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

Ravi Bapna*, Jui Ramaprasad**, Galit Shmueli***, Akhmed Umyarov*[1]
{rbapna@umn.edu, jui.ramaprasad@mcgill.ca, Galit_Shmueli@isb.edu, aumyarov@umn.edu}

July 2013

Abstract

The growing popularity of online dating sites 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, in that they have the ability to check out a potential mate's profile and not leave any observable trail of doing that. While this may decrease search costs and allow users to search without inhibition, it also eliminates a "weak signal" for their potential mates. 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 view profiles of other users anonymously. Compared to the control group, the users treated with anonymity become disinhibited: they view more profiles, and are more likely to engage in viewing 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. The treated users lose the ability to leave a weak-signal in the form of a profile view and, therefore, achieve fewer matches than their non-anonymous counterparts. This effect is significantly stronger for women, reflecting and quantifying the impact of an age-old social norm that prevents them from making the first move.

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

Acknowledgements: 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 SCECR workshop, the 2013 Winter Conference on Business Intelligence, the 2013 International Symposium on IS (Goa, India) as well as the Friday workshop series at the 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.

[1] Author names in alphabetical order. All authors contributed equally to this research.
*Carlson School of Management, University of Minnesota
**Desautels Faculty of Management, McGill University
***Indian School of Business

1. Motivation and Background

According to the United States (US) Census, 46% of the single population in the US uses online dating[2] 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 welfare implications of the new features and capabilities that these new online matching markets bring to an age-old human activity. In this paper, we bridge this gap by studying the causal impact of anonymity, a key feature unique to the online environment, by means of 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 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 extreme scale of population of these websites as well as standardized nature of users' profiles in online world, these capabilities range from extensive search and algorithmic matching

[2] "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?#axzz1TbHiT1Xv

to big-data based recommendations (Gelles 2011), a science perfected for books and movies, 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 have high external validity, we rely on a randomized experiment 'in the wild' (Aral and Walker 2011) to isolate the causal impact of our focal factor. In doing so, we fill a gap in the extant research that has not addressed whether these IT-enabled features impact the search, viewing, message initiation, and matching outcomes of individuals.

In particular, we focus our attention on the impact of an anonymity feature on matching outcomes. This feature allows users to browse profiles of other users anonymously and thus, eliminate a weak signal for their potential mates—anonymous users have the ability to check out a potential mate's profile and not leave a clear and observable trail of doing that. Weak signaling is the ability to visit, or "check out," a potential mate's profile so that the potential mate knows the focal user visited her. It is akin to making a move, through viewing, without actually making a move, by sending a message. Yet, importantly, the counter-party becomes explicitly aware that a 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 players in question. Much less ambiguity exists in the online environment if the focal user views the target user's profile and leaves a visible trail in the target's "Recent Visitors" list.

Based on a novel large-scale randomized trial, similar in spirit to Aral and Walker (2011) and Bapna and Umyarov (2013), and in partnership with one of the largest online dating companies in North America, we causally demonstrate that weak signaling is a key mechanism linked to increased matches. This is especially so for women, helping them overcome social norms that discourage women from making the explicit first move in dating markets (Maccoby and Jacklin 1974).

Our treatment involves gifting one month of anonymous profile viewing to random 50,000 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 $14.95 (value changed for de-identification purposes) per month. In our study, we treat the randomly selected users only with anonymity and not with the other features for the purpose of observing the changes in behavior and outcomes that are induced specifically by anonymity.

Conventional wisdom suggests that anonymous profile viewing, by lowering search costs, should be associated with improved matching. Note that a possible disadvantage of non-anonymous browsing for a user is that in one simultaneous action she is both collecting information or "checking out" another user and leaving a weak signal to that user. This could potentially introduce associated risks such as accidentally making a first move towards an undesired communication because it is hard to learn about the target user without 'provoking' her into communication, or violating societal norms that suggest that there is a stigma associated with an overly active discovery process.

The above scenarios 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. Across genders, social norms 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). These inhibitions on preferences manifest themselves in the search stage of dating, limiting whom one looks for. The anonymity gift, then, may potentially lower this stigma, therefore lowering the search costs, 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 film 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 of the sources of frictions and restrictive social norms and encourage people to express their true preferences.

If the above scenarios dominate, an argument can be made for a positive effect of anonymity on the number of matches. Yet, the advantage of non-anonymous browsing is that it allows a focal user to advertise herself by leaving a 'weak signal' to another user without actually making any unambiguous explicit first move such as sending a personal message. Thus, treating individuals with the ability to anonymously view profiles, in effect, takes away their weak-signaling mechanism. This implies that anonymity might have a role in altering social inhibitions at the contact initiation stage. Individuals in the non-anonymous profile-viewing regime (control group), by virtue of leaving a trail through a visible profile visit, initiate contact by providing a weak signal of interest to the counter-party. In contrast, in the anonymous regime, individuals must send a message to be noticed. 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 is "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 that were 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 relying on leaving a weak signal and waiting for the counter-party to initiate the actual contact (Maccoby and Jacklin 1974). In summary, if weak signaling is an important tool, in particular for women, towards overcoming the effect of social barriers preventing them from making the first move, then an argument can be made for a negative effect of anonymity on the total number of matches. Further, we expect that this negative effect would be stronger for women.

Thus, these competing theories suggest different theoretical directions of the causal effect of anonymity on matches. Establishing the net effect of anonymity, therefore, is 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.

The advent of online dating platforms is increasing 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 any 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 toss these competing forces into a cauldron of a large-scale randomized experiment in the wild, 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 us with 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 mating 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 our 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 some robustness analysis and Section 8 concludes with some 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, while another explanation is that they may just be employed together, or be clustered in educational institutions, leading to romantic liaisons due to spending time together simply because of low search costs and irrespective of preferences.

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. Because search frictions are substantially lower in online dating markets—consider the infeasibility of getting detailed profile and attribute information from even a handful of potential mates at a bar—Hitsch et al (2010) are able to break down the observed sorting outcomes in dating into preferences over mate attributes. The negligible geographic search costs of the online environment rule out the alternative explanation that two individuals would date because they happen to go to the same school or live in the same neighborhood, rather than due to a preference for each other. Preferences can manifest themselves horizontally (men and women may prefer matching with a similar partner) or vertically, wherein each mate ranks all potential mates in an identical manner, and in a frictionless market, the ranks of matched men and women will be perfectly correlated (Hitsch et al 2010). The main finding is that people do have preferences towards those similar to them, along various attributes such as age, income, education, and ethnicity. They find, as expected, that heterosexual users of the online dating service prefer a partner whose age is similar to their own; that women generally avoid divorced men; that attractiveness is important to both men and women; that women place twice as much weight on income than men; and that while women have an overall strong preference for an educated partner, men generally shy away from educated women.

These gender asymmetries in mate selection are also the key findings of Fisman et al. (2006), who obtain mate preference data from a speed dating experiment. Similar to their research, we focus on dating, an activity that usually precedes marriage, and usually 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.

The use of a speed dating experiment with random assignment helps Fisman et al. (2006) overcome the challenge of backing out mate preferences from observational equilibrium outcomes data, where multiple preferences structures would be consistent with a given outcome. Subjects meet between nine and twenty-one potential mates for four minutes each and have the opportunity to accept or reject each partner. If both parties desire a future meeting, each receives the other's email address. Findings from this study indicate that women put greater weight on intelligence than men, while men place more value on physical appearance. Also, they find that women put more emphasis on the partner's race.

Recognizing the gender asymmetries established in the literature and similar to Fisman et al. (2006) and Hitsch et al. (2010), we will report our empirical findings separately for men and women.

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 factors: an underlying search process as well as a post-search contact initiation. That is, one party making a move and another party responding to that move. Given that there are significant social inhibitions in both the search stage and the contact initiation stage of the dating process, and that the contact initiation is gender asymmetric, we believe that it is 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's (2012) idea of a social failure.

The crux of Piskorski's (2012) idea is that while online dating reduces multiple sources of friction that are present in offline dating markets, they do not eliminate them. Piskorski (2012) 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 (in that it would have increased the welfare of both sides), 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 2012).

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 (2002), 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 higher matching outcomes. Yet, social exchange theory, which Piskorski (2012) draws upon, reminds us that while age-old social norms prevent women from making 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, and thereby increasing social frictions.

Thus, in departure from anything considered in 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 trial.

Our work links to the welfare implications of design of large scale matching markets. These markets offer key features and capabilities, such as anonymity. But how these features play out in a multi-faceted real world social process that makes up romantic markets requires careful scientific enquiry 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. In the next section we provide some institutional details about our research site.

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 regular online dating website and offers the following features to its users, which are typical of most other online dating websites:

- Users may set up their own well-structured online profiles where they describe themselves as well as reveal characteristics sought in a desired partner. Users may also place a set of their photos into their profiles.
- Users may view profiles of all other users without limitations.
- 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.

In addition to these features, monCherie.com constitutes a typical *freemium* community: most of the users sign up for free and with that can utilize all the key features of the monCherie.com website listed above. In addition to these free features, users can obtain a premium subscription if they pay $14.95 per month (value changed for de-identification purposes). The premium subscription consists of a fixed bundle of premium features that include anonymous browsing of profiles of others as well as a few other incremental features.

By default, free users of monCherie.com browse in the *non-anonymous* mode such that if the focal user A visits the profile of the target user B, user B knows through her "Recent Visitors" page that user A checked her out. In contrast, premium users browse in the *anonymous* mode such that if the focal user A visits the profile of the target user B, user B does not know that user A checked her out. However if user B were to visit user A’s profile, user A would know it. This feature is the proverbial “one-way mirror” of the online world, the impact of which is the subject of the research of this paper. It is important to highlight that a user’s 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.

4. Data, Empirical Regularities, and Outcomes

Based on the specifics of the agreement with monCherie.com, our experiment was conducted on 100,000 random new users of the website from one geographical area over the period of three months, which we refer to as month 1 (pre-treatment), month 2 (during treatment) and month 3 (post-treatment). For each of the 100,000 users, we know whether they were given the gift of anonymity (manipulation = 1) or whether they were in the control group (manipulation = 0), as well as a set of demographic variables such as gender (gender = 1 for men, 0 otherwise), age, sexual orientation (straight = 1 for straight users, 0 otherwise), whether their race is white (white = 1 for white users, 0 otherwise) and their attractiveness score. Users on monCherie.com can secretly rate each other on attractiveness on a scale of 1 (least attractive) to 5 (most attractive). We define a variable, *AttractScore*, as the average rating reflecting user’s attractiveness as per monCherie’s rating system. This variable can be missing for some individuals if they were not rated by other website users. In addition, we know whether the users are valid (valid = 0 if the 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 has visited at least one profile ten days prior to the manipulation.

In this study we limit our attention only to users who were valid and active prior to our manipulation. Table 1 outlines descriptive statistics of user demographics and their attractiveness score and statistically 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.

Gender	Variable	Mean	Median	Std Dev	Min	Max	t-value	p-value
F	Age	30.6260	27.6249	10.2905	18	75.38	5.14	<0.0001
M	Age	30.0155	27.4634	9.2806	18	78.80		
Combined	Age	30.2486	27.5483	9.6831	18	78.80		
F	Straight	0.8352	1	0.3710	0	1	-17.2	<0.0001
M	Straight	0.9039	1	0.2947	0	1		
Combined	Straight	0.8777	1	0.3277	0	1		
F	White	0.7815	1	0.4133	0	1	7.39	<0.0001
M	White	0.7427	1	0.4372	0	1		
Combined	White	0.7575	1	0.4286	0	1		
F	AttractScore	3.0573	3.1	0.7567	1	5	85.11	<0.0001
M	AttractScore	2.2126	2.1	0.7644	1	5		
Combined	AttractScore	2.5498	2.5	0.8665	1	5		

Table 1. Summary Statistics of User Characteristics

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 profiles that the focal user visited in month 1), *ViewSentCount* (number of profiles the focal user visited in month 2) and *ViewSentCountPost* (number of profiles the focal user visited in month 3), *ViewRcvdCountPre* (number of different users who visited the focal user in month 1), *ViewRcvdCount* (number of different users who visited the focal user in month 2) and *ViewRcvdCountPost* (number of different 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. Likewise, 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 for the statistical significance of differences between the two genders. As is evident from Table 2, women, on average, receive more than five times the viewing attention as compared to men (303 unique visitors versus 59.8) and receive more than ten times more messages (49.9 unique conversations versus 4.7). In addition, note that women are far less likely to initiate explicit contact via sending a message. Although they are 3.6 times less likely than men to initiate explicit contact by sending a message, they are only 1.6 times less likely to leave a weak signal by viewing a profile. 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 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 the subsequent statistics and results separately for the two genders so that we compare not the overall averages but rather the averages by gender: treated women are compared to control women and treated men are compared to control men.

Gender	Variable	Mean	Std Dev	Min	Max	t-value	p-value
F	ViewRcvdCountPre	302.993	3.17987	0	3443	98.19	<0.0001
M	ViewRcvdCountPre	59.821	0.59841	0	1291		
F	ViewSentCountPre	117.291	1.39285	1	3184	-40.6	<0.0001
M	ViewSentCountPre	191.282	1.89665	1	3588		
F	MsgRcvdCountPre	49.936	66.510	1	927	88.81	<0.0001
M	MsgRcvdCountPre	4.747	6.754	1	171		
F	MsgSentCountPre	8.293	18.376	0	662	-30.3	<0.0001
M	MsgSentCountPre	29.831	72.093	0	2167		

Table 2. Descriptive Statistics of User Activity

While defining our outcome of interest, 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 measure could be two users attempting to move their online interactions offline[3]. However, even if we could observe which couples went for an offline date, this measure is far from perfect, as many offline dates turn out to be unsuccessful and do not result in a long-term relationship. Another measure could be whether a couple got married. Again, even if we had the data on actual weddings for our users, such a measure would still hardly constitute a perfect success measure, given that current divorce rates are 40% to 50% according to the American Psychological Association.

Recognizing that any relationship is an ongoing process, and 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

[3] Such approach was taken by Hitsch et al (2010) who used an indicator variable on 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.

‘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 A and user B as a match if user A messaged user B, user B responded, and then user A messaged user B again (with user A possibly responding to that and so on), therefore forming a sequence of *at least* three messages between user A and user B. Communication theorists call this measure a *double interact* (Weick 1979), and it is considered a sense-making process that people use when they organize in a variety of contexts.

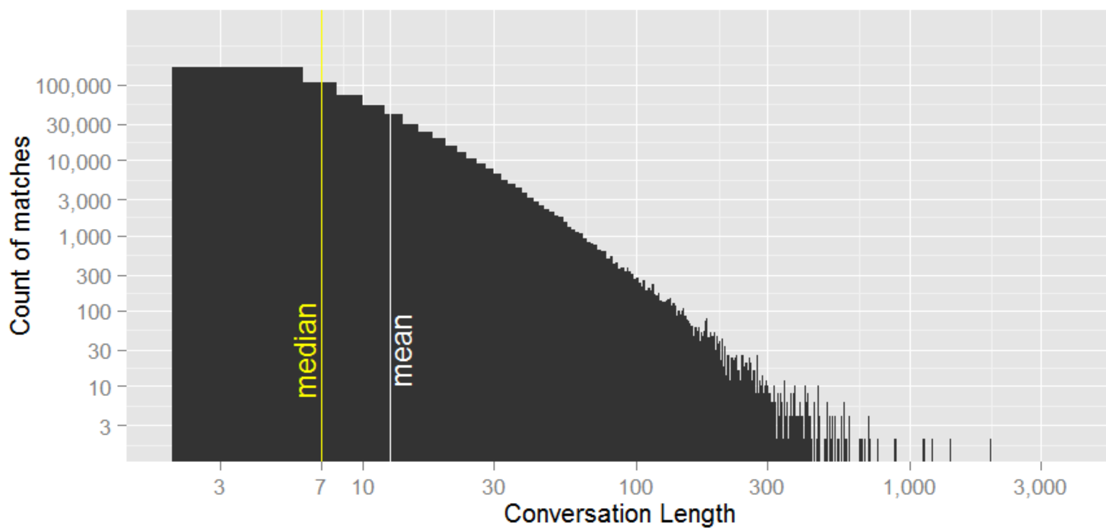

Figure 1. The Typical Number of Messages Exchanged in a Match

As is evident from this definition and demonstrated on Figure 1, a conversation that constitutes a match is typically much longer than three messages. More specifically, in our data, the average number of messages exchanged by matched users 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, monCherie.com is not. Our conversations with the senior executives revealed that they strongly believe that this 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 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 at least five or seven messages. We find that our results do not change even after altering the definition. We present the robustness of our results to different definitions of a match in more detail in Section 7.

Given the above definition of a match, our outcome variable is the number of 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 schemes, we examine some 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, and that 75% of these matches for women are received matches (that is, matches when the man sends the first message and the woman simply responds, thus supporting the conversation and leading to a match that she did not initiate). More than 75% of matches for men are "matches sent," that is, matches initiated by the man himself, while less than 25% of matches for women are "matches sent."

We also explore whether the users browse not just more, but also differently, in a disinhibited manner, indicating the overcoming of social norms under anonymity. We define the following variables for this purpose:

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

We define these constructs 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 constructs:

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

In addition to these measures, we observe more nuanced behavioral patterns that shed additional insight on the effect of anonymity. These behavioral patterns help us understand whether the selectivity levels of users change under the anonymity condition. We operationalize these selectivity constructs as we discuss the results in Section 6.

5. Experimental Design

In order 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 randomly selecting a 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. In other words, our field experiment removes the ability to send a “weak signal” for the treatment group while keeping it for the control group, allowing us to compare the resulting search intensity, search diversity, messaging behavior, and number of matches between the two groups.

As demonstrated in Table 3, the treatment (manip=1) and control (manip=0) groups have statistically indistinguishable properties before manipulation. They are also indistinguishable in terms of their viewing and messaging behavior in the pre-treatment month 1.[4]

Gender	Manip	Variable	Mean	Std Err	Min	Max	t-value	p-value
F	0	Age	30.6528	0.14227	18	75.21	0.27	0.7896
F	1	Age	30.5998	0.13874	18	75.38		
M	0	Age	30.0518	0.09999	18	78.21	0.52	0.6035
M	1	Age	29.9786	0.09916	18	78.80		
F	0	Straight	0.8388	0.00505	0	1	0.99	0.3199
F	1	Straight	0.8316	0.00508	0	1		
M	0	Straight	0.9025	0.00317	0	1	-0.65	0.5184
M	1	Straight	0.9054	0.00315	0	1		
F	0	AttractScore	3.0534	0.01088	1	5	-0.50	0.6183
F	1	AttractScore	3.0611	0.01075	1	5		
M	0	AttractScore	2.2176	0.00895	1	5	0.80	0.4210
M	1	AttractScore	2.2075	0.00886	1	5		
F	0	White	0.7835	0.00566	0	1	0.50	0.6160
F	1	White	0.7795	0.00563	0	1		
M	0	White	0.7417	0.00468	0	1	-0.28	0.7801
M	1	White	0.7436	0.00470	0	1		

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

[4] Tables are available upon request.

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

6. Experimental Results

6.1. Average treatment effects

We start our analysis by exploring changes in profile browsing behavior that were induced by our treatment, using t-tests to compare the treatment and control groups. As demonstrated in Table 4, treated users of both genders viewed significantly more profiles as compared to their non-treated counterparts.

Gender	Manip	Variable	Mean	Std Err	Min	Max	t-value	p-value
F	0	ViewSentCount	44.073	1.07086	0	1414	-3.41	0.0006
F	1	ViewSentCount	49.698	1.24910	0	2475		
M	0	ViewSentCount	74.278	1.64199	0	3216	-2.91	0.0036
M	1	ViewSentCount	81.317	1.77359	0	3634		

Table 4. The Effect of Treatment on Outbound Views

Further, as demonstrated by Table 5, we find that heterosexual individuals of both genders significantly increase their likelihood of viewing profiles of people of the same gender when they are anonymous. In particular, when anonymous, heterosexual men have a 12% higher likelihood of viewing men and heterosexual women have a 19% higher likelihood of viewing other women. We also find that 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 significant for men. This disinhibition effect, however, is only seen in viewing behavior and does not translate to messaging (Table 6), as predicted, since only profile views are hidden by anonymity, not messages.

Gender	Manip	Variable	Mean	Std Err	Min	Max	t-value	p-value
F	0	ViewSameSex	0.0708	0.00385	0	1	-2.40	0.0164
F	1	ViewSameSex	0.0844	0.00414	0	1		
M	0	ViewSameSex	0.0768	0.00300	0	1	-2.05	0.0406
M	1	ViewSameSex	0.0857	0.00317	0	1		
F	0	ViewOtherRace	0.4880	0.00686	0	1	-2.45	0.0142
F	1	ViewOtherRace	0.5117	0.00678	0	1		
M	0	ViewOtherRace	0.6523	0.00509	0	1	-1.40	0.1618
M	1	ViewOtherRace	0.6624	0.00509	0	1		

Table 5. The Effect of Treatment on Same-sex and Inter-racial Viewing

Gender	Manip	Variable	Mean	Std Err	Min	Max	t-value	p-value
F	0	MsgSameSex	0.0092	0.00143	0	1	-0.26	0.7967
F	1	MsgSameSex	0.0097	0.00146	0	1		
M	0	MsgSameSex	0.0106	0.00115	0	1	-0.30	0.7615
M	1	MsgSameSex	0.0111	0.00119	0	1		
F	0	MsgOtherRace	0.1280	0.00459	0	1	-1.35	0.1780
F	1	MsgOtherRace	0.1369	0.00467	0	1		
M	0	MsgOtherRace	0.2572	0.00467	0	1	-0.92	0.3569
M	1	MsgOtherRace	0.2634	0.00474	0	1		

Table 6. The Effect of Treatment 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 number of matches goes in the opposite direction. Table 7 shows that despite apparent disinhibition in browsing, the total number of matches actually decreases both for men and for women. Further, this effect is significantly stronger for women as compared to men. For women, the average match count reduces by a significant 14%[5]. For men, the effect is marginally significant ($p<0.1$), with matches decreasing by 7%, almost half of the effect on women.

[5] Given that our manipulation was exogenously randomized, we do not need to control for any user characteristics in order to establish the average effect of treatment. A regular t-test of observed outcomes is enough to establish statistical significance of our results.

Gender	Manip	Variable	Mean	Std Err	Min	Max	t-value	p-value
F	0	TotalMatchCount	4.30850	0.12207	0	267	3.88	0.0001
F	1	TotalMatchCount	3.71210	0.09398	0	94		
M	0	TotalMatchCount	2.58980	0.07584	0	149	1.72	0.0861
M	1	TotalMatchCount	2.41041	0.07184	0	190		

Table 7. The Effect of Treatment on Total Number of Matches

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 control group and no signaling for treatment group) and messaging (strong signaling) 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 outgoing messages remains statistically the same (Table 9). In order 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. Therefore, our results directly 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. This finding emphasizes the importance of weak-signaling: despite visiting more profiles, the treated users were visited by a smaller number of potential mates.

Gender	Manip	Variable	Mean	Std Err	Min	Max	t-value	p-value
F	0	ViewRcvdCount	129.236	2.17057	0	1842	2.40	0.0165
F	1	ViewRcvdCount	122.156	2.00323	0	1423		
M	0	ViewRcvdCount	26.708	0.49685	0	1710	3.42	0.0006
M	1	ViewRcvdCount	24.485	0.41736	0	785		
F	0	MsgRcvdCount	19.787	0.41282	0	469	2.28	0.0227
F	1	MsgRcvdCount	18.519	0.37398	0	434		
M	0	MsgRcvdCount	1.741	0.04296	0	132	5.00	<0.0001
M	1	MsgRcvdCount	1.475	0.03085	0	57		

Table 8. The Effect of Treatment on Views and Messages Received

Gender	Manip	Variable	Mean	Std Err	Min	Max	t-value	p-value
F	0	MsgSentCount	2.655	0.09832	0	94	-1.16	0.2460
F	1	MsgSentCount	2.833	0.11779	0	262		
M	0	MsgSentCount	11.958	0.55317	0	2116	-0.31	0.7565
M	1	MsgSentCount	12.187	0.48880	0	1181		

Table 9. The Effect of Treatment on Messages Sent

We apply a similar analysis to *TotalMatchCount*, splitting it into two variables that emphasize 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 very differently by our manipulation. *MatchSentCount* remains statistically unchanged for both genders (just like *MsgSentCount*), while *MatchRcvdCount* is reduced significantly with an approximate equal drop of 20%-25% for each gender.

This finding clearly explains the observed gender asymmetry in the effect of anonymity on *TotalMatchCount.* As demonstrated, both genders lose approximately 20-25% of their "matches received" because of anonymity. Yet, unlike women, most of the matches for men are actually "matches sent" (that are unaffected by anonymity), not "matches received." Therefore, the similar scale of 20-25% reduction in the "matches received" induced by anonymity does not have as significant impact on the total number of matches for men as it does for women. This finding demonstrates that removing the weak-signaling capability is especially damaging for women who tend to rely more on the incoming messages and tend not to make a first move. In other words, having weak signaling ability is especially helpful for women.

Gender	Manip	Variable	Mean	Std Err	Min	Max	t-value	p-value
F	0	MatchRcvdCount	3.28588	0.10095	0	265	5.15	<0.0001
F	1	MatchRcvdCount	2.65500	0.07040	0	87		
M	0	MatchRcvdCount	0.59838	0.01889	0	88	6.19	<0.0001
M	1	MatchRcvdCount	0.46121	0.01147	0	19		
F	0	MatchSentCount	1.02263	0.04019	0	48	-0.60	0.5476
F	1	MatchSentCount	1.05710	0.04084	0	58		
M	0	MatchSentCount	1.99143	0.06806	0	148	0.44	0.6607
M	1	MatchSentCount	1.94921	0.06798	0	180		

Table 10. The Effect of Treatment on Matches Received and Matches Sent

6.2. Individual level marginal effects

The results from the Section 6.1 demonstrate the significance and direction of the average effect of the treatment. However, being an average effect, these results do not provide us with any 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 then use this model to explore the marginal effects of our manipulation for each user in our sample.

Given that our dependent variable *TotalMatchCount* is a count variable, we fit a Zero Inflated Poisson (ZIP) model 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, gender, attractiveness, race, and orientation. In addition, we also control for the prior success of the user, as measured by the number of matches the user achieved in the month prior to the treatment (*TotalMatchCountPre*).

We use the standard Zero-Inflated Poisson model setting, modeling a mixture of the following two sub-models: a Poisson model, which accounts for the expected number of matches given a user's demographics, attractiveness, and prior success in matching; and a zero-model, which accounts for zero matches due to a user's decision not to engage in mate seeking behavior during the manipulation month.

The results of the Poisson model are displayed in Table 11a, while the results of the corresponding zero model are displayed in Table 11b. In order to make the coefficients comparable between the models with and without the interaction terms, we normalize all the independent variables (except for *Manipulation*) so that their means are all equal to zero.

Parameter	Estimate	Std Err	p-value	Estimate	Std Err	p-value
(Intercept)	1.02962	0.00720	<.0001	1.011149	0.00885	<.0001
Manipulation	**-0.07179**	**0.00705**	**<.0001**	**-0.03445**	**0.01273**	**0.0068**
Gender (centered)	0.097619	0.00826	<.0001	0.106031	0.01122	<.0001
Age (centered)	-0.00092	0.00039	0.0194	-0.00067	0.00053	0.2035
White (centered)	-0.01014	0.01041	0.3298	-0.08049	0.01427	<.0001
LogMatchCountPre (centered)	0.785506	0.00377	<.0001	0.791542	0.00522	<.0001
AttractScore (centered)	0.024766	0.00520	<.0001	0.07453	0.00716	<.0001
Straight (centered)	0.090662	0.01093	<.0001	0.066446	0.01498	<.0001
Manipulation * Gender				-0.02429	0.01659	0.1430
Manipulation * Age				-0.00079	0.00079	0.3147
Manipulation * White				0.144605	0.02086	<.0001
Manipulation * LogMatchCountPre				-0.01412	0.00754	0.0612
Manipulation * AttractScore				-0.10623	0.01044	<.0001
Manipulation * Straight				0.056388	0.02192	0.0101

Table 11a. ZIP model Coefficients for the Experiment (with and without interactions)

Parameter	Estimate	Std Err	p-value	Estimate	Std Err	p-value
(Intercept)	-0.55744	0.02408	<.0001	-0.5705	0.02480	<.0001
Manipulation	**0.068296**	**0.03343**	**0.0411**	**0.090042**	**0.03482**	**0.0097**
Gender (centered)	0.046807	0.03929	0.2335	0.122526	0.05634	0.0297
Age (centered)	-0.00589	0.00185	0.0015	-0.00548	0.00263	0.0373
White (centered)	-0.39447	0.04219	<.0001	-0.47607	0.05962	<.0001
LogMatchCountPre (centered)	-0.89409	0.01967	<.0001	-0.91525	0.02808	<.0001
AttractScore (centered)	-0.21875	0.02334	<.0001	-0.22244	0.03303	<.0001
Straight (centered)	0.107066	0.05246	0.0413	0.044474	0.07462	0.5512
Manipulation * Gender				-0.15243	0.07871	0.0528
Manipulation * Age				-0.00114	0.00371	0.7595
Manipulation * White				0.166572	0.08452	0.0487
Manipulation * LogMatchCountPre				0.040531	0.03919	0.3010
Manipulation * AttractScore				0.001748	0.04674	0.9702
Manipulation * Straight				0.130307	0.10514	0.2152

Table 11b. Zero model Coefficients for the Experiment (with and without interactions)

As is evident from Tables 11a and 11b, the randomized experiment demonstrates that the average effect of treatment on matches is negative. More specifically, the negative coefficient in the Poisson part of the model (Table 11a) suggests that the treatment has a significant negative effect on the expected match count, while the positive coefficient in the zero part of the model (Table 11b) suggests that the treatment at the same time significantly increases the probability of getting zero matches.

The estimated ZIP model is a complex and non-linear model with interactions. Therefore, the actual marginal effects predicted by this model are hard to infer from the coefficients alone, because they depend on the characteristics of a particular user. In Figure 2, we present a scatter plot of the estimated marginal effects for every user in our sample. More specifically, each point in this plot corresponds to one user in our sample. Assuming user *u* has demographic profile *U* (including age, gender, attractiveness, etc.), the marginal effect and marginal effect ratio are defined as:

$$\text{Marginal Effect}(U) = \mathrm{E}[TotalMatchCnt \mid Manip = 1, U] - \mathrm{E}[TotalMatchCnt \mid Manip = 0, U]$$

$$\text{Marginal Effect Ratio}(U) = \frac{\mathrm{E}[TotalMatchCnt \mid Manip = 1, U] - \mathrm{E}[TotalMatchCnt \mid Manip = 0, U]}{\mathrm{E}[TotalMatchCnt \mid Manip = 0, U]} * 100\%$$

These formulas represent the expected change in the total number of matches for a particular user *u* as predicted by our model, both in absolute terms (the absolute number of matches a user such as *u* is expected to gain) and in relative terms (the percent of the matches a user such as *u* is expected to gain). Also, since the ZIP model contains interactions of manipulation with all user-level variables, it provides us insights into how users of different ages, attractiveness, gender, race and orientation respond differently to our manipulation.

Figure 2. Estimated Marginal Effects for Both Genders in Absolute (left) and Relative (right) Terms

Figure 2 demonstrates the distribution of predicted individual-level marginal effects of treatment with each point on the plot representing the marginal effect of one particular user from our sample. As is evident, the average marginal effects, as represented by blue and red rhombuses, are negative in the absolute number of matches for both genders, indicating an average loss of approximately 0.6 matches for women and 0.18 matches for men, consistent with levels reported in Table 7. The result reiterates the clear gender asymmetry that was initially revealed in Table 7 for average treatment effect, but on an individual-level. The ZIP model suggests that, unlike men, the vast majority of women are expected to lose matches from anonymity, both in absolute terms and in relative terms. Interestingly, the ZIP model also suggests that although the average effect of anonymity on men is negative and marginally significant, a large number of men is expected not to lose matches because of anonymity, an interesting new finding that was not captured by the t-test in Table 7.

6.3. Selectivity

As we discovered in Section 6.1, our treatment causes a decrease in incoming matches, while causing no changes 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 can be the case that because our treatment causes a decrease in the number of incoming messages (see Table 7) a person 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 manipulation causes changes in reply patterns. That is, users may be more likely to anonymously 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, so here we focus on examining the second mechanism by defining a selectivity construct, called *GetCheckRatio*, which is the ratio of the number of message senders the focal user checks out after receiving a message, divided by the total number of message senders that attempted to contacted the focal user. We can extend this logic further by defining another selectivity construct called *GetCheckReplyRatio*, which is the ratio of the number of message senders the focal user decides to reply to, divided by the number of message senders they received a message from and subsequently checked out.

It is evident in Table 12, that once manipulated and under anonymity, both genders indeed prefer to visit more of the profiles of the candidates who messaged them. However, as demonstrated by the *GetCheckReplyRatio* construct, treated women are 10% less likely to reply after visiting the profiles of their candidates as compared to the control group of women. The difference for men is much smaller (3%) and is statistically insignificant. The stronger selectivity effect for women could also contribute to their receiving less incoming matches in total. Therefore, the effect of anonymous browsing is not limited to reducing the incoming communication for the focal user, but also incorporates the changes in the selection behavior of the focal user herself. These findings are suggestive of both disinhibition and selectivity driven by anonymity.

Gender	Manip	Variable	Mean	Std Err	Min	Max	t-value	p-value
F	0	GetCheckRatio	0.416	0.00402	0	1	-5.65	0.0000
F	1	GetCheckRatio	0.449	0.00416	0	1		
M	0	GetCheckRatio	0.519	0.00603	0	1	-6.50	0.0000
M	1	GetCheckRatio	0.575	0.00618	0	1		
F	0	GetCheckReplyRatio	0.229	0.00411	0	1	3.36	0.0008
F	1	GetCheckReplyRatio	0.210	0.00397	0	1		
M	0	GetCheckReplyRatio	0.367	0.00735	0	1	1.08	0.2795
M	1	GetCheckReplyRatio	0.356	0.00739	0	1		

Table 12. The Effect of the Treatment on Selectivity

7. Robustness

We carried out a series of robustness analyses to rule out any alternative explanations for our results. First, we examine if any other factor could lead to the treatment effect that we find. In addition, we specifically test whether the effect that we see is induced simply by the happiness associated with receiving a gift rather than by anonymity itself. Next, we test the robustness of our results to different definitions of a match by considering sequences of five and seven messages as new definitions of a match. Finally, we rule out the possibility that anonymity could have caused users to find their perfect match during the treatment month and quit using the website, thus decreasing the observed number of matches.

7.1 Ruling out alternate explanations of the treatment effect

If the difference between treatment and control groups in response to our manipulation was due to any reason other than our treatment, we would observe the effect to be persistent after our treatment expired. Recall that there was no difference in the characteristics, behaviors, and outcomes between our treatment and control groups in the month prior to manipulation. To rule out alternate explanations of the treatment, we similarly contrast treatment and control groups in the month after the manipulation expired. We find that the observed effect of anonymity disappears in month 3 (post-treatment month). As demonstrated by Tables 13-15, the treatment group reverts to the behavior that is statistically indistinguishable from the control group in the post-treatment month.

Gender	Manip	Variable	Mean	Std Err	Min	Max	t-value	p-value
F	0	ViewSentCountPost	18.3702	0.67783	0	1228	-0.87	0.3828
F	1	ViewSentCountPost	19.2024	0.67056	0	828		
M	0	ViewSentCountPost	50.7441	1.28378	0	2640	-0.84	0.4009
M	1	ViewSentCountPost	52.3459	1.41125	0	3542		
F	0	ViewRcvdCountPost	87.6012	1.63607	0	1896	1.01	0.3131
F	1	ViewRcvdCountPost	85.3288	1.54968	0	1929		
M	0	ViewRcvdCountPost	13.3855	0.24834	0	408	0.09	0.9304
M	1	ViewRcvdCountPost	13.3559	0.23002	0	272		

Table 13. The Effect on Profile Visits Post-Treatment

Gender	Manip	Variable	Mean	Std Err	Min	Max	t-value	p-value
F	0	MsgSentCountPost	1.8825	0.08102	0	104	-0.01	0.9951
F	1	MsgSentCountPost	1.8832	0.08064	0	131		
M	0	MsgSentCountPost	8.9467	0.43456	0	1600	0.11	0.9135
M	1	MsgSentCountPost	8.8815	0.41414	0	1465		
F	0	MsgRcvdCountPost	13.9244	0.32814	0	461	0.79	0.4298
F	1	MsgRcvdCountPost	13.5712	0.30445	0	378		
M	0	MsgRcvdCountPost	1.3237	0.03398	0	119	0.70	0.4818
M	1	MsgRcvdCountPost	1.2931	0.02706	0	39		

Table 14. The Effect on Messaging Post-Treatment

Gender	Manip	Variable	Mean	Std Err	Min	Max	t-value	p-value
F	0	MatchCountPost	2.60041	0.090997	0	227	0.77	0.4428
F	1	MatchCountPost	2.50893	0.077298	0	143		
M	0	MatchCountPost	1.79833	0.057232	0	115	0.25	0.8055
M	1	MatchCountPost	1.77815	0.058660	0	167		
F	0	MatchRcvdCountPost	1.94211	0.076456	0	223	0.98	0.3254
F	1	MatchRcvdCountPost	1.84638	0.060639	0	141		
M	0	MatchRcvdCountPost	0.40848	0.013033	0	32	0.81	0.4165
M	1	MatchRcvdCountPost	0.39464	0.010933	0	20		
F	0	MatchSentCountPost	0.65831	0.029330	0	43	-0.10	0.9189
F	1	MatchSentCountPost	0.66255	0.029629	0	60		
M	0	MatchSentCountPost	1.38985	0.052490	0	115	0.08	0.9333
M	1	MatchSentCountPost	1.38351	0.054678	0	167		

Table 15. Ruling out Alternative Explanations

7.2 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 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. Table 16 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	Std Err	Min	Max	t-value	p-value
F	0	MatchRcvdCountLastW	0.61079	0.025277	0	71	2.74	0.0062
F	1	MatchRcvdCountLastW	0.52441	0.019087	0	25		
M	0	MatchRcvdCountLastW	0.11524	0.004780	0	15	3.81	0.0001
M	1	MatchRcvdCountLastW	0.09173	0.003892	0	5		
F	0	MatchRcvdCountPostFirstW	0.56496	0.023648	0	60	0.77	0.4431
F	1	MatchRcvdCountPostFirstW	0.54006	0.022271	0	54		
M	0	MatchRcvdCountPostFirstW	0.10987	0.004895	0	13	-0.70	0.4838
M	1	MatchRcvdCountPostFirstW	0.11458	0.004611	0	6		

Table 16. Gift Effect Check

7.3 Examining different definitions of a match

As we mentioned earlier, we also test the robustness of our results to different definitions of a match by replicating our analysis with more stringent definitions of a *match*. More specifically, we define that user A and user B achieved a 5-match if user A and B exchanged at least five messages. Similarly, we define a 7-match between user A and user B if user A and B exchanged at least seven messages. We find in Tables 17a and 17b that the effect of anonymity is strongly observed and constitutes the same decrease of approximately 20%-25% in both *Match5RcvdCount* and *Match7RcvdCount* despite having a stricter threshold for declaring a communication as a match. Similarly, we also observe that *Match5SentCount* and *Match7SentCount* remain unchanged by the manipulation, similar to the results reported with our original definition of a match. We decided to limit the robustness check with the seven message threshold following insights learned from Hitsch et al. (2010) who had access to the actual content of the messages in an online dating website. Hitsch et al. (2010) discovered that it takes a median number of six messages for users to reveal their phone number, email address, or for a key phrase like “let’s meet” to appear.

Gender	Manip	Variable	Mean	Std Err	Min	Max	t-value	p-value
F	0	TotalMatch5Count	2.97643	0.091038	0	194	3.87	0.0001
F	1	TotalMatch5Count	2.53601	0.069083	0	79		
M	0	TotalMatch5Count	1.77249	0.053308	0	127	1.53	0.1271
M	1	TotalMatch5Count	1.65940	0.051483	0	162		
F	0	Match5RcvdCount	2.22968	0.074106	0	192	5.23	<0.0001
F	1	Match5RcvdCount	1.76183	0.050746	0	75		
M	0	Match5RcvdCount	0.41706	0.012124	0	35	6.07	<0.0001
M	1	Match5RcvdCount	0.32541	0.008932	0	18		
F	0	Match5SentCount	0.74675	0.031275	0	40	-0.62	0.5366
F	1	Match5SentCount	0.77418	0.031480	0	42		
M	0	Match5SentCount	1.35544	0.048036	0	126	0.31	0.7531
M	1	Match5SentCount	1.33399	0.048388	0	155		

Table 17a. Effect of Manipulation on 5-matches

Gender	Manip	Variable	Mean	Std Err	Min	Max	t-value	p-value
F	0	TotalMatch7Count	2.22120	0.071754	0	142	3.61	0.0003
F	1	TotalMatch7Count	1.89556	0.055007	0	67		
M	0	TotalMatch7Count	1.32617	0.041081	0	105	1.42	0.1548
M	1	TotalMatch7Count	1.24527	0.039269	0	123		
F	0	Match7RcvdCount	1.64473	0.057156	0	140	4.98	<0.0001
F	1	Match7RcvdCount	1.29748	0.040334	0	63		
M	0	Match7RcvdCount	0.31497	0.009542	0	21	5.74	<0.0001
M	1	Match7RcvdCount	0.24609	0.007237	0	13		
F	0	Match7SentCount	0.57647	0.025724	0	38	-0.60	0.5492
F	1	Match7SentCount	0.59808	0.025316	0	36		
M	0	Match7SentCount	1.01120	0.037049	0	104	0.23	0.8180
M	1	Match7SentCount	0.99919	0.036782	0	118		

Table 17b. Effect of Manipulation on 7-matches

7.4 Ruling out decrease of matches due to finding a "perfect" match and exiting the site

It could be the case that anonymity could have caused treated users to find a single highly-desirable match more often and thus exit the website as compared to the control group, and that this could then lead us to observe our result, a lowering of the total match count of the treatment group. To rule out this explanation of our result, we note that the behavior of treated users is indistinguishable from the control group in the post-treatment month as reported by Tables 13-15. In addition to that, we explore whether treated users are still active on the dating website in the last weeks of post-treatment month. Table 18 compares the levels of activity of users in the treatment and control groups in the last weeks of the post-treatment month and finds that treated users are still on the website in the same quantities (approximately 71% of users are still active) and exhibit the same behaviors as the control group.

Gender	Manip	Variable	Mean	Std Err	Min	Max	t-value	p-value
F	0	StillActive	0.7105	0.00623	0	1	-0.52	0.6059
F	1	StillActive	0.7150	0.00613	0	1		
M	0	StillActive	0.7172	0.00482	0	1	-0.27	0.7868
M	1	StillActive	0.7190	0.00484	0	1		

Table 18. Post-Treatment Activity Levels

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 trials. 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 has to like the book but the book does not have to like the consumer, in dating markets the individuals have to like each other) to anonymity-linked features such as weak-signaling, 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 in advance.

Our theoretical understanding of dating and marriage starts with Becker's (1973) exposition of assortative sorting as the equilibrium outcome assuming a frictionless market. Subsequent research (Burdett and Coles 1997, Mortensen and Pissarides 1999, and Smith 2002) 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 2012) 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 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 significant drop in matching outcomes, particularly for women. We demonstrate, by breaking down, measuring, and analyzing the mate-seeking process in detail, that this occurs because of a dominating social friction force of being unable to leave a weak-signal. Women, for whom historically-established social norms make them more reluctant to make the first move, say by messaging a potential mate, are deprived by our treatment of even leaving a profile visit trail, which turns out to be the key source of incoming messages and subsequent matches for them.

Under conditions of weak-signaling, when users browse their potential mates' profiles in non-anonymous mode, they leave a clear, definitive and observable trail to their potential mates without actually messaging them and such a trail plays an important role in their number of matches. Based on our large-scale randomized field experiment, we demonstrate that taking away this weak signaling ability causes a significant decline in the number of matches. Conversely, the presence of a weak-signaling ability improves matching outcomes.

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 crucial to better understanding our main results.

When we break down total matches into incoming and outgoing matches, based on who initiated the process, we find that our manipulation causes a decrease in the number of incoming matches while causing no change in the outgoing match count. We find that because our manipulation decreases the number of incoming messages, a person 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, users explore and anonymously visit more profiles of their prospective mates from whom they have received a message before making a decision to 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 message them, reply to significantly fewer (10% less) of those visited profiles, as compared to the control group of women. The difference for men is much smaller (3% less) and is statistically insignificant. 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 quality of matches and long-term outcomes as they relate to marriage, happiness, and divorce.

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. To some extent, our work is limited by being carried out in the online dating context, especially if there are systematic differences in the underlying populations of the online and offline dating worlds, though recent research (Hitsch et al 2010) finds proximity between online and offline dating markets. Current trends suggest that this concern, while valid, is not threatening as one-third to one-half of the single population of the United States currently 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. Beyond the context of online dating, anonymity is a key feature on a variety of other types of online platforms where people connect with one another including Facebook, where social interactions could be altered with a switch in anonymity settings, and LinkedIn, where anonymity could 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 there are two sets of individual preferences that have to be taken into account in order to produce a successful match. Matching two humans applies not only to dating and marriage, but also to new models of distributed work and crowdsourcing (Burtch et al. 2012). 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 that is interested in 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 2012). These newer platforms reduce many of the frictions that are 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 this. 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, Telang and Acquisti 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 to not just outcomes but also 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.

References

1. Aral, S., D. Walker, (2011), "Creating social contagion through viral product design: A randomized trial of peer influence in networks," *Management Science*, 57(9), 1623.
2. Bapna, R., Umyarov, A., (2013), "Do Your Online Friends Make You Pay? A Randomized Field Experiment in an Online Music Social Network," National Bureau of Economic Research Summer Institute on the Economics of IT and Digitization, 2012, Cambridge, MA, available at http://users.nber.org/~confer/2012/SI2012/PRIT/Bapna_Umyarov.pdf
3. Becker, G. S., (1973), "A Theory of Marriage: Part I." *Journal of Political Economy*, 81(4): 813–46.
4. Burdett, K., and Coles, M., (1997), "Marriage and Class," *Quarterly Journal of Economics*, CXII, 141–168.
5. Burtch, G., A. Ghose and S. Wattal (2012), "An Empirical Examination of the Antecedents and Consequences of Contribution Patterns in Crowdfunded Markets." *Information Systems Research*, Forthcoming.
6. Fisman, R., Iyengar, S., S., Kamenica,E., and Simonson, I. (2006), "Gender Differences in Mate Selection: Evidence from a Speed Dating Experiment." *Quarterly Journal of Economics*, 121(2): 673–97.
7. Frost, J., Chance. Z., Norton, M., Ariely, D. (2008) "People are Experience Goods: Improving Online Dating with Virtual Dates." *Journal of Interactive Marketing*, 22(1): 51-61.
8. Gale, D. and Shapley, L. (1962). College admissions and the stability of marriage. *The American Mathematical Monthly*, Vol. 69, No. 1, pp. 9-15
9. Gelles, D. (2011, July 29) Inside Match.com. *Financial Times.* Accessed here: http://www.ft.com/intl/cms/s/2/f31cae04-b8ca-11e0-8206-00144feabdc0.html?#axzz1TbHiT1Xv
10. Goldfarb, A. and Tucker, C. (2011). Privacy Regulation and Online Advertising. *Management Science*, Vol. 57, No. 1, pp. 57-71
11. Hall, J. A., Carter, S., Cody, M. J., Albright, J. M., (2010), " Individual Differences in the Communication of Romantic Interest: Development of the Flirting Styles Inventory," *Communication Quarterly*, 58:4, 365-393.
12. Holmes, Ronald M., Richard Tewksbury, and Stephen T. Holmes. "Hidden JPGs: A functional alternative to voyeurism." *The Journal of Popular Culture* 32.3 (1998): 17-29.
13. Harris, Tina M., Pamela J. Kalbfleisch, Interracial Dating: The Implications of Race for Initiating a Romantic Relationship. *Howard Journal of Communications*. 11, 49-64 (2000).
14. Henningsen, D. (2004). "Flirting with Meaning: An Examination of Miscommunication in Flirting Interactions," *Sex Roles*, 50 (7/8), 481-489.

15. Hitsch, G. J., Hortaçsu, A., and Ariely, D., (2010), "Matching and Sorting in Online Dating." *American Economic Review*, 100(1): 130-63.
16. Joinson, A. (1998), "Disinhibition and the Internet," in Psychology and the Internet: Intrapersonal, Interpersonal, and Transpersonal Implications, pp. 75 - 92.
17. Kalmijn, M., (1998). "Intermarriage and Homogamy: Causes, Patterns, Trends." *Annual Review of Sociology*, 24: 395–421.
18. Kling, R., Lee, Y., Teich, A. & Frankel, M. (1999), "Assessing anonymous communication on the Internet: Policy deliberations," *The Information Society*, 15(2), pp79 - 90.
19. Maccoby, C.N. Jacklin, **The psychology of sex differences** (Stanford University Press, Stanford, CA, 1974
20. McDevitt, R. C., (2012), "The Internet Lowers Inhibitions: Implications for the Long Tail," University of Rochester working paper, available at https://dl.dropbox.com/u/9082176/mcdevitt_internet_inhibitions_april_2012.pdf
21. Miller A and Tucker C. (2009). Privacy Protection and Technology Diffusion: The Case of Electronic Medical Records. *Management Science*, 55:7, 1077-1093.
22. Mortensen, D. T., and Pissarides, C. A., (1999) "New Developments in Models of Search in the Labor Market," in O. Ashenfelter and D. Card, eds., Handbook of Labor Economics, Vol. 3 (Amsterdam: Elsevier Science), pp 2567–2627.
23. Pachankis, John E., Marvin R. Goldfried, Social anxiety in young gay men. *Journal of Anxiety Disorders*. 20, 996-1015 (2006).
24. Paumgarten, N., (2011), "Looking For Someone: Sex, Love, And Loneliness On The Internet, *New Yorker Magazine*, available at http://www.newyorker.com/reporting/2011/07/04/110704fa_fact_paumgarten#ixzz2eaiykevk
25. Piskorski, M., J., (2012), "Social Failures and Social Solutions: Evidence from OkCupid," adapted from Connect: How Do Online Social Platforms Work and How Can Firms Leverage Them For Success?, forthcoming Princeton University Press.
26. Romanosky S, Telang R, and Acquisti A. (2011). Do data breach disclosure laws reduce identity theft? *Journal of Policy Analysis and Management*, 30:2, 256-286.
27. Roth, A., Sotomayor, M. A. O., (1992), Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis, Econometric Society Monographs.
28. Slater, D., (2013), **Love in the Time of Algorithms: What Technology Does to Meeting and Mating**, Penguin, New York.
29. Smith, L., (2002), "The Marriage Model with Search Frictions," Working Paper.

30. Solove DJ. (2004). **The Digital Person: Technology and Privacy in the Information Age**. New York: New York University Press
31. Suler, J. (2004), "The online disinhibition effect," *CyberPsychology & Behavior*, 7(3), pp 321 - 326.