

REPORT

SOCIAL NETWORKS

A causal test of the strength of weak ties

Karthik Rajkumar¹, Guillaume Saint-Jacques¹, Iavor Bojinov², Erik Brynjolfsson^{3,4}, Sinan Aral^{5*}

The authors analyzed data from multiple large-scale randomized experiments on LinkedIn's People You May Know algorithm, which recommends new connections to LinkedIn members, to test the extent to which weak ties increased job mobility in the world's largest professional social network. The experiments randomly varied the prevalence of weak ties in the networks of over 20 million people over a 5-year period, during which 2 billion new ties and 600,000 new jobs were created. The results provided experimental causal evidence supporting the strength of weak ties and suggested three revisions to the theory. First, the strength of weak ties was nonlinear. Statistical analysis found an inverted U-shaped relationship between tie strength and job transmission such that weaker ties increased job transmission but only to a point, after which there were diminishing marginal returns to tie weakness. Second, weak ties measured by interaction intensity and the number of mutual connections displayed varying effects. Moderately weak ties (measured by mutual connections) and the weakest ties (measured by interaction intensity) created the most job mobility. Third, the strength of weak ties varied by industry. Whereas weak ties increased job mobility in more digital industries, strong ties increased job mobility in less digital industries.

The Strength of Weak Ties (1) is one of the most influential social theories of the past century, underpinning networked theories of information diffusion (2, 3), social contagion (4, 5), social movements (6), industry structure (7), influence maximization (8), and human cooperation (9, 10). It argues that infrequent, arms-length relationships, known as “weak ties,” provide more new employment opportunities (11), promotions and greater wage increases (12), creativity (13), innovation (14, 15), productivity (16), and performance (17) because they deliver more novel information than strong ties. Weak ties are thought to provide access to diverse, novel information because they connect us to disparate and diverse parts of the human social network (18–24). In addition to productivity, performance, innovation, and other benefits, weak ties are thought to be specifically well suited to deliver new employment opportunities because they provide novel labor market information, making job mobility a centerpiece of the original weak tie theory.

Recent large-scale correlational investigations of the weak tie hypothesis, however, have uncovered a seeming “paradox of weak ties,” suggesting that strong ties are more valuable than weak ties in generating job transmissions (25, 26). Though these are the largest, most direct empirical examinations of the weak tie hypothesis to date, because the

work is not experimental the authors rightfully acknowledge that their results “may not be the true causal effect of tie strength on the probability of a sequential job.” More generally, two empirical challenges have prevented robust causal tests of the weak tie theory to date: First, a lack of large-scale data linking human social networks to job transmission makes measurement of the relationship between weak ties and labor market outcomes difficult. Second, network ties and labor market outcomes are endogenous, making the causal link between weak ties and job placement elusive. Individuals' labor market outcomes are likely to be determined by and to simultaneously determine their social networks. The evolution of social networks and job trajectories are also likely correlated with unobserved factors such as effort, ability, and sociability, which confound empirical identification of the link between weak ties and jobs.

We address these two empirical challenges and provide an experimental causal test of the weak tie theory with data from multiple large-scale randomized experiments on LinkedIn, the world's largest professional social network. The experiments randomly varied the prevalence of strong and weak ties in the professional networks of over 20 million LinkedIn members by adjusting the platform's People You May Know (PYMK) algorithm, which recommends new connections to members (Fig. 1A illustrates the experimental design). LinkedIn's PYMK algorithm is an ensemble machine learning model comprising the following: (i) a model for estimating the propensity of an ego (i.e., a focal member) to send a connection invite to an alter (i.e., a member the focal member is not currently connected with), (ii) a model estimating the alter's propensity to accept an

invite from the ego, (iii) a model estimating the engagement between the ego and alter once connected and (iv) weights on each of these models for relative importance. The experiments tuned these components, introduced new data sources, and relied on the number of mutual connections between the ego and a potential tie recommendation as one of the most important features of the ensemble model to randomly vary weak and strong tie recommendations. We performed a retrospective analysis of the randomization created by the PYMK experiments conducted by LinkedIn between 2015 and 2019 in two waves.

The first wave examined a global experiment conducted in 2015 that had over 4 million experimental subjects and created over 19 million new connections. We collected edge-level observations of tie strength and job transmission outcomes for each tie created during this experiment. We then analyzed a larger second wave of node-level PYMK experiments that took place worldwide in 2019. The second wave spanned every continent and US state, had more than 16 million experimental subjects, created ~2 billion new connections and recorded more than 70 million job applications that led to 600,000 new jobs during the experimental period (Fig. 1, B and C). The data were collected both at the node level (in 2019), where each observation corresponds to a unique LinkedIn member, and at the edge level (in 2015), where each observation corresponds to a unique tie between two LinkedIn members (see Fig. 1A for a description of how we compiled the edge- and node-level datasets).

We analyzed labor market mobility by measuring both job applications and job transmissions. Job applications are simply the number of jobs LinkedIn members applied to on the platform in the three months after an experiment. In accordance with the literature (25, 26), we consider a job transmission to have occurred when three criteria are satisfied: First, user *A* reports working at company *c* at date *D*₁. Second, user *B* reports working at that same company *c* at a later date *D*₂, with *D*₂ and *D*₁ being at least one year apart. Third, user *A* and user *B* were friends on the social network at least one full year before *D*₂. In the weak tie literature, when these three criteria are met, a tie is considered a “sequential job” tie, which represents the state of the art in measuring relational job mobility.

We measured tie strength by its two leading indicators: the intensity of the interaction between two people and the number of mutual connections they had in common. We measured interaction intensity by counting the number of interactions LinkedIn members had with one another through bilateral messaging. We measured mutual friendship by counting the number of friends any two connected individuals had in common when their tie was

¹LinkedIn Corporation, 700 E Middlefield Rd, Mountain View, CA 94043, USA. ²Harvard Business School, Soldiers Field, Boston, MA 02163, USA. ³Stanford University, 353 Jane Stanford Way, Stanford, CA 94305, USA. ⁴National Bureau of Economic Research, 1050 Massachusetts Ave., Cambridge, MA 02138, USA. ⁵Massachusetts Institute of Technology, MIT Sloan School of Management, 100 Main St., Cambridge, MA 02142, USA. *Corresponding author. Email: sinan@mit.edu

created. Structural tie strength, based on mutual friendship, was then defined bidirectionally as follows:

$$\text{StructuralTieStrength}_{ij} = \frac{M_{ij}}{D_i + D_j - M_{ij} - 2}$$

where i and j are LinkedIn members, M_{ij} is the number of mutual connections between them, and D_i and D_j are the total number of direct connections of members i and j , respectively. Network diversity is defined as $1 - C_i$, where C_i is the local clustering coefficient (formally defined in the SM).

Because tie strength changes in response to one's own friending behavior and the friending behavior of one's connections, we measured structural tie strength pretreatment and examined the causal effect of adding a new connection whose pretreatment tie strength was either strong or weak depending on whether it was above or below the median of the pre-treatment tie strength distribution (Fig. 2C). Interaction intensity is observed once a new tie is created. We therefore measured interaction intensity during the experimental period after ties were formed. Job seekers have more connections (greater degree) (Fig. 2A) and greater network diversity at higher degrees (Fig. 2B). But because these network variables are endogenously determined in observational data, random variation in LinkedIn members networks is necessary for a robust causal assessment of the relationship between weak ties and job mobility.

We estimated the causal effects of strong and weak ties on job mobility with an instrumental variables (IV) approach (27–29). The IV framework disentangles endogeneity by using random variation created by exogenous treatment assignments as a shock to endogenous counts of newly created weak and strong ties to estimate their causal effect on job mobility. We estimated these effects in a two-stage least squares (2SLS) specification, using the random assignment of members to weak- or strong-tie experimental variants as instruments for identifying the effect of adding weak or strong ties on job applications and job transmissions.

We conducted both node-level and edge-level analyses of the relationship between tie strength and job mobility. Node-level analyses estimated the effect of the number of weak or strong ties created by the experiments on job applications and job mobility. Though the node-level analysis estimates how assignments to weak- or strong-tie-inducing experimental treatments created changes in job mobility, it obfuscates which weak or strong ties led to job transmission. We therefore also conducted edge-level analyses to estimate the marginal effect of adding strong or weak ties to members' networks on their subsequent job mobility.

The first stage regressions estimated the effects of assignments to different experimental

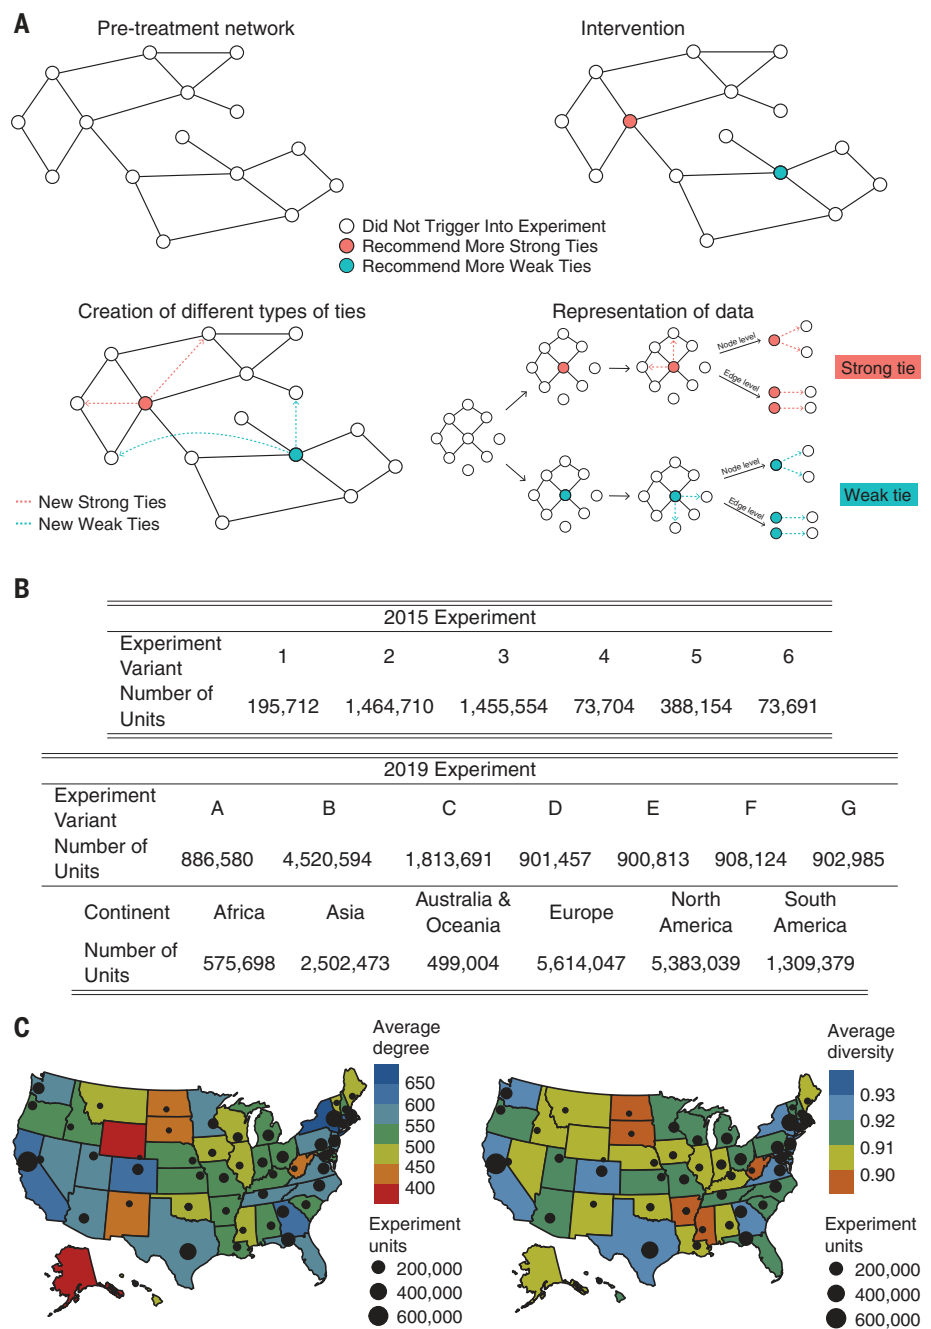


Fig. 1. Experiment design and summary statistics. (A) describes the experimental design and representation of the resultant data in node- and edge-level analyses; (B) displays the number of experimental units in the 2015 and 2019 experiments by continent and experimental variant (98.8% of the 2015 data was from the U.S.); and (C) displays the average degree, network diversity (formally defined in the SM), and number of experimental units by U.S. state in our 2019 experiments.

treatment variants on the creation of weak and strong ties in LinkedIn members' professional networks. The results of these first stage regressions, shown in Fig. 2D, demonstrate the random variation created by our experiments by displaying the effects of the experimental treatment variants on the creation of weak and strong ties between members in the LinkedIn network. As the figure shows,

some treatment variants caused members to form more weak ties (e.g., variants A and E), whereas others caused members to form fewer weak ties (e.g., variants F and G). The different variants also caused members to create more ties (e.g., variants A, C, and E), fewer ties (e.g., variant G), or approximately the same number of ties (e.g., variants B, D, and F), allowing us to distinguish the causal effect of tie strength

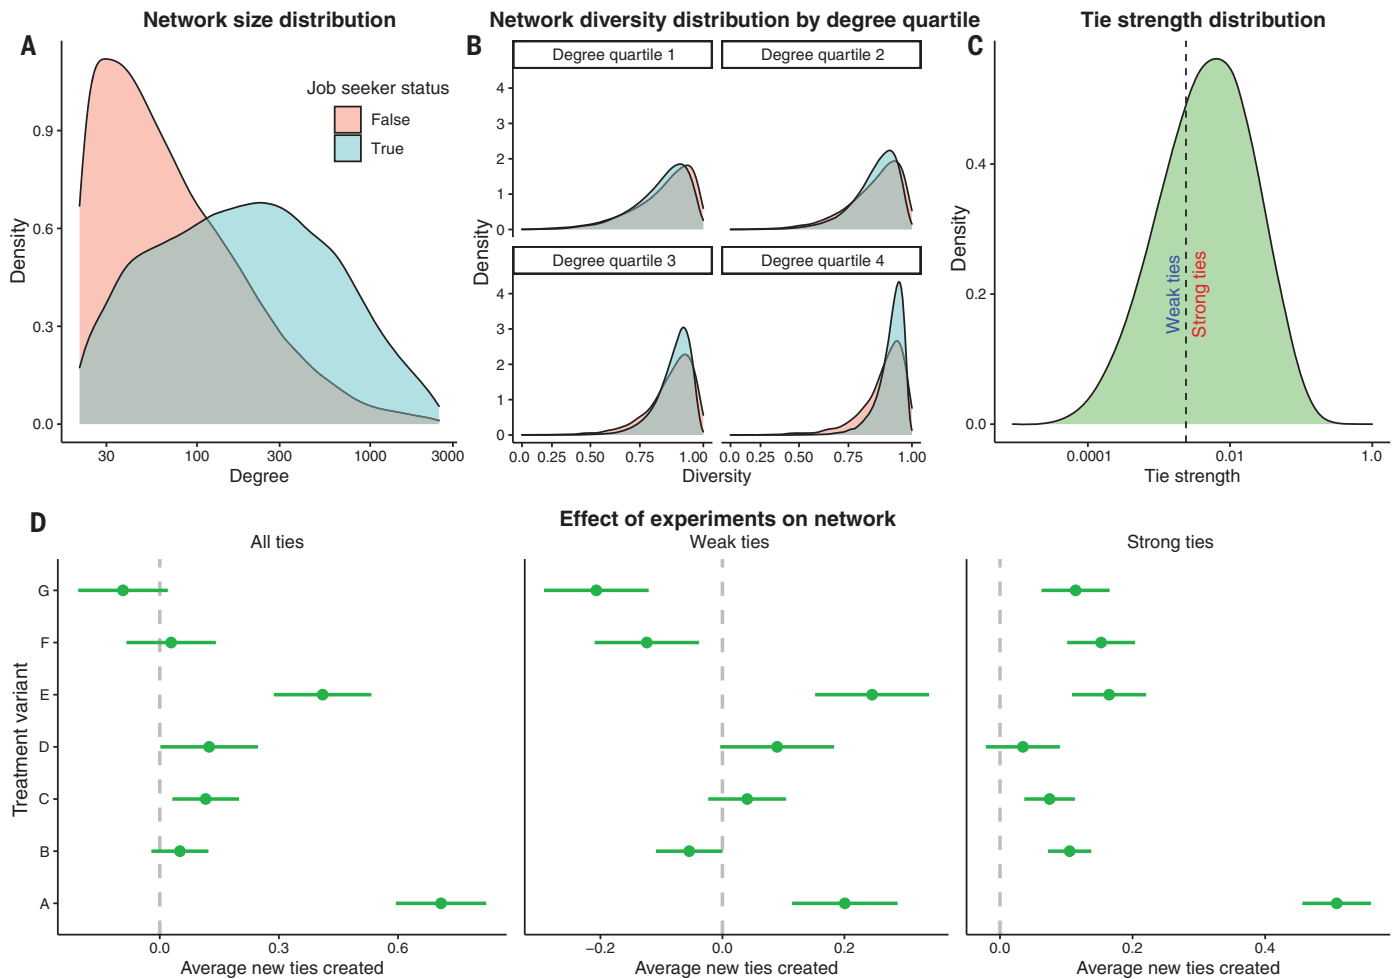


Fig. 2. Network statistics and first stage effects of experimental treatments. (A) displays the node-level degree distributions of job seekers and non-job seekers whereas (B) displays the corresponding distribution of network diversity by job seeker status by quartiles of members' degrees to distinguish diversity and network size, in which job seekers are members who applied for a job in the three months before an experiment; (C) displays the edge-level structural tie strength distribution of all ties created during the 2019

experiments, in which the cutoff for determining weak or strong ties is the median of structural tie strength in the LinkedIn network before the experiments; and (D) displays the "first stage" effects of the experimental treatments on how many new ties are created by members. The point estimates and standard error bars report the number of new ties created and their relative split between strong and weak ties by members assigned to different treatment variants compared to a control variant.

on job outcomes from the causal effects of the number of new ties created.

The second stage regressions estimated the effect of weak and strong ties on job mobility. The fitted values estimated in the first stage captured only those changes in the number of new weak or strong ties caused by our experiments. In the second stage only the variation in the creation of new weak or strong ties caused by our exogenous treatment assignments was used to estimate the effects of weak or strong ties on job applications and job transmissions. In this way the IV approach enabled causal inference by excluding (i) the effects of job mobility on the formation of weak and strong ties, (ii) the effects of strategic network formation behaviors that precede job mobility, and (iii) variation created by observ-

able and unobservable confounding factors that can affect both network formation and job mobility, from estimates of the effects of weak and strong ties on labor market behaviors and outcomes.

For our approach to provide valid causal inference, the treatment assignment to a PYMK algorithm variant should be a valid instrument for the number of weak and strong ties created by experimental subjects and thus should satisfy four assumptions (27). First, the independence assumption, which requires that the instrument was randomly assigned, was satisfied as the LinkedIn experimentation platform used a Bernoulli design to randomly assign all users to different treatment arms. Second, the exclusion restriction, which requires that the instrument did not affect the

outcome through any channels except the treatment channel, held because the experiments uniquely altered tie recommendations without altering any other algorithm related to job outcomes such as job recommendations or job search rankings. It also had no effect on how members interacted or shared social information with each other (e.g., messages, posts, comments, likes, and shares), other than through the new connections formed as a result of the experiments (see Section C.3 and table S24 in the SM). Third, the monotonicity assumption, precluding the existence of "defiers"—egos that initiated more weak tie connections when assigned to an algorithm that promoted strong ties—was satisfied by the design of the PYMK algorithm, which assigned suggested treatment connections to

ranks that were much more likely to be clicked on and initiated. This assumption was also satisfied by the strong suggestive evidence that these assignments worked to create the desired behaviors observed in the new connection outcomes for each treatment variant (Fig. 2D). Fourth, the relevance assumption, which requires that the instrument had an effect on the treatment, was satisfied by the varying numbers of weak, strong, and total ties created by the different treatment arms shown in Fig. 2 (see Section C.3 of the SM for an in-depth discussion of these assumptions and their verification). To estimate the bias in correlational analyses of the weak tie hypothesis, we also specified and estimated standard ordinary least squares (OLS) regressions assessing the correspondence between tie strength and job mobility.

Our main results are summarized in Fig. 3. The first column in each panel displays results from OLS estimation whereas the second column displays results from the experimental instrumental variables (IV) analysis. Dots represent point estimates of the effects and bars represent standard errors. Figure 3A displays the results of our OLS and experimental estimates of the effect of weak and strong ties on job transmissions at the node level. Although the OLS analysis replicated previous findings of an apparent paradox of weak ties in which strong ties were more strongly correlated with job transmissions, the experimental IV analysis reversed this result and suggested a non-linear relationship between tie strength and job transmission in which medium strength ties were the most effective in generating job mobility. As the statistical power of the node-level analysis was not sufficient to confirm individual differences between the effects of strong, medium, and weak ties in our experiments, we also analyzed these relationships at the more granular edge level.

Figure 3B displays the results of our OLS and experimental estimates of the effect of the strength of a newly added tie, measured by the number of mutual connections between LinkedIn members before treatment, on the probability of a job transmission between them after treatment at the edge level. Although the OLS estimates again replicated previous correlational research demonstrating the paradox of weak ties and showed that strong ties were more strongly correlated with an increased probability of a job transmission through the tie, the experimental IV results mirrored the node-level experimental results and revealed a more nuanced correspondence—namely that there was an inverted U-shaped relationship between tie strength and the likelihood of a job transmission. At low levels of mutual friendship adding new ties with more mutual friends caused the probability of a job transmission to go up. However, adding ties with

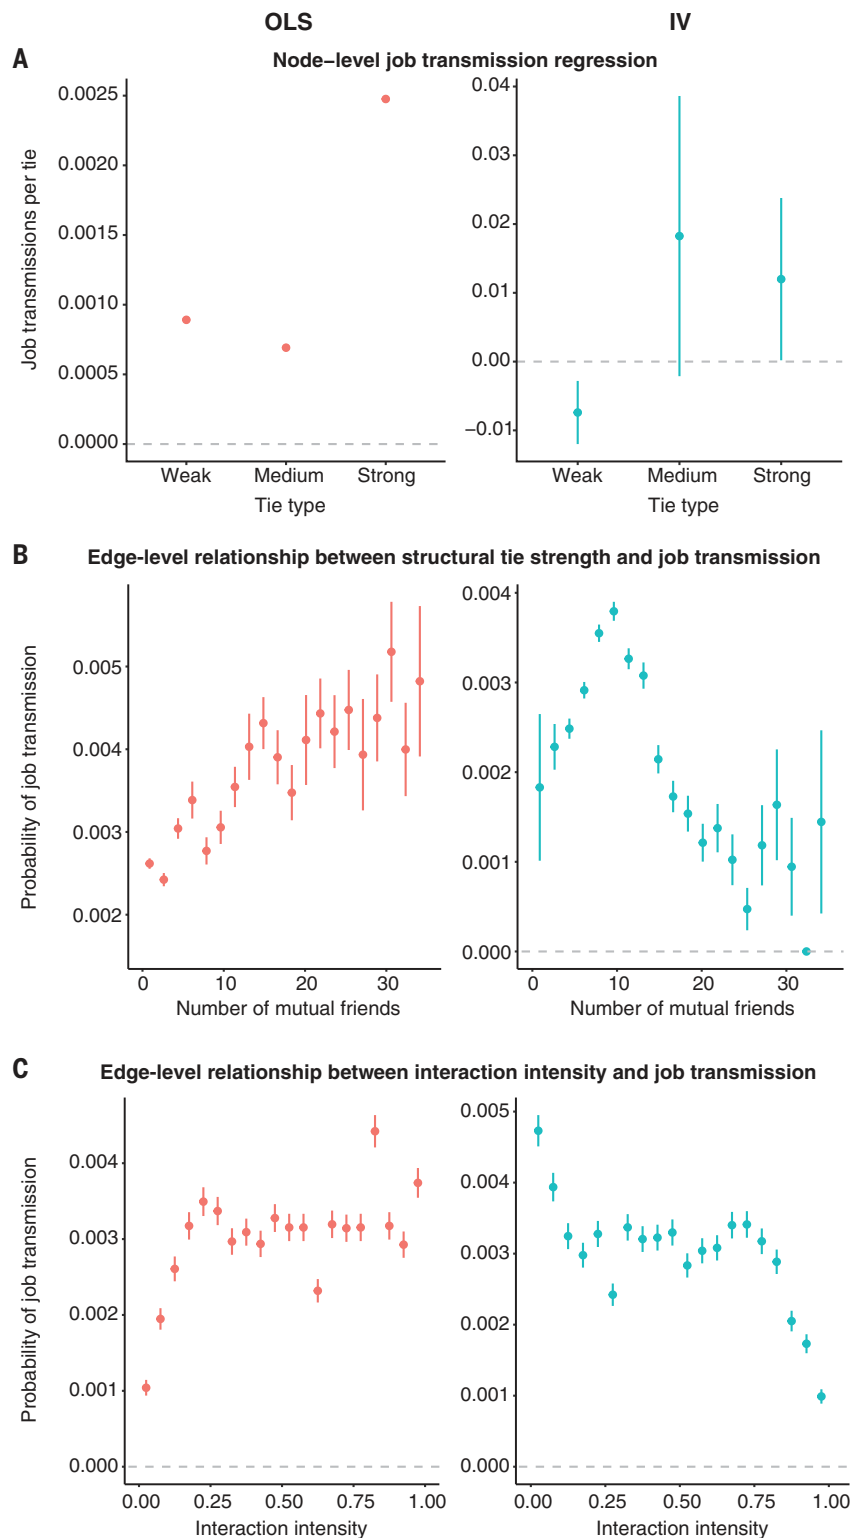


Fig. 3. The causal effect of tie strength on job mobility. The figure displays the estimated effects of tie strength on job transmissions at the node and edge levels. In each panel the left column displays the results of the OLS analysis whereas the right column displays the experimental IV results from first wave experiments conducted in 2015. **(A)** displays the effect of weak, medium, and strong ties (defined by terciles of mutual friends) on job transmissions aggregated to the node level; **(B)** displays the effects of structural tie strength based on the number of mutual friends on job transmissions at the edge level whereas **(C)** displays the effects of tie strength based on interaction intensity on job transmissions at the edge level.

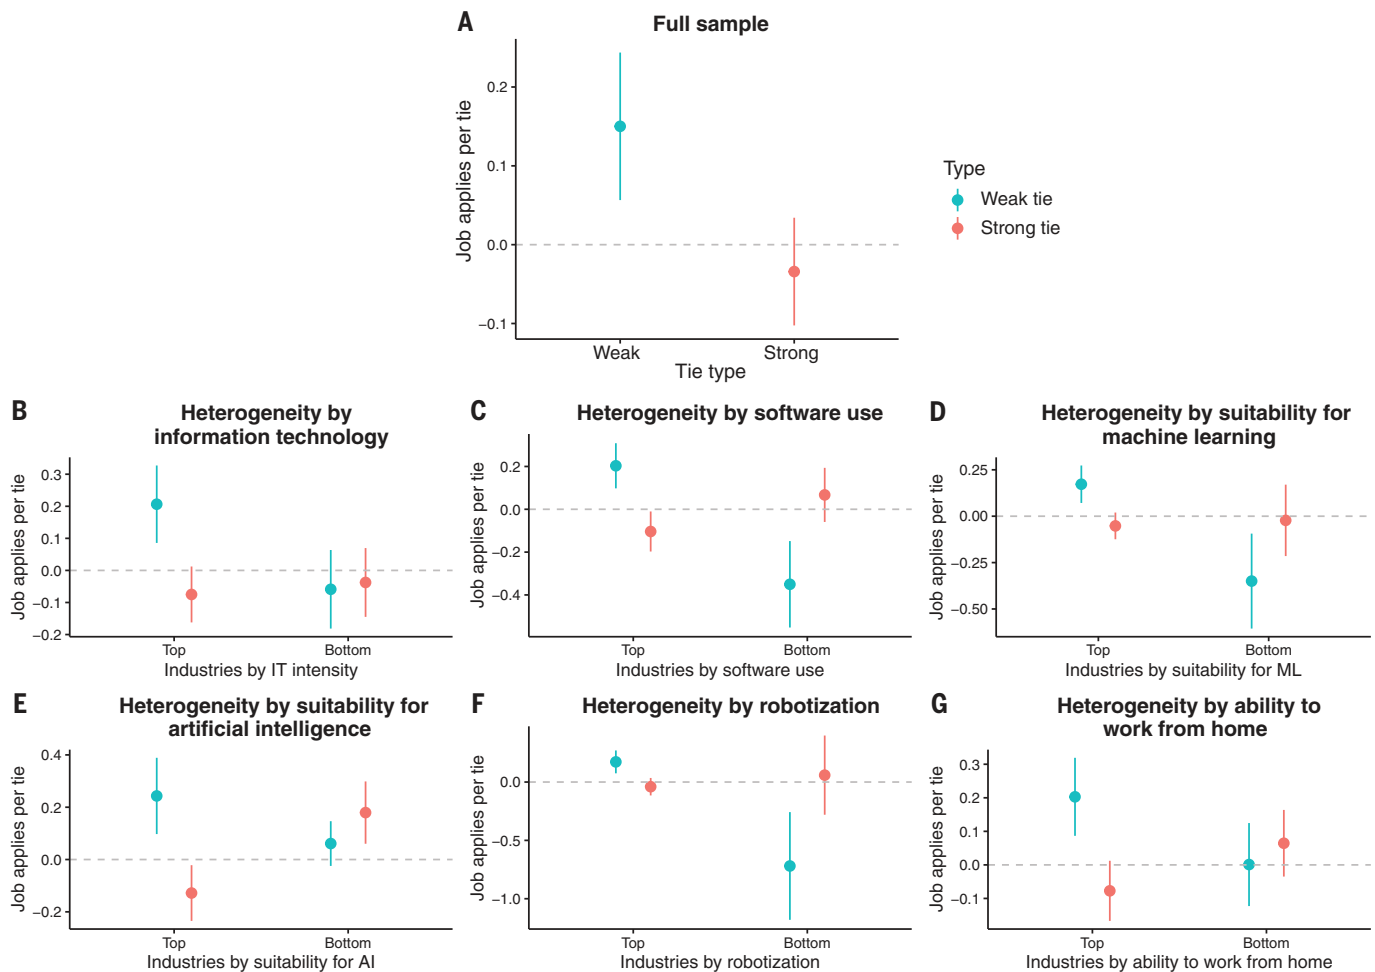


Fig. 4. Heterogeneous causal effects by industry. (A) displays experimental IV analysis of the effects of weak and strong ties on job applications at the node level in the full sample. The rest of the panels display experimental IV analysis of the heterogeneous effects of weak and strong ties on job applications

at the node level to jobs in industries in the top and bottom half of the industry distributions of (B) IT intensity; (C) software intensity; (D) suitability for machine learning; (E) suitability for artificial intelligence; (F) degree of robotization; and (G) suitability for remote work.

more than ten friends in common reduced the probability of a job transmission.

Figure 3C displays the results of our OLS and experimental estimates of the effect of the strength of a newly added tie, measured by the interaction intensity between LinkedIn members and on the probability of a job transmission between them, again at the edge level. Although the OLS estimates showed that stronger ties were correlated with an increased probability of job transmission the experimental IV results revealed the opposite—the stronger the newly added ties the less likely they were to lead to a job transmission. This relationship was also nonlinear. The weakest ties with the least interaction intensity increased the likelihood of a job transmission the most whereas the strongest ties with the greatest interaction intensity increased the likelihood of a job transmission the least; further, the relationship between interaction intensity and job transmission was approximately flat for the

middle quartiles of the interaction intensity distribution.

Three major conclusions emerged from our main results: First, experimental analysis helped resolve the apparent paradox of weak ties in multiple large-scale experiments of job mobility in the world's largest professional social network. Although the correlational analysis supported the seeming importance of strong ties for job mobility, the experimental analyses—conducted over multiple sample populations, numerous years, and in all geographic regions of the world—confirmed that relatively weaker ties increased the likelihood of job mobility the most.

Second, our experiments uncovered a consistent nonlinearity in the relationship between tie strength and job mobility. In contrast to the increasing likelihood of job transmission associated with greater tie strength in correlational analyses, our experiments showed that when considering structural tie strength based

on the number of mutual friends between contacts, an inverted U-shaped relationship exists between tie strength and job mobility—with moderately weak ties increasing job mobility the most and the strongest ties increasing job mobility the least. When considering tie strength based on interaction intensity, in a direct reversal of the correlational evidence, the experimental analysis showed that the weakest ties had the greatest impact on job mobility whereas the strongest ties had the least.

Third, whereas node-level analyses measured the impact of experimental variation in the number of weak or strong ties in one's network on job transmission, our edge-level analysis enabled an assessment of the marginal effects of adding strong or weak ties. The results showed that adding new moderately structurally diverse ties with weak interaction intensity created the greatest marginal increases in the likelihood of job transmissions.

Prior research also suggests that weak and strong ties have different effects across different industries (22). We therefore examined the heterogeneity in the impact of strong and weak ties on job mobility across industry sectors. Although the second wave experimental sample was sufficiently powered to examine this heterogeneity, experiments conducted in this wave in 2019 do not leave sufficient time to examine impacts on longer-term job transmission outcomes. Therefore we limited our analysis of these heterogeneous effects to job applications, which are estimable in the short term.

We classified the industries in which LinkedIn members applied for jobs on the basis of the demand for particular skills listed for those jobs and the counts of occupations in different industries calculated from Burning Glass Technologies (BGT) data and other sources (see supplementary material for details). The industry classifications were created by measuring the weighted skill demands of all job postings within an industry listed in the BGT data and the counts of an industry's hiring for different occupations listed in the job postings of that industry. Based on these metrics, we developed scores that measured each industry's information technology (IT) intensity, software intensity, suitability for machine learning, suitability for artificial intelligence, degree of robotization, and suitability for remote work using known indices for these metrics in the labor economics literature (30–32). We then measured the degree to which experimental variation in the acquisition of new strong or weak ties led to increases or decreases in job applications to industries of these types.

Results of our IV analysis showed that, in the full sample, adding weak ties led to more job applications overall (Fig. 4A), which provides evidence of the mechanism linking weak ties to job transmissions. As members acquired more weak ties through the PYMK algorithm experiments they applied to more jobs and experienced considerably greater job mobility. The heterogeneous treatment effects also reflect a clear trend toward weak ties creating greater job mobility in more digital sectors of the economy. The results showed that weak ties resulted in more job applications than strong ties to industries with greater IT (Fig. 4B) and software intensity (Fig. 4C), as well as industries more suitable for machine learning (Fig. 4D), artificial intelligence (Fig. 4E) and remote work (Fig. 4G), along with those that have experienced a greater degree of robotization (Fig. 4F). By contrast, adding strong ties caused more job applications to industries that relied less on software (Fig. 4C) and were less automated by robots (Fig. 4F).

Although our work presents the first large-scale, longitudinal, experimental evidence on

the causal effects of strong and weak ties on job mobility in a global sample and across multiple industries, it is not without limitations. First, although PYMK experiments provided a robust channel through which to introduce experimental variation into the evolution of human social networks, we could not compel LinkedIn users to take these recommendations. Therefore a degree of self selection exists in who acted on the connection recommendations. For this reason we analyzed our experiments as having an “intent to treat” and compared the population assigned to weak-tie experimental variants to those assigned to strong-tie experimental variants and control groups (for raw intent to treat point estimates, see table S14). Although this approach controls for any bias from self selection, it circumscribes the populations to which our results generalize. Although there were some observable differences between members who took PYMK recommendations on LinkedIn and those who did not, most did, making our results broadly generalizable to the LinkedIn population. However, unsurprisingly, exposure depended on use of the platform and viewing the PYMK recommendations. As we report in the SM, LinkedIn members exposed to our treatments were slightly younger and more active job seekers, clarifying the population to which our results reliably generalize.

Second, LinkedIn is a professional social network and may be different than other online social networks such as Facebook or offline social relationships such as those originally studied by Granovetter. However we do know that certain characteristics such as network clustering, for example, are similar across Facebook, LinkedIn, and Twitter (see SM Section F.2). Furthermore, our OLS results closely mirrored the results of very large global studies of networks and job mobility on Facebook, which suggests that similar processes are occurring in both networks. There are some differences between the population of workers on LinkedIn and those in the US, European, and broader global economies. For example, LinkedIn skews more heavily toward workers in finance, information and professional services, high technology industries, and construction and manufacturing, and less toward wholesale and retail trade work than the US workforce (see SM Section F.2 for a comparison of LinkedIn profiles with the US and EU workforces). However, LinkedIn is also the world's largest professional social network and one of the largest websites for job listings. Many people rely on LinkedIn to find work so this network may be even more representative of how networks affect job mobility in the larger labor market than, for example, friendship networks or the Facebook network.

Third, any networked experiment must pay close attention to the possibility of statistical

interference, in which one unit's treatment assignment affects another unit's outcome, the ignorability of which is known as the stable unit treatment value assumption (33). To minimize such interference we only tracked the edges that each member initiated through PYMK recommendations. Nevertheless, there were still three possible channels through which interference could have occurred in our setting. First, interference could have emerged if an ego's treatment assignment affected their alters through changes in ego's behaviors that were visible to the alters. LinkedIn facilitates some social actions that might have been seen by a member's alters including posting on the news feed, commenting on a post, or sending private messages. However, none of these behaviors were considerably affected by the various treatments, making this channel of interference unlikely to affect our results (see table S24). A second interference channel could have arisen if member i intended to connect with member j but, because of member j 's treatment assignment, j initiated a connection request before i had a chance to send one. If accepted, such ties would be attributed to ego j in our analysis. To account for this possibility we verified that the treatments did not affect the number of connection requests received by members in different treatment arms, allowing us to conclude that any effect from receiving connection requests was small and balanced across treatments and therefore negligible (see table S25). Third, interference could have occurred if new ties generated by members as a result of treatment changed the composition of other members' PYMK recommendation lists. Fortunately, LinkedIn's membership is large enough to ensure a sufficient inventory of new ties to replenish any removed ties with comparable individuals, ensuring that the composition of the potential ties was consistent throughout the experiment. Furthermore, PYMK inputs did not change often enough for new connections to immediately change the types of algorithmic recommendations any member saw, ensuring the experiments' stability. This minimized the risk that connection behaviors instantaneously updated the algorithms and thus changed the types of recommendations other members saw. For more details on the interference assumption, please see SM Section C.2.

Despite these limitations, our analysis of several large-scale experiments on the world's largest professional social network demonstrated that weak ties create job mobility. In contrast to recent large-scale correlational evidence of a paradox of weak ties, we found that moderately weak ties with low interaction intensity—measured by the number of mutual friends between two people—increased job applications and job transmissions the most, whereas strong ties—measured by both the number of mutual friends and interaction intensity—increased job

applications and job transmissions the least. We also found an inverted U-shaped relationship between structural tie strength and job transmissions and a nonmonotonically decreasing correspondence between interaction intensity and job transmissions, demonstrating a consistent nonlinearity in the relationship between tie strength and job mobility, as well as heterogeneity in the impact of weak ties on job applications across industries with varying degrees of digitization. The industry analysis showed that weak ties caused more job applications to high-tech industries, broadly speaking, whereas strong ties caused more job applications to low-tech industries. Together, these results provide some of the first large-scale experimental evidence of the strength of weak ties and suggest the need to revise the theory to incorporate the nonlinear effects of tie strength on job transmissions, differences between the effects of structural tie strength and tie strength measured by interaction intensity, and differences between the effects of weak and strong ties on job mobility across industries.

REFERENCES AND NOTES

- M. S. Granovetter, *Am. J. Sociol.* **78**, 1360–1380 (1973).
- E. Bakshy, I. Rosenn, C. Marlow, L. Adamic, The role of social networks in information diffusion, in *Proceedings of the 21st International Conference on World Wide Web* (2012), pp. 519–528.
- P. S. Park, J. E. Blumenstock, M. W. Macy, *Science* **362**, 1410–1413 (2018).
- J. Ugander, L. Backstrom, C. Marlow, J. Kleinberg, *Proc. Natl. Acad. Sci. U.S.A.* **109**, 5962–5966 (2012).
- D. Centola, *Science* **329**, 1194–1197 (2010).
- D. Centola, M. Macy, *Am. J. Sociol.* **113**, 702–734 (2007).
- G. Walker, B. Kogut, W. Shan, *Organ. Sci.* **8**, 109–125 (1997).
- D. Kempe, J. Kleinberg, É. Tardos, Maximizing the spread of influence through a social network, in *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (2003), pp. 137–146.
- D. G. Rand, S. Arbesman, N. A. Christakis, *Proc. Natl. Acad. Sci. U.S.A.* **108**, 19193–19198 (2011).
- C. L. Apicella, F. W. Marlowe, J. H. Fowler, N. A. Christakis, *Nature* **481**, 497–501 (2012).
- J. D. Montgomery, *Am. Econ. Rev.* **81**, 1408–1418 (1991).
- J. M. Podolny, J. N. Baron, *Am. Sociol. Rev.* **62**, 673–693 (1997).
- R. S. Burt, *Am. J. Sociol.* **110**, 349–399 (2004).
- S. Rodan, C. Galunic, *Strateg. Manage. J.* **25**, 541–562 (2004).
- A. Hargadon, R. I. Sutton, *Adm. Sci. Q.* **42**, 716–749 (1997).
- R. Reagans, E. W. Zuckerman, *Organ. Sci.* **12**, 502–517 (2001).
- N. Eagle, M. Macy, R. Claxton, *Science* **328**, 1029–1031 (2010).
- R. S. Burt, *Structural holes* (Harvard Univ. Press, 1992).
- B. Uzzi, *Am. Sociol. Rev.* **61**, 674–698 (1996).
- D. Lazer, A. Friedman, *Adm. Sci. Q.* **52**, 667–694 (2007).
- M. McPherson, L. Smith-Lovin, J. M. Cook, *Annu. Rev. Sociol.* **27**, 415–444 (2001).
- S. Aral, M. Van Alstyne, *Am. J. Sociol.* **117**, 90–171 (2011).
- D. J. Watts, S. H. Strogatz, *Nature* **393**, 440–442 (1998).
- A.-L. Barabási, R. Albert, *Science* **286**, 509–512 (1999).
- L. K. Gee, J. Jones, M. Burke, *J. Labor Econ.* **35**, 485–518 (2017).
- L. K. Gee, J. J. Jones, C. J. Fariss, M. Burke, J. H. Fowler, *J. Econ. Behav. Organ.* **133**, 362–372 (2017).
- J. D. Angrist, G. W. Imbens, D. B. Rubin, *J. Am. Stat. Assoc.* **91**, 444–455 (1996).
- J. D. Angrist, A. B. Krueger, *J. Econ. Perspect.* **15**, 69–85 (2001).
- J. D. Angrist, J.-S. Pischke, *Mostly Harmless Econometrics* (Princeton Univ. Press, 2008).
- E. Brynjolfsson, T. Mitchell, D. Rock, *AEA Pap. Proc.* **108**, 43–47 (2018).
- M. Webb, The Impact of Artificial Intelligence on the Labor Market (Stanford University, 2020); https://www.michaelwebb.co/webb_ai.pdf
- J. I. Dingel, B. Neiman, *J. Public Econ.* **189**, 104235 (2020).
- D. B. Rubin, *J. Am. Stat. Assoc.* **75**, 591–593 (1980).
- K. Rajkumar, rajkumarkarthik/weak_ties_data_and_code: Replication files for A Causal Test of the Strength of Weak Ties, Version v1.0.1, Zenodo (2022); <https://zenodo.org/record/6533898#.YwkcWnbMI2w>.

ACKNOWLEDGMENTS

We thank D. Eckles and participants of the MIT Conference on Digital Experimentation, the Workshop on Information Systems Economics, and the HBS Junior Faculty Research Seminar for their helpful comments on earlier drafts of this manuscript. We also thank the MIT Initiative on the Digital Economy and the Stanford Digital Economy Lab for their intellectual support, as well as M. Garlinghouse, I. Perisic, Y. Xu, and P. Ahammad at LinkedIn for intellectual and organizational support and data access. The study was approved by the MIT Institutional Review Board (IRB Protocols E-3347 and E-2755). **Funding:** The authors received no external financial support for the research and publication of this work.

Author contributions: Conceptualization: G.S.-J., E.B., and S.A.; Methodology: K.R., G.S.-J., I.B., and S.A.; Data curation, software, and formal analysis: K.R. and G.S.-J.; Investigation: K.R., G.S.-J., I.B., E.B., and S.A.; Writing: K.R., I.B., E.B., and S.A.; Visualization: K.R., I.B., E.B., and S.A.; Project administration and resources: K.R., G.S.-J., E.B., and S.A.; Supervision: G.S.-J., I.B., E.B., and S.A.; Validation: K.R. and G.S.-J. **Competing interests:** K.R. and G.S.-J. are or were employees of LinkedIn and hold a significant financial interest in Microsoft Corporation. All other authors declare no competing interests. **Data and materials availability:** All original analysis code, aggregated data and annotated code sufficient to replicate all correlational regressions are available here: (34). Individual-level data supporting this study are not publicly available due to user privacy and other legal restrictions. Researchers seeking replication should submit a request to mailto:egr_data@linkedin.com clearly stating the work that you would like to reproduce and the names and organizational affiliations of the researchers. LinkedIn reviews all such requests but will grant access at its sole discretion. **License information:** Copyright © 2022 the authors, some rights reserved; exclusive licensee American Association for the Advancement of Science. No claim to original US government works. <https://www.sciencemag.org/about/science-licenses-journal-article-reuse>

SUPPLEMENTARY MATERIALS

science.org/doi/10.1126/science.abl4476
Materials and Methods
Fig. S1
Tables S1 to S28
References (35–52)
MDAR Reproducibility Checklist

Submitted 19 July 2021; accepted 20 July 2022
10.1126/science.abl4476

A causal test of the strength of weak ties

Karthik RajkumarGuillaume Saint-Jacqueslavor BojinovErik BrynjolfssonSinan Aral

Science, 377 (6612), • DOI: 10.1126/science.abl4476

The influence of weak associations

The strength of weak ties is an influential social-scientific theory that stresses the importance of weak associations (e.g., acquaintance versus close friendship) in influencing the transmission of information through social networks. However, causal tests of this paradoxical theory have proved difficult. Rajkumar *et al.* address the question using multiple large-scale, randomized experiments conducted on LinkedIn's "People You May Know" algorithm, which recommends connections to users (see the Perspective by Wang and Uzzi). The experiments showed that weak ties increase job transmissions, but only to a point, after which there are diminishing marginal returns to tie weakness. The authors show that the weakest ties had the greatest impact on job mobility, whereas the strongest ties had the least. Together, these results help to resolve the apparent "paradox of weak ties" and provide evidence of the strength of weak ties theory. —AMS

View the article online

<https://www.science.org/doi/10.1126/science.abl4476>

Permissions

<https://www.science.org/help/reprints-and-permissions>

Use of this article is subject to the [Terms of service](#)

Science (ISSN) is published by the American Association for the Advancement of Science. 1200 New York Avenue NW, Washington, DC 20005. The title *Science* is a registered trademark of AAAS.

Copyright © 2022 The Authors, some rights reserved; exclusive licensee American Association for the Advancement of Science. No claim to original U.S. Government Works