_j^{search} converge.

11.4.7 Simulating Agent Actions

Given the value functions estimated in equilibrium, I then simulate behavior for m men and w women. The quality type of each simulated agent is drawn from the same observed distribution of q for the respective gender. Similarly, their parameters $\theta_i = \{\delta_{1i}, \delta_{2i}, c_i^v, c_i^m, \bar{s}_i, \lambda_i, \alpha_i^l\}$ are drawn from the estimated distributions.

I simulate $T = 500$ time periods. I list the simulation steps from the perspective of an agent i.

1. At time t, i has L_{it} `likes` left and decides whether to search. The probability that he searches is determined by his value functions estimated in the previous equilibrium step. If $L_{it} = 0$, then he cannot search and his session ends.
   (a) If i decides to not search, his session is over and he does not search for the remaining 499 periods.
2. If i searches, he is shown a randomly drawn (without replacement) agent's profile from the set of agents of the opposite gender. Because it is drawn without replacement, he will never see the same profile twice, much `like` the app. I denote this agent by j.
3. He observes q_j.
   (a) If he `likes` j and j has `liked` him at time $t' \leq t$, then i and j match at t.

i. i decides to propose a date with probability $\frac{\exp(\lambda_{1i}q_j)}{\exp(\lambda_{1i}q_j+V^{search}_{i,t+1})}$. j proposes a date with probability $\frac{\exp(\lambda_j q_i)}{\exp(\lambda_j q_i+V^{search}_{j,t+1})}$. If they both propose a date, the date is realized. The search session ends for a user once he pursues a date, regardless of whether the date is realized.

ii. If i does not pursue a date, then $L_{i,t+1} = L_{it} - 1$, and he goes back to step 1 at $t+1$.

(b) If i does not `like` j, he goes back to step 1 at $t + 1$.

I want to highlight a couple differences between the simulation and the model. First, in the simulation, profiles are shown are completely random. However, in the app, the matching algorithm is more likely to show a profile of similar quality. This was done purely to simplify computation. Future work will include incorporating the matching algorithm into the counterfactual simulations. Second, the model considers only matches that were formed instantaneously. For instance, if i `liked` j at time t and j `liked` i at time $t' > t$, then the match would have been formed instantaneously from j's perspective, but not i's. Thus, in the model, j would have matched with i, but i could not have matched with j. Modeling only instantaneous matches greatly simplifies the model, and since the subset of treated users is so small relative to the entire population, agents rarely saw profiles of other agents who were also in the treated sample. However, in the simulation, every agent in the market is modeled. Thus, the model needs a way to resolve matches when they do not occur instantaneously. In the simulation, if agent i views agent j at t but also matches with agent k at t, he considers going on a date with k, and stops search if a date is realized. If i matches with multiple agents at t, then he considers proposing a date to each of them. In this case, i may go on multiple dates. If at least one date is realized, i ends his session at $t + 1$.

Counterfactual	Gender	% Change	Lower 5% CI	Upper 5% CI
Match Quantity				
Increase Market Thickness	Men	-12	-19	-5
Increase Women	Men	44	31	56
Increase Women, Double Like Limit for Women	Men	57	41	70
Increase Market Thickness	Women	1	-5	8
Increase Women	Women	-7	-12	-2
Increase Women, Double Like Limit for Women	Women	-7	-12	0
Match Quality				
Increase Market Thickness	Men	-17	-24	-11
Increase Women	Men	8	-2	17
Increase Women, Double Like Limit for Women	Men	17	6	27
Increase Market Thickness	Women	-1	-3	2
Increase Women	Women	0	-2	3
Increase Women, Double Like Limit for Women	Women	2	-1	5

Table 8: Counterfactual results for users in small markets. The top half shows the changes in the number of matches that a person gets, and the second half shows the change in the quality of users they match with, conditional on matching.

Counterfactual	Gender	% Change	Lower 5% CI	Upper 5% CI
Match Quantity				
Increase Market Thickness	Men	-14	-21	-5
Increase Women	Men	21	12	30
Increase Women, Double Like Limit for Women	Men	61	52	73
Increase Market Thickness	Women	-2	-6	4
Increase Women	Women	-2	-5	2
Increase Women, Double Like Limit for Women	Women	-1	-5	3
Match Quality				
Increase Market Thickness	Men	-19	-26	-11
Increase Women	Men	-9	-16	-2
Increase Women, Double Like Limit for Women	Men	21	14	29
Increase Market Thickness	Women	0	-4	3
Increase Women	Women	1	-4	3
Increase Women, Double Like Limit for Women	Women	-3	-6	0

Table 9: Counterfactual results for users in large markets. The top half shows the changes in the number of matches that a person gets, and the second half shows the change in the quality of users they match with, conditional on matching.

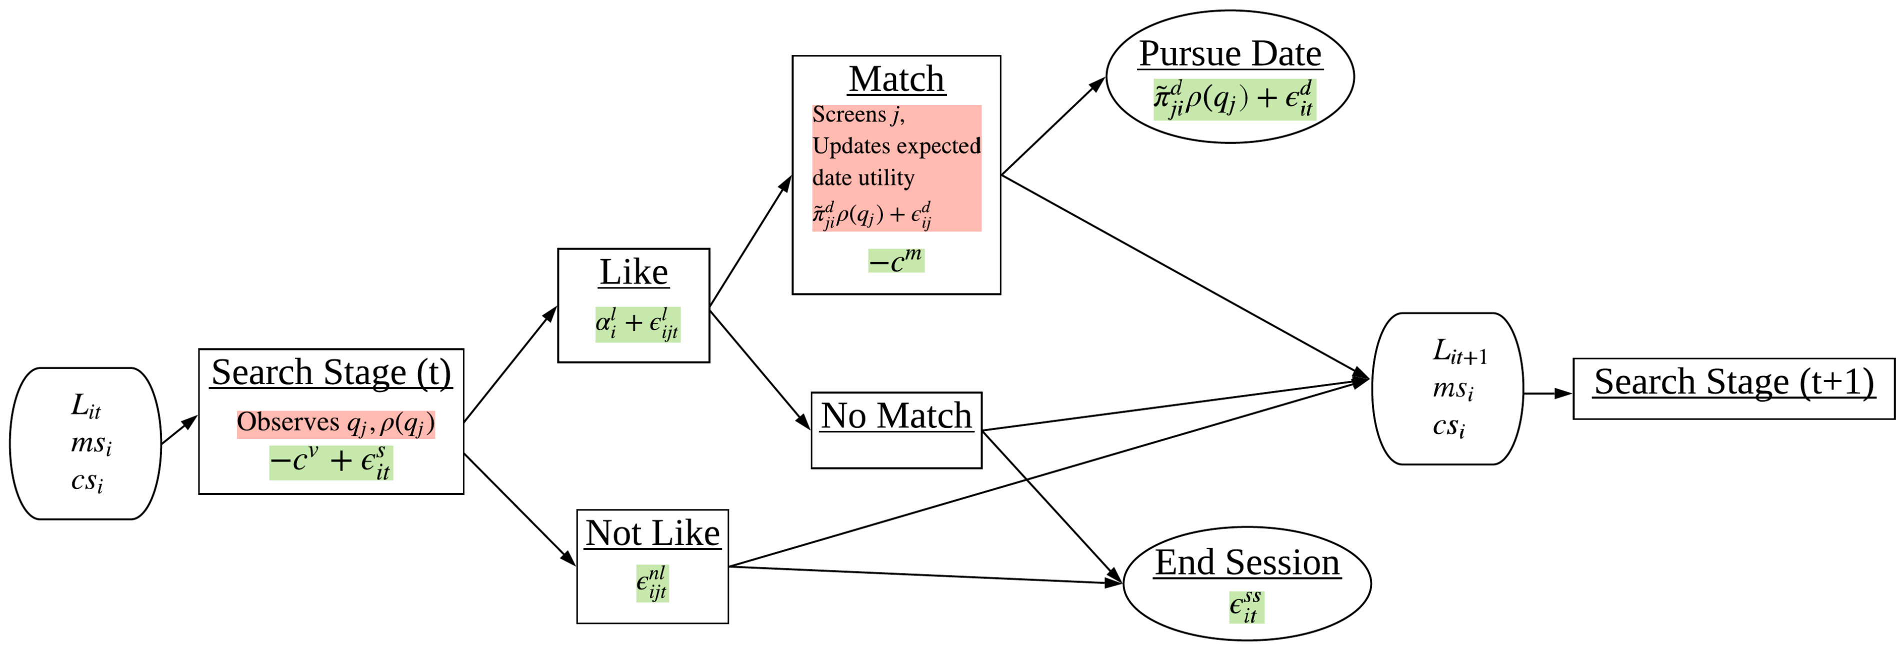

Figure 9: Detailed illustration of what happens during each stage of the model, and the current period utility the agent receives. Ovals are end stages. The text highlighted in green is the utility received, and text highlighted in red is the action that happens in the corresponding stage.

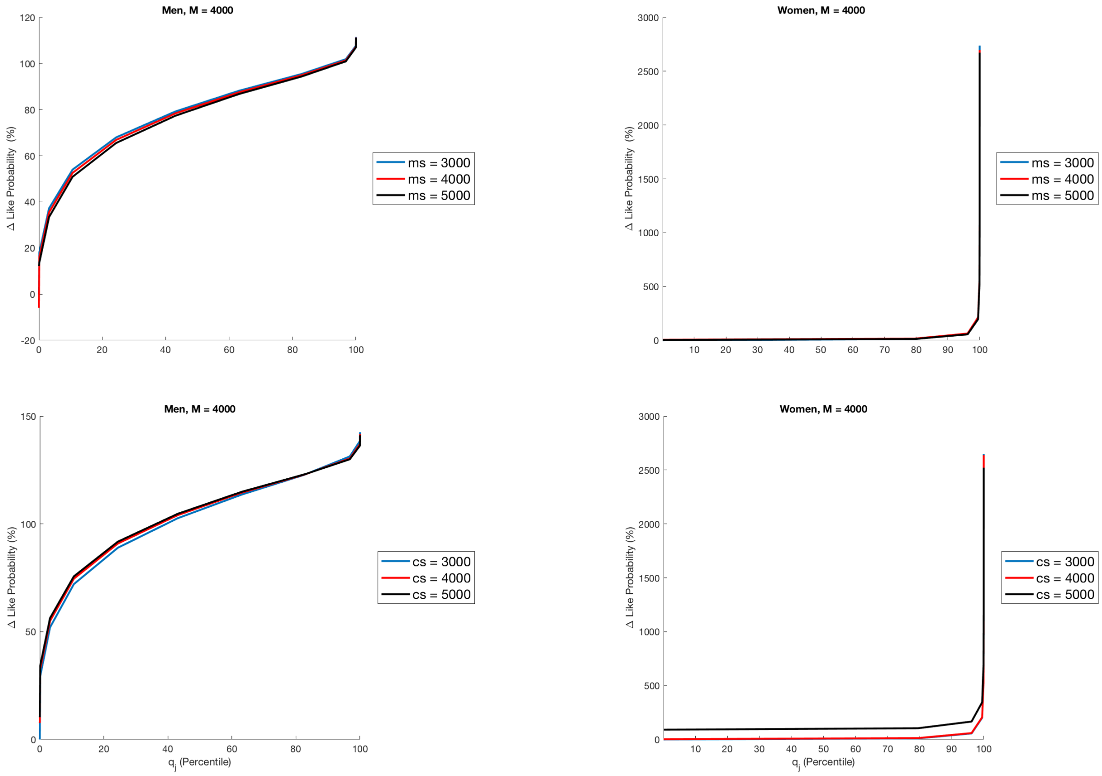

Figure 10: How selectivity changes with market size and competition size for agents in a small market.

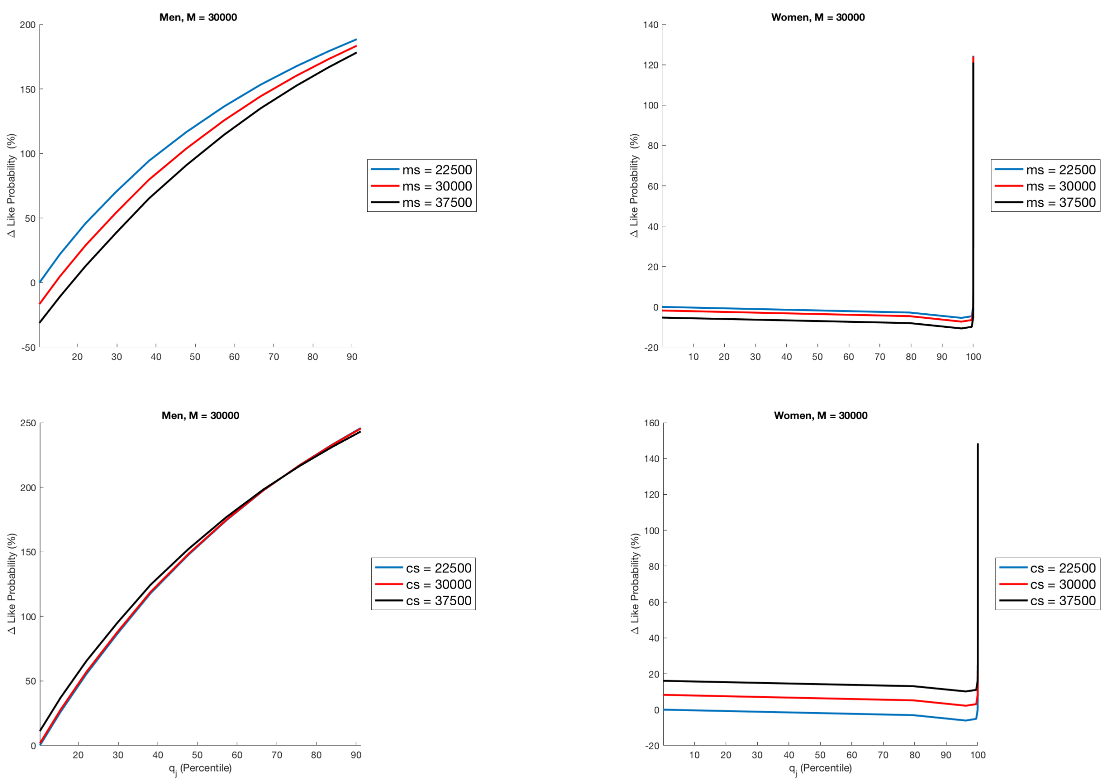

Figure 11: Selectivity and market size in a large market.

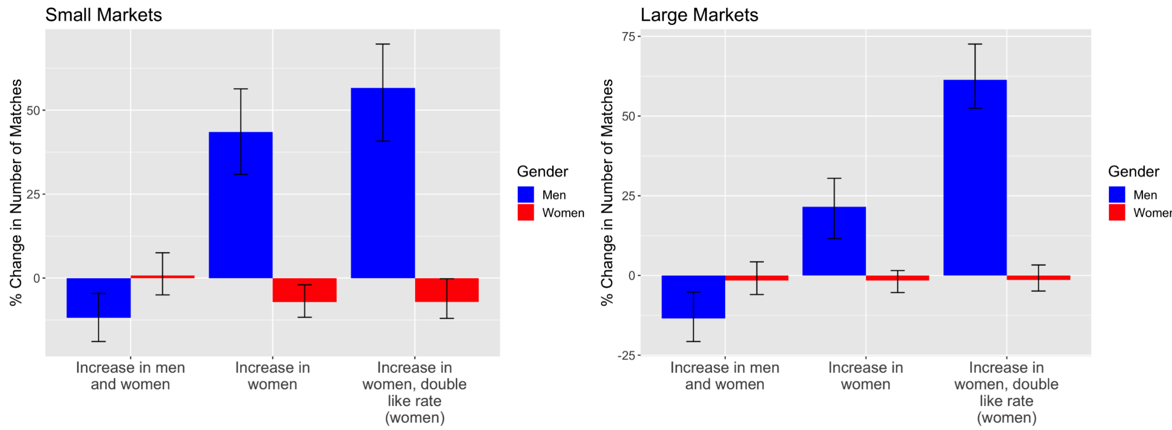

Figure 12: The percentage change in the number of matches that a man (blue bars) or woman (red bars) gets under the different counterfactuals for small (left) and large (right) markets. The error bars represent the 2.5 and 97.5 percentiles over 50 simulations to account for simulation error.

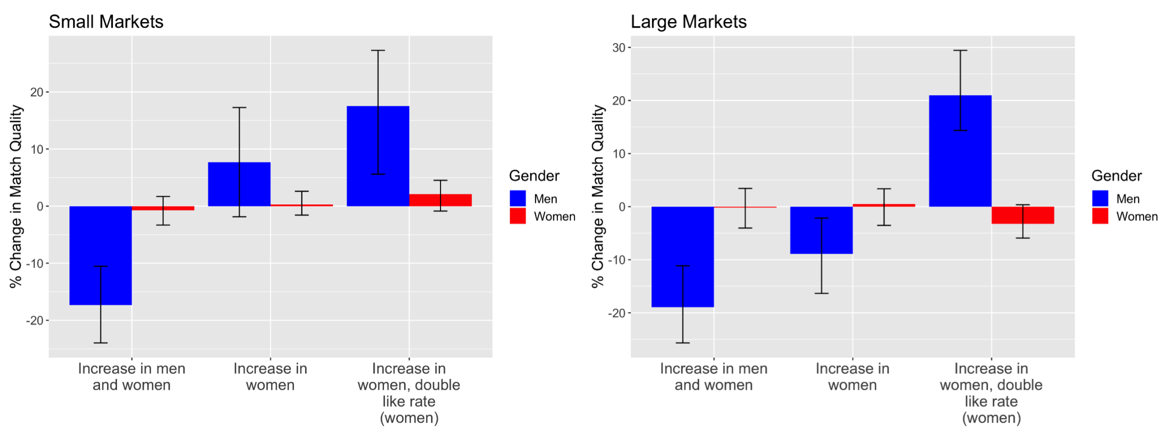

Figure 13: The percentage change in the quality of users that that a man (blue bars) or woman (red bars) matches with, conditional on getting a match, under the different counterfactuals for small (left) and large (right) markets. The error bars represent the 2.5 and 97.5 percentiles over 100 simulations to account for simulation error.

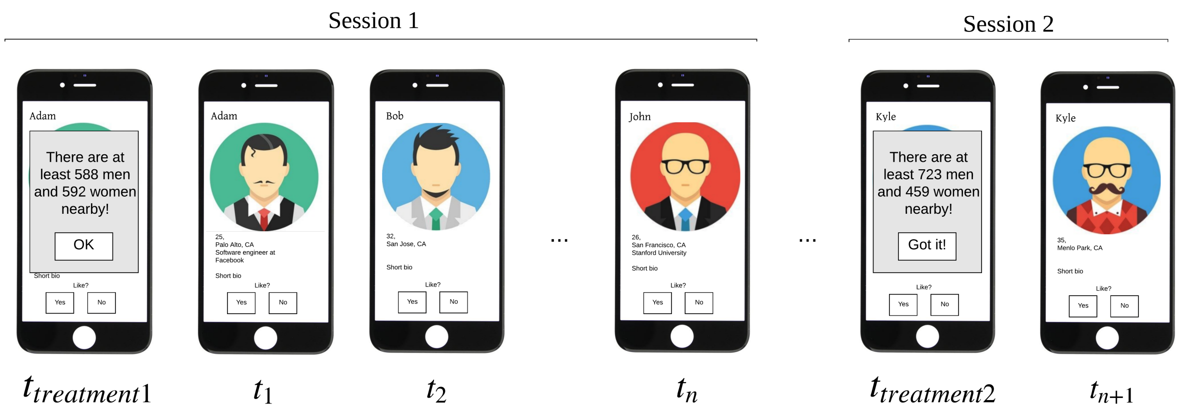

Figure 15: Example of how randomization occurs at the user-session level. The user is first exposed to the treatment at $t_{treatment1}$. She then views some profiles, and closes the app. If she opens the app at a later time and it has been at least 6 hours since $t_{treatment1}$, then he is exposed to the treatment again. The difference between $t_{treatment1}$ and $t_{treatment2}$ is always greater than or equal to 6 hours.

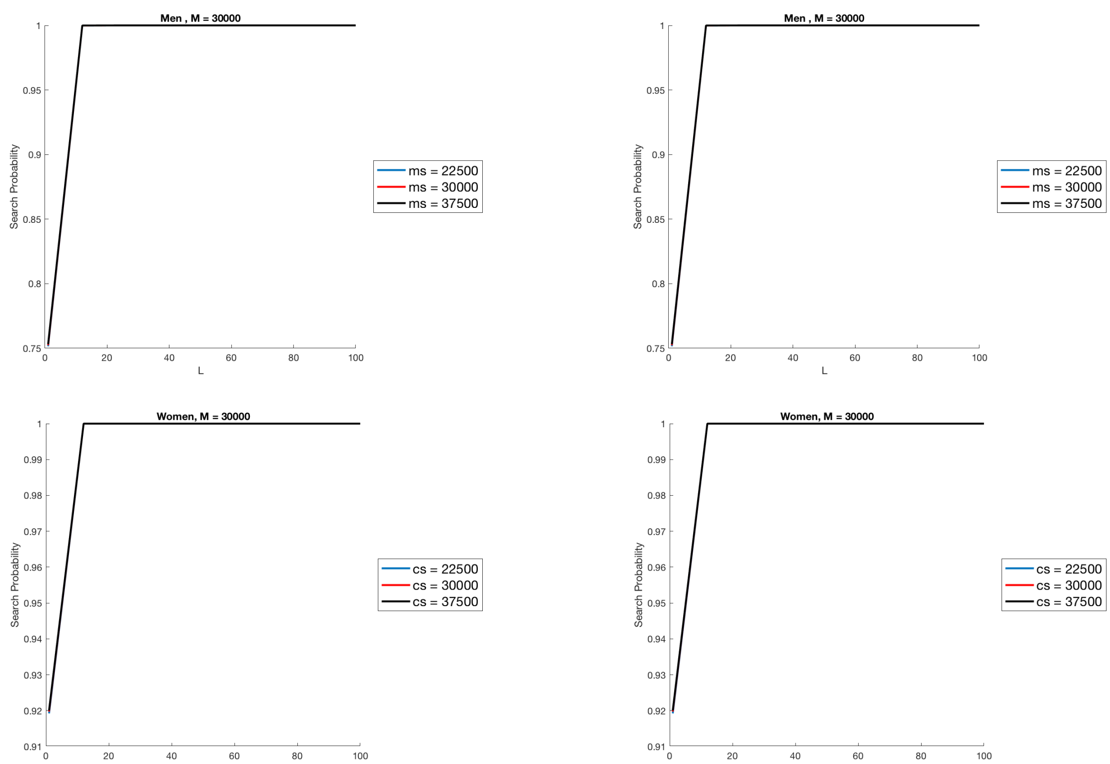

Figure 18: Selectivity and market size in a large market.