

Social Networks and Labor Markets: How Strong Ties Relate to Job Finding on Facebook's Social Network

Laura K. Gee, *Tufts University*

Jason Jones, *Stony Brook University*

Moira Burke, *Facebook*

Social networks are important for finding jobs, but which ties are most useful? Granovetter has suggested that “weak ties” are more valuable than “strong ties,” since strong ties have redundant information, while weak ties have new information. Using 6 million Facebook users’ data, we find evidence for the opposite. We proxy for job help by identifying people who eventually work with a pre-existing friend. Using objective tie strength measures and our job help proxy, we find that most people are helped through one of their numerous weak ties but a single stronger tie is significantly more valuable at the margin.

We would like to thank James Fowler, who has been instrumental throughout this project. We would also like to thank Jim Andreoni, Julie Cullen, Gordon Dahl, Philip Babcock, Shane Greenstein, Leah Lakdawala, Yannis Ioannides, Kelsey Jack, George Norman, Melissa McInerney, Marshall Van Alstyne, and numerous seminar participants for their helpful comments during the writing of this paper. A special thanks to the people at Facebook, especially Cameron Marlow and Dean Eckles. Contact the corresponding author, Laura K. Gee, at laura.gee@tufts.edu. Information concerning access to the data used in this article is available as supplementary material online.

[*Journal of Labor Economics*, 2017, vol. 35, no. 2]

© 2017 by The University of Chicago. All rights reserved. 0734-306X/2017/3502-0002\$10.00

Submitted April 27, 2015; Accepted January 13, 2016; Electronically published January 26, 2017

I. Introduction

Over 50% of jobs are found through a social tie, and social networks help explain observed labor market phenomenon like duration dependence and the socioeconomic, geographic, and racial concentration of unemployment.¹ Additionally, workers who find jobs via a social tie have lower turnover and higher productivity (Burks et al. 2015). Previous theoretical, experimental, and quasi-experimental work finds that if the prospects of the network improve, for example, with increased job vacancy information, then the job finding prospects of all network members improve on average (e.g., Calvo-Armengol and Jackson 2004; Bayer, Ross, and Topa 2008; Beaman 2012). Although there is an average improvement, very little is known about which network members benefit the most. A person's social network is made up of many social ties, and each tie may be of varying tie strength (e.g., a close friend is a stronger tie, and an acquaintance is a weaker tie). If social ties of differing strength benefit more, then this may result in the allocation of a job away from one person and toward another who has the right type of connection (weak or strong) but who is not necessarily a better worker. In this paper, we use data drawn from millions of Facebook users. We are unable to directly observe if these 6 million users were helped by any of their friends, so we create a proxy variable for job help by identifying users who eventually work at the same employer as a pre-existing friend. We find that a person is more likely to get help in obtaining a job at the same workplace as a pre-existing weak tie because collectively there are many weak ties in one's social network. However, when using information about all of a person's social ties, a single stronger tie has a higher probability of a shared workplace than does a weak tie.

This paper contributes to a body of research that seeks to answer the question, which type of social tie is most useful in job finding? Tie strength is a measure of how close two people are to each other. The previous research is mixed. Granovetter (1973) emphasizes the importance of weak ties, but more recent work finds in favor of strong ties (e.g., Kramarz and Skans 2014). In our paper, we measure tie strength either as the number of times two people interact in a year or as the number of mutual friends they share. It follows that a tie is weak if they have very few interactions or very few mutual friends. Using these tie strength measures, our findings reconcile the previous disparate results by showing that a user is less likely to eventually join the same workplace as an individual weak tie but that collectively weak ties are more important than strong ties because they are numerous.

We test two main hypotheses in this work. The first, the Descriptive Weak Ties Hypothesis, mirrors that from the empirical portion of Grano-

¹ See Marsden and Gorman (2001), Ioannides and Loury (2004), Jackson (2011), Munshi (2011), and Topa (2011).

vetter (1973) and asserts that most jobs are found through a weaker tie rather than a stronger tie. Our data support this hypothesis: the majority of job seekers end up working with weak ties. However, the distribution of tie strength in general is also highly skewed toward weak ties. So weak ties are important because people have many more weak ties than they do strong ties.

The second hypothesis we introduce is the Conditional Weak Ties Hypothesis. This states that there is an inverse relationship between the probability of working with a friend and the tie strength with that specific friend. We refute this hypothesis by finding a positive relationship between increased tie strength and increased probability of working together. Using a user-friend pair (“dyad”)–level analysis while controlling for individual-level heterogeneity with a user-level fixed effect and numerous dyad controls to mitigate dyad-level heterogeneity, we find that greater tie strength is associated with a greater probability of working with a friend. Because tie strength is endogenous, we also use a placebo test through which we artificially reverse the direction of who began work at the employer first. This puts a lower bound on the causal effect of increased tie strength on shared workplace. We find that there is a robust positive relationship between tie strength and shared workplace that refutes the Conditional Weak Ties Hypothesis. For example, increasing the amount of contact with a friend by 10 percentage points increases the likelihood that one will eventually work with that friend by at least 20%.

Social network ties are widely used in the job search process, either through formal referral programs or through more informal means, such as telling a friend about a job opening or helping an acquaintance prepare for an interview. In the United States, between 15% and 23% of workers reported using friends or relatives in their job search, over 50% of workers found their job through a network contact, and 70% of firms have programs encouraging referrals (Rees 1966; Granovetter 1973; Ioannides and Loury 2004; Topa 2011; Burks et al. 2015).

Although this paper will concentrate on the United States, the use of social ties to find jobs is pervasive in many countries, as documented by cross-country comparisons for 55 countries (Gee et al. 2017). From the previous work, we can conclude that networks are important in job allocation. However, there is still very little work exploring which specific types of connections in a network are most important. In fact, there is reason to believe that gains accrue differentially depending on the strength of the social tie.

The previous work is mixed on whether stronger or weaker ties are most important for job search.² In his seminal work, Granovetter (1973) empha-

² The evidence is also mixed on whether jobs found through stronger or weaker ties result in higher income (Tassier 2006).

sizes the importance of weak ties when measuring tie strength as self-reported contact between 54 recent job changers and the helping friend.³ Similarly, Yakubovich (2005) finds that self-reported weaker ties are most helpful in Russia, when looking only at connections considered or used during the job search. Granovetter (1973) theoretically models weak ties as bridging between different groups or across structural holes (Burt 2004), thus offering more novel job vacancy information.⁴ Calvo-Armengol and Jackson (2004) capture this same dynamic in a model in which when an unemployed person receives job vacancy information, he/she takes the job, but if an employed person receives the vacancy information, he/she randomly passes it on to an unemployed connection. This treatment results in more diversified information sets and a lower unemployment rate for those weak ties who bridge between disparate groups.⁵

Although Granovetter's "strength of weak ties" claim is prominent, others have pointed out the importance of strong ties. Boorman (1975) presents a theoretical model in which employed individuals communicate job vacancy information to weak ties only if their strong ties are already employed. He finds that if the probability of joblessness is high, it is optimal to have all strong ties. This is in line with the empirical finding that the poor and less educated rely more heavily on stronger ties (Granovetter 1983). This "strength of strong ties" is documented empirically in a number of studies for specific definitions of strong ties, such as geographic distance (Bayer et al. 2008), length of tenure at a previously shared workplace (Cingano and Rosolia 2012), and parental links (Kramarz and Skans 2014).

These differing findings may stem from differences in how tie strength is theoretically modeled and empirically measured. Much of the previous em-

³ Granovetter found that 26.7% of jobs came from a friend a person saw "often" (at least twice a week), 55.6% from someone seen "occasionally" (more than once a year but less than twice a week), and the remaining 27.8% from someone seen "rarely" (once a year or less).

⁴ Granovetter concentrates on the idea of weak ties as being more useful individually because of this novel information, but he also mentions that an alternative explanation for their importance could be "that most of any given person's ties are weak, so that we should expect, on a 'random' model, that most ties through which job information flows should be weak" (Granovetter 1973, 1372). However, because in Granovetter's study "baseline data on acquaintance networks are lacking, this objection remains inconclusive. Even if the premise were correct, however, one might still expect that greater motivation of close friends would overcome their being outnumbered." We interpret Granovetter's comment to mean that he expected weak ties would be most useful both collectively and individually.

⁵ In fact, there are a number of studies looking beyond job finding that support the idea that weak ties are most useful in the diffusion of information, which reinforces Granovetter's original "strength of weak ties" claim (e.g., see Lin, Dayton, and Greenwald 1978; Bakshy et al. 2012; Grabowicz et al. 2012; Ugander et al. 2012).

irical work has been limited by the scope of the data available. Many studies are firm-specific, use survey-reported tie strength that may suffer from self-reporting bias, or only have information for a subset of a person's full network (e.g., just the job helping connection but not all the other connections; see, e.g., Granovetter 1973; Simon and Warner 1992; Castilla 2005; Yakubovich 2005; Loury 2006; Tassier 2006; Beaman and Magruder 2012; Cappellari and Tatsiramos 2015; Brown, Setren, and Topa 2016). More representative samples have often been unable to measure true network ties, so instead researchers have used proxies such as geographic proximity (Topa 2001; Hellerstein, McInerney, and Neumark 2011; Schmutte 2015) or ethnic groups (Beaman 2012; Åslund, Hensvik and Skans 2014; Dustmann, Glitz, and Schönberg 2015). Alternately, studies with measures of true network ties have had to concentrate on only very specific types of ties, such as parental links (Munshi 2003; Babcock 2008; Magruder 2010; Kramarz and Skans 2014), school links (Marmaros and Sacerdote 2002; Shue 2013), or shared previous employer links (Bandiera, Barankay and Rasul 2009; Cingano and Rosolia 2012; Rider 2014). Our paper uses de-identified data from US Facebook users. Mayer (2012) also uses Facebook data, but only from Texas A&M students in 2005, and he performs an individual-level analysis in contrast to our user-friend-level analysis. The Facebook data offer a full map of a job seeker's Facebook social network, include many firms, and have non-self-reported continuous objective measures of tie strength for many types of ties (schoolmates, neighbors, etc.).

Theoretical models generally either make assumptions about ties of different strengths (Boorman 1975; Montgomery 1992; Zenou 2015) or they operationalize tie strength by the network structure of relationships (Calvo-Armengol and Jackson 2004; Ioannides and Soetevent 2006; Calvo-Armengol, Verdier, and Zenou 2007; Bramouille and Saint-Paul 2010). To mirror these modeling differences, we use two broad categories of tie strength: "contact-based" and "structure-based." *Contact-based measures* record the number of interactions between an individual and a friend and are similar to making an assumption about ties of different strengths. Our *structure-based measure* records the number of mutual friends, making it similar to models using the structure of the network to model tie strength and capturing the idea of bridging across different groups. Our paper is the first that we are aware of that uses both types of tie strength measures to explain labor market outcomes. We find that, although the contact-based and structure-based measures are correlated, they also have individual explanatory power.

Using these tie strength measures, we are able to show that collectively weak ties are most likely to result in working with a friend because weak ties are numerous in social networks. However, the probability of working with a friend is higher from a single strong tie than it is from a single weak tie. These results imply that strong ties are important to job finding. So policy

makers hoping to increase employment may want to encourage strengthening existing ties through mentorship programs. In addition, these findings inform future theoretical and empirical modeling by emphasizing the importance of the information used in the model and the differences in how tie strength is measured.

The next section introduces the data. Section III discusses the empirical strategy, Section IV presents the results, and Section V concludes.

II. Data

The data include de-identified information about people and their friends from the social networking website Facebook. A person's Facebook network is not an exact representation of his/her true network, and a large amount of unobservable contact takes place outside of Facebook. However, despite these shortcomings, Facebook interaction is a good predictor of real-world tie strength (Gilbert and Karahalios 2009; Jones et al. 2013).⁶

Facebook users are not a randomly selected sample of the US population. However, in the United States, over 54% of adults have a Facebook account, and 40% of social network users have "friended" their closest friends on social networking websites (Hampton et al. 2011; Bakshy et al. 2012; Burke and Kraut 2013). We restrict our analysis to users and friends who list employer information and who have been on Facebook for at least 1 year. These requirements allow us to measure if a person sequentially begins to work at the same firm as a friend, which we believe is a proxy for receiving help from that friend in the job search process (as explained in Subsecs. II.A and II.B). We only use US Facebook users. We remove common nonemployers that people list in their employer information, such as "stay at home parent" (see Appendix Section 1.1; the appendix is available online). We use those ages 16–64, so individuals are of working age. Finally, we restrict the sample to those who list some education on Facebook, because the lack of education may simply mean that a person has chosen not to self-report this information. These restrictions leave 6 million individuals, or about 4% of the US Facebook population. When we connect these 6 million individuals back to their friends, we have a total of 260 million dyads. Of these 6 million people, about 400,000 sequentially work with a pre-existing social tie, which we believe is a proxy for getting help from that tie (explained in Sec. II.A).

All our analysis involves de-identified data and takes place at the dyad level. This allows us to use a random sample of dyads, making an assump-

⁶ Gilbert and Karahalios (2009) asked 35 people to describe their tie strength with a random subset of all their Facebook friends, and they were able to use Facebook data to predict the survey-reported tie strength with 85% accuracy. More recently, Jones et al. (2013) asked over 700 people to report their closest friend, and they could predict the named friend using interactions on Facebook with 92% accuracy.

tion of independence across dyads more convincing.⁷ We created an approximately 1 million dyad subsample because of computational concerns by randomly selecting 3% of the approximately 400,000 users who end up working at a shared employer. Although we could have looked at both persons who end up working with a friend, as well as those who did not eventually work with a friend, we chose to restrict ourselves to people who end up working with a friend because we are most interested in identifying which ties are most helpful for those who have actually been helped. We are left with a sample of 12,263 individuals who, when connected with their friends, make up 1,438,699 dyads. The sample is weighted so that each user's weights sum to one, ensuring that individuals with many friends are treated similarly to individuals with very few friends. The distribution of tie strength measures for our subsample is very similar to the distribution for the full 260 million dyads (as shown in appendix figure 6). Due to our agreement with Facebook, we cannot report the exact demographics of the total US Facebook population. However, if we compare all US Facebook users to those in our subsample, those in our subsample are about a decade younger. We believe the age gap is driven by lack of employer information for those who joined Facebook earlier, since many people fill out this information when they first sign up, and this functionality was not always available on Facebook. Additionally, if we compare respondents of the Current Population Survey (CPS) to the those in our subsample, we find that Facebook users as a whole are about a decade younger, are more likely to have graduated from high school, and are less likely to be married. Although Facebook is not a nationally representative sample, it is a large and important group of individuals to study because of the increasing use of social media and online job search technologies.

A. Measuring Job Help

Our primary outcome variable is whether a person eventually works at the same employer as a pre-existing friend. We believe that this outcome variable is a proxy for being helped by a connection to find a job. A survey would accurately measure if a person was actually helped by a friend, but surveys often suffer from low response rates leading to a heavily selected sample. So we concentrate our analysis on whether a person works with a friend, a "sequential job," rather than use a survey-based outcome. We validate the sequential job outcome using a permutation test, a survey, and a number of robustness checks (see Sec. IV.C and Appendix Sections 1.2–1.3). People on Facebook can report their employment history, including their current and past employers, start date, and end date for each position

⁷ In our data, each dyad only occurs a single time. For example, we observe the dyad with user A and friend B only a single time. That is, we see only the dyad with A-B but not also the dyad containing B-A.

listed.⁸ Clearly, outcomes like wages would also be useful, but due to our data agreement with Facebook, we are not allowed to link individuals' data to wages by occupation. Additionally, job tenure would be an interesting outcome, but the lack of a self-reported end date does not necessarily mean that a person is still working at that firm; it could just show that the person has not updated his/her employer information recently. So, we concentrate on job help and define a "sequential job" as occurring when the following criteria are met:

1. The user and this friend currently work or previously worked at the same employer.
2. The user began working at the employer at least 1 year after his/her friend started at that employer.
3. The user and the friend were Facebook friends at least 1 year before the user started working at that employer.

The requirement that a person began work at least 1 year after his/her friend should exclude cases in which a person and his/her friend jointly apply to the same employer but the friend starts employment slightly before the person of interest. The requirement that the friend and the person of interest have been friends on Facebook for at least 1 year ensures that there is at least 1 year's worth of Facebook interactions. This requirement excludes most dyads who became friends during the interview process.⁹ While previous work has concentrated on a single helpful connection, our sequential job variable can measure if multiple friends help an individual. This definition of a sequential job requires us to restrict the sample to those who list an employer and have been on Facebook at least 1 year since the user's most recent start date. All Facebook data were anonymized and analyzed in aggregate.

Using this sequential job measure, we found that 7% of the 6 million users, about 400,000 users overall, were helped by at least one friend in finding their most recent job. Our sequential job definition is very conservative and cannot measure job help from non-Facebook friends. Additionally Facebook is not primarily a professional social networking website, so we are likely missing many purely professional connections; thus, we think of our 7% rate as a lower bound on the amount of job help in this popula-

⁸ See appendix figure 1 for a picture of how this information is recorded in Facebook.

⁹ Imagine that a person (who has not been helped by any pre-existing friends) interviews for a position at Tufts in September 2012 and that this person becomes Facebook friends with some of her interviewers at Tufts during the process. If the person begins the new job before September 2013, then this would not be a sequential job. Using this timing, we can exclude people who become friends with their interviewers.

tion. This may explain why our rate is much lower than the 50% rate found by previous work (Granovetter 1973; Topa 2011).

It is possible that our three sequential job criteria may be met accidentally, meaning two ties may eventually work together, but the friend was not helpful in getting their friend that job. To validate our sequential job variable, we ran a permutation test. Each user has an employer/start-date pair (e.g., Tufts and started September 2013). We randomly re-assigned these employer/start-date pairs without replacement to other users in the data and then checked if a sequential job still occurred for each user. The sequential job rate fell to 0.3% when we did this permutation, so we believe a substantial portion of the sequential jobs are nonrandom (see Appendix Section 1.3).

Also, we sent a survey to our subsample of 12,263 Facebook users, asking for the name of their employer and the name of the most helpful friend in their job search. There was a 7.5% response rate, and matching the survey-provided friends' names was possible for 63 respondents, or 7,497 dyads. Comparing the survey-named friend to the sequential job variable, we accurately identified a dyad as helpful or unhelpful 98% of the time (see Appendix Section 1.2). Although job help is a dyad-level variable, we also checked accuracy at the user level. For 63% of users, we identified the survey-named friend using our sequential job variable. Although this is not a very large sample and there are selection issues, we perform the same analysis presented in Section III using the survey-reported outcome variable, and the distributions are similar (appendix figure 7). Additionally, the direction of effects are always larger using the survey-reported outcome, and significance never fell below the 5% level (appendix table 3).

B. Measuring Tie Strength

Tie strength is often measured by amount of contact or by network structure. To closely match the measures used by Granovetter (1973), contact is measured between a user and a friend for the full year before the user started his/her most recent job, and network structure is measured by the number of mutual friends a year before the user's job starts. If a person started a job on June 1, 2011, then contact is measured from June 1, 2010 ($T = -365$) to the day before June 1, 2011 ($T = 0$).¹⁰ And for that same person, mutual friends with each connection are measured on June 1, 2010 ($T = -365$), as illustrated in figure 1.

Photo tags and wall posts are the contact-based tie strength measures. A photo tag occurs when a user marks a photo with a friend's name so the photo is easily located. Photo tags may be evidence of real-world interac-

¹⁰ This is meant to closely match Granovetter (1973), who "used the following categories for frequency of contact: often = at least twice a week; occasionally = more than once a year but less than twice a week; rarely = once a year or less."

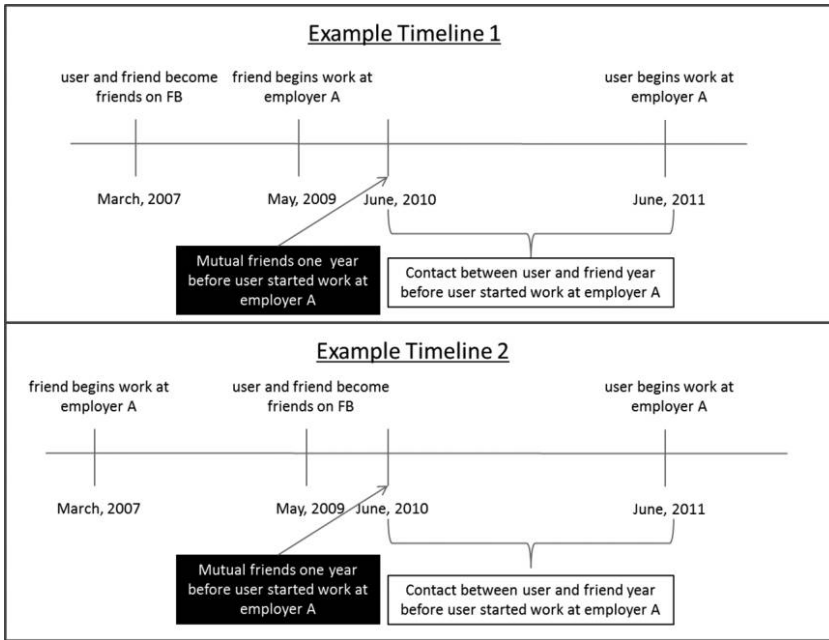


FIG. 1.—Sequential job examples. This figure shows two examples of the timeline of events where a user would be identified as having obtained a sequential job from a friend. We are ambivalent about the timing of Facebook friendship as compared to a friend’s start date, because the time of Facebook friendship is only observable from 2007 onward, whereas the start date can take values before 2007. The figure shows the time period contact-based tie strength is measured over, as well as the point in time mutual friends are measured. A color version of this figure is available online.

tion.¹¹ A wall post occurs when a user posts a message on the Facebook homepage (wall) of a friend. Although there are many measures of contact available for Facebook users, this analysis concentrates on tags and posts because they are good predictors of real-world friendships (Jones et al. 2013). We exclude more commonly used modes of contact (e.g., likes), which on a social network like Facebook may represent inconsequential rather than meaningful contact. We scale photo tags and wall posts to control for Facebook members who tag and post more often than others. Specifically, we use the percentage of tags, meaning the number of tags to a spe-

¹¹ For example, if user A and friend B are together, and A takes a photo of B, then it is likely that A will tag B in the photo when she uploads the photo to Facebook. See appendix figure 2 and appendix figure 3 for a picture of how photo tags and wall posts appear on Facebook.

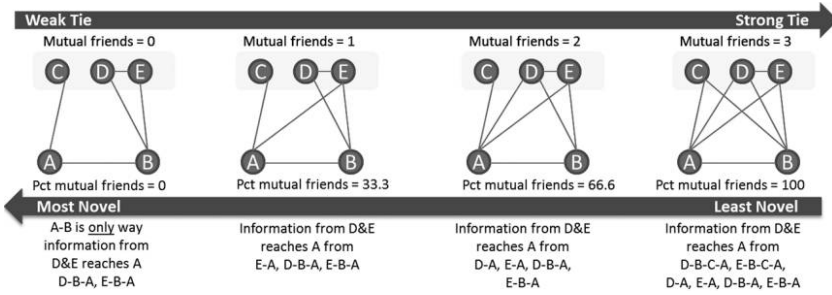


FIG. 2.—Network structure examples. The number of mutual friends shared by a user and a friend is the network structure-based measure of tie strength, and overlap is the scaled version of this measure. As overlap increases, so do the number of paths for information to flow from person A to person B. A color version of this figure is available online.

cific friend divided by total tags sent in the previous year, and we use an analogous definition for percentage of posts.¹²

The number of mutual friends shared by a user and a friend is the network-structure-based measure of tie strength. Our network structure measure is based on the number of friends who are mutual friends of person A and of person B, as illustrated in figure 2. For example, if Anne (A) and Bobby (B) have the exact same group of friends, they have 100% overlap, and if Anne and Bobby share no friends in common, they have 0% overlap. We chose this measure because it closely mirrors the spirit of Granovetter’s theoretical model and the idea of bridging across different groups.¹³

¹² Consider user A with friends B and C. If A tagged B two times last year and C six times last year, we would say that 25% ($2/(2 + 6)$) of user A’s tags were to B and that the remaining 75% were to C. The denominator uses friends we observe in our data rather than all Facebook friends (e.g., only US users, ages 16–64). At the mean start date (November 2010) for the 6 million US users with education and employer information, the modes of communication, from most to least used, were Comment, Like, Message, Wall Post, Tag, Poke, and Chat. In addition to scaled tie strength, we have done the same analysis using the raw number of tags and posts. For this test, the distributions are similar (appendix figure 4), and the regression results have the same sign and significance (appendix table 7).

¹³ This is meant to closely match Granovetter (1973, 1362), which tells the reader to consider “any two arbitrarily selected individuals—call them A and B—and the set, $S = C, D, E, \dots$, of all persons with ties to either or both of them. The hypothesis which enables us to relate dyadic ties to larger structures is: the stronger the tie between A and B, the larger the proportion of individuals in S to whom they will both be tied, that is, connected by a weak or strong tie. This overlap in their friendship circles is predicted to be least when their tie is absent, most when it is strong, and intermediate when it is weak.”

Granovetter also discusses the idea that all weak ties are bridges (where a bridge occurs if a user and a friend are connected by only one path between them).¹⁴ We did not define weak ties as bridges because the Facebook network is so densely connected that there are likely very few bridges. The average distance between US users is 4.3 links, and 99.7% of users are connected by fewer than 6 degrees of separation (Ugander et al. 2011).¹⁵ However, our network structure measure does partially address the idea of a bridge by showing how many paths information can flow through to get from person A to person B (the more mutual friends, the more paths). We present the number of mutual friends as a percentage of possible mutual friends (overlap) to control for Facebook members who are heavier users than others.¹⁶

C. Summary Statistics

Table 1 presents the user-level summary statistics, and table 2 contains the friend-level summary statistics. Both users and their friends are in their mid-20s, and the sample has slightly more women than men. The sample is well educated, with the highest level of education listed on their Facebook page as high school for 9%, college for 72%, and graduate school for 19%.

We characterize the relationship between the tie strength measures and the likelihood of a sequential job. A sequential job J_{ik} is defined as a dummy variable that takes the value $J_{ik} = 1$ if a dyad meets our sequential job criteria and takes the value $J_{ik} = 0$ otherwise. A user i with N friends has information for each friend: $J_{i,k=1}, J_{i,k=2}, \dots, J_{i,k=N}$. A sequential job $J_{ik} = 1$ is a rare

¹⁴ Granovetter pointed out this alternative measure of weak ties defined as a bridge. Specifically, Granovetter (1973, 1364) states: “A bridge between A and B provides the only route along which information or influence can flow from any contact of A to any contact of B, and, consequently, from anyone connected indirectly to A to anyone connected indirectly to B. Thus, in the study of diffusion, we can expect bridges to assume an important role.” He then goes on to state: “Intuitively speaking, this means that whatever is to be diffused can reach a larger number of people, and traverse greater social distance (i.e., path length), when passed through weak ties rather than strong” (1366).

¹⁵ Granovetter (1973, 1364) makes a similar point that in “large networks it probably happens only rarely, in practice, that a specific tie provides the only path between two points.”

¹⁶ Network overlap is defined as $O_{ik} = m_{ik} / (d_i - 1 + d_k - 1 - m_{ik})$, where m_{ik} is the number of mutual friends between i and k , d_i is the number of friends of person of i (degree of i), and d_k is defined analogously. The friend counts m_{ik} , d_i , and d_k are measured for all the friends on Facebook, not only those with our data restrictions (e.g., US, aged 16–64). Then $O_{ik} = 1$ if i and k have all the same friends, and $O_{ik} = 0$ if they have no friends in common. Granovetter (1973, 1362) states: “The stronger the tie between A and B, the larger proportion of individuals S to whom they will both be tied, that is, connected by a weak or strong tie. This overlap in their friendship circles is predicated to be least when their tie is absent, most when it is strong, and intermediate when it is weak.”

Table 1
User-Level Summary Statistics

Variable	Mean	SD	Minimum	Maximum
User's age	24.58	5.418	16	65
User is male	.475	.499	0	1
User has some high school	.091	.287	0	1
User has some college	.719	.45	0	1
User has some post-BA	.191	.393	0	1
User is married	.179	.384	0	1
User's friend count	521.039	345.134	1	4,653
All tags (1 year)	34.587	54.159	0	1,468
All posts (1 year)	42.962	55.513	0	945

NOTE.—Number of dyads = 1,438,699. Number of users = 12,263. Each dyad-level observation is weighted by 1/(number of times user is in data). User's friend count includes all Facebook friends, not just those meeting our data restrictions (e.g., United States, ages 16–64, with employer and education listed); in our sample, a user is connected to an average of 117 friends who met the data restrictions.

Table 2
Friend-Level Summary Statistics

Variable	Mean	SD	Minimum	Maximum
Friend's age	25.029	5.516	16	65
Friend is male	.472	.499	0	1
Friend has some high school	.089	.285	0	1
Friend has some college	.727	.446	0	1
Friend has some post-BA	.184	.388	0	1
Friend is married	.202	.402	0	1
Friend's friend count	499.386	380.163	1	6,329

NOTE.—Number of dyads = 1,438,699. Number of friends = 1,149,562. Each dyad-level observation is weighted by 1/(number of times friend is in data).

occurrence between dyads. Conditional on the user of interest i having at least one occurrence of $J_{ik} = 1$, the average level of sequential jobs is about 2% between these users and all their friends.¹⁷

Table 3 summarizes the statistics for sequential jobs and tie strength in the primary subsample. The average level of the sequential job rate J_{ik} is 2.038%. Only 4.2% of dyads have any tags between them, while many more, 16.1%, have some posts. On average, a dyad has 55 mutual friends. Almost all dyads have some friend overlap, and on average, a dyad shares about 5% of their friends, as measured by percentage of friend overlap. The average level of “Percentage tags” between a dyad (including zeros) is 0.621%, and this rises to 14.672% when excluding zeros. That is, if a user made 100 tags in

¹⁷ A dyad-level sequential job rate of 2% means that if a user i has a sequential job, then 2% of his/her friends, on average, meet the sequential job criteria. For example, if he/she had 100 friends, then on average he/she would eventually work with two friends at their shared workplace. The average number of friends a person is connected to in our sample is 117.

Table 3
Dyad-Level Job Help and Tie Strength Summary Statistics

Variable	Mean	SD	Minimum	Maximum	N
Sequential job rate	2.038	14.129	0	100	1,438,699
Any dyad tag	4.200	20.100	0	100	1,438,699
Count of tags	.280	3.744	0	974	1,438,699
Count of tags > 0 ^a	6.621	17.002	1	974	60,919
Percentage tags	.621	5.347	0	100	1,438,699
Percentage tags > 0 ^a	14.672	21.659	.082	100	60,919
Any dyad post	16.100	36.800	0	100	1,438,699
Count of posts	.366	2.19	0	627	1,438,699
Count of posts > 0 ^a	2.268	5.039	1	627	232,301
Percentage posts	.823	3.92	0	100	1,438,699
Percentage posts > 0 ^a	5.097	8.566	.106	100	232,301
Any friend overlap	98.600	11.700	0	100	1,438,699
Mutual friends	54.957	58.792	0	1303	1,438,699
Percentage mutual friends	5.127	5.108	0	100	1,438,699

NOTE.—Number of dyads = 1,438,699. All dummy variables are multiplied by 100 for ease of readability.

^a Average of tie strength, excluding dyads with zero tags or zero posts.

a year, on average, she would tag each friend about 0.621 times (0.621% of 100). The average level of “Percentage post” (including zeros) is 0.823%, and this rises to 5.097% when excluding zeros. Table 4 summarizes how similar, or homophilous, the dyads are. A correlation table for all the dyad-level variables is available in appendix table 1.

III. Empirical Strategy

To find the causal impact of tie strength on the likelihood of a sequential job from that friend, one would ideally observe an experiment that randomly assigns first networks and then tie strength. We have found no such experiment, so exploring the causal relationship between tie strength and a sequential job is difficult. The network itself is endogenously determined, and the level of tie strength between each dyad is also endogenous. We are most interested in the question of how tie strength affects sequential jobs, so we take the network as given.¹⁸ Our empirical strategy aims to address the second type of endogeneity; therefore, the results should be interpreted as how best to strengthen or weaken ties in a pre-existing network.

Our identification issues are closely related to those in the peer effects literature, so we use the terms “reflection,” “correlated effects,” and “context-

¹⁸ The previous work on exogenous variations of networks has concentrated on the effect of the network as a whole rather than the effect of tie strength between people in the network (Graham 2008; Sacerdote 2011; Beaman 2012; Mayer 2013). The network alone, ignoring tie strength, has an effect on outcomes. We believe that examining the effect of exogenous networks as well as the effect of exogenous tie strength are important next steps in this line of research.

Table 4
Dyad-Level Demographic Summary Statistics

Variable	Mean	SD	Minimum	Maximum
Friend years older (10)	.044	.453	-4.4	4.8
Both male	.251	.434	0	1
Both female	.289	.453	0	1
Friend more educated	.153	.36	0	1
Friend less educated	.192	.394	0	1
Both married	.056	.23	0	1
Same state	.477	.499	0	1
Same city	.177	.382	0	1
Same high school	.293	.455	0	1
Same college	.308	.462	0	1
Same grad school	.016	.126	0	1
Friend's tenure at firm (years)	1.121	1.773	-4.088	42.027

NOTE.—Number of dyads = 1,438,699. “Friend years older” is the number of years the friend is older than the user. The (10) on “Friends years older” means measured in 10 year increments, so average age difference is 4.4 years “Friend more educated” takes the value one if the friend has more education than the user. “Friend less educated” takes the value one if the friend has less education than the user. “Same city” takes the value one if the user and friend live in the same city within the same state. “Same high school,” “Same college,” and “Same grad school” take the value one if a user and friend list the same school on their Facebook profiles. “Friend's tenure at firm” is the number of years the friend had been employed at his firm before the user began her most recent job.

tual effects” to highlight this parallel, even though it does not map one-to-one (Manski 1993; Sacerdote 2011). The first candidate problem is “reflection,” which occurs if an individual's outcome is a function of the outcome of his/her network connections, and vice-versa. However, our sequential job measure does not suffer from reflection, because user A has a sequential job from friend B only if B has worked at their shared employer for at least 1 year. Thus, it is not possible for user A to have a sequential job from friend B and simultaneously for user B to have a sequential job from user A.¹⁹ Second, “correlated effects” occur when individuals select into peer groups in a way that is unobservable. Third, “contextual effects” point to the issue that, for both the individual and the peer, background characteristics of the peer group may affect the outcome variables.

We attempt to control for correlated and contextual effects using user-level and dyad-level controls and a number of robustness checks (see Sec. IV.C). We will also exploit the fact that a sequential job can only go from A to B or from B to A to construct a placebo test in which we reverse the direction of job assistance, which will allow us to put a lower bound on the causal effect of tie strength on likelihood of job help.

¹⁹ It is possible for user A's sequential job from B to affect future sequential jobs he/she may obtain or help others obtain. But, we take a random subsample at a given point in time, so this will not be recorded in our cross-sectional data. Additionally, the use of a nonlinear model mitigates the reflection issue (Blume et al. 2011). So, the fact that our results hold in nonlinear models (see appendix table 7) is further evidence that reflection is not an issue in our setting.

Keeping in mind the endogenous nature of our tie strength variables, we begin by explaining our empirical strategy for testing the Descriptive Weak Ties Hypothesis, which mirrors the analysis from the seminal paper (Granovetter 1973), and we then proceed beyond descriptive analysis to the empirical strategy for testing the Conditional Weak Ties Hypothesis.

Recall that Granovetter (1973) asked 54 recent job changers who found their job through a contact how close that tie was. He found that 26.7% of jobs came from a friend a person saw “often” (at least twice a week), 55.6% from someone seen “occasionally” (more than once a year but less than twice a week), and the remaining 27.8% from someone seen “rarely” (once a year or less). Thus, he concludes that most jobs collectively come from weak ties. This is in spite of the intuition that strong ties may be more motivated to help a friend. And Granovetter conjectures that the novelty of information available from weak ties makes an individual weak tie more useful than an individual strong tie.²⁰

A. Empirical Strategy for Descriptive Weak Ties Hypothesis

This section is purely descriptive in nature, to replicate the previous descriptive results. To test the Descriptive Weak Ties Hypothesis, we need to explore if most sequential jobs collectively come from weak ties rather than stronger ties. So, the empirical strategy for testing the Descriptive Weak Ties Hypothesis is to create a histogram of the proportion of sequential jobs that are transmitted from weaker versus stronger ties (recall our tie strength measure is continuous). To confirm the Descriptive Weak Ties Hypothesis, we would expect the histograms for each tie strength measure to be skewed toward the weak end of the continuum. The distribution of tie strength is endogenously determined. So we cannot make any causal inference about the effect of a change in this distribution on the level of the sequential job rate.

B. Empirical Strategy for Conditional Weak Ties Hypothesis

To test the Conditional Weak Ties Hypothesis, we need to explore if the probability of a sequential job from a single weak tie is higher than the probability of a sequential job from a single strong tie. So we will explore the relationship between the propensity that user i eventually works with a specific friend k (J_{ik}), and we will measure tie strength with that specific friend (T_{ik}).

1. Empirical Strategy for Conditional Weak Ties Hypothesis: Simple Specification

A simple specification would be

$$J_{ik} = \beta T_{ik} + c + \epsilon_{ik}. \quad (1)$$

²⁰ Granovetter does not have information about the total number of weak ties in a person’s network, but he posits that one might still expect that greater motivation of close friends would overcome their being outnumbered.

A user i with N friends has a sequential job dummy for each of those friends: $J_{i,k} = 1, J_{i,k} = 2 \dots J_{i,k} = N$. The dummy variable J_{ik} takes the value 1 if person i eventually works with a friend k , and T_{ik} represents our tie strength variables. The standard errors are clustered at the user i level.²¹

One might believe that there is a nonmonotonic relationship between the likelihood of a sequential job and tie strength; however, nonparametric models result in a roughly linear relationship (see appendix figure 9). Additionally, the dependent variable takes the values zero or one, so a logit model would be appropriate. However, we are most interested in the average probability of a sequential job, and the coefficients from the linear model are more easily interpreted. So, we ignore the special nature of the dependent variable and report the results from the linear model in text.²²

To confirm the Conditional Weak Ties Hypothesis, we would need the coefficients on the tie strength measures, β in equation (1), to be negative. However, even if we take the network as given, this is likely not a causal relationship because tie strength is endogenous, $E[T_{ik}\epsilon_{ik}] \neq 0$. So we move on to a more complex model to deal with the endogenous nature of tie strength.

2. Empirical Strategy for Conditional Weak Ties
Hypothesis: Improved Specification

In equation (2), J_{ik} is a sequential job, T_{ik} is tie strength, X_{ik} is a vector of dyad-level control variables, and E_i is the user fixed effect.

$$J_{ik} = \beta T_{ik} + \alpha X_{ik} + E_i + \epsilon_{ik}. \tag{2}$$

The user fixed effect, E_i , controls for all observable and unobservable attributes about the individual. For example, an extroverted individual may be more likely to have a sequential job and have higher levels of tie strength. With the inclusion of E_i , the variation in tie strength, T_{ik} , comes from variations within a user’s friendships instead of across all dyads.

Additionally, dyad-level variables may affect tie strength and the likelihood of a sequential job. Dyads with higher tie strength may be more likely to be similar (homophily), and so these dyads may be more likely to work at the same employers, even in the absence of actual help through their so-

²¹ Clearly ϵ_{ik} is not independent and identically distributed because we have multiple observations within each individual i that are not independent of each other, so we have clustered our standard errors at the user level. There is the additional worry that ϵ_{ik} for individual i may be correlated with $\epsilon_{j\ell}$; however, the use of a random subsample of only 1 million dyads mitigates this concern. The Facebook network is highly connected in general, so it is very possible for a random person i and another random person j to interact with each other. However, because we are only looking at a random selection of 12,000 individuals out of 400,000 who end up working at a shared employer (who in turn are from a total of 6 million individuals), we believe it is reasonably safe to assume that in our subsample ϵ_{ik} is uncorrelated with $\epsilon_{j\ell}$.

²² The sign and significance of the logit models are the same as those from the linear model (see appendix table 7).

cial networks. Ideally, we would use a dyad fixed effect to control for this source of omitted variable bias, but we only observe each dyad a single time. The Facebook data, however, have a rich set of dyad-level control variables that are predictive in decision making (Holzer 1987; Leicht and Marx 1997; Lin 1999; Aral and Walker 2012), X_{ik} , such as differences in gender, age, and education. We are able to control for all of these observable variables and the following additional variables: friend's tenure at firm at time of user's most recent start date, both married, same state, and same city. The full set of dyad-level controls are summarized in table 4.

Also, friend-specific unobservable attributes may affect the likelihood of a sequential job. For example, a friend may be the unofficial recruiter for his/her firm. This would increase the likelihood of a sequential job from that friend, but it is generally unobservable. Because we use a random sample of dyads, most friends only occur one time in the data. So we cannot use a friend fixed effect to control for friend-level unobservables.²³ We include a user fixed effect and dyad-level controls, which are computed from differences between the user- and friend-level variables.

To confirm the Conditional Weak Ties Hypothesis, we would need the coefficients on the tie strength measures, β in equation (2), to be negative. We would like to emphasize that these results may not be the true causal effects of tie strength on the likelihood of a sequential job. The principal concern is that tie strength T_{ik} is endogenous, $E[T_{ik}\epsilon_{ik}] \neq 0$. Even with this endogeneity, we can confidently make statements about the correlation between tie strength and the likelihood of a sequential job after controlling for unobservable individual heterogeneity and observable dyad heterogeneity.

It is not clear which direction our estimates of β will be biased. For example, a dyad made up of two economists who work in the same subfield might be more likely to have a sequential job and to be stronger ties with one another. Recall that we have controlled for the user's subfield by including a fixed effect, but if a same-subfield friend simultaneously increases tie strength and the likelihood of a sequential job, then the coefficient on tie strength, β , will be biased upward from the true causal β . Another scenario is that unproductive friends may be less likely to help obtain a sequential job but more likely to be strong ties. Again, we have controlled for the user being unproductive with a fixed effect, but β would be biased downward from the true causal β if unproductive friends simultaneously lower the likelihood of a sequential job and increase tie strength. If ϵ_{ik} is catching these types of confounding factors, our estimates of β may be biased away from the true causal effect of tie strength, and the direction of this bias is unclear. We cannot include a dyad fixed effect, since we observe each dyad a single time, but we can construct a placebo test.

²³ We have replicated the same analysis with a friend fixed effect for those friends who occur two or more times in the data. The results are of the same magnitude, sign, and significance.

3. Empirical Strategy for Conditional Weak Ties Hypothesis: Placebo Test

In our placebo test, we exploit the fact that a sequential job can only be transmitted in one direction because the transmitting friend must have worked at the firm for at least 1 year before the user of interest joins the firm. Our data record a sequential job from a pre-existing friend L to a user A if A joins L 's employer ($J_{A,L} = 1$). Suppose the user A had a total of three friends B , C , and L . We would use data from all three of these friends to test our Conditional Weak Ties Hypothesis ($J_{A,L}$, $J_{A,B}$, and $J_{A,C}$). We know that friend L must have started his/her job before user A because by the definition of a sequential job, L started at the firm at least a year before A . In our original data, $J_{A,L} = 1$. So our original estimating equation would be:

$$J_{i=A, k=L, B, C} = \beta_{\text{original}} T_{i=A, k=L, B, C} + \alpha_{\text{original}} X_{i=A, k=L, B, C} + E_{i=A} + \epsilon_{i=A, k=L, B, C} \tag{3}$$

Next, to create our placebo data, we artificially code $J_{L,A} = 1$; that is, we switch the order of sequential job transmission so that L received his/her job from A (even though this is impossible given the true timing of events). Then we connect the placebo sequential job transmittee (L) to all his/her friends, so if L had five friends, then we connect L to his/her friends A , M , N , O , and P ($J_{L,A}$, $J_{L,M}$, $J_{L,N}$, $J_{L,O}$, and $J_{L,P}$). We measure tie strength for L to his/her friends over the same time period, as in the original data. So our placebo estimating equation would be:

$$J_{i=L, k=A, M, N, O, P} = \beta_{\text{placebo}} T_{i=L, k=A, M, N, O, P} + \alpha_{\text{placebo}} X_{i=L, k=A, M, N, O, P} + E_{i=L} + \epsilon_{i=L, k=A, M, N, O, P} \tag{4}$$

In short, if, in the original data, user A has a sequential job from friend L ($J_{A,L} = 1$), in the placebo data, we artificially record the sequential job as coming from A to L instead ($J_{L,A} = 1$), even though that is impossible given the true timing of events. Then we perform the analysis using this placebo data.²⁴

²⁴ In this example, the original data could look like the following:

Original Data		
User of Interest	Friends	Sequential Job Dummy
A	L	1
A	B	0
A	C	0
Original rate = 33.3%		

And the placebo data could look like the following:

We compare the results using the placebo data to the results using the original data. If the coefficients using the original data are not statistically significantly different from the results using the placebo data, then unobservable dyad-level heterogeneity is driving the relationship between tie strength and the likelihood of a sequential job. If the coefficients using the placebo data are zero, this suggests that there is no effect of unobservable dyad-level heterogeneity. Finally, if the coefficients using the placebo data are nonzero but statistically significantly different from the coefficients using the original data, then that difference is attributable to dyad-level unobservable variables. One way to interpret the placebo coefficients is as the portion of the relationship between tie strength and the likelihood of a sequential job that is driven by unobservable dyad level variables. The difference between the placebo coefficients and the original coefficients is essentially like the coefficients from a specification including a dyad fixed effect.

The scaled difference between β_{original} and β_{placebo} is a lower bound on the causal portion of an attribute’s effect on the likelihood of a sequential job.²⁵ If β_{original} is statistically significantly smaller than β_{placebo} , then we will confirm the Conditional Weak Ties Hypothesis because this relationship implies that there is some negative causal effect of increased tie strength on the likelihood of a sequential job.

IV. Results

We begin by presenting results that show that collectively most people end up working with their weaker ties, so we confirm the Descriptive Weak

Placebo Data		
Placebo User of Interest	Placebo’s Friend	Placebo Sequential Job Dummy
L	A	1
L	M	0
L	N	0
L	O	0
L	P	0
Placebo rate = 20%		

²⁵ The most compelling way to control for endogeneity would be to experimentally assign tie strength to dyads. Secondly, we could use an exogenous shock to tie strength in an instrumental variables approach. We did find a possible instrument for number of mutual friends in our data in the form of a randomized experiment run by Facebook. We need an instrument to be both valid and relevant, so as a first step, we tested for relevance. That is, we tested if $E[T_{ik}Z_{ik}] \neq 0$, where T_{ik} is our measure of tie strength and Z_{ik} represents our instrumental variable. Unfortunately, the relationship between random assignment into the experimental treatment and tie strength was very weak, so we did not pursue this analysis any further.

Ties Hypothesis. Then we proceed with results that refute the Conditional Weak Ties Hypothesis by showing that individually stronger ties are consistently associated with a higher, rather than lower, probability of a sequential job from that specific tie.

A. Results for the Descriptive Weak Ties Hypothesis

In figure 3, we present the proportion of dyads with $N\%$ tagging, $N\%$ posting, and $N\%$ friends overlapping. For convenience, let us say that if the user of interest begins to work with a pre-existing friend, then the user is the “transmittee” and the friend is the “transmitter” of that specific sequential job. Then, the shaded bars represent dyads between a sequential job transmittee and the sequential job transmitter. For all three tie strength measures, the majority of sequential jobs came from a dyad with below-average tie strength, and this is clearly shown by the leftward skew of the distributions.²⁶

Also pictured in figure 3 is the distribution of tie strength between a sequential job transmittee and all his/her friends, both helpful and nonhelpful. The proportions for all friends are shown by the clear bars with a black outline. What is most striking about all these tie strength distributions is their similarity. Most sequential jobs come from a weak tie, but most ties in the population are weak. This means that the weak ties are collectively important because people have many weak ties.²⁷

There is support for the Descriptive Weak Ties Hypothesis: most sequential jobs will be found through a weaker tie rather than a stronger tie. In our data, the majority of sequential jobs are transmitted through a weaker tie. However, this is largely driven by the fact that most ties are weak. So weak ties are very useful collectively, but to see if they are also

²⁶ Ninety percent of sequential jobs were between a dyad with less than average percent tagging (average = 0.62%). Seventy-seven percent of sequential jobs were between a dyad with less than average percent posting (average = 0.82%). Sixty percent of sequential jobs were between a dyad with less than average percent friends overlapping (average = 5.12%).

²⁷ This type of tie distribution is not unique to Facebook or to our measure of tie strength. Networks are often characterized by many more weak ties than strong ties, whether measured by mobile phone usage (Onnela et al. 2007), academic coauthorship (van der Leij and Goyal 2011), Twitter usage (Grabowicz et al. 2012; Harrigan, Achananuparp, and Lim 2012), or Facebook usage (Bakshy et al. 2012; Ferrara et al. 2012). Although the distributions in fig. 3 are statistically significantly different from each other at the 15% level using a Wilcoxon signed-rank test, it is clear from a visual inspection that all distributions are characterized by many more weak than strong ties. For tagging: $z = -8.251$ $\text{Prob} > |z| = 0.000$. For posting: $z = -7.641$ $\text{Prob} > |z| = 0.000$. For friend overlap: $z = -1.519$ $\text{Prob} > |z| = 0.128$. All use a Wilcoxon signed-rank test. In the appendix, we show that the distribution of tie strength is similar for all 6 million users and 6 million dyads.

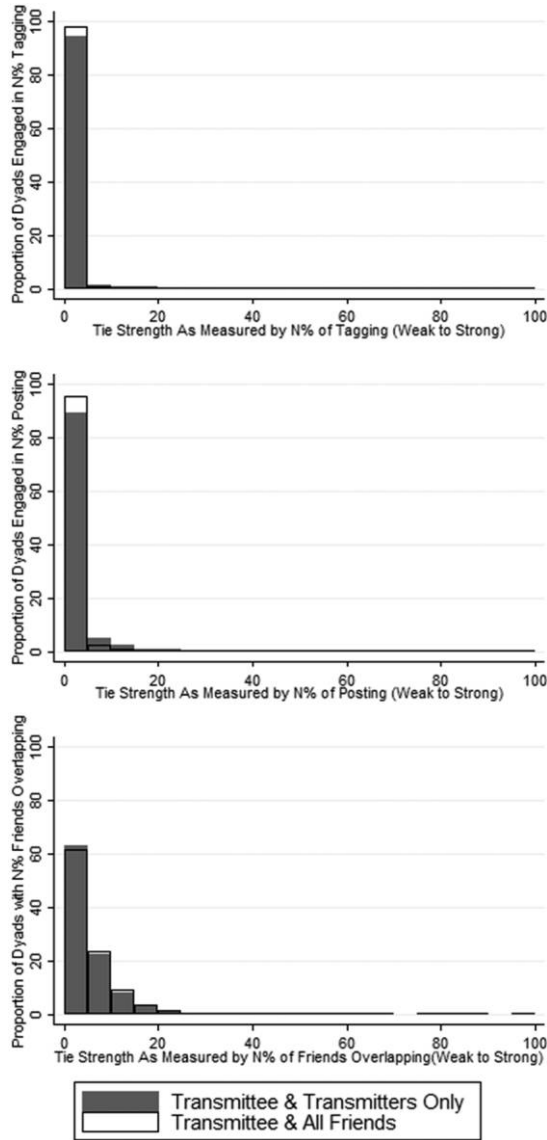


FIG. 3.—Distribution of tie strength (including 0). This figure shows the distribution of weak to strong ties. The shaded bars are for dyads between a sequential job transmittee (12,263) and the sequential job transmitters (29,319). The clear bars with black outlines are for dyads between the sequential job transmittee (12,263) and all their friends (1,438,699). The first panel uses percentage of tags from a user to a friend the year before the user began her most recent job as the measure of tie strength. The second panel uses percentage of posts from a user to a friend the year before the user began her most recent job as the measure of tie strength. The third panel uses the number of mutual friends as a percentage of possible mutual friends (overlapping friends) as the measure of tie strength. A color version of this figure is available online.

useful individually, we will next present the results of the tests of Conditional Weak Ties Hypothesis.

B. Results for the Conditional Weak Ties Hypothesis

The original theoretical motivation for the importance of weak ties is that they act as bridges that convey novel information. This suggests that weak ties should be individually more helpful than strong ties. To test this Conditional Weak Ties Hypothesis, we estimate the relationship between the likelihood of a sequential job and tie strength.

The mean level of the sequential job rate is 2.038%, so a person who has a sequential job usually shares that shared workplace with 2.038% of her friends in the data set (United States only, ages 16–64, employer/education listed). To confirm the Conditional Weak Ties Hypothesis, we expect a negative and significant coefficient, β , using either our simple model from equation (1) or our improved specification from equation (2). Table 5 shows the coefficients for the simple model in columns 1–3 and for the improved model in columns 4–6. The coefficients are all positive and statistically significant. This refutes the Conditional Weak Ties Hypothesis. These coefficients show that a single weak tie is actually less helpful than a single strong tie. In columns 1–6 of table 5, the coefficient β is the average percentage point difference in the likelihood of a sequential job attributable to a unit increase in T_{ik} , tie strength. A 1 unit increase would be a 1 percentage point increase in, for example, the amount of tags from i to k , which on average would be an increase of only one-third of a tag. Given that this is quite small, we will interpret the coefficients for a 10 percentage point increase. Because the simple specification is likely heavily influenced by omitted variable bias, let us begin by interpreting the coefficients from the improved specification as reported in columns 4–6. Column 4 implies that a dyad with a 10 unit increase in tags has a probability of a sequential job about .94 percentage points higher than a similar dyad, that is, a 46% proportional increase over the mean sequential job rate (2.038%). A 10 unit increase in posts is associated with a 60% proportional increase over the mean sequential job rate. And a 10 unit increase in friend overlap is associated with a 46% proportional increase over the mean sequential job rate.²⁸

²⁸ Although our tie strength measures are positively correlated, each contributes something distinct to our understanding. Tags may measure real-world contact, while posts measure online contact and mutual friends measure the structure of the network. Tags and posts measure contact in the past year, whereas the number of mutual friends is a long-term level of closeness, which is difficult for an individual to fully control. In appendix table 7, we show that a model with all three tie strength measures still has the same positive relationship between tie strength and the probability of a sequential job but that the magnitude of the effect decreases by about 23%.

Table 5
Dependent Variable Is Probability of a Sequential Job

	Simple Specification 1			Improved Specification 2			Placebo Test		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Tag	.126*** (.009)			.094*** (.008)			.021*** (.001)		
Post		.203*** (.012)			.124*** (.002)			.036*** (.001)	
Friend overlap			.001 (.011)			.093*** (.003)			.036*** (.002)
R ²	.002	.004	.000	.055	.056	.054	.006	.006	.006
N	1,438,699	1,438,699	1,438,699	1,438,699	1,438,699	1,438,699	2,241,023	2,241,023	2,241,023
Outcome mean \bar{J}_{it}	2.038	2.038	2.038	2.038	2.038	2.038	.839	.839	.839
Proportional ↑ if 10 unit ↑ (%)	62	100	.5	46	60	46	25	43	43
Original - Placebo							.073	.088	.057
Proportional ↑ if 10 unit ↑ Original - Placebo (%)							21	17	3

NOTE.—All coefficients are multiplied by 100 for ease of readability. Standard errors (in parentheses) are clustered at the user level and are weighted so that each user's weights sum to 1. Columns 4-9 include a user fixed effect and control variables: friend x years older, both male/female, friend more/less educated, both married, same state, same city, friend's tenure at firm (years), and same high school/college/graduate school. Full results for cols. 4-9 are in appendix table 2. The coefficients in col. 4 are different from those in col. 7 ($\chi^2(1) = 64.78$ Prob > $\chi^2 = .0000$). The coefficients in col. 5 are different from those in col. 8 ($\chi^2(1) = 27.29$ Prob > $\chi^2 = .0000$). The coefficients in col. 6 are different from those in col. 9 ($\chi^2(1) = 47.25$ Prob > $\chi^2 = .0000$). "Proportional ↑ if 10 unit ↑" shows the proportional increase in the probability of a sequential job associated with a 10 unit increase in tie strength. "Proportional ↑ if 10 unit ↑ Original - Placebo" shows the difference between the proportional increase using Improved Specification 2 on the original data versus using the Improved Specification 2 on the placebo data.

* $p < .05$.
 *** $p < .01$.
 *** $p < .001$.

Because we are not able to control for dyad-level heterogeneity, the results reported in columns 1–6 of table 5 may not be the true causal effect of tie strength on the probability of a sequential job. In an attempt to put a lower bound on the causal portion of the effect, we present the results of a placebo test for which we have artificially switched who worked at the employer first and rerun the same analysis using equation (2) on these placebo data. The scaled difference between β_{original} and β_{placebo} is a lower bound on the causal portion of an tie strength measures' effect on the likelihood of a sequential job. We have already refuted the Conditional Weak Ties Hypothesis by showing that individually weak ties are less helpful than strong ties—as evidenced by the positive and significant coefficients in columns 4–6 of table 5. If some portion of this is causal rather than simply a correlation, we expect the original coefficients (β_{original}) in columns 4–6 to be larger and statistically significantly different than those using the placebo data (β_{placebo}) in columns 7–9, which is exactly what we find. Column 7 implies that the causal effect of a 10 unit increase in tags is at minimum a 21% increase in the likelihood of a sequential job from that friend. A 10 unit increase in posts increases the likelihood of a sequential job by at least 17%. And a 10 unit increase in friend overlap increases the likelihood of a sequential job by at least 3%. So all our tie strength measures pass the placebo test, though the effect of increasing mutual friends on the likelihood of a sequential job becomes much smaller.

We conclude that there is strong evidence against the Conditional Weak Ties Hypothesis. Increasing tie strength as measured by contact (tags, posts) or network structure (friend overlap) is associated with increasing the propensity of eventually working with that specific friend. Furthermore, our placebo test shows that this is not totally driven by unobservable dyad-level attributes.

C. Robustness Checks

The previously presented results hold under a number of robustness checks.

1. Robustness Checks: Functional Form

In testing our Conditional Weak Ties Hypothesis, we use linear models. But, to explore more complex relationships in the data, we divide each tie strength measure into roughly equally sized bins (excluding zero) and estimate the coefficients from a linear model using these bins, a user fixed effect, and the control variables. When we plot the coefficients on these bins against tie strength, we find that the relationship is positive and generally linear.²⁹ Also, when we use a conditional logit model, the coefficients are

²⁹ See appendix figure 9. The underlying regression results are quite long, so they are available from Laura K. Gee (laura.gee@tufts.edu) by request.

of the same sign and significance as those presented from the linear model (see appendix table 7).

2. Robustness Checks: Mismeasurement of Sequential Jobs

The sequential job variable may include incidental occurrences of two people working at the same workplace when no actual job help took place, so we present the results from a number of robustness checks. As previously mentioned, we ran a permutation test and found that very few sequential jobs could be attributed to the chance occurrence, based on our sequential job criteria (Appendix Section 1.3). And when we test the Conditional Weak Ties Hypothesis using survey-reported job help, the coefficients on tie strength are still positive and significant, and they are actually larger than those from when we use the sequential job dependent variable (Appendix Section 1.2).³⁰

Incidental sequential jobs may be especially likely for large employers, but excluding large employers, we find that the distribution of friendships is still heavily skewed toward weaker ties.³¹ Also we find positive and significant coefficients with larger proportional magnitudes when we exclude large employers, that is, the top 25% of employers by size (appendix table 5).³²

Individuals who work for the same firm in different cities may be more likely to have an incidental sequential job. When we redefine our dependent variable J_{ik} to only take the value 1 if a user and a friend live in the same city, we find that most sequential jobs are still transmitted by a very weak tie. There is still a positive and significant relationship between tie strength and the likelihood of a sequential job, and the proportional magnitude of the coefficients is either larger or does not decline by more than 69%. Additionally, when we include distance between friends in the model, there is no change in the magnitude, sign, or significance of the coefficients on tie strength (appendix table 5).

We observe many sequential jobs at employers who do not generally require a college education, and one may believe that these jobs are more prone to incidental sequential jobs. If we restrict the analysis to only users with some college, we find that most sequential jobs still come from from a

³⁰ The average dyad-level survey-reported job help rate is only 0.8%. Using this as our dependent variable, we find that the coefficient (standard error) on “Percentage tag” is 0.279* (0.119), “Percentage post” is 0.528** (0.161), and “Percentage friend overlap” is 0.361*** (0.104).

³¹ The largest employer in our primary subsample was listed by 232 of the 12,263 users, which is only 1.8% of this sample.

³² When we exclude large employers, the average rate of sequential jobs falls from 2.038% to 1.75%, so even though the raw coefficients on tie strength fall, the proportional magnitudes are larger (e.g., the original model has a coefficient of .069 on tags, but that is a 3.3% rise over 2.038%, while for the smaller employer model, the coefficient is 0.07 on tags, but that is a 4.1% rise over 1.75%).

weaker tie. Additionally, if we restrict the analysis to only dyads for which both the user and the friend have some college, the sign and the significance of the coefficients on tie strength remain the same as the models in the text, with the proportional magnitude never decreasing by more than 7% (appendix table 5).³³

Certain industries may be especially prone to incidental sequential jobs, so we match the self-reported employer names from Facebook to industries. When we include either a dummy for a user and a friend being in the same industry or industry-specific fixed effects, all the tie strength coefficients remain positive and significant, and the proportional magnitude does not decrease by more than 17% (appendix table 5).

3. *Robustness Checks: Definition of Tie Existence and Tie Strength*

In this section, we show that the results are robust to varying definitions of tie existence and tie strength. A major concern is that dyads without any contact are too weak to be considered friends. So, we redefine a tie as only existing if some contact has occurred. Using this definition, we find that most users still get their sequential job through a very weak tie. And if we limit the sample to dyads with at least one tag or post in the previous year, we find positive and significant coefficients on tie strength, although they are only about half the size of the original coefficients (appendix figure 5 and appendix table 6).

A person may increase contact and tie strength in the hopes of obtaining a job from a friend. If this is the case, then we capture the effect of strategic tie strength rather than underlying tie strength on the likelihood of working together. In choosing the time frame for measuring contact-based tie strength, there is a trade-off between how current the measure of tie strength is and how likely the tie strength is strategically motivated. If we vary the time frame and the direction of contact measured, we find that the coefficients remain of the same sign and significance and the proportional magnitude never decreases by more than 38% (appendix table 6).³⁴

³³ It is also possible that high-skill jobs may be more prone to incidental sequential jobs because large employers hire at elite university campuses, so that two people who both went to elite colleges and were both recruited in this manner in subsequent years might appear to have exchanged information when in reality there was no such exchange. If we limit the random subsample to only users and friends who have only high school, then the sample falls to about 1,000 users and 21,000 dyads. With this small sample, it is difficult to make inferences. We find that most jobs still come from from a weaker tie and that the coefficients on the tie strength measures are still positive but are only significantly different from zero for the mutual friends measure.

³⁴ In the text, tags and posts are measured from a user to a friend during the year previous to the user's most recent start date. If a person started a job in September 1, 2012, then tags are measured from September 1, 2011 ($T = -365$) to the day before September 1, 2012 ($T = 0$). We have replicated the analysis using a 1 month lag. For

It is interesting that we find such robust support for a positive relationship between tie strength and the likelihood of a sequential job because there is a strong intuitive argument for the importance of novel information from a weak tie. As another proxy for novel information, we use whether a friend was formed through a pre-existing friend, and we find that non-friend-of-friends are generally more positively correlated with the likelihood of a sequential job but that the relationship is not statistically significant (Appendix Section 1.6 and appendix table 7).

V. Concluding Remarks

One of the most influential claims in the literature about social networks and labor markets has been Granovetter's "strength of weak ties" result. Granovetter's empirical work found that most jobs came from a weak rather than a stronger tie, but his data did not allow him to disentangle if this meant that weak ties are more useful collectively or individually or both. We re-examine that result in further depth using two different hypotheses based on the information that is available for analysis. We use a proxy for job help by identifying users who eventually work at the same employer as a pre-existing friend. And we use two objective tie strength measures: amount of contact and number of mutual friends. With these data, we find, like the original paper, support for the Descriptive Weak Ties Hypothesis, which finds that collectively most jobs came from weak rather than strong ties. We also find that the majority of job seekers began working with a weaker rather than a stronger tie. However, the distribution of tie strength in the population at large is also highly skewed toward weak ties, so weak ties are collectively important because weak ties are numerous in social networks.

We test our second hypothesis, the Conditional Weak Ties Hypothesis, that individually weak ties are more useful than strong ties. This is suggested by theories by which bridges provide access to better job vacancy information. We reject the hypothesis that a single weak tie is more helpful than a single strong tie by finding that an increase in tie strength is associated with an increase in the probability of working with a pre-existing friend at a shared workplace. This relationship is not driven by user-level unobservable variables or observable dyad-level variables, and it remains after a number of robustness checks. Furthermore, after using a placebo test to put a lower bound on the causal portion of this positive relationship, we find that

the 1 month lag, then tags are measured from August 1, 2011 ($T = -390 = -365 - 30$) to the day before August 1, 2012 ($T = -30$). Additionally, we compute tags from a friend to a user, and bi-directional tags. We have also run the same models measuring tags from 2 years to 1 year before the person began to work at the shared employer on a different subsample, and the results were still positive and significant.

strengthening a single weak tie will result in a higher probability of working with that specific friend.

In short, a person is most likely to eventually work with a weak tie because weak ties collectively make up most of a person's social network. However, strengthening an existing tie should increase the probability that one will work with that specific friend. This has implications for the formation of job finding networks and programs (Babcock et al. 2012). For example, if a university has an alumni network and the costs of pairing any two people together is the same, then the university should match dyads that they expect to be strong rather than weak ties. The ultimate research goal suggested by our study is to find the causal effect of network formation and tie strength on job finding. We believe that both laboratory experiments and large-scale field experiments that exogenously affect the structure of the network and the level of contact between friends are a natural extension of our work.

This paper can only speak about whether or not a job was obtained; due to the scope of our data, it cannot talk about the quality of that match. A natural extension would be to find a data set with both tie strength and long-term outcome variables such as wages or tenure to see if jobs found through weaker or stronger ties result in better matches.³⁵ Additionally, one may wonder if jobs found through different types of ties are more likely to be within one's current career path or pull one toward the profession of a friend. Unfortunately, the current data only provide a snapshot and so cannot identify job switching of this type. However, we believe that these are all important next steps in this area of research.

Previous work has shown that a majority of jobs are found through social ties and that those who find a job via social ties have higher productivity and longer tenure. This paper illustrates that whether strong or weak ties are more valuable in job search is a very nuanced question. The answer depends on the scope of the data used in the analysis. Contact and network structure-based measures need to be accounted for both empirically and theoretically. When looking at the collective power of weak ties, we find that weak ties matter most. But when we look individually at all of a person's social connections, we find that a single strong tie is more influential than a single weak tie. Weak

³⁵ Karlan et al. (2009) suggest that persons hired through the use of network connections should earn higher wages than those not using networks and that this wage gap should intensify with the skill intensity of the work. Unfortunately, we do not have data on wages. However, we can use highest level of schooling as a proxy for skill intensity of the vacancy. If we compare the results for the subsample with only a high school education to those with college or above, we find that that indeed the coefficients are larger for those with a college education than they are for those with only a high school education. But the high school-only sample is quite small, so the point estimates are very noisy.

ties are important in aggregate because they are numerous, while single strong ties are scarce but are associated with a higher probability of job help.

References

- Aral, Sinan, and Dylan Walker. 2012. Identifying influential and susceptible members of social networks. *Science* 337, no. 6092:337–41.
- Åslund, Olof, Lena Hensvik, and Oskar Nordström Skans. 2014. Seeking similarity: How immigrants and natives manage in the labor market. *Journal of Labor Economics* 32, no. 3:405–41.
- Babcock, Linda, William J., Congdon, Lawrence F. Katz, and Sendhil Mullainathan. 2012. Notes on behavioral economics and labor market policy. *IZA Journal of Labor Policy* 1, no. 1:1–14.
- Babcock, Philip. 2008. From ties to gains? Evidence on connectedness and human capital acquisition. *Journal of Human Capital* 2, no. 4:379–409.
- Bakshy, Eytan, Itamar Rosenn, Cameron Marlow, and Lada Adamic. 2012. The role of social networks in information diffusion. In *Proceedings of the ACM Conference on the World Wide Web*, Lyon, France.
- Bandiera, Oriana, Iwan Barankay, and Imran Rasul. 2009. Social connections and incentives in the workplace: Evidence from personnel data. *Econometrica* 77, no. 4:1047–94.
- Bayer, Patrick, Stephen L. Ross, and Giorgio Topa. 2008. Place of work and place of residence: Informal hiring networks and labor market outcomes. *Journal of Political Economy* 116, no. 6:1150–96.
- Beaman, Lori A. 2012. Social networks and the dynamics of labour market outcomes: Evidence from refugees resettled in the U.S. *Review of Economic Studies* 79, no. 1:128–61.
- Beaman, Lori, and Jeremy Magruder. 2012. Who gets the job referral? Evidence from a social networks experiment. *American Economic Review* 102, no. 7:3574–93.
- Blume, Lawrence E., William A. Brock, Steven N. Durlauf, and Yannis M. Ioannides. 2011. Identification of social interactions, vol. 1 of *Handbook of social economics*, ed. Jess Benhabib, Alberto Bisin, and Matthew O. Jackson, 853–964. Amsterdam: Elsevier.
- Boorman, Scott A. 1975. A combinatorial optimization model for transmission of job information through contact networks. *Bell Journal of Economics* 6, no. 1:216–49.
- Bramoulle, Yann, and Gilles Saint-Paul. 2010. Social networks and labor market transitions. *Labour Economics* 17, no. 1:188–95.
- Brown, Meta, Elizabeth Setren, and Giorgio Topa. 2016. Do informal referrals lead to better matches? Evidence from a firm's employee referral system. *Journal of Labor Economics* 34, no. 1:161–209.
- Burke, Moira, and Robert Kraut. 2013. Using Facebook after losing a job: Differential benefits of strong and weak ties. In *Proceedings of the 2013*

- Conference on computer supported cooperative work*, 1419–30. New York: ACM.
- Burks, Stephen V., Bo Cowgill, Mitchell Hoffman, and Michael Housman. 2015. The value of hiring through employee referrals. *Quarterly Journal of Economics* 130, no. 2:805–39.
- Burt, Ronald S. 2004. Structural holes and good ideas. *American Journal of Sociology* 110, no. 2:349–99.
- Calvo-Armengol, Antoni, and Matthew O. Jackson. 2004. The effects of social networks on employment and inequality. *American Economic Review* 94, no. 3:426–54.
- Calvo-Armengol, Antoni, Thierry Verdier, and Yves Zenou. 2007. Strong and weak ties in employment and crime. *Journal of Public Economics* 91, no. 12:203–33.
- Cappellari, Lorenzo, and Konstantinos Tatsiramos. 2015. With a little help from my friends? Quality of social networks, job finding and job match quality. *European Economic Review* 78, no. C:55–75.
- Castilla, Emilio J. 2005. Social networks and employee performance in a call center. *American Journal of Sociology* 110, no. 5:1243–83.
- Cingano, Federico, and Alfonso Rosolia. 2012. People I know: Job search and social networks. *Journal of Labor Economics* 30, no. 2:291–332.
- Dustmann, Christian, Albrecht Glitz, Uta Schönberg, and Herbert Brücker. 2015. Referral-based job search networks. *Review of Economic Studies* 83, no. 2:514–46.
- Ferrara, Emilio, Pasquale De Meo, Giacomo Fiumara, and Alessandro Provetti. 2012. The role of strong and weak ties in Facebook: A community structure perspective. *International Journal of Social Network Mining* 1, no. 1:67–90.
- Gee, Laura K., Jason Jones, Christopher J. Fariss, Moira Burke, and James Folwer. 2017. The paradox of weak ties in 55 countries. *Journal of Economic Behavior and Organization* 133 (January): 362–72.
- Gilbert, Eric, and Karrie Karahalios. 2009. Predicting tie strength with social media. *CHI '09 Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 211–20. New York: ACM.
- Grabowicz, Przemyslaw A., José J. Ramasco, Esteban Moro, Josep M. Pujol, and Victor M. Eguiluz. 2012. Social features of online networks: The strength of intermediary ties in online social media. *PLoS ONE* 7 (1), e29358.
- Graham, Bryan S. 2008. Identifying social interactions through conditional variance restrictions. *Econometrica* 76, no. 3:643–60.
- Granovetter, Mark S. 1973. The strength of weak ties. *American Journal of Sociology* 78, no. 6:1360–80.
- . 1983. The strength of weak ties: A network theory revisited. *Sociological Theory* 1, no. 1:201–33.

- Hampton, Keith, Lauren Sessions Goulet, Lee Rainie, and Kristen Purcell. 2011. Social networking sites and our lives. Pew Internet and American Life Project, Pew Research Center.
- Harrigan, Nicholas, Palakorn Achananuparp, and Ee-Peng Lim. 2012. Influentials, novelty, and social contagion: The viral power of average friends, close communities, and old news. *Social Networks* 34 (4):470–80.
- Hellerstein, Judith K., Melissa McInerney, and David Neumark. 2011. Neighbors and coworkers: The importance of residential labor market networks. *Journal of Labor Economics* 29, no. 4:659–95.
- Holzer, Harry J. 1987. Informal job search and black youth unemployment. *American Economic Review* 77, no. 3:446–52.
- Ioannides, Yannis M., and Linda Datcher Lounsbury. 2004. Job information networks, neighborhood effects, and inequality. *Journal of Economic Literature* 42, no. 4:1056–93.
- Ioannides, Yannis M., and Adriaan R. Soetevent. 2006. Wages and employment in a random social network with arbitrary degree distribution. *American Economic Review* 96, no. 2:270–74.
- Jackson, Matthew O. 2011. An overview of social networks and economic applications. In *Handbook of social economics*, vol. 1, ed. Jess Benhabib, Alberto Bisin, and Matthew O. Jackson, 511–85. Amsterdam: Elsevier.
- Jones, Jason J., Jaime E. Settle, Robert M. Bond, Christopher J. Fariss, Cameron Marlow, and James H. Fowler. 2013. Inferring tie strength from online directed behavior. *PLOS ONE*, January.
- Karlan, Dean, Markus Mobius, Tanya Rosenblat, and Adam Szeidl. 2009. Trust and social collateral. *Quarterly Journal of Economics* 124, no. 3: 1307–61.
- Kramarz, Francis, and Oskar Nordström Skans. 2014. When strong ties are strong: Networks and youth labour market entry. *Review of Economic Studies* 81, no. 3:1164–1200.
- Leicht, Kevin T., and Jonathan Marx. 1997. The consequences of informal job finding for men and women. *Academy of Management Journal* 40, no. 4:967–87.
- Lin, Nan. 1999. Social networks and status attainment. *Annual Review of Sociology* 25:467–87.
- Lin, Nan, Paul W. Dayton, and Peter Greenwald. 1978. Analyzing the instrumental use of relations in the context of social structure. *Sociological Methods and Research* 7, no. 2:149–66.
- Lounsbury, Linda Datcher. 2006. Some contacts are more equal than others: Informal networks, job tenure, and wages. *Journal of Labor Economics* 24, no. 2:299–318.
- Magruder, Jeremy R. 2010. Intergenerational networks, unemployment, and persistent inequality in South Africa. *American Economic Journal: Applied Economics* 2, no. 1:62–85.

- Manski, Charles F. 1993. Identification of endogenous social effects: The reflection problem. *Review of Economic Studies* 60, no. 3:531–42.
- Marmaros, David, and Bruce Sacerdote. 2002. Peer and social networks in job search. *European Economic Review* 46, nos. 4–5:870–79.
- Marsden, Peter V., and Elizabeth H. Gorman. 2001. Social networks, job changes, and recruitment. In *Sourcebook of labor markets*, ed. Ivar Berg, Arne L. Kalleberg, Ivar Berg, and Arne L. Kalleberg. Plenum Studies in Work and Industry. New York: Springer.
- Mayer, Adalbert. 2012. The structure of social networks and labour market success. *Applied Economics Letters* 19, no. 13:1271–74.
- . 2013. First friends, then co-workers: An empirical investigation of social connections and the initial employment of college graduates. Working paper (January), Washington College.
- Montgomery, James D. 1992. Job search and network composition: Implications of the strength-of-weak-ties hypothesis. *American Sociological Review* 57, no. 5:586–96.
- Munshi, Kaivan. 2003. Networks in the modern economy: Mexican migrants in the US labor market. *Quarterly Journal of Economics* 118, no. 2: 549–99.
- . 2011. Labor and credit networks in developing economies. In *Handbook of social economics*, vol. 1, ed. Jess Benhabib, Matthew O. Jackson, and Alberto Bisin, 1224–54. Amsterdam: Elsevier.
- Onnela, Jukka-Pekka., Jari Saramški, J. Hyvšnen, György Szabó, David Lazer, Kimmo Kaski, János Kertész, and Albert-Laszlo Barabási. 2007. Structure and tie strengths in mobile communication networks. *Proceedings of the National Academy of Sciences* 104, no. 18:7332–36.
- Rees, Albert. 1966. Information networks in labor markets. *American Economic Review* 56, nos. 1–2:559–66.
- Rider, Christopher I. 2014. Educational credentials, hiring, and intra-occupational inequality: Evidence from law firm dissolutions. Paper presented at the annual meeting of the American Sociological Association, San Francisco, August 15.
- Sacerdote, Bruce. 2011. Peer effects in education: How might they work, how big are they, and how much do we know thus far? In *Handbook of the economics of education*, vol. 3, ed. Eric A. Hanushek, Stephen Machin, and Ludger Woessmann, 249–77. Amsterdam: Elsevier.
- Schmutte, Ian M. 2015. Job referral networks and the determination of earnings in local labor markets. *Journal of Labor Economics* 33, no. 1:1–32.
- Shue, Kelly. 2013. Executive networks and firm policies: Evidence from the random assignment of MBA peers. *Review of Financial Studies* 26, no. 6: 1401–42.
- Simon, Curtis J., and John T. Warner. 1992. Matchmaker, matchmaker: The effect of old boy networks on job match quality, earnings, and tenure. *Journal of Labor Economics* 10, no. 3:306–30.

- Tassier, Troy. 2006. Labor market implications of weak ties. *Southern Economic Journal* 72, no. 3:704–19.
- Topa, Giorgio. 2001. Social interactions, local spillovers and unemployment. *Review of Economic Studies* 68, no. 2:261–95.
- . 2011. Labor markets and referrals. In *Handbook of social economics*, vol. 1, ed. Jess Benahbib, Matthew O. Jackson, and Alberto Bisin, 1193–1221. Amsterdam: Elsevier.
- Ugander, Johan, Lars Backstrom, Cameron Marlow, and Jon Kleinberg. 2012. Structural diversity in social contagion. *Proceedings of the National Academy of Sciences* 109, no. 16:5962–66.
- Ugander, Johan, Brian Karrer, Lars Backstrom, and Cameron Marlow. 2011. The anatomy of the Facebook social graph. *arXiv preprint*, arXiv:1111.4503.
- van der Leij, Marco, and Sanjeev Goyal. 2011. Strong ties in a small world. *Review of Network Economics* 10, no. 2: article 1.
- Yakubovich, Valery. 2005. Weak ties, information, and influence: How workers find jobs in a local Russian labor market. *American Sociological Review* 70, no. 3:408–21.
- Zenou, Yves. 2015. A dynamic model of weak and strong ties in the labor market. *Journal of Labor Economics* 33, no. 4:891–932.