NET Institute*

www.NETinst.org

Working Paper #16-05

September 2016

Assessing and Quantifying Local Network Effects in an Online Dating Market

Gordon Burtch
University of Minnesota

Jui Ramaprasad
McGill University

* The Networks, Electronic Commerce, and Telecommunications (“NET”) Institute, <http://www.NETinst.org>, is a non-profit institution devoted to research on network industries, electronic commerce, telecommunications, the Internet, “virtual networks” comprised of computers that share the same technical standard or operating system, and on network issues in general.

Assessing and Quantifying Local Network Effects in an Online Dating Market

Gordon Burtch*

Carlson School of Management, University of Minnesota, gburtch@umn.edu,

Jui Ramaprasad*

Information Systems, Desautels Faculty of Management, McGill University, jui.ramaprasad@mcgill.ca,

We empirically examine and quantify network effects on a large online dating platform in Brazil. We consider the effects of a seeding intervention by the platform operator, wherein it acquired its primary competitor and subsequently imported the competitor’s 150,000+ user accounts over a 3-day period. The acquisition thus constitutes a large exogenous shock the composition of the acquiring platform’s user base. We estimate the effect of the shock on the rate of subsequent enrollments and exits amongst heterosexual users across 120 cities. Bearing in mind that the purchased users were exclusively heterosexual, we employ a difference-in-differences specification in which homosexual enrollment and exit rates serve as plausible controls. Our estimates indicate that the treatment increased the rates of both enrollment and exit, for both genders, with a net positive effect that translated to a 22% increase in short-term revenue for the platform. Further, we find that the response amongst male users was significantly stronger. When we consider that female participation was being fully subsidized by the acquiring platform, this result is consistent with the idea that subsidies and seeding strategies are substitutes, rather than complements. Finally, we explore nuances of the observed effects, quantifying *local* features. In particular, we show that the treatment effect varied significantly, depending on age differences and the degree of co-location between new and existing users.

Key words: online dating, network effects, platforms, natural experiment, difference in differences

1. Introduction

Two-sided markets (or multi-sided markets more generally) are roughly defined as platforms that enable transaction and interaction amongst end-users of different types (Rochet and Tirole 2006). Examples of these markets are now quite prevalent, e.g., Hulu matches video content providers

* The authors gratefully acknowledge financial support from the NET Institute, www.NETinst.org.

with end-consumers and advertisers, eBay connects buyers and sellers, and AirBNB links renters and hosts. Perhaps the oldest example, however, is the market for dating, which we consider here.

The success of two-sided markets hinges on network effects. Most commonly, individuals experience greater utility when the number of potential transaction partners is large (cross-side positive effects). At the same time, individuals experience lower utility when the number of competitors is large (same-side negative effects, or congestion). Platform operators employ a variety of techniques to stimulate and take advantage of network effects. In this work, we offer a case study of an online dating platform, which implemented one such technique: a seeding strategy. We do so with an eye toward identifying and quantifying network effects in this setting. We also explore the interaction between multiple interventions that a platform operator might consider employing in tandem. We are able to do this because, in addition to the observed intervention that we focus upon, the platform operator has maintained a consistent policy of subsidizing female users. By exploring heterogeneity in the response between the male and female population, we obtain some insight into the nature of the interaction between these two interventions. Finally, we seek to understand the importance of local network structure in the market (i.e., the fit between the seeding action and the preferences and interests of targeted users), and how this can determine the efficacy of the seeding strategy in stimulating network effects. Empirically understanding the nature of network effects is becoming increasingly important as the number of platforms competing in the online space grows. This is particularly true in the dating space, which has grown extremely competitive in recent years.¹ More generally, our efforts contribute to the broader platforms literature, which has largely focused on aspects of platform competition and pricing, instead of the factors that determine platform growth, and in which empirical work has typically been situated in offline markets (Chu and Manchanda 2016).

Similar to more traditional markets of sale and purchase (e.g., eBay), we might expect to observe positive cross-side and negative same-side network effects. However, we will argue that the story

¹Online dating magazine estimates that there are 2,500 online dating sites in the US alone: <http://www.forbes.com/sites/martinzwilling/2013/03/01/how-many-more-online-dating-sites-do-we-need/>

is somewhat more nuanced when it comes to online dating, for a number of reasons. First, broadly speaking, these nuances and differences derive from the fact that these markets facilitate social transactions and not economic ones, and thus individuals preferences for a transaction partner are likely to be highly heterogeneous. Second, by construction, participants in these markets generally have no direct visibility into levels of congestion or competition, because users do not typically receive information about other users who would not provide a viable relationship match.² Third, the online dating context, despite digital intermediation, nonetheless continues to depend heavily on offline word-of-mouth. Together, these differences highlight that the role and impacts of network effects may be highly contextual and should be considered in a case-by-case basis.

We seek to address the following three research questions: *1) To what degree does seeding new users stimulate network effects in online dating markets; 2) Is seeding complementary to subsidies, or substitutable?; and 3) How does the impact of the seeding strategy vary with differences between the characteristics of new (seeded) and existing users (i.e., matching potential)?* To address these questions, we draw upon anonymized, aggregated data from the Ashley Madison leak of 2015. We follow an approach similar to that of other researchers who have used this dataset for academic research purposes (Griffin et al. 2016, Grieser et al. 2016, Parsons et al. 2016), employing anonymized, aggregated data. Our interest in this dataset comes from one key event that enables us to study seeding in online dating: Ashley Madison purchased a Brazilian competitor, Ohhtel, and thereby acquired its user base, an action which resulted in the sudden injection of more than 150,000 new user accounts into the platform.

We exploit three features of this study context to identify the effects deriving from the shock on subsequent patterns of enrollment and exit on the dating site. First, whereas the competitor only allowed heterosexual matching, the focal site allowed both heterosexual and homosexual matching. Hence, the exogenous shock to the number of user accounts on our focal site was only applied

² Admittedly, this observation applies only to heterosexual users; homosexual users would of course serve, simultaneously, as possible matches and as competition for one another

to heterosexual users. Second, prior to the treatment, the platform had a standing policy of fully subsidizing female participation in the marketplace. Third, the treatment varied in a number of ways across cities, in that the characteristics of purchased user accounts (e.g., physical location within the city) differed to varying degrees from those of the existing base of users.

We leverage the first feature to obtain a plausible control group that was not subject to the exogenous shock, to estimate a difference-in-differences (DD) specification, contrasting the enrollment and exit rates between heterosexual and homosexual users. The second feature allows us to explore the degree to which subsidies and seeding strategies are complementary or substitutable, by examining heterogeneity in the response to treatment between males and females. Finally, we leverage the third feature to explore the notion of *local* network effects; the idea that the strength of network effects will vary based on the degree to which the new entrants can serve as a possible match for other users in the market.

We find, first, that the seeding action did deliver a positive return for enrollments, amongst both genders. Noting that site design precludes the manifestation of direct or indirect network effects on enrollment (because new users are incapable of viewing potential matches before they create an account), we attribute this finding to a mechanism of social proof, manifest via offline word-of-mouth. We also find that the treatment increased the rate of exit amongst both genders, suggesting a countervailing result. However, our calculations suggest an overall net benefit to the platform equivalent to an approximate 22% increase in the rate of acquisition of male users, and thus revenue. Second, we find evidence of a statistically significant, stronger effect amongst males. This result is consistent with the idea that seeding and subsidies are in fact substitutes; in short, it seems that females, because they are fully subsidized, demonstrated significantly less sensitivity to the seeding intervention. Third, exploring local structure in the market, we observe that the treatment effect was indeed heterogeneous between cities. In particular, we find that the treatment was significantly stronger when purchased users were i) co-located with, ii) younger than, or iii) healthier than existing users.

This paper extends the current work on network effects in online platforms. Our main contribution lies in identifying and quantifying network effects in the context of online dating, in contrast to the current literature that mainly focuses on platforms that facilitate economic exchanges. Importantly, we look at how one seeding strategy (seeding the platform with additional users) can stimulate network effects, addressing the critical “chicken-and-egg” problem faced by two-sided platforms. Additionally, we provide preliminary evidence that two common approaches to engineering network effects, namely seeding and subsidies, are substitutes, rather than complements. From a practical perspective, we are able to provide a measure of the economic effects of the seeding strategy by estimating the proportional revenue returns. In exploring these questions, we consider that the nature of network effects can vary across contexts and market types. One significant reason that network effects may vary, as we demonstrate here, is based on local structure in the market, and how that structure aligns with the preferences and interests of users. We demonstrate that the strength of network effects varies in our setting, depending on users’ relative ages, body types and locations. We argue that this occurs because these three features are important determinants of matching.

The remainder of this paper is structured as follows. We review the literature on network effects, two-sided markets, and online dating in the following section. We then describe our study context, the natural experiment and the data at our disposal, before detailing our empirical approach. We then report our results, discuss the implications and limitations of our work, and propose a number of avenues for future work.

2. Related Literature

In the literature review below, we highlight literature relevant to our paper, focusing on work that looks at network effects in two-sided platforms as well as the existing literature examining online dating platforms.

2.1. Network Effects

Two-sided markets (or multi-sided markets more generally) are roughly defined as platforms that enable transaction and interaction amongst end-users of different types (Parker and Van Alstyne

2005, Rochet and Tirole 2006). Examples of these markets are now quite prevalent, e.g., Hulu matches video content providers with end-consumers and advertisers, eBay connects buyers and sellers, and AirBNB links renters and hosts. Perhaps the oldest example, however, is the market for dating, which we consider here.

Any two-sided network is generally faced with the challenge of how best to attract and retain participants on each side of the market, to achieve and maintain critical mass, and ultimately turn a profit (Rochet and Tirole 2006). Achieving and maintaining participation on either side is a difficult proposition because platforms face a cold-start problem; they must attract membership, yet members have a significantly reduced interest in joining if no other members exist to transact with. This notion, which speaks to cross-side network effects (Katz and Shapiro 1985), implies that the platform often requires a specific strategy to achieve growth. That strategy often hinges on pricing. The platform operator may choose to subsidize participants on one side of the market, offering free or discounted access and features (Parker and Van Alstyne 2005). Indeed, this has long been a standard practice in dating markets, wherein free access is often granted to female users, as noted in the introduction. This is also the case in the market we study here, and long before the emergence of online dating, bars were using promotions such as “Ladies Night,” letting ladies in for free in order to attract men (in the case of heterosexual pairs), i.e. to stimulate cross-side network effects.

However, online platforms have experimented with a variety of other, non-pricing strategies to stimulate cross-side network effects, many of which have been shown to be quite effective. For example, past work has demonstrated that if a platform can institute a perception of adoption by members on one side of the market, this can have a profound effect on enrollment on the opposite side. Tucker and Zhang (2010) report on a field experiment in which they demonstrate the benefit of advertising existing user volumes to potential users. They find that informing potential sellers of the volume of existing sellers in a B2B marketplace has a positive effect on the probability of sign-up. However, informing potential sellers of existing sellers and buyers, in tandem, eliminates

the effect. The authors take this as evidence of an indirect network effect, wherein the absence of information about buyer volumes causes the potential seller to infer the number based on known seller volumes.

Anecdotal evidence also speaks to the success of platforms that have initially “faked” membership and usage (an example of a seeding strategy).³ In 2012, Steve Huffman, co-founder of Reddit, revealed that the site stimulated early growth in its user base via the creation of fake accounts. Similar practices have been reported at PayPal, which reportedly grew its early base of eBay sellers by creating a bot that would automatically purchase goods and then insist that the payment be completed via PayPal. This strategy caused eBay sellers to flock to the Paypal platform, under the perception that it was heavily populated by consumers (Jackson 2004).

While a variety of interventions have been explored, both in practice, and in the academic literature, to our knowledge, no prior work has explored the effects of combining different interventions in tandem. In particular, it is unclear how subsidies may interact with alternative strategies, such as seeding. We anticipate that these two types of interventions are likely to be substitutes, for two reasons. First, subsidies will be most effective for price sensitive consumers. Past work has observed that price sensitivity is positively associated with search costs (Lynch and Ariely 2000). As new users are seeded into the market, we might expect a decline in search costs, because users will have an easier time finding a transaction partner that closely meets their preferences. Thus, the seeding action should serve to reduce price sensitivity, weakening the efficacy of subsidies. Second, paying users are likely to be more sensitive to any changes in the market that improve their return on investment. As matching probabilities increase with the seeding action, paying consumers are likely to respond particularly strongly, as they see a greater relative increase in the value of the service.

Just as positive cross-side effects manifest with a growth in potential transaction partners, negative same-side effects can arise with the enrollment of competitors. Sun and Tse (2009) explain

³ <http://venturebeat.com/2012/06/22/reddit-fake-users/>

this phenomenon as follows: “Even if there are infinite potential participants of the network, the exponential growth of the network may be limited by other factors such as competition among participants. This is called a congestion effect in the two-sided market literature (Rysman 2004) and it represents negative network externalities among the participants on the same side of the market. Sun (2007) finds that when such negative externalities exist, the growth of a two-sided network has a finite limit.”

Network effects can manifest in a heterogeneous fashion in many settings, when individuals prefer to transact with specific subsets of the population. Sundararajan (2007) formalizes this idea, presenting a model in which users prefer adoption of a product by their social network neighbors. Other work has frequently acknowledged the importance of local network structure in the manifestation of network effects, considering, for example, local network density (Katana et al. 2011) and strong tie subsets (Suarez 2005).

Given the highly personal nature of relationships, we would argue that a similar pattern of localized preference can play out in ways that need not depend strictly on social network structure. That is, we propose a broader definition of what is considered “local”. In particular, we note that individuals may prefer adoption by a specific subset of the population for a variety of other reasons; individuals may prefer adoption by people who speak the same language (Meer and Rigbi 2013), who come from a similar cultural background (Burtch et al. 2014), or who have a particular set of personal features.

2.2. Online Dating

Online dating itself is a fairly new phenomenon that has just recently achieved social acceptability in the mainstream. The related online dating literature is nascent but growing. While the bulk of the work in online dating has been from the psychology literature, with a focus on individual self-presentation (Ellison et al. 2006), deception (Toma et al. 2008), and image, the economics literature has also examined online dating, looking specifically at mate preferences and sorting in matching (Hitsch et al. 2010).

Most relevant to our work is the literature on platform design of online dating sites. Recent literature in Information Systems has looked at the impact of different features (e.g. anonymous browsing) on matching outcomes (Bapna et al. 2016). Other work has examined the design of recommendation systems for human-to-human matching which, because it must take into account two sets of preferences instead of one, is different from human-to-product matching that we see on sites like Amazon.com (Pizzato et al. 2010), and the inclusion of search features to narrow down the choices of potential matches (Fiore and Donath 2004). Though critically important to the success of an online dating site, few papers have looked at network effects in the online dating context, and in particular little work has examined impacts on enrollment of new users and attrition of existing users—critical factors in the success of such a site. The objective of this project is to address this gap.

Similar to more traditional markets of sale and purchase (e.g., eBay), which have received most of the consideration in the network effects literature to date, we might expect to observe positive cross-side and negative same-side network effects. However, we will argue that the story is quite different in the online dating context, for a number of reasons, most notably because these markets do not facilitate economic transactions; instead, they facilitate social transactions. Moreover, by construction, despite advertisements that tout the "millions of possibilities" available on online dating sites (Heino et al. 2010), participants in these markets generally have no direct visibility into levels of congestion or competition, because users do not typically receive information about other users who would not provide a viable relationship match. Finally, a successful transaction on an online dating platform may result in the loss of two participants, thus requiring the platform to continually seek new enrollments in order to sustain itself. Together, these differences highlight that the role and impacts of network effects may be highly contextual and should be considered on a context-by-context basis.

Online dating markets provide perhaps the best opportunity to explore the importance of local preferences as a determinant of network effects. In online dating, individuals vary quite a bit in

their statements about the characteristics they deem to be desirable in a relationship partner (i.e., “turn-ons” vs. “turn-offs”). Accordingly, an ideal seeding strategy for an online dating market might be to focus on seeding new users who most clearly meet the preferences of the existing user base. In the online dating context, taking this approach makes sense, as this would be most likely to stimulate word of mouth on the part of existing users, who might then draw in new users.

Of course, when it comes to platform growth strategies in the general sense, it likely behooves platform operators to explore the preference distributions of users in a given market, and to then optimize based on what is found. That is, firms may wish to incentivize participation amongst *specific* subsets of the population or they may wish to seed specific types of content, to maximize the resulting demand.

3. Methodology

3.1. The Natural Experiment

On April 10th of 2012, Ashley Madison’s acquisition of Brazilian competitor Ohhtel was announced, for an undisclosed dollar amount⁴. With the announcement, the media reported that the Ohhtel user data was transferred directly to Ashley Madison’s servers. Given that this was the only media announcement that was made about this acquisition and it occurred on the day of the account transfer, we are not concerned that the media buzz around this event is driving our results. Ohhtel users, upon attempting to login to their user accounts after that transfer, would be redirected to the Ashley Madison website and asked to accept a set of terms and conditions before proceeding. This purchase and user transfer reportedly resulted in the addition of approximately 150,000 new users to Ashley Madison’s Brazilian website.

The user account transfer described by the media is readily apparent in our sample. Figure 1 depicts the volume of Brazilian user account creations over time in our sample, in the two months surrounding April 10th. Here, we clearly see the creation of approximately 150,000 new user accounts between April 9th and April 11th. When we look more closely at these accounts, we observe that

⁴ Globo News: <http://g1.globo.com/economia/negocios/noticia/2012/04/site-de-traicao-ashley-madison-compra-base-de-usuarios-do-ohhtel.html>

they were all created from the same IP address. A *whois* lookup on that IP reveals that it is maintained by OnX Enterprise Solutions (US) Ltd., a hosting services provider that also hosts AvidLifeMedia.com (ALM is the umbrella company that owns Ashley Madison). These observations are consistent with the occurrence of a large database import of user accounts having taken place at the time the acquisition was announced in the media.

A notable feature of these purchased accounts is that they were almost exclusively heterosexual; just 78 of the 150,000+ users indicated an interest in same-sex partners. This implies that the natural experiment presents an exogenous shock to the volume of heterosexual user accounts on Ashley Madison’s Brazilian platform, yet virtually no shock to the volume of homosexual user accounts.

Figure 1 User Account Creation Over Time.

3.2. Data

We employ aggregate, anonymized data from the Ashley Madison data leak of 2015. In anonymizing the data, we follow the approach of other recent academic research that has drawn on the same

data set (Griffin et al. 2016, Grieser et al. 2016). The natural experiment takes place from April 9th to April 11th of 2012, and our sample spans a 2-month window around these dates (March through May of 2012). We focus on the platform’s 120 largest Brazilian markets (cities). Figure 2 depicts the geographic distribution of users in our sample who had registered prior to the treatment date. As can be seen, user volumes follow a power distribution; São Paulo and Rio de Janeiro jointly account for more than 30% of registered users, or approximately 80,000 individuals, while the smallest city in our sample, Rio Grande, hosts slightly more than 400 users.

Figure 2 Geographic Distribution of Users.

We aggregate user accounts to construct daily counts enrollment and exit, i.e., last date of profile modification. We are relatively confident that the last date of account activity in our data is a reliable indicator of exit, because although we do not study data post-2012, we draw upon all data through 2015 to identify a user’s last action. Thus, a user whose last date of activity was observed in 2012 was inactive for at least 3 years afterward. Accordingly, right censoring is not a serious

concern in our data. We construct four panels for each city in our data, for a total of 480 panels. These panels reflect the cross-product of gender and sexual orientation; one panel for heterosexual males, one for heterosexual females, one for homosexual males and one for homosexual females.

In Table 1, we report descriptive statistics related to daily enrollments and exits in each of the cities over the month prior to the natural experiment, for each combination of gender and sexual orientation.

Table 1 Average Daily City Enrollment & Exit in Two Weeks Before Treatment.

<i>Variable</i>	Heterosexual						Homosexual					
	Male			Female			Male			Female		
	Max	Mean	STDev	Max	Mean	STDev	Max	Mean	STDev	Max	Mean	STDev
<i>Enrollment</i>	288.00	12.721	29.537	37.000	1.441	3.750	14.571	0.651	1.491	12.286	0.589	1.360
<i>Exit</i>	8.00	0.502	0.847	34.643	1.364	3.482	14.357	0.636	1.466	12.286	0.568	1.334
<i>Cities</i>	120	120	120	120	120	120	120	120	120	120	120	120

3.3. Econometric Specification

All purchased user accounts were heterosexual. As such, homosexual enrollment serves as our control. Because our dependent variable represents a strictly positive count measure and exhibits a power (right skewed) distribution, we draw upon the Conditional Fixed Effect Poisson Quasi-Maximum Likelihood (PQML) estimator proposed by Wooldridge (1999), and we subsequently evaluate the robustness of our primary results to the use of a Log-OLS specification.

Bertrand et al. (2004) advocate a number of strategies to avoid issues of false significance deriving from the use of a difference-in-differences estimator on data with an outcome variable that exhibits serial correlation. One such strategy, easily implemented, is to ignore time-series data and average observations across the pre period, as well as observations across the post period, collapsing the panel into pairs of observations for each unit. We take that approach here. Thus, we begin by constructing a panel of observation pairs, reflecting enrollment volumes for each gender and sexual orientation, in each city, in the pre and post periods. Our panel is thus short ($T = 2$), yet relatively wide. In our estimations, we consider two alternative time windows; 2 weeks before and after the treatment (i.e., 1-month), and 4 weeks before and after the treatment (i.e., 2-months). We execute

separate estimations by gender, in order to compare the coefficient estimates, and thereby assess the relative strength of response between paying and subsidized users (males vs. females).

Equation 1 reflects the log-linear specification, where i indexes cities, j indexes sexual orientation and t indexes time. Our outcome of interest, Y is the count of enrollments or exits, which we model as a function of *Orientation*, an indicator that equals 1 if the panel pertains to heterosexual users, *Post*, an indicator of whether the observation takes place after the natural experiment date, and the interaction between the two. Additionally, we incorporate a city fixed effect, α . Here, we are particularly interested in the coefficient β_3 , the difference-in-differences (DD) estimate. In performing these estimations, we exclude the purchased user accounts from our enrollment and exit measures. Moreover, in the specific case of our attrition (exit) estimations, we exclude any users who registered on or after the treatment date. This is necessary in order to rule out the possibility that users who arrive due to the treatment may be systematically different than users who arrived “organically.” For example, it may be the case that treatment-induced enrollments are systematically less interested or committed to the platform, and thus that they may exit more rapidly. Accordingly, were we to include post treatment entrants in our calculation of exit rates over time, we might observe an uninformative increase in the exit rate following treatment.

$$\log(Y_{ijt}) = \alpha_i + \beta_1 \cdot \text{Orientation}_j + \beta_2 \cdot \text{Post}_t + \beta_3 \cdot \text{Orientation}_j * \text{Post}_t + \varepsilon_{ijt} \quad (1)$$

Before progressing to our primary estimations, it is also useful to first explore the validity of our assumption that the homosexual population can serve as a valid control for the heterosexual population. To determine this, we first estimate two relative time difference-in-differences models (Autor 2003, Angrist and Pischke 2009), pooling the gender panels, separately considering enrollments and exits. We estimate the specification presented in Equation 2, treating April 8th as the reference period (the day prior to the natural experiment). If the homosexual population is a valid control group, we would expect to find no statistically significant coefficients associated with pre-treatment relative time dummies, yet significant coefficients associated with post-treatment relative

time dummies. We estimate this model on our un-collapsed panels for the 28 days around April 9th. We plot our estimates of γ from each regression in Figures 3 and 4, with bands representing the 95% confidence intervals. As can be seen in the two figures, there are no significant differences in enrollment or exit between the heterosexual and homosexual populations before the natural experiment, yet we begin to see differences emerge shortly after. The lack of significant differences prior to the experiment suggests that we have no reason to believe that the homosexual population cannot serve as a valid control in our setting.

$$\log(Y_{ijt}) = \alpha_i + Orientation_j + \sum_q \tau_{qt} + \gamma \cdot \sum_q \tau_{qt} * Orientation_j + \varepsilon_{ijt} \quad (2)$$

Figure 3 Poisson Relative Time Estimates for Enrollments; t-1 Omitted.

Figure 4 Poisson Relative Time Estimates for Exits; $t-1$ Omitted.

4. Results

We next proceed with our primary analyses. Table 2 reports our findings for new user enrollments, by gender, for both the 4-week and 2-week pre/post windows. Table 3 reports the results using a Log-OLS specification. We observe consistent results in both cases. Referring to Table 3, we see that female enrollments increased by approximately 21.1% in the two weeks after treatment, whereas male enrollments rose by approximately 23.8%. The latter result is of particular note, given that only male users pay to use the Ashley Madison website.

Next, we consider the natural experiment's impact on exit rates. We observe significant positive effects across all estimations. Table 4 reports our findings for user exits, by gender, for both the 4-week and 2-week pre/post windows, while Table 5 reports the results using a Log-OLS specification. Referring to Table 5, we see that female exits increased by approximately 24.8% in the two weeks after treatment, whereas male exits rose by approximately 56.3%. The latter result is once again of particular note, because only male users pay to use the Ashley Madison website.

Table 2 Enrollments: Poisson FE.

<i>DV = Enrolls</i>	Male		Female	
	4-Week	2-Week	4-Week	2-Week
<i>Post</i>	-0.799*** (0.035)	-0.970*** (0.043)	-0.697*** (0.042)	-0.795*** (0.050)
<i>Orientation</i>	2.995*** (0.032)	2.970*** (0.034)	0.920*** (0.034)	0.895*** (0.054)
<i>Post * Orientation</i>	0.288*** (0.041)	0.467*** (0.048)	0.211*** (0.051)	0.309*** (0.050)
<i>Observations</i>	480	476	476	476
<i>Cities</i>	120	119	119	119
<i>Wald χ^2</i>	10,923.64 (3)	14,645.41 (3)	1,272.02 (3)	1,461.96 (3)

Notes: Robust standard errors in brackets; *** p<0.001, ** p<0.01, * p<0.05; Some observations dropped because no variation is observed in the DV over time

Table 3 Enrollments: Log-OLS FE.

<i>DV = Log(Enrolls)</i>	Male		Female	
	4-Week	2-Week	4-Week	2-Week
<i>Post</i>	-0.759*** (0.059)	-0.804*** (0.064)	-0.738*** (0.062)	-0.806*** (0.069)
<i>Orientation</i>	3.032*** (0.040)	3.018*** (0.051)	0.926*** (0.043)	0.834*** (0.057)
<i>Post * Orientation</i>	0.207*** (0.062)	0.238*** (0.067)	0.169*** (0.075)	0.211*** (0.086)
<i>Observations</i>	469	454	463	445
<i>Cities</i>	120	119	119	119
<i>R²</i>	0.967	0.960	0.735	0.664
<i>F statistic</i>	3,042.41 (3, 119)	2,103.36 (3, 118)	359.25 (3, 118)	239.78 (3, 118)

Notes: Robust standard errors in brackets; *** p<0.001, ** p<0.01, * p<0.05; Some observations dropped because Log(0) is undefined

Table 4 Exits: Poisson FE.

<i>DV = Exits</i>	Male		Female	
	4-Week	2-Week	4-Week	2-Week
<i>Post</i>	-0.721*** (0.039)	-0.570*** (0.040)	-0.606*** (0.043)	-0.333*** (0.048)
<i>Orientation</i>	-0.268*** (0.093)	-0.238* (0.092)	0.906*** (0.033)	0.838*** (0.051)
<i>Post * Orientation</i>	0.688*** (0.051)	0.880*** (0.069)	0.212*** (0.046)	0.199*** (0.051)
<i>Observations</i>	476	476	476	476
<i>Cities</i>	119	119	119	119
<i>Wald χ^2</i>	395.76 (3)	251.87 (3)	994.72 (3)	695.95 (3)

Notes: Robust standard errors in brackets; *** p<0.001, ** p<0.01, * p<0.05; Some observations dropped because no variation is observed in the DV over time

We assess the relative effect on paying and subsidized users by comparing the DD estimates between our male and female regressions, in both the enrollment and exit case. We find, as anticipated, that the male response was significantly larger in both cases ($p < 0.05$ and $p < 0.001$, respectively). This finding is consistent with the notion that subsidies and seeding strategies are substitutable interventions. Again, males, as paying users are likely to respond more strongly because they see a larger increase in value for money. At the same time, the decline in search costs

Table 5 Exits: Log-OLS FE.

<i>DV = Log(Exits)</i>	Male		Female	
	4-Week	2-Week	4-Week	2-Week
<i>Post</i>	-0.608*** (0.059)	-0.649*** (0.066)	-0.619*** (0.068)	-0.715*** (0.071)
<i>Orientation</i>	-0.183*** (0.089)	-0.001 (0.093)	0.929*** (0.044)	0.838*** (0.060)
<i>Post * Orientation</i>	0.679*** (0.084)	0.563*** (0.103)	0.148* (0.074)	0.248*** (0.079)
<i>Observations</i>	459	434	461	445
<i>Cities</i>	119	119	119	119
<i>R²</i>	0.178	0.214	0.700	0.642
<i>F statistic</i>	39.21 (3, 118)	38.11 (3, 118)	306.79 (3, 118)	171.79 (3, 118)

Notes: Robust standard errors in brackets; *** p<0.001, ** p<0.01, * p<0.05; Some observations dropped because Log(0) is undefined

might be expected to decrease the price sensitivity of female users, making the subsidy less effective (Lynch and Ariely 2000).

Having obtained estimates of the impact of the treatment on both entry and exit rates, we can now assess the net effect on revenue by calculating the shift in the net acquisition of paying (male) users, accounting for both entry and exit. Again considering the two-week pre-post window, where the effects were strongest, the elasticity of enrollment, equal to 24%, equates to an approximate increase of approximately 3 new male users per day in the average city, or 69 new male users per day in the largest city, São Paulo. The elasticity of exit, equal to 56%, equates to an approximately increase of 0.28 additional male user exits per day in the average city, or 4.5 additional male exits per day in São Paulo. Accordingly, in the average city, we see an approximate net daily acquisition of 2.72 new male (paying) user per day in the average city, over and above the baseline net acquisition rate of approximately 12.22 new male (paying) users, or an estimated 22% increase in short term revenue.

4.1. Local Network Effects

To explore the role of local network effects in this marketplace, i.e., the degree to which our results depend on the distribution of preference in the existing user base, in a given city, we consider user separation in various dimensions, including location, age and health (as measured by body mass index).

4.1.1. Geographic Separation Geographic distance has been highlighted as a particularly important factor in users selection of online dating partners (Couch and Liamputtong 2008). This

is because physical distance translates to increased travel time. As such, intuitively, if our results are indeed driven by network effects, we would expect to observe weaker effects in cities where the average location of pre-existing users deviates from the average location of purchased users (e.g., analogous to the notion of cluster separation) . To determine average location, we draw upon latitude and longitude data recorded for each user account at the time of creation, which are likely based on the user’s IP address at registration, or a self-reported address. For each city, we calculate the average geographic coordinates of pre-existing heterosexual users, and another set of average coordinates for newly entered (purchased) users. Based on these coordinate pairs, we then calculate our measure of geographic separation, Geo , as the Euclidean distance between the two (e.g., analogous to a centroid distance). This measure then enters our baseline specification as a moderator, along with all lower-order interactions, as reflected in Equation 3. In this model, we are thus particularly interested in β_6 and β_7 , our difference in differences estimate and its moderation by geographic distance. In this setup, we would expect to once again observe a positive coefficient for β_6 and a negative coefficient on the moderator, β_7 .

$$\begin{aligned} \log(Y_{ijt}) = & \alpha_i + \beta_1 \cdot Orientation_j + \beta_2 \cdot Post_t + \beta_3 \cdot Geo_i + \\ & \beta_4 \cdot Orientation_j * Geo_i + \beta_5 \cdot Post_t * Geo_i + \beta_6 \cdot Orientation_j * Post_t + \\ & \beta_7 \cdot Orientation_j * Post_t * Geo_i + \varepsilon_{ijt} \end{aligned} \quad (3)$$

It should be noted that our measure of geographic separation is not strictly accurate, because degrees of latitude and longitude vary in their relationship to traditional measures of physical distance (e.g., miles), depending upon the one’s location on the earth. Fortunately, however, Brazil is situated on the equator, where a degree of latitude and a degree of longitude both translate to roughly 68 miles or 110km. Accordingly, a simple Euclidean distance measure can provide a rough approximation of actual physical distance in a localized area (e.g., within a city). We incorporate our distance measure as a moderator on the interaction term in each of our 2-week, pre-post regressions, for both enrollments and exits.

The results of the enrollment regressions are reported in Table 6, and those for the exit regressions are reported in Table 7. Broadly speaking, as we expected, we find that a greater average geographic distance between pre-existing users and purchased users results in weaker treatment effects. The one exception to this is that we observe no significant moderating effect when it comes to male exits.

Table 6 Geographic Distances Moderate Enrollment: Poisson FE.

<i>DV = Enrolls</i>	Male	Female
<i>Post</i>	-0.980*** (0.045)	-0.832*** (0.054)
<i>Orientation</i>	2.973*** (0.035)	0.887*** (0.056)
<i>Post * Geo</i>	0.019** (0.007)	0.060*** (0.017)
<i>Orientation * Geo</i>	0.001 (0.014)	0.027 (0.020)
<i>Post * Orientation</i>	0.479*** (0.049)	0.350*** (0.049)
<i>Post * Orientation * Geo</i>	-0.018* (0.008)	-0.060* (0.025)
<i>Observations</i>	456	456
<i>Cities</i>	114	114
<i>Wald χ^2</i>	23,172.01 (6)	1,400.89 (6)

Notes: Robust standard errors in brackets; *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$; Some cities dropped because no variation was observed in the DV within the group or because some height or weight data was not available for some users

Table 7 Geographic Distances Moderate Exit: Poisson FE.

<i>DV = Exits</i>	Male	Female
<i>Post</i>	-0.888*** (0.048)	-0.726*** (0.050)
<i>Orientation</i>	-0.242* (0.096)	0.869*** (0.053)
<i>Post * Geo</i>	0.021** (0.008)	0.059** (0.022)
<i>Orientation * Geo</i>	-0.013 (0.028)	0.025 (0.021)
<i>Post * Orientation</i>	0.875*** (0.067)	0.340*** (0.049)
<i>Post * Orientation * Geo</i>	-0.004 (0.010)	-0.044* (0.019)
<i>Observations</i>	456	456
<i>Cities</i>	114	114
<i>Wald χ^2</i>	697.27 (6)	867.96 (6)

Notes: Robust standard errors in brackets; *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$; Some cities dropped because no variation was observed in the DV within the group or because some height or weight data was not available for some users

Considering the point we had made above, that a single unit change in our average geographic separation measure equates to approximately 68 miles, with the fact that a number of the largest cities in our sample are quite large, in some cases spanning hundreds of miles, we observe what

appears to be a relatively significant attenuation of the treatment effect from locational separation amongst users.

4.1.2. Age Differences Broadly speaking, we expect that users will hold preferences when it comes to the relative age of other users. Accordingly, we constructed a directed measure of the difference in average ages between existing and new users, by city (existing user age minus new user age). Thus, if existing users in a city were, on average, older than new users, our measure would take on a positive value. Once again, this measure enters our baseline specification as a moderator, along with all lower-order interactions, as reflected in Equation 4. Again, we are interested in β_6 and β_7 , our difference in differences estimate and its moderation by age differences. In this setup, we once again expect a positive coefficient for β_6 . It is not immediately apparent what coefficient we might expect for, β_7 , and, indeed, the coefficient may vary in direction depending on gender.

$$\begin{aligned} \log(Y_{ijt}) = & \alpha_i + \beta_1 \cdot Orientation_j + \beta_2 \cdot Post_t + \beta_3 \cdot Age_i + \\ & \beta_4 \cdot Orientation_j * Age_i + \beta_5 \cdot Post_t * Age_i + \beta_6 \cdot Orientation_j * Post_t + \\ & \beta_7 \cdot Orientation_j * Post_t * Age_i + \varepsilon_{ijt} \end{aligned} \quad (4)$$

The results of our estimations are presented below in Tables 8 and 9. We observe, for females in particular, that the positive effect on enrollment rates from the treatment is magnified when the purchased male users are relatively older. Alternatively, we might interpret this finding as evidence that the treatment is magnified when new male users encounter younger existing female users. We observe no significant moderating effects in our exit regressions.

4.1.3. Body Type Differences Finally, we consider differences in body type. As in the case of Age, we generally expect that users will hold preferences about others' relative body type. Accordingly, we constructed a measure of Body Mass Index (BMI) based on users reported weight and height information. Because our weight data is recorded in kilograms, and height data is recorded in meters, we employed the metric formula. Here, BMI is calculated by dividing individuals' weight by the square of their height. Once again, we then subtracted the average BMI of

Table 8 Age Differences Moderate Enrollment: Poisson FE.

<i>DV = Enrolls</i>	Male	Female
<i>Post</i>	-0.967*** (0.043)	-0.794*** (0.058)
<i>Orientation</i>	2.975*** (0.036)	0.905*** (0.057)
<i>Post * Age</i>	0.024* (0.011)	0.021 (0.011)
<i>Orientation * Age</i>	0.021* (0.011)	-0.003 (0.011)
<i>Post * Orientation</i>	0.468*** (0.048)	0.310*** (0.055)
<i>Post * Orientation * Age</i>	-0.023 (0.012)	-0.025* (0.012)
<i>Observations</i>	448	444
<i>Cities</i>	112	111
<i>Wald χ^2</i>	16,142.19 (6)	1,443.63 (6)

Notes: Robust standard errors in brackets; *** p<0.001, ** p<0.01, * p<0.05; Some cities dropped because no variation was observed in the DV within the group or because some height or weight data was not available for some users

Table 9 Age Differences Moderate Exit: Poisson FE.

<i>DV = Exits</i>	Male	Female
<i>Post</i>	-0.873*** (0.047)	-0.710*** (0.057)
<i>Orientation</i>	-0.245** (0.093)	0.885*** (0.053)
<i>Post * Age</i>	0.019 (0.012)	0.007 (0.011)
<i>Orientation * Age</i>	0.014 (0.028)	-0.006 (0.012)
<i>Post * Orientation</i>	0.863*** (0.068)	0.322*** (0.056)
<i>Post * Orientation * Age</i>	-0.002 (0.022)	-0.010 (0.013)
<i>Observations</i>	448	444
<i>Cities</i>	112	111
<i>Wald χ^2</i>	397.46 (6)	970.67 (6)

Notes: Robust standard errors in brackets; *** p<0.001, ** p<0.01, * p<0.05; Some cities dropped because no variation was observed in the DV within the group or because some height or weight data was not available for some users

purchased users in a city from the average BMI of existing users, to obtain a directed measure of body type differences. This measure once again enters our baseline specification as a moderator, as per Equation 5.

$$\begin{aligned}
\log(Y_{ijt}) = & \alpha_i + \beta_1 \cdot \text{Orientation}_j + \beta_2 \cdot \text{Post}_t + \beta_3 \cdot \text{BMI}_i + \\
& \beta_4 \cdot \text{Orientation}_j * \text{BMI}_i + \beta_5 \cdot \text{Post}_t * \text{BMI}_i + \beta_6 \cdot \text{Orientation}_j * \text{Post}_t + \\
& \beta_7 \cdot \text{Orientation}_j * \text{Post}_t * \text{BMI}_i + \varepsilon_{ijt}
\end{aligned} \tag{5}$$

The results of our enrollment and exit regressions are reported below, in Tables 10 and 11.

Unlike with geographic separation and age differences, we find no evidence indicating that the

treatment effect is significantly moderated by differences in average BMI between purchased and existing users.

Table 10 Body Type Differences Moderate Enrollment: Poisson FE.

<i>DV = Enrolls</i>	Male	Female
<i>Post</i>	-1.049*** (0.057)	-0.832*** (0.072)
<i>Orientation</i>	2.975*** (0.034)	0.822*** (0.052)
<i>Post * BMI</i>	0.050 (.033)	0.026 (0.034)
<i>Orientation * BMI</i>	-0.004 (0.018)	0.071** (0.026)
<i>Post * Orientation</i>	0.561*** (0.054)	0.393*** (0.084)
<i>Post * Orientation * BMI</i>	-0.058 (0.031)	-0.060 (0.040)
<i>Observations</i>	448	444
<i>Cities</i>	112	111
<i>Wald χ^2</i>	16,177.06 (6)	1,797.11 (6)

Notes: Robust standard errors in brackets; *** p<0.001, ** p<0.01, * p<0.05; Some cities dropped because no variation was observed in the DV within the group or because some height or weight data was not available for some users

Table 11 Body Type Differences Moderate Exit: Poisson FE.

<i>DV = Exits</i>	Male	Female
<i>Post</i>	-0.990*** (0.047)	-0.718*** (0.073)
<i>Orientation</i>	-0.409** (0.142)	0.831*** (0.058)
<i>Post * BMI</i>	0.074** (0.027)	0.004 (0.032)
<i>Orientation * BMI</i>	0.103 (0.068)	0.048 (0.028)
<i>Post * Orientation</i>	0.985*** (0.104)	0.368*** (0.074)
<i>Post * Orientation * BMI</i>	-0.078 (0.050)	-0.034 (0.038)
<i>Observations</i>	448	444
<i>Cities</i>	112	111
<i>Wald χ^2</i>	615.73 (6)	1,208.70 (6)

Notes: Robust standard errors in brackets; *** p<0.001, ** p<0.01, * p<0.05; Some cities dropped because no variation was observed in the DV within the group or because some height or weight data was not available for some users

5. Discussion & Conclusion

This work offers a first look at network effects in online dating. We consider a case study in which a Brazilian online dating platform purchased its primary competitor, and subsequently imported the entirety of the competitor’s heterosexual user base over a three day period. This action, i.e., seeding new users in the market, imposed an exogenous shock on the volume of heterosexual users in the market, with little to no effect on the homosexual population. We demonstrate through

a series of difference-in-difference estimations (treating the homosexual population as a control) that this strategy was indeed effective, increasing the rate of new user entry as well as the rate of attrition, with a net positive effect that translated to an approximate 22% increase in short term revenue. Further, we demonstrate that the strength of the treatment, and thus the resultant network effects, were moderated by the degree of co-location amongst new and existing users, as well as age differences.

We thus contribute to the literature by identifying and quantifying network effects in the context of online dating, in contrast to the current literature that mainly focuses on platforms that facilitate economic exchanges. In doing this, we suggest that given the differences between the different types of platforms, such as in the visibility of network size, the length of the transaction between the two parties, the importance of offline interactions, and importantly, the individual preferences of the users, the nature of network effects can vary across contexts and market types. Moreover, network effects can manifest to varying degrees within a market, depending on local preferences or network structure. We demonstrate here how two particular features of users, geography and age, influence the manifestation of network effects in an online dating market.

It is readily apparent that relevant factors will vary by transaction-type and context. As noted previously, there are numerous online platforms in which the geography of users would be unlikely to play an important role in determining network effects, namely those platforms which fully mediate matching and transaction (e.g., video game consoles and online multiplayer games). At the same time, other factors that bear no relevance to the dating context may play a particularly important role in other markets. Returning to the example of online games, players may, for example, exhibit a preference for opponents of a comparable skill level, in order to ensure a “challenge.”

Interestingly, these ideas suggest that more traditional launch strategies for highly innovative products, such as targeted launch (Lee and O’Connor 2003), remain somewhat relevant for products and services characterized by network effects. Although the volume of adopting users may be of paramount importance in the presence of network effects, our findings suggest that, under a

resource constraint, it may be to the platform operator’s advantage to allocate it’s resources toward the acquisition of specific individuals, namely those who are likely to be of greatest interest to others, and thus most likely to stimulate subsequent adoption.

Thus, the results of this work can help inform strategies for online dating platforms, whose number and popularity are growing. As most platform operators are now aware of the importance of network effects, it is important for firms not only to achieve a large installed base, but to do so as quickly as possible. By optimizing platform launch and growth strategies, our findings suggest that firms can accelerate their acquisition of new users and increase their likelihood of succeeding in a winner-take-all scenario. At the same time, platforms need to remain cognizant of congestion and search costs. As the installed base of an online dating platform grows large, users will face greater difficulty sorting through suggested matches to identify those with real potential. By limiting growth to users who are quite likely to match with others in the market, the platform can maintain efficiency and ensure an ideal user experience. Interestingly, some relatively new online dating applications appear to have struck upon this same conclusion. Newer online dating services like *Coffee Meets Bagel*⁵ place a greater focus on relevant matches, providing users with just one recommendation each day, based on user-specified filters. In taking this approach, the platform operators seek to reduce noise and congestion while simultaneously ensuring the relevance of potential matches.

Our work is of course subject to a number of limitations. First, we are unable to identify the exact mechanism of the observed treatment effect on exit rates and thus our main contribution focuses on enrollments, where the mechanism is more easily understood. That being said, as noted above, the immediacy of the response and the fact that heterosexual users are unable to directly observe their competition implies that the increase in exits is due to successful matches. Future work might look to build on our findings employing more granular, nuanced data, enabling clearer identification of this mechanism.

⁵ <https://coffeemeetsbagel.com>

Second, our data is limited to a single geography (Brazil), which could limit the generalizability of our results in other countries, given that dating norms vary across cultures (Hatfield and Rapson 1996). Future work could examine online dating seeding strategies in other countries. Third, we note that we are only able to identify the short-term impact of this influx of new users to the focal website on enrollments and exits. While indirect network effects could take some time to develop, there are other factors that also confound drivers of enrollments and exits as the time window increases making it difficult to isolate the impact of this seeding strategy.

Further, our conclusions regarding the interaction between subsidies and seeding are not as well identified as we might wish. It is possible that the heterogeneous responses we observe between males and females are in fact due to other factors that we do not observe, such as differences in patterns of website use between the two genders. Future work can improve upon our understanding of the interaction between subsidies and other approaches to engineering network effects, e.g., via randomized experimentation.

Finally, to further pursue this line of work, it would be particularly useful for researchers to explore the features and characteristics of individuals and markets that may drive matching outcomes in various settings characterized by network effects. We have demonstrated three examples of user characteristics that impact matching in this particular context. Going forward, researchers might seek to develop a theoretical typology (Bailey 1994) of features that can result in local network effects, based on transaction and market characteristics.

References

- Angrist J, Pischke JS (2009) *Mostly Harmless Econometrics: An Empiricists Companion* (Princeton, NJ: Princeton University Press).
- Autor DH (2003) Outsourcing at will: The contribution of unjust dismissal doctrine to the growth of employment outsourcing. *Journal of Labor Economics* 21(1):1–42.
- Bailey K (1994) *Typologies and Taxonomies: An Introduction to Classification Techniques* (SAGE Publications).
- Bapna R, Ramaprasad J, Shmueli G, Umyarov A (2016) One way mirrors in online dating: A randomized field experiment. *Management Science* Articles in Advance.
- Bertrand M, Duflo E, Mullainathan S (2004) How much should we trust differences-in-differences estimates? *The Quarterly Journal of Economics* 119(1):249–275.
- Burtch G, Ghose A, Wattal S (2014) Cultural differences and geography as determinants of online prosocial lending. *MIS Quarterly* 38(3):773–794.
- Chu J, Manchanda P (2016) Quantifying cross and direct network effects in online consumer-to-consumer platforms. *Marketing Science* Articles in Advance.
- Couch D, Liamputtong P (2008) Online dating and mating: The use of the internet to meet sexual partners. *Qualitative Health Research* 18(2):268–279.
- Ellison N, Heino R, Gibbs J (2006) Managing impressions online: Self-presentation processes in the online dating environment. *Journal of Computer-Mediated Communication* 11(2):415–441.
- Fiore A, Donath J (2004) Online personals: An overview. *CHI'04 extended abstracts on Human factors in computing systems*, 1395–1398 (ACM).
- Grieser WD, Kapadia N, Li R, Simonov A (2016) Fifty shades of corporate culture. *Available at SSRN* <http://ssrn.com/abstract=2741049>.
- Griffin JM, Kruger SA, Maturana G (2016) Do personal ethics influence corporate ethics. *Available at SSRN* <http://ssrn.com/abstract=2745062>.
- Hatfield E, Rapson RL (1996) *Love and Sex: Cross-cultural Perspectives* (Allyn & Bacon).

- Heino RD, Ellison NB, Gibbs JL (2010) Relationshopping: Investigating the market metaphor in online dating. *Journal of Social and Personal Relationships* 27(4):427–447.
- Hitsch GJ, Hortacsu A, Ariely D (2010) Matching and sorting in online dating. *The American Economic Review* 100(1):130–163.
- Jackson E (2004) *The PayPal Wars: Battles with eBay, the Media, the Mafia, and the Rest of Planet Earth* (Washington, D.C.: World Ahead Publishing).
- Katana Z, Zubcsek P, Sarvary M (2011) Network effects and personal influences: The diffusion of an online social network. *Journal of Marketing Research* 48(3):425–443.
- Katz M, Shapiro C (1985) Network externalities, competition, and compatibility. *American Economic Review* 75(3):424–440.
- Lee Y, O'Connor GC (2003) New product launch strategy for network effects products. *Journal of the Academy of Marketing Science* 31(3):241–255.
- Lynch J, Ariely D (2000) Wine online: Search costs affect competition on price, quality, and distribution. *Marketing Science* 19(1):83–103.
- Meer J, Rigbi O (2013) The effects of transaction costs and social distance: Evidence from a field experiment. *The B.E. Journal of Economic Analysis & Policy* 13(1).
- Parker G, Van Alstyne M (2005) Two-sided network effects: A theory of information product design. *Management Science* 51(10):1494–1504.
- Parsons CA, Sulaeman J, Titman S (2016) The geography of financial misconduct. *Available at SSRN* <http://ssrn.com/abstract=2412970>.
- Pizzato L, Rej T, Chung T, Koprinska I, Kay J (2010) Recon: A reciprocal recommender for online dating. *Proceedings of the Fourth ACM Conference on Recommender Systems*, 207–214 (ACM).
- Rochet JC, Tirole J (2006) Two-sided markets: A progress report. *The RAND Journal of Economics* 37(3):645–667.
- Rysman M (2004) Competition between networks: A study of the market for yellow pages. *The Review of Economic Studies* 71(2):483–512.

-
- Suarez F (2005) Network effects revisited: The role of strong ties in technology selection. *Academy of Management Journal* 48(4):710–720.
- Sun M (2007) *Weaving a Two-Sided Network: Winning Strategies for Network Platform Competition* (VDM Publishing).
- Sun M, Tse E (2009) The resource-based view of competitive advantage in two-sided markets. *Journal of Management Studies* 46(1):45–64.
- Sundararajan A (2007) Local network effects and complex network structure. *The BE Journal of Theoretical Economics* 7(1):Article 46.
- Toma CL, Hancock JT, Ellison NB (2008) Separating fact from fiction: An examination of deceptive self-presentation in online dating profiles. *Personal and Social Psychology Bulletin* 34(8):1023–1036.
- Tucker C, Zhang J (2010) Growing two-sided networks by advertising the user base: A field experiment. *Marketing Science* 29(5):805–814.
- Wooldridge J (1999) Distribution-free estimation of some nonlinear panel data models. *The Journal of Econometrics* 90(1):77–97.