One-Way Mirrors in Online Dating: A Randomized Field Experiment

Ravi Bapna
Carlson School of Management, University of Minnesota, rbapna@umn.edu

Jui Ramaprasad
Desautels Faculty of Management, McGill University, jui.ramaprasad@mcgill.ca

Galit Shmueli
Institute of Service Science, National Tsing Hua University, galit.shmueli@iss.nthu.edu.tw

Akhmed Umyarov*
Carlson School of Management, University of Minnesota, aumyarov@umn.edu

The growing popularity of online dating websites is altering one of the most fundamental human activities: finding a date or a marriage partner. Online dating platforms offer new capabilities, such as extensive search, big-data based mate recommendations and varying levels of anonymity, whose parallels do not exist in the physical world. Yet, little is known about the causal effects of these new features. In this study we examine the impact of a particular anonymity feature, which is unique to online environments, on matching outcomes. This feature allows users to browse profiles of other users anonymously, by being able to check out a potential mate's profile while not leaving any visible online record of the visit. While this feature may decrease search costs and allow users to search without inhibition, it also eliminates "weak signals" of interest for their potential mates that may play an important role in establishing successful communication. We run a randomized field experiment on a major North American online dating website, where 50,000 of 100,000 randomly selected new users are gifted the ability to anonymously view profiles of other users. Compared to the control group, the users treated with anonymity become disinhibited, in that they view more profiles, and are more likely to view same-sex and inter-racial mates. However, based on our analysis, we demonstrate causally that weak signaling is a key mechanism in achieving higher levels of matching outcomes. Anonymous users, who lose the ability to leave a weak signal, end up having fewer matches as compared to their non-anonymous counterparts. This effect of anonymity is particularly strong for women who tend not to make the first move and instead rely on the counter-party to initiate the communication. Further, the reduction in quantity of matches by anonymous users is not compensated by a corresponding increase in quality of matches.

Key words: online dating, anonymity, weak-signaling, randomized trial, field experiment

---

* Author names are in alphabetical order. All authors contributed equally to this research.

1. Motivation and Background

According to the United States (US) Census, 46% of the single population in the US uses online dating[1] to initiate and engage in the process of selecting a partner for reasons ranging from finding companionship to marrying and conceiving children, and everything in between. Finding the optimal dating and ultimately marriage partner is one of the most important socio-economic decisions made by humans. Yet, such dating markets are fraught with frictions and inefficiencies, often leading people to rely on choices made through happenstance — an offhand referral, or perhaps a late night at the office (Paumgarten 2011). Interestingly, this primal human activity is being reshaped with the advent of big data and algorithmic match-making (Slater 2013). The continued growth of online dating despite the presence of a close substitute, the physical world, reflects the presence of significant frictions in the offline dating and marriage markets. Yet, the underlying processes, dynamics, and implications of mate seeking in the online world are largely unstudied. Also unknown are the implications of the new features and capabilities that these new online matching markets bring to an age-old human activity. In this paper, we address this gap by studying the causal impact of anonymity, a key feature unique to the online environment, through a randomized experiment in partnership with a major online dating site. We study anonymity because the design of this particular feature is a critical decision in the major emerging social platforms, ranging from Facebook to LinkedIn to our own context of online dating. One can easily imagine how users' behavior would change if Facebook suddenly required them to browse non-anonymously so that all their visits were visible to the visited person.

As is often the case, the Internet not only replicates the physical world processes of human interaction, but also extends them, supporting a variety of features that afford new capabilities that are next to inconceivable in the physical world, and that can vary the search costs for individuals looking for prospective dates. Given the extremely large user base of these websites as well as the standardized nature of users' profiles in the online world, these extra capabilities range from extensive search and algorithmic matching to big-data based recommendations (Gelles 2011), a science perfected for books and movies and now being deployed to what might be the ultimate experience good — human partners (Frost et al. 2008). However, certain features of these websites, such as completely-anonymous browsing of user profiles, have no direct analogies in the offline world. Thus, existing theories may not be adequate in explaining these online phenomena. Further, human behavior in the context of matters of the heart is inherently primal and complex. Studies that test existing theories based on purely observational data are likely to suffer from incompleteness due to key variables being unobserved or even unanticipated. To overcome this limitation and to achieve

[1] Of the 87 million singles in the US, nearly half of them, or 40 million, have tried online dating, according to the US Census. `www.ft.com/intl/cms/s/2/f31cae04-b8ca-11e0-8206-00144feabdc0.html`

high external and internal validity, we rely on an in-vivo randomized experiment to isolate the causal impact of our focal factor, anonymity. In doing so, we fill a gap in the extant research that has not addressed whether such IT-enabled features impact the search, viewing, message initiation, and matching outcomes of individuals.

The anonymity feature of online dating enables us to examine an interesting horse race between social frictions and search costs in online dating markets. Social frictions in such contexts arise from existing social norms that govern the contact initiation process in romantic situations (Maccoby and Jacklin 1974). At the same, it is well known that online environments reduce search costs in a variety of markets (Bakos 1997), including matching markets. In particular, we argue that anonymity may impact user behavior through two distinct causal mechanisms. On one hand, anonymity is likely to lower search costs and therefore lead to *disinhibition*. Users need not worry about how others interpret or perceive their visits, possibly even repeated visits that may otherwise be considered "stalking." On the other hand, with respect to social frictions, anonymity may impact the matching process via *signaling* related mechanisms. In particular, by hiding the focal user's actions from others, anonymity may influence signaling protocols that are necessary to establish successful communication with a potential mate. Thus, broadly speaking, our research objective is to examine the net effect of *disinhibition* and *signaling* in online dating.

Conventional economic wisdom suggests that anonymous profile viewing, by lowering search costs, should be associated with improved matching due to users being able to express their true preferences while browsing for mates. This may imply that, in a world of non-anonymous browsing, the focal user may search sub-optimally — by not searching enough or not searching in a way that reflects her true preferences — thereby limiting the options available to her and resulting in weaker matching outcomes. Another way to look at this is through the lens of information asymmetry where anonymous and non-anonymous users differ in their ability to gather information about users they are interested in. An anonymous user has uninhibited access to information as compared to a non-anonymous user, who may not visit a profile or visit a profile to resolve information asymmetry and later regret this decision. The anonymous user would have no such cause for inhibitions or subsequent regret.

In addition, across genders, social norms may inhibit the expression of what are considered taboo preferences, such as tendencies towards inter-racial or same-sex mate-seeking (Harris and Kalbfleisch 2000, Pachankis and Goldfried 2006). Normally, these inhibitions on preferences manifest themselves in the search stage of dating, limiting whom one looks for. An anonymity feature may potentially lower this stigma, thereby lowering the search costs and resulting in improved search and ultimately improved matching.

Support for this argument also comes from the growing literature on the disinhibition effect of the Internet, where a user's behavior changes once she is anonymous. The disinhibition literature has its roots in social psychology (Joinson 1998, Suler 2004). Kling et al. (1999) review social behavior on the Web, and state, "people say or write things under the cloak of anonymity that they might not otherwise say or write." Such anonymity-induced changes have been observed in settings ranging from adult films and books (Holmes et al. 1998) to pizza orders (McDevitt 2012). In the context of dating, the reduction in search costs due to the ability to view profiles anonymously, combined with internet-induced disinhibition, could overcome some sources of frictions and restrictive social norms and encourage people to express their true preferences. Therefore, if the disinhibition scenario indeed dominates, an argument can be made for a positive effect of anonymity on the number of matches due to a reduction in search costs.

In contrast to the disinhibition argument, the signaling advantage of non-anonymous browsing is that it allows a focal user to advertise herself by leaving a "weak signal" for another user without actually making any unambiguous explicit first move such as sending a personal message. We define a *weak signal* as a visit from the focal user i to a potential mate, that we call user j, such that user j knows user i visited her and thus showed some interest. It is akin to making an implicit move through viewing, without making an definitive explicit move by sending a message. Yet, importantly, the counter-party becomes aware that some move was made. Weak signaling is an important market feature that is unique to the online environment, and next to impossible to implement reliably in the physical world, at least with anywhere close to the level of definitiveness that can be done online. The offline "flirting" equivalents, at best, would be a suggestive look or a preening bodily gesture such as a hair toss to one side or an over-the-shoulder glance (Hall et al. 2010), each subject to myriad interpretations and possible misinterpretations (Henningsen 2004) contingent on the perceptiveness of the people involved. Much less ambiguity exists in the online environment if the focal user views another user's profile and leaves a visible trail in his "Recent Visitors" list. Thus, anonymously profile browsing, in effect, takes away a user's weak-signaling mechanism.

Individuals in the non-anonymous profile-viewing regime (control group) can provoke new conversations by virtue of leaving a visible profile visit and thus by providing a weak signal of interest to the counter-party. In contrast, in the anonymous regime (treatment group), individuals cannot provoke communications by a simple visit to a profile. Thus, the social inhibition of making the initial contact is higher towards the treated anonymous group relative to the non-anonymous group that has the weak-signaling capability. In a sense, it may be "easier" for other users to make contact with non-anonymous users than with anonymous users, since anonymous users leave no trail of their interest. In this scenario, we would expect the anonymity treatment to decrease the total number of matches and especially the number of matches initiated by a counter-party. Further, we would expect a significantly higher impact for women, given that social norms inhibit them from making the first contact through explicit messaging and who are thus waiting for the counter-party to initiate the actual contact (Maccoby and Jacklin 1974). In summary, if weak signaling is an important tool in initiating communication, in particular for women, then an argument can be made for a negative effect of anonymity on the total number of matches. Further, we would expect that this negative effect would be stronger for women.

These competing theories suggest different theoretical directions of the causal effect of anonymity on matches. Establishing the net effect of anonymity is therefore an empirical question, which we address using a randomized trial. These opposing forces reflect the fact that human behavior in the context of dating is incredibly complex. Studying such complex phenomena requires data on dating behavior that is challenging to obtain, particularly in the traditional offline world, and therefore such interactions are largely unmeasured and scientifically untested at the micro-level.

Based on a novel large-scale randomized experiment, similar in spirit to Aral and Walker (2011) and Bapna and Umyarov (2015), and in partnership with one of the largest online dating companies in North America, we examine these competing forces — lowering of search costs (through anonymity) vs. weak signaling (through non-anonymity). Our treatment involves gifting one month of anonymous profile viewing to 50,000 random users from a pool of 100,000 randomly selected new users of the site, while leaving the other 50,000 users untreated in order to serve as a control group. On this website, the anonymity feature is bundled with other advanced features and is available for purchase to any user of the dating site for \$15 (value changed for de-identification purposes) per month. In our study, we make sure to treat the randomly selected users only with anonymity and not with any other features of the premium subscription bundle for the purpose of observing the changes in behavior and outcomes that are induced specifically by anonymity.

The results of our experiment suggest that weak signaling is a key mechanism in increasing the number of matches. This is especially important for women, helping them overcome social norms that discourage them from making the explicit first move in dating markets (Maccoby and Jacklin 1974). Specifically, we observe that those treated with the ability to browse anonymously lose their ability to send a weak signal, thereby reducing the number of matches they achieve without any increase in the quality of their fewer matches. In other words, the quality of the match does not compensate for the reduction in quantity of matches.

The advent of online dating platforms has increased measurability, while also introducing new modalities of behaviors that do not have offline parallels. Thus, our approach in examining these opposing forces is positivist in nature. We refrain from making any broad welfare claims, but provide a thorough analysis of the impact of anonymity on both the quantity and quality of the matches. We also refrain from a priori judgments about the relative efficacy of the competing hypotheses: lower search costs improving matching versus the absence of weak-signaling hurting matching. Instead, we examine these competing forces using a large-scale randomized experiment, measure the effect, and then analyze the sub-processes to understand the observed outcome. A key aspect here is that the online dating platforms provide an environment where participants' choices at sub-stages of the dating process are available to the researcher in unprecedented detail, which is not observable in the offline world. We exploit this rich micro-level data to explain our key finding in a detailed and nuanced manner.

In summary, we seek to answer the following research questions in a causal manner:

1. Does anonymity change the searching behavior of individuals in dating markets?

2. Does anonymity change the number of matches achieved by individuals in dating markets?

3. Given known gender asymmetries in dating markets (Fisman et al. 2006), does the effect of anonymity differ across genders?

4. How does anonymity and its counterpart weak-signaling manifest itself in the overall dating process, which begins with viewing, is followed by messaging, and ends (potentially) in matching?

The remainder of the paper proceeds as follows. Section 2 explores the current state of the literature in more detail. In Section 3 we provide institutional details of the online dating site we partner with as well as share some empirical regularities in our data. In Sections 4, 5, and 6 we describe our experimental data, design, and results, respectively. Section 7 presents robustness analysis and Section 8 concludes with directions for future research.

2. Literature Review

Our work builds upon and contributes to three streams of literature: the economics literature on marriage markets and subsequent empirical studies on dating, the related literature on two-sided marriage markets, and more recent work on social frictions in dating markets.

The stream of economics literature on marriage markets, starting with Gale and Shapley (1962) and Becker (1973) and related work across multiple disciplines, establishes the theoretical basis for the sorting patterns that are exhibited in marriages. The literature establishes that marriage partners are similar in age, education levels, and physical traits such as looks (Kalmijn 1998) with the sorting being attributed to either search frictions or preferences. For instance, one explanation for the observed sorting based on education is that people of certain education levels may prefer their partners to be of the same educational level. Another explanation is that these people may be employed together, or be educated in the same educational institutions, leading to romantic liaisons due to spending time together simply because of lower search costs and not because of preferences. Similar sorting patterns based on equivalence in social desirability have been established in the online dating context (Taylor et al. 2011).

Existing approaches, based on analyzing observational data from online dating markets (Hitsch et al. 2010), assume equilibrium outcomes, and thus, are able to tease out peoples' preferences conditional on the observed outcome. As Hitsch et al. (2010) note, because search frictions are substantially lower in online dating markets – considering the infeasibility of getting detailed profile and attribute information from even a handful of potential mates at a bar – they are able to break down the observed sorting outcomes in dating into preferences over mate attributes and find gender asymmetries in matching preferences. Fisman et al. (2006), who obtain mate preference data from a speed dating experiment, also find significant gender asymmetries in mate selection. Recognizing these gender asymmetries established in the literature, we report our empirical findings separately for men and women. Similar to Fisman et al. (2006), we focus on dating, an activity that usually precedes marriage, and often manifests itself in the form of a long learning period during which people engage in more informal and often polygamous relationships. That said, in discussing the related literature, we use dating and marriage interchangeably so as to be expansive in our coverage of the various streams of thought that can possibly influence our work.

Our point of departure from what is already known regarding preferences in heterosexual matching rests on the idea that observed preferences are conditional on two initial stages in the mate selection process: an underlying search process, i.e. "checking people out" as well as a post-search contact initiation stage, i.e. "making a move." In the online world, the first stage involves searching for and viewing multiple user profiles, allowing a user to explore the potential market of date-able users – similar to checking people out at a bar or a party in the offline world. The second stage involves making the move with the user(s) identified through the search stage – through initiating contact by sending a message in the online world, similar to approaching and talking to a potential mate in the offline world. Given the significant social inhibitions in both the search stage and the contact initiation stage of the dating process, and given that contact initiation is gender asymmetric, we believe it a worthy reprise to examine the causal impact of each of the two aforementioned social inhibitions on matching outcomes.

Our research also relates to the economics literature on two-sided matching markets, e.g., Gale and Shapley (1962) and Roth and Sotomayor (1992), who formulate marriage as a two-sided matching problem given the differences between women and men. They model preference orderings in the matching process and, importantly for this research, introduce the idea of unstable matching, an outcome wherein people would have been better off having different partners. This idea of unstable matching is intricately linked to Piskorski (2014)'s idea of a social failure.

The crux of Piskorski (2014)'s idea is that while online dating reduces multiple sources of friction that are present in offline dating markets, it does not eliminate them. Piskorski (2014) documents that dating markets are fraught with frictions ranging from high search costs to asymmetric societal norms that often lead to social failures. Akin to a market failure, which implies an economic exchange that did not take place but had it taken place would have made everybody better off, a social failure is a human connection that should have taken place, but did not. In the context of heterosexual dating, these matching inefficiencies arise due to social frictions such as physical constraints of time and space, the costliness of the initial information acquisition, and societal norms, such as those inhibiting women from making the first move (Piskorski 2014).

In this paper, we contribute to the literature on dating and marriage markets by causally examining the role of such frictions on matching. This view is somewhat distant from the early thinking of the economic modeling of marriage markets as being frictionless (Becker 1973), and even broader than the more recent developments by Burdett and Coles (1997), Mortensen and Pissarides (1999) and Smith (2006), who account for search frictions but do not account for social frictions. Our research is motivated by taking into account these well-documented frictions and examining whether the newer capabilities afforded by the online environment can mitigate them. Our random assignment of the anonymity feature to a subset of users in the online dating site can, at one level, be interpreted as an exogenous shock that lowers search frictions. Anonymous users can uninhibitedly search for potential mates (McDevitt 2012) and, if search frictions are the only force at play, this should naturally lead to better matching outcomes. Yet, social exchange theory, which Piskorski (2014) draws upon, stipulates that while women tend not to make the first move, say by messaging a potential partner, the online dating markets give women an opportunity to leave a weak signal. This "trail" of a profile visit can then serve as an implicit move that could trigger a response and possibly lead to a match. When we gift anonymity to our treatment group, we are in effect taking away this ability to leave a weak signal, thereby increasing social frictions.

Thus, in departure from the extant literature, our treatment is in effect a horse race between search frictions, which decrease with anonymity and should result in more matches, and social frictions, which rise when we take away weak-signaling and therefore should result in fewer matches. Again, while the economics literature has extended the original frictionless matching models to account for search frictions, no one has looked at social frictions and compared the two in the setting of a randomized controlled experiment.

Our work has implications for the design of large scale matching markets. In designing these markets, important decisions about features and capabilities, such as anonymity, must be made. But how these features play out in a multi-faceted real world social process that makes up romantic markets requires careful scientific inquiry and experimentation. We contribute to prior work by rigorously and causally investigating the impact of the new capabilities afforded by the online dating environment on the underlying process and resulting outcomes of this fundamental human activity of mating.

3. Institutional Details

To conduct the experiment, we partnered with one of the largest online dating websites in North America, which we call monCherie.com (name disguised). MonCherie.com constitutes a typical online dating website and offers the following features to its users, which are common to most online dating websites:

- Users set up their own well-structured online profiles where they describe themselves as well as reveal characteristics sought in a desired partner. User profiles typically also include a set of their photos.
- Users may view profiles of all other users without limitations and for no cost.
- Users may search for profiles of other users using an advanced search engine that allows filtering by age, location, religion, and a large number of other demographic variables. Users may also discover partners using a proprietary recommendation engine that is provided by the website.
- Users may send private messages to any other user without limitations and for no cost.

In addition to these features, monCherie.com constitutes a typical freemium community: most of the users sign up for a free account ("free users"), which allows them to utilize all the key features listed above. In addition to these free features, users can buy a premium subscription for approximately $15 per month (exact value changed for de-identification purposes). The premium subscription consists of a fixed bundle of premium features[2] that include, among other incremental features, the ability to anonymously browse profiles of other users.

By default, free users of monCherie.com browse using the *non-anonymous* mode, where if focal user i visits the profile of user j, user j knows through her "Recent Visitors" page that user i checked her out. In contrast, premium users browse in the anonymous mode, where if the focal user i visits the profile of user j, user j does not know that user i checked her out. However if free user j were to visit premium user i's profile, user i would know about this visit. This feature is the proverbial "one-way mirror" of the online world, the impact of which is the research subject of this paper. It is important to highlight that a user's public profile does not reflect whether the user is anonymous or not and therefore, it is impossible to distinguish premium users from non-premium users by looking at their profile or by messaging them.

4. Data, Empirical Regularities, and Outcomes

Based on the specifics of the agreement with monCherie.com, our experiment was conducted on a random sample of 100,000 new users of the website from one geographical area (unknown to us

[2] In addition to anonymity, key premium features may include extra search and messaging options, no advertisements and so on. A complete list of premium features is not provided due to the confidentiality agreement with the online dating site.

for privacy reasons) over a period of three months, which we refer to as month 1 (pre-treatment), month 2 (during treatment) and month 3 (post-treatment). A randomly selected sub-sample of 50,000 users was given the gift of anonymity (manipulation = 1), and the remaining 50,000 users served as the control group (manipulation = 0). The treated users received a message from the site indicating that they had been gifted this premium feature for one month and were automatically defaulted to anonymous browsing. No opt-in was required, thus there are no selection issues with respect to who was treated. While the treated users had the option of switching back to non-anonymous browsing, empirically this rarely occurred in our experiment and 98% of users remained with the default setting of anonymity. This behavior is not surprising as anonymity is marketed as a key premium feature.

For all users in our experiment we had information on a set of demographic variables such as gender (gender = 1 for men, 0 otherwise), age, sexual orientation (straight = 1 for straight users, 0 otherwise), race (white = 1 for white users, 0 otherwise, black = 1 for black users, 0 otherwise and so on). In addition, MonCherie.com users can secretly rate each other on attractiveness on a scale of 1 (least attractive) to 5 (most attractive). Based on these scores, we define each user's attractiveness, AttractScore, as the average rating that the user received[3].

In our analysis, AttractScore is meant to capture the "dating market value" of the focal user to prospective dating partners taking into consideration photos and other profile characteristics, as per monCherie's rating system. In Section 6, we also use AttractScore as a proxy measure in analyzing the quality of a match. While attractiveness does not necessarily capture the complete picture with regards to quality of a long-term relationship, our focus is on the initial stages of a dating process. The evidence suggests that attractiveness score is what users mainly care about in online dating context (Rudder 2014). This choice of attractiveness measure is also supported by Ellison et al. (2006) who conduct a qualitative analysis and find that users invest significantly in creating profiles to represent themselves authentically and reflect their "ideal self."

Finally, we know whether the users are valid (valid = 0 if the focal user is a spammer or a bot as determined by internal algorithms at monCherie.com) and we know whether users are active or not. A user is defined as active if s/he visited at least one profile ten days prior to our manipulation.

In this study we limit our attention only to users who are not premium subscribers, are straight, and were valid and active prior to our manipulation. Table 1 outlines descriptive statistics of user demographics and their attractiveness scores and compares men to women. As is evident from Table 1, men are statistically different from women in every single demographic attribute and with large differences in attractiveness scores.

[3] The attractiveness variable can be missing in our data for users who were not rated by other website users

Gender	Variable	Mean	StdErr	Median	Min	Max	t-value	p-value
F	Age	31.375	0.113	28.296	18.045	75.381		
M	Age	30.179	0.075	27.625	18.045	78.795	8.849	<0.0001
F	AttractScore	3.097	0.008	3.126	1.000	5.000		
M	AttractScore	2.179	0.007	2.053	1.000	5.000	86.353	<0.0001

(a) Demographic characteristics

Gender	Variable	Mean	StdErr	Median	Min	Max	t-value	p-value
F	Asian	0.065	0.003	0	0	1		
M	Asian	0.048	0.002	0	0	1	5.453	<0.0001
F	Black	0.066	0.003	0	0	1		
M	Black	0.052	0.002	0	0	1	4.168	<0.0001
F	Indian	0.011	0.001	0	0	1		
M	Indian	0.015	0.001	0	0	1	-2.998	0.0027
F	Latino	0.069	0.003	0	0	1		
M	Latino	0.068	0.002	0	0	1	0.269	0.7879
F	White	0.577	0.005	1	0	1		
M	White	0.561	0.004	1	0	1	2.418	0.0156

(b) Racial composition

Table 1 Summary Statistics of User Characteristics in Study Data

In addition to the demographic variables, we collected all profile viewing and messaging activity for the users in our sample for the same three months. We name our variables as follows: ViewSentCountPre (number of unique profiles the focal user visited in month 1), ViewSentCount (number of unique profiles the focal user visited in month 2) and ViewSentCountPost (number of unique profiles the focal user visited in month 3), ViewRcvdCountPre (number of unique users who visited the focal user in month 1), ViewRcvdCount (number of unique users who visited the focal user in month 2) and ViewRcvdCountPost (number of unique users who visited the focal user in month 3). In other words, a "sent" view reflects a focal user viewing another user's profile, whereas a "received" view reflects another user viewing a focal user's profile. We follow a similar naming convention for messages and matches.

Table 2 outlines the statistics of user activity for our sample of users in month 1 with t-tests indicating the statistical significance of differences of the various measures of activity between the two genders. As is evident from Table 2, women, on average, receive almost six times the viewing attention as compared to men and receive more than sixteen times the number of messages. In addition, note that women are far less likely to initiate explicit contact via sending a message, consistent with the findings of Fiore et al. (2010). Although women, on average, visit only 1.6 times fewer profiles than men, they send four times fewer messages. These extreme gender asymmetries in user behavior with respect to viewing and messaging play out in a significant way in our findings. To the best of our knowledge, this study provides the first quantification of an age-old social norm, the extent to which women are not likely to make the first move. Recognizing these large differences between the two genders as presented in Tables 1 and 2, we report all subsequent statistics and results separately for the two genders.

Gender	Variable	Mean	StdErr	Median	Min	Max	t-value	p-value
F	ViewRcvdCountPre	314.927	3.408	223	0	3030		
M	ViewRcvdCountPre	54.222	0.580	32	0	1291	75.421	<0.0001
F	ViewSentCountPre	115.670	1.500	75	1	3184		
M	ViewSentCountPre	190.871	1.991	112	1	3588	-30.170	<0.0001
F	MsgRcvdCountPre	51.066	0.681	30	0	926		
M	MsgRcvdCountPre	3.092	0.044	1	0	170	70.309	<0.0001
F	MsgSentCountPre	7.892	0.196	3	0	662		
M	MsgSentCountPre	31.366	0.604	10	0	2167	-36.939	<0.0001

Table 2 Descriptive Statistics of User Activity, By Gender

While defining our outcome of interest, a match, we recognize that there is an inherent challenge in creating a perfect and all-encompassing measure of success in the online dating scenario. For example, one success measure could be two users attempting to move their online interactions offline, or getting married. However, even if we could observe which couples went for an offline date or got married, these measures are far from perfect, as many offline dates are ultimately unsuccessful in the long-run and current divorce rates indicate that a marriage is not a perfect measure either.

Recognizing that any relationship is an ongoing process, and that significant difficulty exists in learning ex-ante the ultimate success of any observed relationship unless it is observed for the entire lifetime of both partners, we refrain from defining a measure of 'ultimate success.' Instead, we define 'success' in online dating as a successful outcome at a certain initial and critical step: successful online communication. Without successful initial online communication, no further steps are possible in the online dating process: there will be no offline date, no relationship, and no marriage.

More specifically, we define the communication of user i with user j as a match if there's a sequence of *three or more* messages exchanged between user i and user j. Communication theorists call this measure a double interact (Weick and Kiesler 1979), and it is considered a sense-making process that people use when they organize in a variety of contexts. We can further distinguish between a "sent match" and a "received match" for user i: A "sent match" occurs when user i messaged user j, user j responded, and then user i messaged user j again (with user i possibly responding to that and so on). A "received match" for user i is similarly obtained when user j initiates the chain of messaging (j messages i, then i responds, then j messages i again, with possible further messaging).

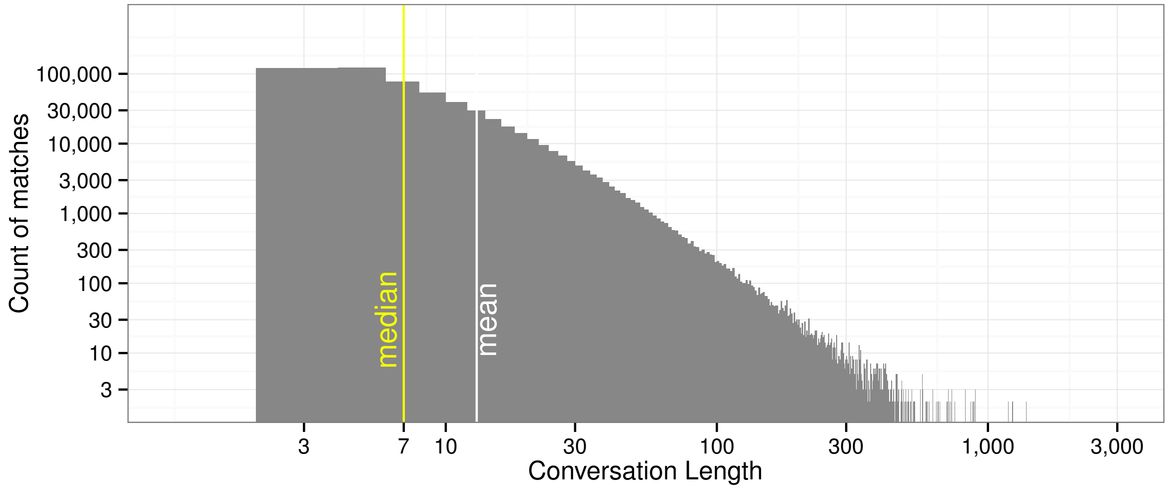

Figure 1 Distribution of the Number of Messages Exchanged in a Conversation

As is evident from this definition and as demonstrated in Figure 1, a typical message exchange sequence is much longer than the minimum threshold of three messages[4]. More specifically, in our data, the average number of messages exchanged in a match is 12.6, while the median is seven messages. These statistics are particularly encouraging given almost identical results reported by Hitsch et al. (2010) who had access to the actual content of the messages exchanged on a different online dating website. As reported by Hitsch et al. (2010), it took users on average 12.6 messages for women and 11.6 for men (with an overall median of six messages) to reveal their phone number, email address, or to say a key phrase like "get together" or "let's meet." Also, while we are blind to the actual content of these messages[5], monCherie.com is not. Our conversations with the senior executives revealed that they strongly believe that the double-interact measure of a match is an accurate predictor of an offline date and that it is used as an industry-standard measure of success. Indeed, despite knowing the content of users' messages, monCherie.com uses this double-interact metric as a measure of matching for their own internal recommendation engine, a key component of their value proposition to the users. We also test the robustness of our results to different definitions of a match, by redefining a match to consist of an exchange of three to twelve messages. By including match sequences of an even number of messages exchanged, we rule out the possibility that our results are driven by the combination of men being more likely to initiate messages and

[4] In Section 6.3 we demonstrate that our results do not change when we adopt a more stringent definition of a match as we increase the threshold from three messages to 4, 5 and so on up to 12. Also, as we discuss in this section, three messages is a minimum threshold and most conversations are longer than three messages.

[5] Such an approach was taken by Hitsch et al. (2010) who used an indicator variable for whether users exchanged phone numbers or email addresses somewhere in their messages on an online dating website. For reasons of privacy and sensitivity to the user base, particularly because our study involves a randomized experiment, our research partner could not provide us access to the actual content of the messages.

therefore more likely to achieve a "sent match" initiated by themselves than a "received match" by a women. We find that our results do not change even after altering the definition to include longer message exchanges (see Section 6).

Given the above definition of a match, our main outcome variable of interest is the number of matches (sent and received matches) achieved by a focal user in the treatment month, month 2. This choice reflects the dating context of mate seeking that we study. Dating is defined as a prolonged period of polygamous learning that eventually leads to a long-term relationship such as marriage. In that spirit, we posit that there is positive expected utility in each additional date. Individuals return to the dating market and search until the value of any expected improvement in the date they can find is no greater than the cost of their time and other inputs into the additional search.

Using this definition of a match, and consistent with our prior variable naming convention, we examine pre-treatment data to give the reader a feel for the site's efficacy in matching. We find that women achieve a significantly higher number of matches than men on average[6], and that 75% of these matches for women are received matches. That is, 75% of women's matches are when the man initiates contact by sending the first message and the woman responds, thus supporting the conversation and leading to a match that she did not initiate. More than 75% of matches for men are "sent matches," that is, matches initiated by the man himself, while less than 25% of matches for women are "sent matches."

In addition to statistics on profile viewing, messaging and matching, we also explore whether, under anonymity, users not only increase the quantity of their views, but also change the nature of their viewing. In particular, we ask whether anonymity reduces inhibitions along the dimensions of sexual orientation and race, potentially indicating users overcoming existing social norms. We define the following variables for this purpose:

- ViewSameSex: a binary variable indicating whether the focal user initiated a visit to a user of the same sex at least once during the treatment month (ViewSameSex $= 1$) or did not (ViewSameSex $= 0$).
- MsgSameSex: a binary variable indicating whether the focal user initiated a message to a user of the same sex at least once during the treatment month (MsgSameSex $= 1$) or did not (MsgSameSex $= 0$).

We define the above variables only for users who are valid, active and straight, since for users who reported themselves as non-straight these concepts would not capture disinhibition. In addition to same-sex browsing, we also explore changes in inter-racial browsing patterns by defining the following variables:

[6] This is mathematically possible since there are more men on this site than women.

- ViewOtherRace: a binary variable indicating whether the focal user initiated a visit to a user of a different race at least once during the treatment month (ViewOtherRace $= 1$) or did not (ViewOtherRace $= 0$).
- MsgOtherRace: a binary variable indicating whether the focal user initiated a message to a user of a different race at least once during the treatment month (MsgOtherRace $= 1$) or did not (MsgOtherRace $= 0$).

In addition to these measures, we observe more nuanced behavioral patterns that shed additional light on the effect of anonymity. These behavioral patterns help us understand whether the selectivity levels of users change under the anonymity condition. Finally, we also examine whether anonymity changes the quality of the achieved matches. We operationalize selectivity and quality constructs while discussing the results in Section 6.

5. Experimental Design

To test the impact of anonymity on user behavior in online dating markets, we collaborated with a large online dating website monCherie.com. Our experimental design involves selecting a random subset of 100,000 users from an undisclosed geographical area in North America and treating 50,000 of them with a gift of one month of anonymous browsing (treatment group), while keeping the remaining 50,000 in the default non-anonymous setting as our control group of untreated users. Our field experiment removes the ability to send a "weak signal" for the treatment group while maintaining it for the control group, allowing us to compare the two groups in terms of the resulting search intensity, search diversity, messaging behavior, and number of matches.

We derived our sample size based on the statistical power calculations based on Pocock (2013) and institutional constraints of our partner website. More specifically, based on a sample of pre-experimental observational data, we observe that women receive on average 2.95 matches per month (with a standard deviation of 7.97). We assumed a 10% effect size for our treatment, which would imply a mean of 2.68 matches per month. Running the power analysis with α of 5% and power of 90% yields a sample size calculation indicating that we should have at least 30,680 women in our sample. The same analysis for men with 2.18 matches per month (6.76) leads to 1.98 matches under the treatment and the sample size of 41,168 men. This leads to a total sample size of 71,848 users. Given this calculation along with institutional constraints of our partner company, we agreed to set the sample size at 100,000 users. Increasing the sample size to 100,000 is necessitated by the anticipation that we would need to filter out some data generated by bots, and wanting to estimate effects for sub-populations such as gender-specific differences that are motivated by the theories described in Section 3. Broadly speaking, in determining sample size for in-vivo randomized trials it is important to consider that contextual data cleaning and comparing sub-samples can reduce the final sample. Further, for researchers desiring to examine multiple outcomes of interest, sample size calculations also need to take into account the noisiest outcome.

One concern in designing a field experiment is that a very large treatment group could raise questions about whether the treatment and control groups interact with each other and that the treatment effect spills over to the control group, i.e. we are concerned about contamination between the treatment and control groups. We test for this empirically and find that the intersection — that is the users from the treatment group who match with users of the control group — comprises of only 3% of the matches, which we exclude from the analysis. Re-running the analysis including the 3% yields no qualitative change in the results.

As demonstrated in Table 3, the treatment ($\mathsf{Manipulation} = 1$) and control ($\mathsf{Manipulation} = 0$) groups have statistically indistinguishable demographic properties before manipulation. They are also indistinguishable in terms of their viewing and messaging behavior in the pre-treatment month 1.

The exogenous random assignment of the treatment rules out myriad problems of endogeneity and alternative explanations that could confound any analysis of such a causal question based on observational data. Our treatment is carefully implemented in that we do not ask for anything in return from users who receive the gift and no action is needed on their side. Users are also unaware of being part of an experiment, so observer bias is not applicable. As mentioned earlier, we limit our sample only to valid, active, straight users.

6. Experimental Results

6.1. Treatment effects

We start our analysis by exploring changes in profile browsing behavior that were induced by our treatment, comparing the average behaviors in the treatment and control groups, by gender. As demonstrated in Table 4, treated users of both genders viewed significantly more profiles as compared to their non-treated counterparts (14.1% increase for women and 8.35% for men).

Further, as demonstrated by Table 5, we find that straight individuals of both genders significantly increase their likelihood of viewing profiles of users of the same gender when they are anonymous. In particular, anonymous heterosexual women increase their likelihood of viewing other women by 19.8% . Additionally, anonymous heterosexual men, on average, increase their propensity to view other men by 12.2% . We also find that, on average, anonymity induces white women to have a 5% higher likelihood of viewing a race other than their own, while this inter-racial effect is not statistically significant for men. This disinhibition effect, however, is only seen in viewing behavior and does not translate to initiating a message (Table 6), as predicted, since only profile views are hidden by anonymity, but not messages.

Gender	Manip	Variable	Mean	StdErr	Median	Min	Max	t-value	p-value
F	0	Age	31.365	0.161	28.214	18.045	75.214		
F	1	Age	31.386	0.158	28.296	18.045	75.381	-0.092	0.9267
M	0	Age	30.214	0.107	27.625	18.045	78.215		
M	1	Age	30.144	0.105	27.625	18.130	78.795	0.469	0.6388
F	0	AttractScore	3.093	0.012	3.115	1.000	5.000		
F	1	AttractScore	3.100	0.012	3.140	1.000	5.000	-0.403	0.6869
M	0	AttractScore	2.182	0.010	2.059	1.000	5.000		
M	1	AttractScore	2.177	0.009	2.047	1.000	5.000	0.421	0.6737

(a) Demographic characteristics

Gender	Manip	Variable	Mean	StdErr	Median	Min	Max	t-value	p-value
F	0	Asian	0.064	0.004	0	0	1		
F	1	Asian	0.065	0.004	0	0	1	-0.216	0.8287
M	0	Asian	0.045	0.002	0	0	1		
M	1	Asian	0.050	0.003	0	0	1	-1.522	0.1279
F	0	Black	0.065	0.004	0	0	1		
F	1	Black	0.066	0.004	0	0	1	-0.212	0.8321
M	0	Black	0.052	0.003	0	0	1		
M	1	Black	0.053	0.003	0	0	1	-0.529	0.5965
F	0	Indian	0.012	0.002	0	0	1		
F	1	Indian	0.009	0.001	0	0	1	1.121	0.2623
M	0	Indian	0.014	0.001	0	0	1		
M	1	Indian	0.016	0.001	0	0	1	-0.684	0.4937
F	0	Latino	0.071	0.004	0	0	1		
F	1	Latino	0.066	0.004	0	0	1	0.981	0.3267
M	0	Latino	0.067	0.003	0	0	1		
M	1	Latino	0.069	0.003	0	0	1	-0.449	0.6536
F	0	White	0.579	0.007	1	0	1		
F	1	White	0.574	0.007	1	0	1	0.526	0.5990
M	0	White	0.563	0.006	1	0	1		
M	1	White	0.559	0.006	1	0	1	0.499	0.6177

(b) Racial composition

Table 3 Randomization Check: Comparison of Treatment and Control Groups Before Manipulation

Gender	Manip	Variable	Mean	StdErr	Median	Min	Max	t-value	p-value
F	0	ViewSentCount	43.664	1.139	18	0	1414		
F	1	ViewSentCount	49.830	1.378	21	0	2475	-3.449	0.0006
M	0	ViewSentCount	73.520	1.735	23	0	3216		
M	1	ViewSentCount	79.661	1.785	25	0	2777	-2.467	0.0136

Table 4 The Effect of Treatment on Outbound Views

It should be noted that there is possible ambiguity in the intent of same-sex browsing, where it could be argued that users are not necessarily expressing the relationship interest in the users of the same sex but are rather checking out their competition. While we cannot distinguish between these two interpretations, we should note that the choice of viewing same sex users is very deliberate –

Gender	Manip	Variable	Mean	StdErr	Median	Min	Max	t-value	p-value
F	0	ViewSameSex	0.068	0.004	0	0	1		
F	1	ViewSameSex	0.081	0.004	0	0	1	-2.406	0.0162
M	0	ViewSameSex	0.074	0.003	0	0	1		
M	1	ViewSameSex	0.083	0.003	0	0	1	-2.073	0.0381
F	0	ViewOtherRace	0.581	0.007	1	0	1		
F	1	ViewOtherRace	0.611	0.007	1	0	1	-2.781	0.0054
M	0	ViewOtherRace	0.723	0.005	1	0	1		
M	1	ViewOtherRace	0.729	0.005	1	0	1	-0.886	0.3757

Table 5 **Treatment Effect on Same-sex and Inter-racial Viewing**

straight users normally cannot see the profiles of users of the same gender. In order to visit a profile of the same sex a user must make an explicit decision to switch the sexual orientation preference in their settings.

Gender	Manip	Variable	Mean	StdErr	Median	Min	Max	t-value	p-value
F	0	MsgSameSex	0.008	0.001	0	0	1		
F	1	MsgSameSex	0.009	0.001	0	0	1	-0.487	0.6261
M	0	MsgSameSex	0.011	0.001	0	0	1		
M	1	MsgSameSex	0.011	0.001	0	0	1	-0.152	0.8789
F	0	MsgOtherRace	0.195	0.006	0	0	1		
F	1	MsgOtherRace	0.205	0.006	0	0	1	-1.081	0.2798
M	0	MsgOtherRace	0.339	0.005	0	0	1		
M	1	MsgOtherRace	0.342	0.005	0	0	1	-0.306	0.7595

Table 6 **Treatment Effect on Same-sex and Inter-racial Messaging**

Interestingly, despite the observed reduction in social inhibition on preferences and the lowering of search frictions, where individuals not only view more profiles but also view a broader range of profiles, the impact of anonymity on the number of matches goes in the opposite direction. Table 7 shows that despite apparent disinhibition in browsing, the total number of matches decreases significantly (by 14.1%) for women and decreases (but not significantly) for men. Note that our manipulation was exogenously randomized, and thus we do not need to control for any user characteristics to establish the average treatment effect. A t-test of observed outcomes is sufficient to establish causality of our results, and the magnitude of the means and the difference between them provides practical significance.

Gender	Manip	Variable	Mean	StdErr	Median	Min	Max	t-value	p-value
F	0	TotalMatchCount	4.089	0.116	1	0	160		
F	1	TotalMatchCount	3.513	0.096	1	0	90	3.832	0.0001
M	0	TotalMatchCount	2.419	0.078	0	0	142		
M	1	TotalMatchCount	2.268	0.075	0	0	180	1.392	0.1639

Table 7 **Treatment Effect on Total Number of Matches**

Gender	Manip	Variable	Mean	StdErr	Median	Min	Max	t-value	p-value
F	0	ViewRcvdCount	134.864	2.318	89	0	1701		
F	1	ViewRcvdCount	126.503	2.135	85	0	1423	2.653	0.0080
M	0	ViewRcvdCount	24.454	0.506	11	0	1710		
M	1	ViewRcvdCount	22.197	0.421	11	0	785	3.429	0.0006
F	0	MsgRcvdCount	20.662	0.437	12	0	469		
F	1	MsgRcvdCount	19.189	0.397	11	0	434	2.495	0.0126
M	0	MsgRcvdCount	1.444	0.040	0	0	132		
M	1	MsgRcvdCount	1.214	0.027	0	0	57	4.737	<0.0001

Table 8 Treatment Effect on Views and Messages Received

Gender	Manip	Variable	Mean	StdErr	Median	Min	Max	t-value	p-value
F	0	MsgSentCount	2.663	0.112	0	0	94		
F	1	MsgSentCount	2.760	0.117	0	0	156	-0.599	0.5490
M	0	MsgSentCount	12.711	0.624	1	0	2116		
M	1	MsgSentCount	12.725	0.536	1	0	1181	-0.017	0.9861

Table 9 Treatment Effect on Messages Sent

In order to explain the direction of the effect as well as this apparent gender asymmetry, we utilize our micro-level data and break down the initial steps of the dating process, namely viewing (weak signaling for the control group and no signaling for the treatment group) and messaging (strong signaling for both groups), by gender.

As demonstrated in Table 8, both incoming views and messages decrease significantly for both men and women because of anonymity, while the number of initiated outgoing messages remains statistically unchanged (Table 9). To explain this, recall that the only difference between an anonymous user and a non-anonymous user, from the point of view of other website users, is that the anonymous user does not leave a trace when viewing profiles. Therefore, our results suggest that a focal user's inability to leave a weak signal results in a lack of other users viewing that focal user, i.e., a user loses incoming views.

To provide more insights into this finding, we split TotalMatchCount into two variables that indicate whether the match was initiated by the focal user, MatchSentCount, or by the counter-party, MatchRcvdCount. Based on the results reported in Table 10, we can clearly see that MatchSentCount and MatchRcvdCount are indeed affected differently by our manipulation. MatchSentCount remains statistically unchanged for both genders (just like MsgSentCount), while MatchRcvdCount is reduced significantly with a drop of 24.7% for men and 18.5% for women.

This finding clearly explains the observed gender asymmetry in the effect of anonymity on TotalMatchCount. As demonstrated, both genders lose close to one of out every five of their "matches received" because of anonymity. Yet, unlike women, most of the matches for men are "matches sent" (that are unaffected by anonymity), not "matches received." Therefore, the similar scale of reduction in the "matches received" induced by anonymity does not have as significant an

Gender	Manip	Variable	Mean	StdErr	Median	Min	Max	t-value	p-value
F	0	MatchRcvdCount	3.199	0.093	1	0	130		
F	1	MatchRcvdCount	2.607	0.075	1	0	86	4.937	<0.0001
M	0	MatchRcvdCount	0.504	0.017	0	0	66		
M	1	MatchRcvdCount	0.379	0.010	0	0	10	6.381	<0.0001
F	0	MatchSentCount	0.891	0.041	0	0	42		
F	1	MatchSentCount	0.906	0.040	0	0	49	-0.275	0.7835
M	0	MatchSentCount	1.916	0.072	0	0	141		
M	1	MatchSentCount	1.889	0.072	0	0	170	0.260	0.7951

Table 10 Treatment Effect on Matches Received and Matches Sent

impact on the total number of matches for men as it does for women. This finding demonstrates that the ability to send a weak signal is especially helpful for women who tend to rely on incoming messages and not make the first move.

6.2. Weak-signaling efficacy

We can further see the efficacy of the weak signal in generating a response by looking at the response of the potential mate to a profile view left by focal user. Specifically, we define a new measure, MsgReplyToWeakSignal, which is constructed as follows. Consider an event such as user i viewed user j but user i did not message her, then user j messaged user i for the first time[7]. Assume Y_i denotes the number of such events that occurred to user i during the treatment month. We then compute the average of this measure for all treated users and separately for all control users. The measure is well defined for every user (if user i didn't send any views then $Y_i = 0$), so no selection bias occurs. The results of this analysis (Table 11) show a significant difference between treatment and control users' ability to generate a message in response to visiting another user's profile.

This finding emphasizes the importance of weak-signaling: despite visiting more profiles, the treated users were visited by a smaller number of potential mates. Furthermore, the average number of received messages triggered by a sent view for men is 0.17 in the control group, which is significantly higher than that of the treatment group at 0.08 ($p < 0.0001$). This also holds for women, where the average number of received messages triggered by a view in the control group is 0.63 and in the treatment group is 0.25 ($p < 0.0001$). Note that the treated users have non-zero values of this variable since this particular sequence of events can happen by chance as users can be found through other means on the site including search and recommendations.

Further, we can also empirically demonstrate a particular type of disinhibition that we call "lack of regret" that anonymous users experience, as alluded to in Section 1. We examine the

[7] Note that the site does not supply an explicit figure showing the total number of visits from a given user, thus our context does not allow us to quantify the effect of multiple weak signals. Future research in a different context that does reveal the total number of visits could causally examine this question with a blocking-style research design (Aral and Walker 2011).

Gender	Manip	Variable	Mean	StdErr	Median	Min	Max	t-value	p-value
F	0	MsgReplyToWeakSignal	0.632	0.032	0	0	36		
F	1	MsgReplyToWeakSignal	0.251	0.019	0	0	37	10.369	<0.0001
M	0	MsgReplyToWeakSignal	0.169	0.008	0	0	13		
M	1	MsgReplyToWeakSignal	0.076	0.004	0	0	14	10.329	<0.0001

Table 11 Messages Triggered by a View

"regret effect" by constructing a measure of RegretSent. This measure captures a sequence of events where the focal user views but does not message the target user, and the target user subsequently messages the focal user but does not receive a reply back. Given that the focal user collected the information and decided neither to initiate the message towards the target user nor reply to the message from the target user, we assume the focal user collected the necessary information and is no longer interested in the target user. From this chain of events, we infer that the initial profile view ultimately ended up being "regret-laden," and as expected we find that the control group users exhibit this in significantly higher proportions as compared to the treatment group. We have included the table of these results in Online Appendix D.

We treat weak signaling as one specific case of many different types of mechanisms for reducing inward search costs, the search costs that other users incur to find a focal user. Other potential ways of reducing inward search costs, other than weak signaling, may include, say, buying a promotion to temporarily boost the user's ranking in the search results, or any other activity that can promote the user's profile. It should be noted that of all the incoming views that a user receives, 95.4% come from sources other than leaving a weak signal through a profile view visit, most notably from search and recommendations. Note that the decrease in incoming views from the anonymity treatment (Table 8) is approximately 6.2% for women (9.23% for men), largely corresponding to the overall percent of views received by a typical user via weak-signals. However, this 6.2% (and 9.23% respectively) reduction in received views is linked to a 18.5% (24.7%) reduction in received matches. This is important because it highlights the prominence of the weak signaling component. In other words, the profile view leaves a signal that has important information for triggering a social exchange that leads to a match. Blocking a small percent of these views causes a disproportionately large decline in matches.

Being an average effect, the above results do not provide insights on the scale of the marginal effects experienced by individual users. To explore this further, we model each user's matches as a function of a rich set of his/her demographic and other user-specific information. We focus our analysis on MatchRcvdCount as this is where the role of weak signaling manifests itself. Given that MatchRcvdCount, our dependent variable, is a count variable, we fit a negative binomial regression model[8] (Online Appendix A has the full details and results) using manipulation (our treatment) as an independent variable that is uncorrelated with the residual (because of exogenous randomization of this variable), while controlling for observed characteristics namely age, attractiveness, race[9], and orientation. The negative binomial model suggests that the average individual loses close to 1.2 matches per month due to anonymity with significant heterogeneity across attractiveness and gender, especially for White men as well as Asian women.

6.3. Examining different definitions of a match

While we describe a series of robustness checks in Section 7, we seek to firmly establish here that our overall results are not sensitive to the exact definition of match with respect to the length of the message chain. As mentioned in Section 4, Hitsch et al. (2010) who had access to the actual content of the messages in an online dating website, found that it takes a median number of six messages for users to reveal their phone number, email address, or for a key phrase such as "let's meet" to appear. Drawing on this, we define an n-match between users i and j if user i and j exchanged at least n messages. To be conservative, we repeat our analysis for n ranging from four to twelve (Figure 2).

Looking at Figure 2, the top left plot represents the number of n-matches received by women (y-axis) based on the definition of an n-match (x-axis), which varies from $n = 4$ up to[10] to $n = 12$. We find that the effect of anonymity is strongly observed and constitutes the same decrease of approximately 20%-25% in matches irrespective of what threshold n is used as the definition of n-match. Consistent with our main results, we also observe that MatchSentCount has no significant difference between treatment and control users, for all n.

6.4. Selectivity

As shown in Section 6.1, our treatment causes a decrease in incoming matches, while causing no change in outgoing matches. Given that a match is an outcome of a sequence of messaging events, there are multiple mechanisms that can cause a decrease in the number of incoming matches in the manipulated group. For instance, it is possible that because our treatment causes a decrease in the number of incoming messages (see Table 8) a user has fewer messages to respond to and therefore fewer chances to establish a match. At the same time, it could be the case that our

[8] We also fit a Zero-Inflated Negative Binomial (ZINB) and Zero-Inflated Poisson (ZIP) model. The AIC and BIC fit statistics indicated that the negative binomial model had the fit that is comparable to zero-inflated models. Negative binomial model has the additional advantage of providing coefficients that are easier to interpret in terms of estimated marginal effects as compared to ZIP and ZINB. Therefore, we pick negative-binomial regression for our analysis.

[9] We expand the non-white race category here to obtain finer grained marginal effects.

[10] The value 12 on the x-axis corresponds to $n = 12$ which means that we are looking at 12-matches as our outcome variable. A conversation only qualifies as a 12-match if it consists of at least 12 messages exchanged or more.

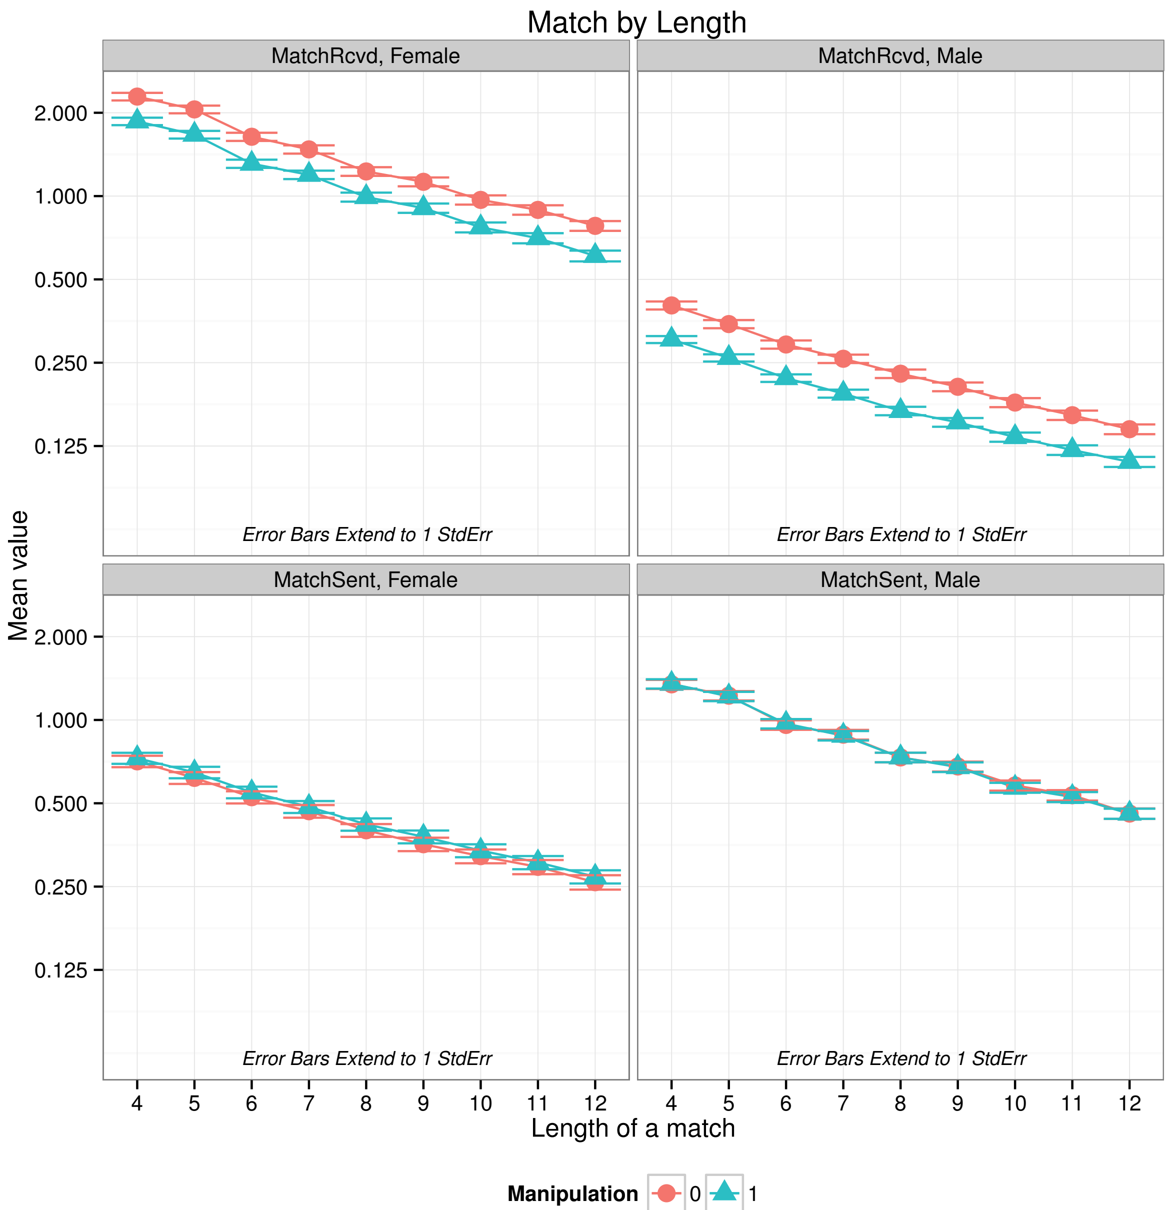

Figure 2 Robustness of Anonymity Effect to Message Exchange Length. Length Ranges from 4 to 12.

manipulation causes changes in reply patterns. That is, anonymous users may be more likely to explore prospective candidates who have messaged them, as they are more disinhibited, before choosing whether or not to reply.

We have already demonstrated that incoming messages do decrease due to anonymity. We now focus on examining the second mechanism by defining a selectivity variable, called GetCheckRatio, which is the ratio of the number of message senders the focal user viewed after receiving their initial message, divided by the total number of message senders that contacted the focal user. We extend this logic by defining another selectivity construct called GetReplyRatio, which is the ratio of the number of message senders the focal user replied to, divided by the total number of message senders that contacted the focal user[11]. Thus, our conceptualization of selectivity focuses on one key construct: the response rate to messages received. This conceptualization of selectivity assumes no differences in the incoming pool of potential mates. We verify that this is the case empirically in Section 6.5 where we show that the average attractiveness rating of users sending messages to the focal users is not significantly different in the treatment and the control groups.

It is evident in Table 12 that under anonymity both genders indeed prefer to visit a larger proportion of the profiles of those who messaged them. However, as demonstrated by GetReplyRatio variable, treated women are 10.5% less likely to reply to people who messaged them compared to the control group of women. The effect for men is 4.42% which is smaller but still marginally significant. The stronger selectivity effect could also contribute to receiving less incoming matches in total. Therefore, the mechanism of the effect of anonymous browsing is not limited to simply reducing the incoming communication for the focal user, but also incorporates the changes in the subsequent selection behavior of the focal user herself. Our findings are suggestive of both weak signaling and selectivity as an explanation for the decline in the number of received matches.

Gender	Manip	Variable	Mean	StdErr	Median	Min	Max	t-value	p-value
F	0	GetCheckRatio	0.416	0.004	0.433	0	1		
F	1	GetCheckRatio	0.455	0.004	0.500	0	1	-6.238	<0.0001
M	0	GetCheckRatio	0.532	0.007	0.500	0	1		
M	1	GetCheckRatio	0.587	0.007	0.667	0	1	-5.744	<0.0001
F	0	GetReplyRatio	0.178	0.003	0.111	0	1		
F	1	GetReplyRatio	0.160	0.003	0.087	0	1	4.011	<0.0001
M	0	GetReplyRatio	0.411	0.007	0.333	0	1		
M	1	GetReplyRatio	0.393	0.007	0.286	0	1	1.840	0.0658

Table 12 Effect of Anonymity on Selectivity

6.5. The impact of anonymity on match quality

The prior section demonstrates that women are more selective in responding to messages under anonymity, which helps explain the reduction in their number of matches. This motivates us to examine whether the quality of these fewer matches that they do get is higher in the treatment group as compared to the control group. As explained in Section 4, in some of the analysis that follows we use attractiveness ratings as a proxy measure for overall quality of the match, which arguably may not be the perfect measure of quality of a long-term relationship. However, there

[11] Intuitively, GetCheckRatio can be phrased as “how often you check out profiles of those who just messaged you”, while GetReplyRatio can be phrased as “how often you reply to those who just messaged you”.

exists experimental evidence suggesting that the attractiveness score is what users mainly care about in the context of online dating (Rudder 2014). Given that we do not know individuals' ultimate objective functions, we approach the question of quality by postulating several preference axioms that a typical online dating user might consider unambiguous and self-evident:

1. It is generally preferable to find at least one match as compared to none;
2. Given similar attractiveness, it is generally preferable to find the match faster;
3. It is generally preferable to match to a more attractive person[12].

We address Axiom 1 using a logit model to predict whether a user gets at least one match or not based on the treatment and control variables. Our analysis indicates that the odds of receiving at least one match decreases by 24.2% for women and 17.5% for men under anonymity (see Online Appendix B Tables B.2 and B.3).

In order to address Axiom 2, we model the hazard rate of getting the first match and find that the time to receiving the first match is significantly longer under anonymity (Online Appendix B has the full model). For ease of interpretation of the Cox proportional hazard model, we present a chart in Figure 3 that breaks down the hazard of receiving the first match by gender, by who initiated contact, and by the different definitions of n-match (the meaning of the x-axis is the same as in Figure 2). Figure 3 shows the estimated hazard ratio of the treatment coefficient with whiskers indicating standard errors. The negative effect means that the treatment group (Manipulation=1) takes longer, on average, than the control group (Manipulation=0) to receive a match. Consistent with our prior results, we find that while there is a significant increase in the time to receiving the first match for the treated group, there is no change in the time to sending a first match. This result is consistent for n-matches for any n ranging from four to twelve messages.

With respect to Axiom 3, a key variable that sheds light on the quality of the match is the average rating that each user of the site receives. Generally, a rating by one user (rater) for another (ratee) reflects the rater's perception of the overall attractiveness of the ratee. This "attractiveness" signal is a composite of the rater's perception of the ratee's entire persona (characteristics, pictures, answers to a series of questions that the site asks, etc.) reflected in the ratee's profile. For each user, we compute their average rating, if they received at least one rating. A natural starting point to measure the impact of our treatment on match quality is to see whether anonymity causes individuals to match with users of a higher attractiveness score than they would have without anonymity. This could indicate an improvement in the quality of the match. To compare the attractiveness of users' matches under treatment versus control, one has to control for the likelihood of getting a match in the first place, which we know is impacted by the treatment. We model

[12] See (Rudder 2014)

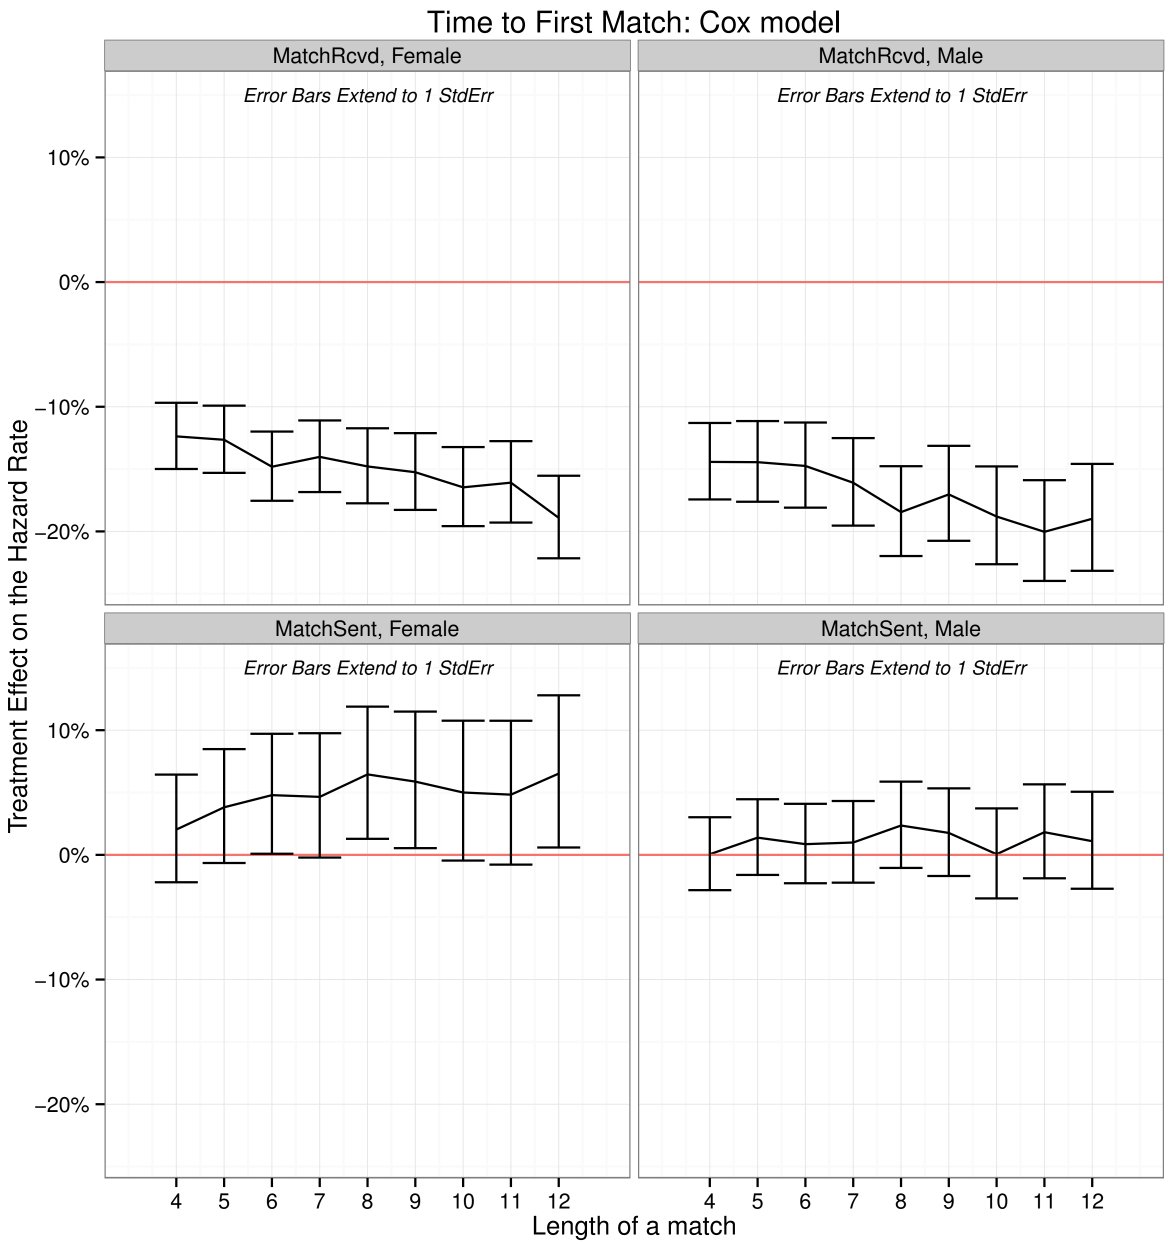

Figure 3 **Treatment Effect on the "Hazard" of Finding a Match**

this using a two-stage Heckman selection model where the first stage reflects the likelihood of getting matched (which is a function of the treatment) and the second stage reflects the average attractiveness score of the users given that at least one match occurred. For illustrative purposes, we provide the full model and results for attractiveness of received matches for women in Online Appendix B (Table B.1). Each row in Table 13 reports the relevant treatment coefficient from the second stage of a Heckman model with average attractiveness as the dependent variable. We find no significant differences in attractiveness of the users who were viewed, messaged and matched under treatment as compared to control, i.e. the treatment coefficient is insignificant in each of these cases. For sake of completeness, it should be noted that for weak signaling to have an overall negative impact, the increase in the number of matches must be sufficiently outweighed by significant decreases in match quality.

Gender	Variable	Estimate	StdErr	t-value	p-value
F	ViewSentAvgAttract	0.007	0.012	0.569	0.5693
M	ViewSentAvgAttract	0.002	0.008	0.280	0.7794
F	ViewRcvdAvgAttract	-0.005	0.009	-0.616	0.5378
M	ViewRcvdAvgAttract	-0.017	0.010	-1.690	0.0911
F	MsgSentAvgAttract	-0.016	0.020	-0.811	0.4176
M	MsgSentAvgAttract	0.003	0.014	0.240	0.8101
F	MsgRcvdAvgAttract	0.011	0.010	1.086	0.2773
M	MsgRcvdAvgAttract	-0.025	0.024	-1.055	0.2915
F	MatchSentAvgAttract	-0.050	0.032	-1.561	0.1185
M	MatchSentAvgAttract	0.011	0.023	0.473	0.6365
F	MatchRcvdAvgAttract	-0.036	0.022	-1.594	0.1109
M	MatchRcvdAvgAttract	-0.033	0.083	-0.395	0.6929

Table 13 Coefficients of the Second Stage of a Heckman model of Average Attractiveness on Viewing, Messaging, and Matching.

Overall, it is evident that under anonymity: 1) users are less likely to receive a match; 2) they take longer to receive the first match; 3) conditional on receiving a match, there is no significant increase (nor a decrease) in the attractiveness of the match.

7. Robustness

We carried out a series of robustness analyses to rule out alternative explanations for our results. First, we examine whether treated users quit the site earlier due to securing a highly attractive match. Second, we explore matches on a weekly level and test whether our monthly granularity is adequate to explain the effect. Third, we explicitly test whether the observed effect is induced simply by the happiness associated with receiving a gift rather than by anonymity itself. Finally, we conclude by providing evidence that users revert to their original behavior in month 3 immediately once anonymity expires demonstrating no observable learning effects in the average population that we analyze.

7.1. Exiting due to a highly attractive match

One potential explanation for the decrease in the monthly number of matches is due to users finding an initial really attractive (high quality) match under anonymity and therefore modifying their subsequent behavior. To test this, we examine the time and the propensity for treated users to obtain matches in the top 20th percentile of attractiveness. Tables B.5 and B.6 in Online

Appendix B reveal that while there is no difference in the duration and propensity of achieving outgoing matches, treated users of both genders not only take longer (i.e. have a lower hazard) to find a highly attractive incoming match, but also get fewer of these highly attractive matches. These results are consistent with our general analysis, and rule out the potential explanation that anonymity causes users to find highly attractive partners.

7.2. Weekly analysis

In the same vein, we explore another alternate explanation related to obtaining a match under treatment by examining the temporal variation in activity. That is, it is possible that the granularity of our analysis is not precise enough and that users under treatment find their partner, say, in the first week of the month and thus decrease their activity in the remainder of the month leading us to observe the overall monthly decline. More specifically, it could be the case that during the treatment month, treated users actually initially search and message more, but as they find more potential matches, they reduce their messaging activity with the other users. To test this possibility, we look at the weekly data instead of monthly and compare weekly messaging activity for each of the four weeks of the treatment month. As shown in Figure 4, we continue to find a consistent difference in the behavior in the treatment and control groups with respect to matches (and messages; results omitted for brevity). We find that our weekly results replicate exactly what we see with the monthly data: while there is no difference in the average number of sent matches between treatment and control, the average number of received matches is significantly lower under treatment. These results are consistent with our monthly-level analysis, and rule out the potential explanation that anonymity causes users to quickly find their partners and reduce subsequent activity on the site.

Finally, if users find their ideal match during the treatment month, we would expect a decrease in activity after the month of treatment. To examine this, we compare the activity levels of users in the treatment and control groups in the three last weeks of the post-treatment month 3. We find that the percent of active users is the same for the treatment and control groups (Online Appendix C Table C.1).

Treated users also exhibit the same behaviors as the control group: the average duration between login times, the maximum time between logins, and the average “conversation” length, are all statistically equivalent for the treatment and control groups and for both genders during the treatment month (Online Appendix C Table C.2).

7.3. Ruling out a gift effect

We want to rule out the possibility that users are simply acting in response to the initial impulse associated with the act of receiving a gift in the beginning of month 2, rather than to anonymity itself. To do this we examine whether the treatment effect was present in the last week of the

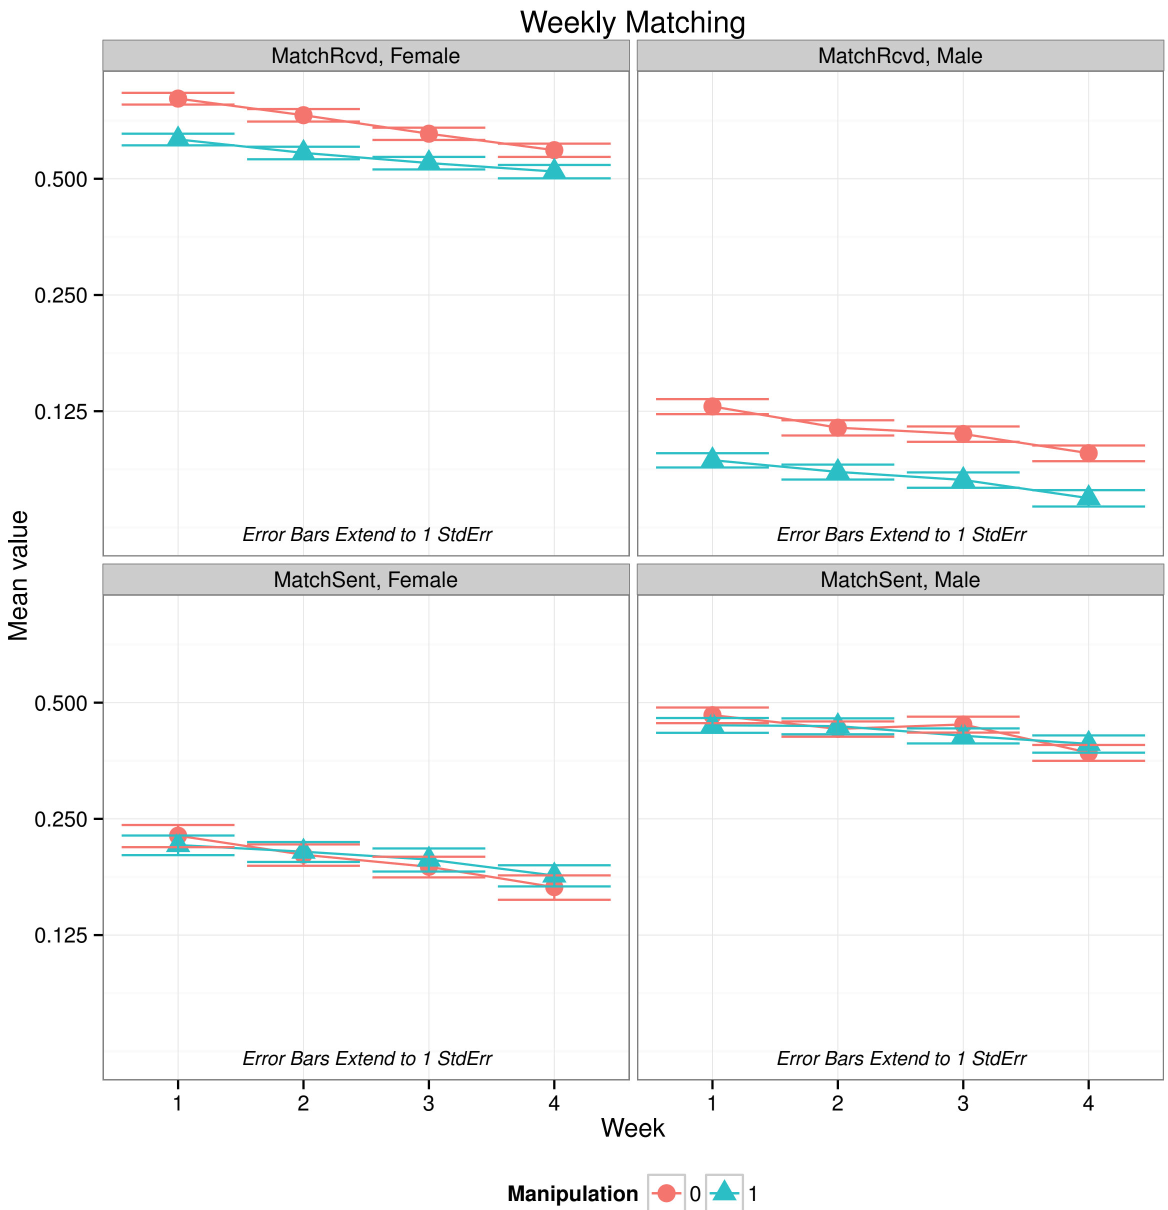

Figure 4 **Average Number of Sent/Received Messages: Weekly Breakdown for the Treatment Month**

treatment month, and compare it to the (adjacent) first week of the post-treatment month to see if the effect disappears. If the effect persists in the last week of the treatment month and immediately disappears in the first week of post-treatment month, we can rule out the gift effect since the gift effect should not demonstrate a sharp decline between weeks 4 and 5 after the gift was awarded. Table 14 shows that the effect of anonymity is strongly observed in the last week of the treatment month, yet disappears as early as the first week of the post-treatment month, indicating the salience of the manipulation and ruling out the gift effect.

Gender	Manip	Variable	Mean	StdErr	Median	Min	Max	t-value	p-value
F	0	MatchRcvdLastW	0.593	0.023	0	0	31		
F	1	MatchRcvdLastW	0.522	0.021	0	0	25	2.264	0.0236
M	0	MatchRcvdLastW	0.097	0.005	0	0	12		
M	1	MatchRcvdLastW	0.074	0.004	0	0	5	3.932	<0.0001
F	0	MatchRcvdPostFirstW	0.540	0.022	0	0	30		
F	1	MatchRcvdPostFirstW	0.522	0.024	0	0	53	0.552	0.5811
M	0	MatchRcvdPostFirstW	0.096	0.005	0	0	11		
M	1	MatchRcvdPostFirstW	0.100	0.004	0	0	5	-0.618	0.5363

Table 14 Gift Effect Check

7.4. Ruling out other alternate explanations

If the differences between the treatment and control groups in response to our manipulation were due to any reason other than our treatment, we would expect the effect to persist after our treatment expired. Recall that there was no difference in the characteristics, behaviors, and outcomes between the treatment and control groups in the month prior to treatment. To rule out alternate explanations of the treatment, we compare the treatment and control groups in the month after the manipulation expired. We find that the effect of anonymity completely disappears in month 3 (post-treatment month). As demonstrated by Tables 15-17, the treatment group reverts to the behavior that is statistically indistinguishable from the control group in the post-treatment month.

Gender	Manip	Variable	Mean	StdErr	Median	Min	Max	t-value	p-value
F	0	ViewSentCountPost	29.447	0.908	9	0	1251		
F	1	ViewSentCountPost	30.729	0.903	10	0	855	-1.001	0.3169
M	0	ViewSentCountPost	55.293	1.412	14	0	2652		
M	1	ViewSentCountPost	56.377	1.493	14	0	3566	-0.528	0.5977
F	0	ViewRcvdCountPost	96.903	1.799	59	0	1359		
F	1	ViewRcvdCountPost	94.849	1.768	60	0	2017	0.814	0.4154
M	0	ViewRcvdCountPost	18.871	0.377	8	0	498		
M	1	ViewRcvdCountPost	18.933	0.381	8	0	617	-0.116	0.9074

Table 15 Profile Visits in Post-Treatment Month

Gender	Manip	Variable	Mean	StdErr	Median	Min	Max	t-value	p-value
F	0	MsgSentCountPost	1.855	0.091	0	0	104		
F	1	MsgSentCountPost	1.794	0.086	0	0	131	0.492	0.6228
M	0	MsgSentCountPost	9.446	0.486	0	0	1600		
M	1	MsgSentCountPost	9.385	0.464	0	0	1465	0.090	0.9281
F	0	MsgRcvdCountPost	14.393	0.346	8	0	461		
F	1	MsgRcvdCountPost	14.060	0.319	8	0	343	0.708	0.4790
M	0	MsgRcvdCountPost	1.096	0.031	0	0	119		
M	1	MsgRcvdCountPost	1.077	0.025	0	0	39	0.491	0.6235

Table 16 Messaging in Post-Treatment Month

Gender	Manip	Variable	Mean	StdErr	Median	Min	Max	t-value	p-value
F	0	MatchSentCountPost	0.574	0.031	0	0	42		
F	1	MatchSentCountPost	0.570	0.030	0	0	60	0.079	0.9374
M	0	MatchSentCountPost	1.349	0.055	0	0	111		
M	1	MatchSentCountPost	1.345	0.058	0	0	159	0.048	0.9618
F	0	MatchRcvdCountPost	1.876	0.070	0	0	140		
F	1	MatchRcvdCountPost	1.820	0.065	0	0	137	0.584	0.5595
M	0	MatchRcvdCountPost	0.348	0.012	0	0	32		
M	1	MatchRcvdCountPost	0.336	0.010	0	0	20	0.716	0.4738

Table 17 Matches in Post-Treatment Month

The astute reader could also interpret the treatment group reverting to behavior that is statistically indistinguishable from the control group as an indication that the effect of anonymity is a short-term effect. In the extreme case that a match on this site, which we used as an indication of successful online communication, is indicative of the beginning of a long-term relationship, we would expect users in the control group to stop visiting the site. At the same time, the lack of matches for users in the treatment group could also result in their absence or departure from the site due to frustration. Therefore, the lack of changes in month 3 is not necessarily evidence of the lack of long-term differences between the treatment and the control group. Unfortunately, our data and time-span examined do not allow us to draw conclusions on whether there are long-term differences between the treatment and control groups.

8. Discussion and Conclusions

Online dating platforms are rapidly growing worldwide. Our work is motivated by the fact that today's IT-enabled online dating and matching platforms introduce new capabilities and features, the causal impact of which are hard to analyze meaningfully in the absence of randomized control experiments. These capabilities range from enhanced search to big-data based mate recommendations (much like Amazon recommends a book or Netflix recommends a movie, with the added nuance that while for books and movies the consumer must like the book but the book need not like the consumer, in dating markets both individuals must like each other) to anonymity-linked features such as weak-signaling, which is the focus of this paper. As we demonstrate, these new technological capabilities affect user behavior and outcomes in a number of different ways that are not always easy to anticipate. Emerging research from scholars in psychology is examining questions of whether online dating is fundamentally different from dating in the offline world and whether online dating leads to better romantic outcomes (Finkel et al. 2012). This research is being examined in an emerging stream of work called "relationship science," where psychologists are going beyond their emphasis on individuals' motivations and actions to a focus on the dyadic nature of romantic liaisons (Reis et al. 2013). While this stream of literature contrasts online versus offline dating, we focus on key features enabled through the online environment.

Our theoretical understanding of dating and marriage starts with Becker (1973)'s exposition of assortative matching as the equilibrium outcome assuming a frictionless market. Subsequent research (Burdett and Coles 1997, Mortensen and Pissarides 1999, Smith 2006) has theoretically accounted for search frictions in the characterization of sorting equilibriums. While the extant research has limited its attention to search frictions, we develop the idea that social frictions, as imposed by long-standing social norms such as women not making the first move in dating markets, are an important and, prior to this study, largely unstudied (with the exception of Piskorski (2014)) source of inefficiencies in such contexts. Thus, in departure from the extant literature, our treatment is in effect a horse race between search frictions, which anonymity lowers and should result in more matches, and social frictions, which rise when we take away weak-signaling and therefore should result in fewer matches. Again, while the economics literature has extended the original frictionless matching models to account for search frictions, no one has looked at social frictions and compared the two types of frictions in the setting of a randomized field experiment.

In particular, in this paper, we explore the effect of one such anonymity-linked feature that we refer to as weak-signaling. The users who are treated with anonymity are able to browse potential mates' profiles anonymously, without leaving a trail in the "Recent Visitors" list of the target user. While conventional economic wisdom suggests that such anonymity of profile viewing should be associated with improved matching outcomes by reducing search costs and allowing users to explore their options freely, our study results demonstrate that, to the contrary, there is a significant drop in the number of matches, particularly for women. Based on data available to us, it appears that this decrease in match quantity is not compensated by an increase in match quality. We demonstrate, by breaking down, measuring, and analyzing the mate-seeking process in detail that the smaller number of matches occurs because of a dominating social friction force of being unable to leave a weak signal. Women, who are more reluctant to make the first move, such as by messaging a potential mate, are deprived by our treatment of leaving a profile visit trail, which turns out to be the import trigger to provoke incoming messages and subsequent matches.

Online dating platforms provide us with an environment where participants' choices at sub-stages of the dating process are available to the researcher in unprecedented detail. We exploit this rich micro-level sub-process dating data to explain our key findings in a detailed and nuanced manner. In particular, we recognize that there are sub-processes of viewing and messaging that precede the final matching outcome (which has been the outcome considered by the existing literature, e.g., Hitsch et al. (2010)). Further, each of these sub-processes can be initiated by the focal user or the target user, giving rise to many possible permutations of arriving at a match. While prior literature has not considered this microscopic view, we find that it is crucial for better understanding our main results.

When we break down total matches into received and sent matches, based on who initiated the process, we find that our manipulation lowers the number of received matches while causing no change in the count of sent matches. Because our manipulation decreases the number of incoming messages, a treated user has fewer messages to respond to, and therefore fewer chances to establish a match. We also find that our manipulation causes changes in reply patterns. Anonymity lowers search frictions — as expected, anonymous users explore and visit more profiles of users who messaged them before making a decision to return a message. Interestingly, in addition to these baseline mechanisms, we find evidence of selectivity under the anonymity condition. In particular, treated women, while visiting more profiles of the men who messaged them, reply at a significantly lower rate (10.5% less) of those visited profiles, as compared to the control group of women. The effect for men is much smaller (4.42% less) and is only marginally significant. This finding is consistent with the disinhibition effect: under the cloak of anonymity, users are more compelled to visit the profile of the other user before deciding to reply. Needless to say, this also contributes to receiving less incoming matches by women in total, but these matches are a result of a more selective process exercised by women. We expect future research to examine in more depth the issue of match quality and long-term outcomes as they relate to marriage, happiness, long-term relationships, and divorce.

Users of online dating websites may have additional components of utility in their objective function that are beyond online dating itself. For example, users may still prefer their privacy even if they know they lose some successful communications because of it. Thus, some users might be willing to trade off online dating success for privacy. In other words, we generally position our paper as providing experimental evidence of the effect of anonymity on success in online dating communications — looking at both quantity and quality of matches – and we refrain from making any claims about the overall impact on the user's welfare, which may have an unobserved utility component from privacy.

In exploring social inhibitions in dating markets, we used the anonymity feature of online dating platforms as our manipulation. This feature is a double-edged sword, lowering social inhibitions along the preferences dimension and increasing social inhibitions along the contact initiation dimension. In this study, we are unable to estimate the absolute strength of each of these inhibitions separately, but rather quantify the total overall effect that these inhibitions cause in real life when induced by an external feature such as anonymity. We expect future work to study each mechanism separately. Future research can also more explicitly examine these trade-offs between the quality of the match and the quantity and time-to-match. For instance, would users be indifferent between matching with a three-star (attractiveness rated) partner in one day versus a four-star partner in fourteen days.

Our work is intended to be directly generalized only to the online dating context, considering the possibility of systematic differences between the underlying populations of the online and offline dating worlds. That said, recent research such as Hitsch et al. (2010), Rudder (2014) finds that the distance between the populations of online and offline dating markets is much closer than what was previously thought. Current trends suggest that about one-third to one-half of the single population of the United States now uses online dating websites in their search for a romantic partner (Slater 2013, Gelles 2011). One main reason cited for this growth is the reduction in search costs (Slater 2013, Paumgarten 2011) as individuals can view multiple in-depth profiles of potential matches in a short amount of time, with further help coming from algorithmic recommendations by the websites.

Another aspect of our work is that our inference is limited to those users who do not tend to purchase an anonymity feature and thus, were not subscribed to the premium package. However, this does not restrict our generalizability noticeably since approximately 99% of all users fall into that category. Premium subscribers who did purchase anonymity on their own and were excluded from our experiment comprise of only 1% of this online dating site, as is typical of freemium communities.

Beyond the context of online dating, anonymity is a key feature on a variety of online platforms where people connect with one another including Facebook, where social interactions might be altered with a switch in anonymity settings, and LinkedIn, where anonymity can play a key role in job market outcomes. Matching two individuals is a complex task, relative to, say, matching a buyer with a product in product markets. In dating, as well as in job search, there are two sets of preferences that must be taken into account to produce a successful match. Matching two humans applies not only to dating, marriage and job search, but also to new models of crowdsourcing (Burtch et al. 2013). Thus, we expect this study and our associated methodology to be the basis for a stream of work on how the Internet and social media are changing some of the fundamental activities we carry out as humans.

Our work fits under the broader umbrella of emerging research examining the societal impact of the new generation of big-data enabled online social platforms that connect people who either know each other (e.g., Facebook) or would like to know each other (e.g., eHarmony and Match.com) (Piskorski 2014). These newer platforms reduce many of the frictions present in the offline world that always exists as a close substitute for them. Many replicate the prior social processes that they digitize; others extend and expand these social processes with new capabilities and at the same time have larger implications. For instance, viewed from another lens, anonymous browsing can be thought of as an increased level of privacy, much like Facebook offers to its users who can check out their friends' profiles any number of times without their friends being aware of it. Thus, our research fits into the emerging stream of work that evaluates the impact of various levels of privacy protection on individuals' outcomes (Solove 2004, Goldfarb and Tucker 2011, Romanosky et al. 2011, Miller and Tucker 2009).

What is consistent across these platforms is that the very act of digitization of these social processes gives us unprecedented micro-level data and access not only to outcomes but also to the underlying sub-processes of getting to the outcomes. This, we argue, is a revolutionary research opening that awaits the broader scientific community. It is an opportunity to understand human behavior around fundamental social, economic, and emotional decisions we make at a level we have been unable to achieve in the past.

Acknowledgments

The paper benefited from feedback received from the participants of the 2013 NBER Summer Institute on the Economics of IT and Digitization, the 2012 Workshop on Information Systems and Economics, the 2013 Statistical Challenges in E-Commerce Research symposium, the 2013 Winter Conference on Business Intelligence, the 2013 International Symposium on IS, as well as the Friday workshop series at Carnegie Mellon University, University of Minnesota, Georgia Institute of Technology, and Michigan State University. In particular, we would like to thank our NBER and WISE discussants Ali Hortacsu and Sarah Rice, respectively, for their invaluable inputs in shaping this paper

References

Aral, Sinan, Dylan Walker. 2011. Creating social contagion through viral product design: A randomized trial of peer influence in networks. *Management Science* **57**(9) 1623–1639.

Bakos, J Yannis. 1997. Reducing buyer search costs: Implications for electronic marketplaces. *Management science* **43**(12) 1676–1692.

Bapna, Ravi, Akhmed Umyarov. 2015. Do your online friends make you pay? a randomized field experiment on peer influence in online social networks. *Management Science* .

Becker, Gary S. 1973. A theory of marriage: Part i. *The Journal of Political Economy* 813–846.

Burdett, Ken, Melvyn G Coles. 1997. Marriage and class. *The Quarterly Journal of Economics* 141–168.

Burtch, Gordon, Anindya Ghose, Sunil Wattal. 2013. An empirical examination of the antecedents and consequences of contribution patterns in crowd-funded markets. *Information Systems Research* **24**(3) 499–519.

Ellison, Nicole, Rebecca Heino, Jennifer Gibbs. 2006. Managing impressions online: Self-presentation processes in the online dating environment. *Journal of Computer-Mediated Communication* **11**(2) 415–441.

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

Fiore, Andrew T, Lindsay Shaw Taylor, Xiaomeng Zhong, Gerald A Mendelsohn, Coye Cheshire. 2010. Who's right and who writes: People, profiles, contacts, and replies in online dating. *HICSS*. 1–10.

Fisman, Raymond, Sheena S Iyengar, Emir Kamenica, Itamar Simonson. 2006. Gender differences in mate selection: Evidence from a speed dating experiment. *The Quarterly Journal of Economics* 673–697.

Frost, Jeana H, Zoe Chance, Michael I Norton, Dan Ariely. 2008. People are experience goods: Improving online dating with virtual dates. *Journal of Interactive Marketing* **22**(1) 51–61.

Gale, David, Lloyd S Shapley. 1962. College admissions and the stability of marriage. *American mathematical monthly* 9–15.

Gelles, David. 2011. Inside Match.com. `http://on.ft.com/mZSt7K`. [Online; accessed 05-May-2015].

Goldfarb, Avi, Catherine E Tucker. 2011. Privacy regulation and online advertising. *Management Science* **57**(1) 57–71.

Hall, Jeffrey A, Steve Carter, Michael J Cody, Julie M Albright. 2010. Individual differences in the communication of romantic interest: Development of the flirting styles inventory. *Communication Quarterly* **58**(4) 365–393.

Harris, Tina M, Pamela J Kalbfleisch. 2000. Interracial dating: The implications of race for initiating a romantic relationship. *Howard journal of Communication* **11**(1) 49–64.

Henningsen, David Dryden. 2004. Flirting with meaning: An examination of miscommunication in flirting interactions. *Sex roles* **50**(7-8) 481–489.

Hitsch, Günter J, Ali Hortaçsu, Dan Ariely. 2010. Matching and sorting in online dating. *The American Economic Review* 130–163.

Holmes, Ronald M, Richard Tewksbury, Stephen T Holmes. 1998. Hidden jpgs: A functional alternative to voyeurism. *The Journal of Popular Culture* **32**(3) 17–29.

Joinson, Adam. 1998. Causes and implications of disinhibited behavior on the internet. .

Kalmijn, Matthijs. 1998. Intermarriage and homogamy: Causes, patterns, trends. *Annual review of sociology* 395–421.

Kling, Rob, Ya-ching Lee, Al Teich, Mark S Frankel. 1999. Assessing anonymous communication on the internet: policy deliberations. *The Information Society* **15**(2) 79–90.

Maccoby, Eleanor E, Carol Nagy Jacklin. 1974. *The psychology of sex differences*, vol. 1. Stanford University Press.

McDevitt, Ryan C. 2012. The internet lowers inhibitions: Implications for the long tail. *University of Rochester working paper* .

Miller, Amalia R, Catherine Tucker. 2009. Privacy protection and technology diffusion: The case of electronic medical records. *Management Science* **55**(7) 1077–1093.

Mortensen, Dale T, Christopher A Pissarides. 1999. New developments in models of search in the labor market. *Handbook of labor economics* **3** 2567–2627.

Pachankis, John E, Marvin R Goldfried. 2006. Social anxiety in young gay men. *Journal of Anxiety Disorders* **20**(8) 996–1015.

Paumgarten, Nick. 2011. Looking For Someone: Sex, Love, And Loneliness On The Internet. `www.newyorker.com/reporting/2011/07/04/110704fa_fact_paumgarten`. [Online; accessed 05-May-2015].

Piskorski, Mikolaj Jan. 2014. *A Social Strategy: How We Profit from Social Media*. Princeton University Press.

Pocock, Stuart J. 2013. *Clinical trials: a practical approach*. John Wiley & Sons.

Reis, Harry T, Arthur Aron, Margaret S Clark, Eli J Finkel. 2013. Ellen berscheid, elaine hatfield, and the emergence of relationship science. *Perspectives on Psychological Science* **8**(5) 558–572.

Romanosky, Sasha, Rahul Telang, Alessandro Acquisti. 2011. Do data breach disclosure laws reduce identity theft? *Journal of Policy Analysis and Management* **30**(2) 256–286.

Roth, Alvin E, Marilda A Oliveira Sotomayor. 1992. *Two-sided matching: A study in game-theoretic modeling and analysis*. 18, Cambridge University Press.

Rudder, Christian. 2014. *Dataclysm: Who We Are (when we think no one's looking)*. Crown.

Slater, Dan. 2013. *Love in the Time of Algorithms: What Technology Does to Meeting and Mating*. Penguin.

Smith, Lones. 2006. The marriage model with search frictions. *Journal of Political Economy* **114**(6) 1124–1144.

Solove, Daniel J. 2004. *The digital person: Technology and privacy in the information age*. NYU Press.

Suler, John. 2004. The online disinhibition effect. *Cyberpsychology & behavior* **7**(3) 321–326.

Taylor, Lindsay Shaw, Andrew T Fiore, GA Mendelsohn, Coye Cheshire. 2011. out of my league: A real-world test of the matching hypothesis. *Personality and Social Psychology Bulletin* **37**(7) 942–954.

Weick, Karl E, Charles A Kiesler. 1979. *The social psychology of organizing*, vol. 2. Random House New York.