

Geography, Ties, and Knowledge Flows: Evidence from Citations in Mathematics

Keith Head, Yao Amber Li, Asier Minondo

Head: Sauder School of Business, University of British Columbia; Li: Hong Kong University of Science and Technology; Minondo: Deusto Business School, University of Deusto

Acknowledgement: The authors thank Mitch Keller at the Mathematics Genealogy Project (MGP) and Nicolas Roy from zentralblatt-math.org for data. Yao Amber Li acknowledges financial support from the Research Grants Council of Hong Kong, China (General Research Funds Project no. 643311), and Asier Minondo from the Spanish Ministry of Economy and Competitiveness (MINECO ECO2015-68057-R and ECO2016-79650-P, co-financed with FEDER) and the Basque Government Department of Education, Language policy and Culture (IT885-16). Seminar participants at Dartmouth, LSE, Oxford, UBC, Glasgow, Birmingham, and Nottingham made helpful suggestions. We also thank Andrew Bernard, Teresa Fort, Joshua Gottlieb, Bob Staiger, Bronwyn Hall, Wolfgang Keller, Anthony J. Venables, Quoc-Anh Do, Edwin Lai, Jim MacGee, Andrés Rodríguez-Clare, Daniel Sturm, Dave Donaldson, Ben Faber, and Pablo Fajgelbaum for valuable discussions. We thank Michal Fabinger for alerting us to the importance of arXiv.org for knowledge dissemination over the internet. We are particularly grateful to Andrei Levchenko for raising issues that led to the results in section 4.2. Ho Yin Tsoi, Bo Jiang, Yiye Cui, and Song Liu provided excellent research assistance during this project. Finally, we thank three anonymous referees and Amit Khandelwal for very helpful suggestions.

Abstract: Combining data on locations with career and educational histories of mathematicians, we study how distance and ties affect citation patterns. The ties considered include coauthorship, past colocation, and relationships mediated by advisors and the *alma mater*. With fixed effects capturing subject similarity and article quality, we find linkages are strongly as-

sociated with citation. Controlling for ties generally *halves* the negative impact of geographic barriers on citations. Ties matter more for less prominent and more recent papers and retain their quantitative importance in recent years. The impact of distance—controlling for ties—has fallen and is statistically insignificant after 2004.

Keywords: knowledge diffusion, distance, borders, networks, academic genealogy

JEL: F1, O3, R1

1 Introduction

Mounting evidence points to the importance of geographic barriers to knowledge flows. Arguing that citations provide the “paper trail” for knowledge flows, Jaffe et al. (1993) establish that cites to patents are geographically localized.¹ Keller (2002) shows that research spillovers on productivity decay with distance and Comin et al. (2012) find the likelihood of adopting new technologies declines with distance to the origin of the invention. Ellison et al. (2010) show that industries that share ideas (proxied by R&D and patent citation flows) have a stronger tendency to coagglomerate in space. While much of the literature focuses on technology diffusion, spatial separation impedes the spread of many other types of information. For example, information frictions account for half of the distance effect in the Allen’s (2014) study of differences in rice prices between Philippine islands. Urbanization continues to increase despite rising land prices and congestion, a fact Glaeser (2011) attributes to the spread of innovations “from person to person across crowded city streets.”

All the above evidence notwithstanding, the notion that borders or distance could prove to

¹Peri (2005), Belenzon and Schankerman (2013), Singh and Marx (2013), and Li (2014) estimate robust negative distance effects on patent citation propensities.

be practical obstacles to flows of knowledge seems hard to square with the fact that information can move anywhere without incurring either tariffs or freight costs. As Keller and Yeaple (2013) put it “Knowledge, as an intangible, seems ideally suited to overcoming spatial frictions ...” Especially in the age of Google, whose self-described mission is to “Organize the world’s information and make it universally accessible and useful,” the microfoundations for geographic knowledge frictions are far from obvious. To the extent there is a standard explanation, it is often mentioned that tacit knowledge is easier to communicate face to face. However, one study shows that even the transmission of highly codified information benefits from proximity. Lissoni (2001) examined a cluster of mechanical firms in Brescia, Italy and found they engaged primarily in the transfer of CAD encoded designs.

In this paper we hypothesize that distance’s impact on knowledge arises in large part due to spatially concentrated personal ties. Proximity facilitates tie formation and those ties foster knowledge flows. The general mechanism we envision is that an agent trying to solve a problem becomes aware of potential solutions by tapping the knowledge residing in their network of personal relationships. This hypothesis can only be tested in a specific context where interpersonal ties, geography, and knowledge flows can all be tracked in a systematic way. We argue that the rich data available on mathematicians makes them, despite their idiosyncrasies, an insightful group to study for this purpose. Our first key finding is that adding controls for a comprehensive set of career and educational linkages between authors of mathematics papers, leads to a halving of estimated geography effects. The role of ties in attenuating the negative effect of distance on citations echoes Keller’s (2001) finding that including trade flows and FDI in the equation for technological knowledge spillovers shrinks the estimated negative effect of distance. The paper proceeds to combine additional results to establish the microfoundations for why ties matter so much.

Prior work on patent citation has already pointed towards ties as an important determinant of knowledge flows. Invoking the idea of “social proximity” Agrawal et al. (2008) and Kerr (2008) show that inventors have a higher propensity to cite patents by those who share their ethnic origins (as revealed by their surnames). While social connections are known to be richer within ethnic groups, sharing surnames with the same ethnic origin does not imply a personal connection between citing and cited inventors. Co-ethnicity can reflect cultural similarities between inventors who do not know each other personally. In order to capture the effect of person-to-person ties on knowledge flows, we need data sources from which we can extract the histories of personal relationships. Patent applications provide enough information to determine past collaboration; Singh (2005) and Breschi and Lissoni (2009) find this type of tie increases citation. Agrawal et al. (2006) investigate a second tie, past colocation. They find that inventors who move institutions are still disproportionately cited in patent applications by their former colleagues.

To capture a richer set of social ties between individuals who potentially transmit knowledge to each other, we believe it useful to consider academics, for whom it is possible to identify ties based on educational histories. We take advantage of the fact that in mathematics, Ph.D. institutions and advisors have been tracked globally for a long time by the Mathematics Genealogy Project (MGP).² There is strong evidence from Waldinger (2010) that the quality of mathematics faculty causally increases subsequent academic success of their doctoral students. The MGP allows us trace the patterns of citation between advisors and advisees, classmates, and the academic “extended family.”

The process through which mathematicians (or anyone else) form ties is not, of course,

²Borjas and Doran (2012) use the MGP to identify immigrant mathematicians who received Ph.D.s from Soviet institutions.

entirely random. A concern for estimating the effect of ties is that the same unobservables that promote scholars to form ties with each other also affect the likelihood of them citing each other. The educational ties we focus on have the advantage of being predetermined with respect to the citation process, since it is rare for an academic to cite or be cited prior to obtaining doctoral education. Unlike colocation and collaboration, educational linkages do not change over time in response to shocks to the interests of citing authors. While there is substantial randomness involved in determining classmates, the matching between advisors and advisees is likely to be shaped by common interests. The worry is that author A may be more likely to cite a paper by a tied author B than author C who has no tie with B because A and B write on the same topics. The way we respond to this concern is to compare citation probabilities only between authors A and C who have written papers in the same 3-digit field of mathematics. We show that this control for article subject is essential. Without it, estimates of ties are substantially inflated. Controlling for 3 or 5 digit fields or even keywords, ties have reduced—but still large—estimated effects on citation. With the controls that give the lowest magnitude, 5-digit subject and a cocitation indicator, a single tie on average boosts the odds of citation by 46%.

In addition to the strength of its academic genealogy data, mathematics offers two additional advantages relative to other academic fields. First, mathematics employs a common language of communication. This suggests transmission of mathematics knowledge would be less influenced by linguistic and cultural factors. In many social sciences and humanities fields, there are journals that focus on certain regions or countries. For example, in the fields of history and literature, there are obvious reasons to expect national borders and language to influence citation patterns. A second advantage of studying mathematics comes from the citation norms of the discipline. New theorems build upon previous theorems, which must be cited. There also appears to be a norm against gratuitous citation, as evidenced by the relatively low num-

ber of references in each paper. Althouse et al. (2009) report that math papers cite 18 papers on average, compared to 30 in economics and 45–51 in sociology, psychology, business and marketing.

Our first set of results establish that ties are an important mechanism underlying estimated geography effects on citations. But what is the mechanism underlying the importance of ties? We present two lines of evidence to argue that ties matter because they transmit information. The first follows from the idea of Arrow (1969) that knowledge flows can be generally thought of as interactions between a teacher (sender) and a student (receiver). We find evidence that citations are stronger to the authors who are more likely to be senders of information. The odds of citation are seven times higher if a paper is written by the advisor of the citing author. The impact of the author being a former advisee is weaker, albeit still very large. Moving one step further apart in the advisor network, we find advisors of advisors have three times the normal odds of being cited, but there are no significant differences in their propensity to cite their advisees' advisees. The second line of evidence is that ties matter more for the types of papers where information is harder to acquire. Our estimates show that ties (and geographic separation) have stronger impacts for papers that were only recently published, or not heavily cited, or just in a different field.

The role of distance—after controlling for ties—even becomes statistically insignificant in recent years. This finding of declining geographic barriers extends the results of two earlier studies using very different methodologies. Keller (2002) estimates the rate of distance decay in the benefits that one country receives from R&D conducted in another country. He finds that the distance decay rate fell by two thirds from the period 1970–1982 to 1983–1995. Griffith et al. (2011) analyze the number of days until the first citation of a newly granted patent. They find home inventors take fewer days on average to be the first to cite home-invented patents

than foreign-based inventors. This home-bias declined substantially between 1975–1989 and 1990–1999. Our study shows that distance effects have fallen by two thirds from the early 1990s to the late 2000s. This extends the evidence from the previous literature to the decade in which internet usage becomes pervasive. Our investigation of time-varying coefficients also reveals that, despite the advances in scholar’s ability to search for information over the internet, the impact of personal ties remains as strong as ever.

While we do not wish to draw conclusions that stray too far from the context of our estimation, the whole rationale for studying citations in mathematics is to obtain insights with broader applicability to knowledge flows. The extent that ties facilitate transfer of valuable knowledge in one context (math) provides a *prima facie* case for their potential importance in all cumulative, collaborative discovery processes. Collaboration in mathematics often takes the form of tied researchers making suggestions of previously proven theorems that could help prove new theorems. In other research contexts, from drug invention to financial engineering, there would be analogous ways that lessons learned by one person could help a tied person to solve a new problem.

Going beyond research, there is a wealth of suggestive evidence that entrepreneurs learn about potential business opportunities from their web of connections. For example, Kerr and Mandorff (2015) explain the remarkable concentration of ethnic groups in certain occupations (Gujarati-speaking Indians are over-represented in the motel industry by a factor of 108) by invoking knowledge acquired through social interactions. Learning from ties might also explain the robust empirical association between bilateral immigration stocks and trade flows.³ Such

³Gould (1994) is the seminal paper. Rauch and Trindade (2002) show that countries with larger ethnic Chinese populations trade more with each other. Combes et al. (2005) use migration and investment data to infer that social and business networks create trade within France.

work generally lacks individual-level evidence on the relevant social ties. Using our person-to-person measures of ties provides insight into the processes underlying the patterns seen in aggregated data.

The remainder of the paper is organized as follows. Section 2 posits a simple citation model to serve as the estimating framework for relating a paper-to-paper citation indicator to the ties and geography variables measured at the author level. Section 3 describes our data on citations, geography and ties and explains how we construct the estimating sample. Section 4 presents the results of our regressions. In the final section we re-interpret other research findings in light of our results. We also suggest the potential policy implications.

2 Specification of citation probability equation

To guide estimation and interpretation, we provide a simple model of the citation process, leading to a reduced-form estimating equation for the probability of one article citing another. We then specify the observed determinants of citation and a method for controlling for key unobservables.

At the article level, citation is a binary choice and we therefore follow the standard approach of defining a latent variable C_{id}^* which leads to a realized citation, $C_{id} = 1$ of paper d by paper i when a threshold κ is exceeded. Thus the probability of citation is $\mathbb{P}(C_{id}^* > \kappa)$. Articles should cite the relevant preceding work. However, author teams can only cite papers if they are aware of them. These truisms suggest that citation probabilities should be increasing in the product of relevance and awareness. We therefore model $C_{id}^* = A_{id}R_{id}$ where A_{id} denotes the level of awareness of citing team i of paper d and R_{id} scores the relevance of the content of paper d for paper i . The marginal effect of awareness is zero for irrelevant ($R_{id} = 0$) papers and the

marginal effect of relevance is zero under the condition of ignorance ($A_{id} = 0$).

We model awareness as an exponential function of a vector of indicators of geographic separation, \mathbf{G}_{id} , and of the educational and career linkages, \mathbf{L}_{id} , between members of the two author teams. Geographic proximity matters because it increases the frequency of face-to-face interactions (from “water-cooler” conversations to conference meetings). Information flows can overcome geographic barriers if authors of papers i and d are connected via overlapping career and/or educational histories. Past colocation or just indirect linkages such as having the same advisor at different times create a kind of connective tissue that facilitates knowledge flows. In summary we hypothesize that $\partial A / \partial \mathbf{G}^{(k)} < 0$ for all k elements of geographic separation and $\partial A / \partial \mathbf{L}^{(k)} > 0$ for all k indicators of ties between author teams.

We model relevance to depend on an article- d specific function of the subject area of the citing article, $s(i)$, the year the citing article is published, $t(i)$, and a random term, ε_{id} , representing idiosyncratic factors operating between the article pair. Thus, we have

$$R_{id} = \exp(\alpha_{s(i)t(i)d} + \varepsilon_{id}).$$

The d component of $\alpha_{s(i)t(i)d}$ embodies the *general* importance of article d to all mathematics articles. The “intellectual distance” between the subject of article i and article d enters via the $s(i)d$ component of α . The $t(i)d$ component captures the idea that relevance of article d to all subjects may decrease over time due to obsolescence of older ideas. The particular usefulness of the combined fixed effect is that it allows article d to have time-varying patterns of relevance that differ across subject areas.

We can take monotonic transformations of C^* and the threshold without affecting probabilit-

ities so we take logs, leading to

$$\ln C_{id}^* = \mathbf{G}'_{id}\boldsymbol{\gamma} + \mathbf{L}'_{id}\boldsymbol{\lambda} + \alpha_{s(i)t(i)d} + \varepsilon_{id}, \quad (1)$$

The probability of citation is the probability $C_{id}^* > \kappa$ and is given by

$$\mathbb{P}(C_{id} = 1) = \mathbb{P}(-\varepsilon_{id} < \mathbf{G}'_{id}\boldsymbol{\gamma} + \mathbf{L}'_{id}\boldsymbol{\lambda} + \alpha_{s(i)t(i)d} - \ln \kappa) \quad (2)$$

For ε distributed logistically with parameters μ and σ the probability of citation takes the familiar logit form:

$$\mathbb{P}(C_{id} = 1) = \Lambda[(\mathbf{G}'_{id}\boldsymbol{\gamma} + \mathbf{L}'_{id}\boldsymbol{\lambda} + \alpha_{s(i)t(i)d} - \ln \kappa - \mu)/\sigma], \quad (3)$$

where $\Lambda(x) = (1 + \exp(-x))^{-1}$.

We use logit as the primary estimator (and discuss linear probability model results in the Appendix robustness section D) since it constrains predicted citation probabilities to be non-negative. Logit coefficients provide the marginal effect on changes in the log odds. In the context of rare events such as citations, marginal effects on probabilities can be tiny. Singh (2005) multiplies his marginal effects by one million for reporting purposes. We find odds ratios are more intuitive, but as with rare diseases, one must keep in mind that a large odds ratio does not imply a large change in the probability of a positive outcome.

The $\alpha_{s(i)t(i)d}$ fixed effects are a critical part of our estimation strategy since there is no reason to expect the geography and ties variables to be orthogonal to the triadic relevance term. Indeed, it is likely that authors of more important articles would be better connected. Moreover, authors who tend to work on similar subjects are more likely to be connected. That

is, intellectual separation between $s(i)$ and article d may be negatively related to \mathbf{L}_{id} . We therefore estimate our model controlling for $\alpha_{s(i)t(i)d}$, the triad of subject of i , year of i , and article d . While we have modeled awareness as a function of geography and ties only, we could easily introduce $s(i)t(i)d$ effects and random article-pair effects. They would simply be incorporated into α and ϵ . This means, for example, that we allow for a completely general pattern of diffusion of awareness of article d on different subjects s .

Estimating $\alpha_{s(i)t(i)d}$ with a large number of articles is computationally difficult and raises concerns over the incidental parameters problem. Instead we take advantage of the logit feature that the total number of cites received by each triad is a sufficient statistic for $\alpha_{s(i)t(i)d}$. This permits estimation in terms of a conditional density to obtain consistent estimators of the γ and λ parameters. Prior work has included fixed effects for time lags (Singh, 2005), cited patents (Thompson, 2006), and cited institutions (Belenzon and Schankerman, 2013). This is the first study to control for the triad of citing article subject, citing article publication year, and cited article.

The unit of observation for citations is the *article* pair. However, the geography and ties variables underlying \mathbf{G}_{id} and \mathbf{L}_{id} are measured at the *author*-pair level. For multiple-author article pairs, we aggregate geography and ties of coauthors in the citing and cited author teams under the assumption of perfect information flow *within* teams. Specifically, the distance between an article pair is the minimum of the author-pair distances and the ties between an article pair are the maximum of the author-pair ties. Appendix A.5 provides greater detail and Appendix D shows that averaging geographic barriers and ties leads to similar results.

3 Data

Our data set combines four main sources. The Web of Science (WOS) provides citations, author affiliations, and keywords. The Mathematics Genealogy Project (MGP) tracks the place and time of doctoral education, as well as the names of the dissertation supervisor(s). Zentralblatt MATH classifies mathematical articles at the 5-digit level. Finally, the longitudes and latitudes for 1000 mathematics institutions used to calculate distance data between citing and cited author teams come from Google Maps. Further information on data sources is provided in Appendix A.1.

Insert Figure 1 here

Citations and distance data

Figure 1 provides a first look at the patterns in the WOS citation data and how they relate to geographic distance between authors. It graphs survival functions for citation flows as a function of distance between the authors of the citing and cited papers. $S^c(D)$ is the share of all cites that occur with distance $\geq D$. Citations from one nation to another are calculated by summing the citations from papers written by authors affiliated with institutions in country j to papers written by authors in country n .⁴

The benchmark for cites is a dartboard model that takes as given each country's outward citations, $C_j \equiv \sum_n C_{jn}$ and inward citations, $C_n \equiv \sum_j C_{jn}$. The international allocation of

⁴For papers with multiple authors from different countries, citations are allocated fractionally. Thus, a paper co-authored by two scholars from countries A and B to a paper written by two other authors from countries C and D would generate four international citation flows of 0.25 each. This fractional accounting of citations ensures that the sum of all citations in the world, C_w , is the same regardless of whether one sums across paper dyads or country dyads; that is $C_w = \sum_{jn} C_{jn} = \sum_{id} C_{id}$.

these citations is completely random; that is, each paper is equally likely to cite any other paper regardless of distance. Randomness implies that outgoing cites from j go to country n with probability given by n 's share of all received cites. Thus, the aggregate flow of benchmark cites from j to n is given by $F_{jn}^c \equiv C_j(C_n/C_w)$, where C_w sums all cites in the world. F_{jn}^c can be thought of as the “frictionless” flow of citations from j to n . The survival curve for the benchmark is $\bar{S}^c(D) = \left(\sum_{\text{dist}_{jn} \geq D} F_{jn}^c \right) / C_w$. Figure 1 displays $S^c(D)$ and $\bar{S}^c(D)$ using solid and dashed black lines. The vertical gap between $\bar{S}^c(D)$ and $S^c(D)$ measures the frictions that divert citations away from the dartboard benchmark.

The blue lines in Figure 1 permit comparison with actual and benchmark flows of trade in goods. Research using gravity equations has established that distance is a major friction impeding trade in goods.⁵ To facilitate comparisons with the citation data, we employ trade data sets that measure each origin's aggregate flows including those that remain within that origin. Thus, trade flows to self are value-added minus exports of value-added to the rest of the world.

Panel (a) displays flows of manufacturing value-added between and within 63 countries derived from the Trade in Value Added (TiVA) dataset made available through a joint effort of the World Trade Organization (WTO) and the Organization for Economic Cooperation and Development (OECD). One prominent aspect of Panel (a) is that a very large share of trade takes place within countries. The precipitous drops in the survival functions for both cites and goods seen at 1854km correspond to the CEPII internal distance of the United States (the average distance between 20 major cities).

⁵See Head and Mayer (2014) for explanation of the gravity methodology and results and Head and Mayer (2013) for a version of the distance distribution figure that considers only gross trade flows between countries.

To see what is happening within this important set of intra-national flows, we display the survival functions for state to state citations and trade in Panel (b). Citations are aggregated up to the state level just as they were for countries in Panel (a). The value of goods transported in 2007 between and within the 50 states and Washington, DC come from the Freight Analysis Framework (FAF) database. We display distances up to 5000km (excluding some Hawaii and Alaska dyads) because by that distance both benchmarks and actual flows are indistinguishable visually from zero.

What we learn from Figure 1 is that distance attenuates knowledge flows in mathematics leading them to occur over shorter ranges than one would expect in a frictionless world. This is true at international scale and also true within the United States. The gap between actual and benchmark citation flows is much smaller than what we observe for goods flows. This is consistent with the hypothesis that trade flows are attenuated by *both* transport costs and information decay. Moreover, distance decay effects in commercial activities may be larger than those that apply to researchers.

Author ties based on career and educational histories

The WOS contributes three indicators of ties based on past coauthors and past affiliations. Each tie variable is based on actions taken prior to the publication year of the relevant citing article. “Coauthors” indicates whether author pairs have collaborated on a paper published in one of the 255 math journals included in WOS since 1975. “Coincided past” requires colocation at the same institution in the *same year* but the authors no longer work at the same place. “Worked same place” indicates that two authors worked at the same institution in *different years* in the past.

The MGP data allow us to construct eleven additional binary ties based on three types of relationships. “Share Ph.D.” denotes author pairs who graduated from the same Ph.D. program

within a 5-year period and who are therefore assumed to have overlapped. There are four types of academic “relatives.” The first type are academic parents: “Advisor citing” which takes the value of 1 if the author of the citing article was the PhD advisor of the author of the cited article. For “Advisor cited” the citing author was the advisee. Academic siblings were both supervised by the same professor. Academic grandparents are the advisors of the citing or cited authors’ advisors. Academic cousins are authors that share a grandparent. Academic uncles are the advisees of one’s academic grandparent. The final category of educational links are via the Alma Mater. These indicate when the citing or cited author is affiliated to the institution where the other author received her PhD. For example, “Alma Mater cited” takes a value of 1 when a Princeton alumnus cites a professor currently affiliated with Princeton. The tie dummy variables are additive: if author i has co-authored with author d who is also i ’s Ph.D. advisor, there would be ones for both co-author and advisor cited.

Realized citations and control citations

The MGP sample we use in most estimations has 29,404 realized citations. The complete set of steps leading to this sample are described in Appendix A.2. To estimate the regressions, we need to combine the realized-citations set with a non-realized citation set. A standard “exogenous sampling” approach would entail picking a set of citing articles and constructing the universe of papers they might cite and predicting which potential cites are actually realized. Applying such an approach in the case of citations creates both conceptual and practical problems. First, it is hard to determine the appropriate “universe.” Should we consider the applied math papers that might have cited a given paper, the physics papers, the economics papers? The data gathering challenge for a true universe of potential citing papers would be formidable. There would also be computational difficulties with incorporating so many non-citation observations. Citations are an example of a rare event problem. In the Web of Science

sample (before imposing the requirement of MGP data on all authors), there are approximately 3 billion potential cites and about 269,000 realized cites. Thus, the rate of citation is only 9 per 100,000. In response to this problem, the patent citation literature has generally adopted a choice-based sampling approach following the matching methodology of Jaffe et al. (1993). For each realized citation (case), a single non-realized citation (control) is selected at random from a larger set of matched potential controls.⁶

We adopt the one case per control approach when using the whole WOS sample. However, the sample featuring our full set of ties has a small enough number of realized citations that we can incorporate *all* potentially cited papers that meet certain criteria. Our baseline matching criteria is that controls be published in the same year and the same 3-digit field as the original citing paper (case). The union of the realized citations and the control group constitutes the sample that is used in the econometric analysis.⁷ The presence of triadic fixed effects means that we have effectively the full set of control observations. To see this imagine another field A in which none of the papers cite a given paper d . Then the A - d part of the triadic fixed effect would be a perfect predictor for non-citation so all such observations would be automatically dropped from the fixed effects logit estimation. Appendix A.2 shows that the differences between the realized citation set and the control set are in line with our expectations.

⁶Singh (2005) uses five controls per realized citation in his weighted estimator.

⁷Kerr and Kominers (2015) use an alternative method that randomly samples patent distances to calculate *expected* citations within a fixed ring.

4 Regression results

This section presents the main results regarding the effect of geography and ties on knowledge flows. All regressions are logits with fixed effects for each group defined by citing field (3-digit subject codes), citing year, and cited article. Conditional on these fixed effects, variation in geography and ties is assumed to be random, allowing for a causal interpretation of the estimates. We recognize this is a strong assumption. Appendix B provides evidence that the 3-digit subject controls are effective at reducing bias due to endogenous ties. A large, highly significant association between ties and citations holds up with even the most stringent measure of subject (using the same keywords).

There are four key findings. First, the effects of distance, borders, and language differences are about half as strong once educational and career links are taken into account. Second, 13 of the 14 measures of ties have positive effects that are significant at the 5% level in our final specification. On average the effect of adding a tie raise the odds of citation by 80%, with some ties having much bigger effects. Third, ties and geography affect different types of papers differently. In particular, less prominent and more recently published papers exhibit stronger effects. Finally, while the importance of distance has declined to the point of statistical insignificance in recent years, ties remain as valuable as ever. Appendix D shows the robustness of our main results to using alternative subsamples and specifications.

4.1 Baseline

Insert Table 1 here

Table 1 reports the result of baseline logit coefficients, which have the interpretation of marginal effects on the log odds. Statistical significance is calculated using standard errors that

are clustered at the cited article level to allow for correlations in the errors across potentially citing articles for the same cited article.⁸

The first specification includes only the four geographic explanatory variables: an indicator for distance greater than zero (not being at the same institution), log distance (interacted with the positive distance indicator), and indicators for residing in different countries and from countries that have different official languages. The two-part distance function is necessary because there is no good way to directly measure the distance between two scholars at the same institution. The first of the two parts implicitly estimates this distance. The indicator for distance greater than zero is equivalent to a “different university” dummy. The two-part formulation has a jump from zero to positive distances, but thereafter the elasticity of citations odds with respect to distance is constant. While a constant elasticity of distance in trade equations is the standard assumption underlying gravity equations, there is little *a priori* reason to expect this relationship to carry over to citations. Therefore we re-estimate specifications (3) and (5) with distance-interval step functions in columns (4) and (6).

The second specification adds ties constructed from the WOS database. The third to sixth specifications restrict the sample to articles with full information from the MGP database. The overall estimating sample does not decline much because the MGP sample uses all available controls (non-citations in the same subject-year), whereas the WOS sample has just one control per case. As in the first two columns, we show the effects of geography without ties (columns 3 and 4) and then add the full set of ties available in the MGP data (columns 5 and 6).

Specification (1) presents significantly negative coefficients on distance and borders, sug-

⁸The table is replicated with standard errors in parentheses in Appendix Table E.1. The entire table is re-estimated using a linear probability model in Appendix Table D.3, with the results compared in Appendix section D.

gesting that physical distance and borders indeed impede knowledge flows. We estimate smaller border and distance effects than those obtained by Singh and Marx (2013) using citations of US patents. Whereas we find that crossing a national border reduces citation odds by $\exp(-0.198) - 1 = -18\%$ they find a 41% reduction (specification (6) of Table 5). Our distance elasticity is -0.073 whereas theirs is -0.137 . While it is tempting to attribute this halving of geography effects to differences between academic and commercial diffusion of ideas, other evidence on patent effects obtains similar magnitudes to our column (1). As the coefficients show the marginal effects on the log odds of citation and citation is rare, the dependent variable approximates the log probability and should therefore be proportionate to the log citation flow in aggregated data. This means we can compare our estimates directly to the results from the gravity-type regressions on patent citations estimated by Peri (2005) and Li (2014). The different country (border) effect we estimate is -0.198 , whereas the baseline estimate of Peri (2005) is -0.19 . Li (2014), also estimating a patent citation gravity equation, reports distance elasticities (after controlling for subnational borders) from -0.03 to -0.067 , which are slightly weaker than those reported in our column (1).

All these results support the conclusion that border and distance decay of citations are considerably smaller than the effects typically estimated for trade in goods. Nevertheless, it may be surprising to many that geography has a significant impact on academic citations at all. We now show that the estimated effects are substantially reduced by controlling for ties.

The second specification shows that the three measures of career ties (past coauthorship, past colocation, and past work at the same institution) all have strong positive associations with citation. As exponentiating the coefficients in a logit expresses the effects in terms of the change in citation odds ratios the 0.712 coefficient on past colocation implies that even after colleagues have moved to separate institutions, they have 104% higher odds of citing each other

$(\exp(0.712) - 1 = 104\%)$. Prior coauthors are even more likely to cite each other. We also see that the inclusion of career ties lowers geography effects somewhat.

Comparing columns (1) and (3) we see that estimating the same specification on the MGP-restricted sample does not change the geography coefficients by more than one would expect given the standard errors.⁹ Comparing columns (3) and (5) we see one of the headline results of this paper: Controlling for ties shrinks the negative effects of geographic separation by about 50%. The ratios of the four geography coefficients in column (5) to the corresponding coefficients of column (3) are 0.46, 0.54, 0.39, and 0.30. The omitted variable bias formula tells us that this means that ties and geography are correlated and that the pure partial effect of being far away or in a foreign country is overestimated in regressions that omit controls for ties.

Table 1, columns (5) and (6) show that, with just one exception, ties have systematically positive effects on citation probability. All of the estimates are statistically significant at the 5% level except “grandparent citing” which has an imprecisely measured negative effect and “uncle citing” which has a borderline significant result in column (5). The average over all fourteen ties coefficients is 0.59, implying that the average tie raises the odds of citation by 80%. The addition of the full set of ties in column (5) dramatically increases the fit of the logit to the data: the pseudo R^2 nearly triples from 0.033 to 0.091.¹⁰

⁹Additional investigation of the possibility of MGP sample selection bias is reported in the robustness section D in the Appendix.

¹⁰Pseudo R^2 is measured as $1 - \mathcal{L}_1/\mathcal{L}_0$ where \mathcal{L}_0 is the likelihood of the constant-only model. Hence it rises with the number of estimated parameters. It is therefore worth noting that the inclusion of ties reduces the Akaike Information Criterion (AIC) by 7995 points compared to column (3), indicating that the rise in the likelihood from adding ties is large enough to offset

There are three tie relationships where one can identify the more “senior” of the two authors: advisors, uncles, and grandparents. In each case we observe the author in the teaching role is more likely to be cited than to cite. Advisees massively over-cite their advisor’s papers by a factor of four (the second largest impact of the 14 types of ties). In the reverse direction, we find advisors over-cite their advisees’ articles by a factor of three. The academic “nephew” overcites his “uncle” by 86% but the reverse direction features a bias of just 26%. The most pronounced asymmetry emerges when we skip a generation. Authors over-cite their advisor’s advisors (academic grandparents) by a factor of three. Yet this intergenerational flow is not reciprocated; the grandparents’ propensity to cite advisees of their advisees is not significantly different from zero. These vertical patterns support the hypothesis that citations transmit knowledge.

Insert Figure 2 here

Figure 2 illustrates the coefficients on each of the 12 steps in the non-parametric estimation of distance effects conducted in specification (6) of Table 1, represented with black circles. The vertical axis depicts the reduction in the log odds of citation associated with each step, relative to working at the same institution. We also show with blue squares the corresponding estimates for the 12-step specification omitting ties. For each set of steps, we overlay the implied reduction in the log odds of citation based on the 2-part coefficients from specifications (3) and (4). The key finding illustrated in the figure is that after the dramatic fall associated with positive distance, the subsequent declines are consistent with a constant elasticity decay rate. Controlling for ties moves the decay function up (lower effect of being at different institutions) and flattens it. After controlling for ties, the two-part prediction lies within two standard errors for 11 out of 12 steps.¹¹ Clearly there is a big discontinuity between zero and positive distances

the penalty AIC imposes for adding 14 parameters.

¹¹The exceptional case is the 25–50km bin, which is driven by the dyad Rutgers-CUNY

corresponding to a same-university effect. Conditional on positive distance, the figure shows that it is hard to distinguish empirically between a decay function that is flat after 1000 kilometers and one that exhibits regular decay with a constant elasticity of -0.037 . Since the 2-part approach adequately captures distance effects, we use it for all the subsequent estimations.

The negative effect of geographic barriers on citation probabilities is presumed to arise because these barriers reduce the frequency of face-to-face interactions. In academics (as well as other areas) co-attendance at conferences provides one of the most important opportunities to meet in person with scholars doing related work. We collected data on papers presented between 1990 and 2009 at one of the most important conferences, the Joint Mathematics Meetings (JMM). Held annually in the United States, an average of 1459 participants present 1037 papers.

The first exercise we conduct, reported in Table C.1, is to show that there is a strong and precisely estimated negative effect of distance on the probability of attending a conference. Because the conference venue moves each year, the data exhibit substantial variation in distance for a given scholar. This permits estimation of the logit with author-specific fixed effects. The distance elasticity in this specification is -0.136 with a 0.016 standard error (clustered by author).¹²

(45km apart). Both of these math departments are very active in the Set Theory 3-digit code but they do not cite each other's papers. The apparent cause is that while Rutgers papers span the field, CUNY authors specialize in two sub-fields, Consistency and Independence Results and Large Cardinals, which comprise 52 out of CUNY's 58 papers.

¹²This estimation includes only those authors who attended at least one meeting but not every one (no perfect predictors). Appendix C also presents an estimation without author fixed effects that includes all potential attendees. The distance effect in this estimation is not as strong (-0.05) but negative and significant border and language effects show up in this specification.

The second exercise, reported in Table C.3, uses the conference data to show the impact of attendance on citation. Using the Table 1 column 5 specification, we add indicators of coinciding at the same conference (as presenters or session organizers). While just coinciding has a negligible effect on citation, coinciding when the (potentially) cited paper is presented increases the odds of citation by a factor of 8.3. Table C.3 also shows a positive effect of presenting at the same session, regardless of whether it was the citing or cited paper. Contrary to our own observation of presenters being encouraged to cite the work of co-attendees at a session, we find no significant evidence of a citing paper effect in this data. While these two exercises are confined to the one conference for which we could obtain long-term conference participation data, they illustrate a broader mechanism that we view as underlying distance effects on citation. Proximate authors are more likely to present at the same conferences and, when they do so, this makes the citing authors aware of new relevant research which they build upon in their own work.¹³

Insert Figure 3 here

Why does controlling for academic linkages lead to the large reduction in distance effects shown in Table 1? It must be that ties are negatively correlated with geographic barriers. We illustrate this in Figure 3(a), which shows that linked authors tend to be closer to each other than authors who have no ties. For example, about 33% of tied authors are more than 5000 kilometers apart, compared to almost 60% of non-tied authors. Similarly, tied authors are much more likely than non-tied authors to reside in the same country (51% vs 16%) or countries that share a common language (65% vs 32%).

¹³Our results align with the finding of Iaria et al. (2018) that the ban on Central scientists from participating at international conferences during and after World War I was associated with a drop in citations between Allied and Central scientists.

Figure 3(b) reveals the phenomenon that helps to understand our baseline results: Mathematicians tend to remain close to the university where they obtained their doctorates. Thirty percent either do not leave or have returned and only 18% move more than 5,000km away.¹⁴ Proximity to the *alma mater* is likely to beget proximity to one's advisor (and his advisor), former classmates, etc. The story underlying our baseline results is a simple but important illustration of omitted variable bias. Ties are very important for citation but ties are negatively correlated with distance. Thus a failure to control for ties leads to the inference that distance has a greater *direct* impact on knowledge flows than is truly the case. Authors are unlikely to cite papers written by faraway authors partly because they are less likely to have interacted at conferences, but an equally important factor is that they are less likely to have an academic or career tie with each other.

4.2 Evidence for information mechanisms

The results we have obtained so far point to an important role for educational and career ties in fostering citations. The underlying mechanism we imagine is one of communication along the network of ties that causes one set of authors to become aware of useful theorems and conjectures provided by other authors. This information transfer mechanism predicts that the presence of ties should matter more for certain types of papers than others. Specifically, we conjecture that authors rely more on their ties to find out about work that is less widely known, more recently written, and further from the author's expertise. To the extent that face-to-face interactions matter more for such papers, geographic barriers should be stronger as well.

¹⁴There is substantial heterogeneity in the tendency to work at the Ph.D. granting institution, with the just 20% of US-educated authors staying/returning compared to 46% in Spain. The sample comprises 2213 MGP authors who published in pure mathematics journals in 2009.

Insert Table 2 here

We develop three proxies for papers that researchers are less likely to know about. First, we categorize papers as “obscure” if they receive less than or equal to the median number of cites (three). Our second proxy for low awareness is the gap in time between when the citing and potentially cited papers were published. A paper is “recent” if the gap is less than or equal to the median gap in our data (nine years). The third awareness measure follows from the observation that authors are more familiar with work in their own fields than in other subject areas. We classify papers as different field if their 2-digit mathematical subject classifications (MSC) differ (for example 11 Number Theory vs 14 Algebraic Geometry). As we show in Table D.4 in the Appendix, these specific rules for categorizing obscure, recent, and different field are not critical for the results.

Table E.2 in the Appendix provides summary statistics on these variables. Not surprisingly, there are lower average number of cites for obscure papers and recent papers. We see approximate balance between the average number of cites to the same and to different fields. There are more observations in total featuring cites within the same field so this suggests that cross-field citations go mainly to more prominent papers. In terms of ties, on average the differences between obscure and recent papers are small. The fact that ties are higher for same-field papers probably reflects greater ties within the same field. This is an important reason why our fixed effects control for the citing paper’s 3-digit subject code.

Table 2 reports the detailed results for the three awareness proxies. To reduce the number of parameters to be displayed and discussed, we report the average of 13 ties indicators, followed in the next column by the averages over 13 interaction terms.¹⁵ Column (1) reports the

¹⁵We drop “grandparent citing” in this table because of a logit perfect predictor problem. In the different-field specification, there were only 9 grandparent citing instances and all of them

corresponding regression *without* interactions for comparison purposes.

The first set of interactions in Table 2 shows the results of interacting geography and ties with an indicator for obscure papers. Column (2) shows the base effects corresponding to non-obscure papers and column (3) shows the coefficient on each corresponding interaction. We find that the more prominent papers (> 3 cites) have a 18% ($= 0.135/(0.619 + 0.135)$) smaller coefficient on the average of ties than the lesser known papers. This is consistent with the interpretation that ties facilitate awareness. Papers that are big successes require less help from networks to promote transmission. The coefficient on log distance is about 1.2 times as large for obscure papers.

When the interaction is changed to distinguish recent versus older papers, the results are similar as shown in columns (4) and (5). Recent papers have a 24% higher coefficient on the average effects of ties. Distance decays are estimated at $-0.021 - 0.028 = -0.049$ for papers in their first nine years after publication (the median age of papers in our sample) and -0.021 thereafter. These numbers are remarkably similar to those reported by Li (2014) in a gravity-style study of inter-city patent citation flows. She finds that the distance elasticity declines monotonically with age from a -0.028 in the first five years to -0.014 for patents granted 20 or more years before. These combined findings of significantly higher geographic concentration of “new knowledge” are intuitively appealing and provide some guidance for models of knowledge diffusion.¹⁶

were for control observations, rather than realized cites. We reinstate grandparent-citing back in a robustness check (Table D.4) where it is a component in a sum of ties variable.

¹⁶A recent paper studying patents finds corroborating results. Packalen and Bhattacharya (2015) show that denser cities are responsible for patents that make use of newer knowledge, as measured by textual analysis of the patent applications.

Ties also have larger impacts for papers in different fields, with a coefficient, reported in column (7) that is 28% larger than for same-field papers. None of the different-field geographic interactions are statistically significant, suggesting that face-to-face communication matters more for obscure and recent papers than for different fields.

All three sets of interactions therefore support the premise that *scholars draw more heavily on their connections when obtaining less familiar information*. The positive interactions between ties and information proxies have similar magnitudes and strong statistical significance in five alternative specifications described in the Appendix section D. These robustness regressions also find statistically significant (10% or better) negative effects for the geography-information interactions in 11 out of 30 estimates. The remaining estimates are mainly negative but not statistically different from zero.

Tables 2 and D.4 show strong and robust evidence that ties matter more for three types of papers where awareness poses a more serious challenge. We also find that geographic barriers pose a greater impediment to citation for recent papers in nearly every specification. This evidence supports the interpretation of ties as facilitating information transfer rather than an alternative mechanism involving “citation cliques.” Under this alternative, scholars have perfect awareness of the relevant research in their field but choose to cite specific prior work because it was written by the scholars for whom they have some kind of social affiliation. If ties are just proxies for intra-group loyalties, it is not obvious why such forces should be relatively more important specifically for the types of papers where the awareness gap is predictably larger.¹⁷

There is a third mechanism, combining elements of information and affiliation, which is also consistent with our results. In this story, mathematicians are aware of relevant work but

¹⁷The fixed effects control for the overall tendency to cite each article d so the interactions measure how ties boost the *relative* tendency to cite specific types of papers.

uncertain of whether the proofs those papers contain are all correct. Since the validity of one's own results hinges on the correctness of the proofs of the cited theorems, the mathematicians we have spoken to claim to check all proofs, regardless of the author. In practice, this may not always occur. There could be cases where, for example, an author would cite her advisors papers because she knows his proofs have always stood up to scrutiny. This trust mechanism would likely be stronger for lesser known and more recent papers because they are less likely to have been thoroughly checked by others. Trust could also matter more for papers outside one's field because those involve unfamiliar techniques that make it difficult for an outsider to verify the proof.

We see the awareness and trust mechanisms as both emphasizing ties as conduits of information. In the first case, the information is about the *existence* of a useful theorem; in the second case the information is about the *reliability* of the theorem. This echoes the situation in international trade where Rauch (2001) summarizes a number of studies showing that “transnational business and social networks promote international trade by alleviating problems of contract enforcement and providing information about trading opportunities.” Thus, networks help exporters by making them aware of the specific needs of foreign buyers, while also promoting trust that buyer and seller will comply with the terms of their contract with each other.

4.3 Time-varying effects of distance and ties

The estimates presented so far pool citations made from 1980 to 2009. This section investigates whether the effects of distance and ties on more recent citations differ from the past. The results of Keller (2002) and Griffith et al. (2011) show a decline in the importance of geographic separation between the 1980s and late 1990s. We extend the investigation these authors initiated by including more recent data and also estimating the time-varying effect of ties. We examine

changes since 1990 because our 1980s citation data are too sparse. Estimation of time-varying coefficients from 1990 to 2009 is of great interest given the many relevant advances observed over this period.

Insert Figure 4 here

To investigate whether the impact of distance and ties have been changing, we estimate regressions based on a moving sample window. We construct the estimation windows by first restricting the citing papers to be published within a 5-year period centered around year t . This implies citing years, t_c , in the interval $t + 2 \geq t_c \geq t - 2$. To make the sample size in later years comparable to that of earlier years, we impose a fixed maximum citation lag L set equal to 5 or 10 years. This implies cited years, t_d in the interval $t_c \geq t_d \geq t_c - L$. The first mid-year t we use is 1990 and the last is 2007 (since our data set runs to 2009).

Figure 4 shows the effects of distance in panel (a) and ties in panel (b). In both panels we use blue squares to depict the point estimates for 10 year maximum citation lags. A solid LOWESS smoother passes through the point estimates. The dashed smoother line depicts the results for a 5-year citation lag and the dot-dash line corresponds to an estimation with no restriction on citation lag (all years). The points in panel (a) are estimated distance elasticities, that is, the marginal effect on the log odds of citation of increasing log distance between citing and cited authors. The time-pattern of distance effects depends on whether the regressions controls for ties or not. We depict these differing results using blue for estimates that control for the sum of 14 ties and red for those that do not. 95% confidence intervals (as before standard errors are clustered at the cited article level) are shaded blue and red for estimates that do and do not (respectively) control for the sum of ties.¹⁸

¹⁸The purple area corresponds to the intersection of the two intervals.

All the specifications plotted in Figure 4(a) show absolute distance elasticities becoming much smaller over time since the early 1990s. In the geography-only specification shown in red, distance remains a statistically significant impediment to citation up to and including the final interval, 2005–2009, when its elasticity is -0.06 (standard error: 0.013). However, the magnitude falls by two thirds from its 1990 value of -0.18 . The confidence intervals also shrink over time, since increasing numbers of digitalized articles raise \sqrt{N} in the standard error calculation. Controlling for the sum of ties, we see the absolute elasticities are uniformly smaller in all periods, with the largest gap between the smoother lines appearing in the last estimation windows. Starting around 2005, the confidence intervals mainly include zero. The final estimated distance elasticity controlling for ties is -0.017 (standard error: 0.012).

Panel (b) of Figure 4 shows the evolution of the coefficient on the sum of ties. The impact of ties on citation has been mainly rising over the 1990s and 2000s. The increase in citation odds from adding a tie rises from 72% in 1990 to 94% in 2007.¹⁹

In all the results presented to this point we have used a world-wide sample. This contrasts with much of the work we cited in the introduction on the geography of knowledge flows that uses citations *within the United States*. It is therefore worth investigating whether the patterns shown in Figure 4 reflect global phenomena or whether the US is special.

Panels (c) and (d) of Figure 4 graph the results for a moving-window specification similar to those depicted in panels (a) and (b) except they estimate separate distance > 0 , and (sum of) ties coefficients for pairs where both i and d are US residents and others ($1 - \text{both US}_{id}$). Panel (c) reveals that the shrinking distance effects depicted in panel (a) derive from author pairs where at least one set is not US-based. Distance effects between US pairs have not been

¹⁹Exponentiate the 10-year lag coefficients shown in Figure 4 and subtract one to obtain these amounts.

significantly different from zero throughout the period of study. Another difference between US pairs and others is that the latter exhibit rising effects of ties, becoming significantly larger than those between US pairs since the 2000s.

It is obviously tempting to try to explain the temporal patterns in the coefficients with reference to the technological advances we have observed since the 1990s. However, it is not possible to identify one cause or the other with so many trends at work during this period. Advances affecting information flows include, but are not limited to, the rise of web browsers in the mid 1990s, and the introduction of the Google search engine in 1998 and Google Scholar in 2004. Of particular importance to scientists was the creation of arXiv.org, a repository of pre-prints, which has included mathematics since 1992.

Insert Figure 5 here

Figure 5(a) plots the growth of the number of arXiv papers in mathematics over time and compares (it on a second scale) with the spectacular increase in Google searches in the 2000s. Panel (a) also depicts the introductions of Skype and Google Scholar. The combination of all these technologies would be expected to have reduced the importance of face-to-face interactions, implying declining geographic separation effects since 1990.

The smoother line for the distance elasticity in Figure 4(a) begins to trend up in the late 1990s, coinciding with the rise of arXiv shown in Figure 5(a). However, the stable importance of ties between US authors and the growing role of ties elsewhere shown in Figure 4(d) is not consistent with the view that arXiv and Google searches have been making all information universally accessible. Furthermore, the rise of internet article depositories and search engines cannot explain why distance effects between US author teams have been insignificantly different from zero during the whole period.

While internet advances capture the most attention, other contemporaneous changes could

reasonably affect the importance of geography and ties in knowledge transmission. Figure 5(b) shows the dramatic decline in the costs of making international calls to and from the US. In real terms international calls fell by 95% between 1990 and 2007, compared to a 75% decline over the same period for interstate calls. As we cannot find comparable series on domestic and international air fares, we use data on the volume of travel as a proxy. Figure 5(c) shows that air travel has been rising relative to the population size.²⁰ The rate of growth outside the US has been much larger, partly because the US started from a much higher base. In 1990 the US had 1.3 air passengers per capita compared to 0.11 for the nine other countries that comprise the top 10 countries in mathematics (measured by number of citing authors in 2009). Two decades later the US ratio had risen by 20% whereas the other countries rose by 140%.

The data shown in Figures 5(b) and (c) suggest an alternative interpretation of advances since the 1990s. Perhaps cheaper phone calls and improved air travel makes it easier for scholars to stay in touch with their ties. Improved contact allows them to share the kind of complex knowledge that is hard to procure via Google searches. Thus communication cost reductions lower the need for face-to-face interactions but raise the opportunities for drawing upon one's ties.²¹ Similarly lower costs for flying to conferences or visiting collaborators could also contribute to the explanation of why distance matters less, but ties matter more.

The greater drop in the effect of distance and the larger increase in the effect of ties for non-US based authors is in line with the big decline in call costs between the US and other countries shown in Figure 5(b) and the rise of air travellers per capita in the rest of countries relative to the US shown in Figure 5(c). While we find this story linking the coefficient patterns

²⁰Data from World Development Indicators series "Air transport, passengers carried."

²¹This story is consistent with the model of complementarity between proximity and communication technology in Gaspar and Glaeser (1998).

in Figure 4 to the trends in Figure 5 to be plausible, future work with different identification strategies would be needed to confirm it.

5 Conclusion

Our results add further evidence to the diverse strands of the literature finding geographic separation impedes knowledge flows. Geography matters in large part because of its role in shaping the personal ties between citing and cited scholars. In the full sample, including 14 linkages based on career and educational histories as controls cuts geography coefficients approximately in half. For the subsample where both citing and cited authors reside in the US, a region where communication and travel costs have long been relatively low, the marginal effect of greater distance between institutions is insignificantly different from zero. The distance effect also disappears in the most recent five years of the world-wide sample. These “zero” partial effects of distance are obtained only in those regressions that control for ties.

Despite the increase in global access to knowledge provided by the internet, the strength of the impact of ties on citation probabilities has not been declining. Because ties matter most for papers where awareness gaps are most acute (recent, obscure, and different-field articles), we infer that ties matter because connected scholars transmit knowledge to each other. This view is further supported by the finding that scholars whose formal role is to impart knowledge (advisors and the academic parents and siblings of advisors) have larger impacts on subsequent academic generations than vice versa. In sum, the evidence suggests that “what you know” depends a great deal on “whom you know.” It is increasingly unrelated to “where you work”—except insofar as where you work influences whom you know.

To the extent that we can generalize from the study of mathematicians, our study suggests

novel interpretations of existing empirical findings. Cities may be valuable not just because of daily face-to-face interactions, but also because they are good places to build networks. Such a view points to a different takeaway from the De La Roca and Puga's (2017) finding that wages rise with experience in big cities but retain much of this growth even when the individual returns to a smaller city. While the authors attribute the wage premium to increased ability, our framework suggests it might also have arisen via an expanded set of professional ties. Since ties created at short distances can be maintained over longer distances, the ties explanation is also consistent with the continued prosperity of city leavers.

In trade, Feyrer (2009) estimates that changes in distance caused by the Suez canal closure have much lower impacts than cross-sectional differences in distance. Our interpretation would be that the lengthening of the shipping route has no impact on the ties between importer and exporter that predated the closure. The puzzle posed by Head and Mayer (2013) calculations that observable barriers such as tariffs and freight charges can only explain less than half the estimated magnitudes of border and distance effects has a simple resolution in light of our results. Traders depend on their networks and those networks are nationally and spatially biased.

It bears repeating that the broader lessons we draw from observing mathematicians are necessarily tentative; they beg for corroboration in other contexts. This is especially true when it comes to policy implications. However, a ties-centered view of knowledge flows does suggest certain types of government actions could be fruitful. To promote more geographically dispersed networks, universities could be strongly discouraged from hiring their own students straight out of graduate school. Another policy to broaden ties of researchers is for the government to fund and promote doctoral study abroad. Invitations to foreign faculty for short and long term visits often lead to the formation of new collaborative ties. Analogous versions of these policies can expand the networks for non-academics. For example, easing visa require-

ments to facilitate medium-run stays by employees of multinationals should thicken the set of connections between foreign and domestic knowledge creators.

References

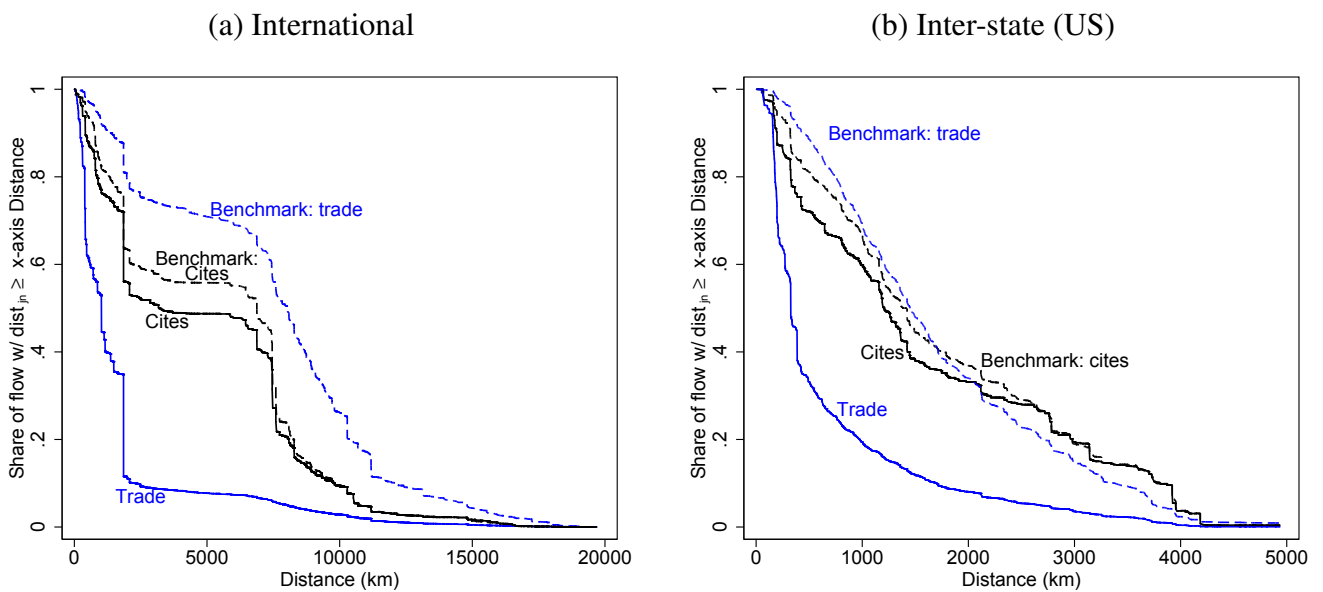
- Agrawal, A., Cockburn, I., and McHale, J. (2006). Gone but not forgotten: Knowledge flows, labor mobility, and enduring social relationships. *Journal of Economic Geography*, 6(5):571–591.
- Agrawal, A., Kapur, D., and McHale, J. (2008). How do spatial and social proximity influence knowledge flows? Evidence from patent data. *Journal of Urban Economics*, 64(2):258–269.
- Allen, T. (2014). Information frictions in trade. *Econometrica*, 82(6):2041–2083.
- Althouse, B. M., West, J. D., Bergstrom, C. T., and Bergstrom, T. (2009). Differences in impact factor across fields and over time. *Journal of the American Society for Information Science and Technology*, 60(1):27–34.
- Arrow, K. J. (1969). Classificatory Notes on the Production and Transmission of Technological Knowledge. *American Economic Review*, 59(2):29–35.
- Belenzon, S. and Schankerman, M. (2013). Spreading the word: Geography, policy, and knowledge spillovers. *The Review of Economics and Statistics*, 95(3):884–903.
- Borjas, G. J. and Doran, K. B. (2012). The collapse of the Soviet Union and the productivity of American mathematicians. *The Quarterly Journal of Economics*, 127(3):1143–1203.
- Breschi, S. and Lissoni, F. (2009). Mobility of skilled workers and co-invention networks: An anatomy of localized knowledge flows. *Journal of Economic Geography*, 9(4):439–468.

- Combes, P.-P., Lafourcade, M., and Mayer, T. (2005). The trade-creating effects of business and social networks: Evidence from France. *Journal of International Economics*, 66(1):1–29.
- Comin, D. A., Dmitriev, M., and Rossi-Hansberg, E. (2012). The spatial diffusion of technology. Working Paper 18534, National Bureau of Economic Research.
- De La Roca, J. and Puga, D. (2017). Learning by working in big cities. *The Review of Economic Studies*, 84(1):106–142.
- Ellison, G., Glaeser, E. L., and Kerr, W. R. (2010). What causes industry agglomeration? Evidence from coagglomeration patterns. *American Economic Review*, 100(3):1195–1213.
- Federal Communications Commission (2010). Trends in telephone service. Statistical report, Wireline Competition Bureau, Washington, D.C.
- Feyrer, J. (2009). Distance, trade, and income—the 1967 to 1975 closing of the Suez Canal as a natural experiment. Technical Report 15557, National Bureau of Economic Research.
- Gaspar, J. and Glaeser, E. L. (1998). Information technology and the future of cities. *Journal of Urban Economics*, 43(1):136–156.
- Glaeser, E. (2011). *Triumph of the city: How our greatest invention makes US richer, smarter, greener, healthier and happier*. Pan Macmillan.
- Gould, D. M. (1994). Immigrant Links to the Home Country: Empirical Implications for U.S. Bilateral Trade Flows. *The Review of Economics and Statistics*, 76(2):302–316.
- Griffith, R., Lee, S., and Van Reenen, J. (2011). Is distance dying at last? Falling home bias in fixed-effects models of patent citations. *Quantitative Economics*, 2(2):211–249.

- Head, K. and Mayer, T. (2013). What separates us? Sources of resistance to globalization. *Canadian Journal of Economics*, 46(4):1196–1231.
- Head, K. and Mayer, T. (2014). Gravity equations: Workhorse, toolkit, and cookbook. In Helpman, E., Gopinath, G., and Rogoff, K., editors, *Handbook of International Economics*, volume 4. Elsevier.
- Iaria, A., Schwarz, C., and Waldinger, F. (2018). Frontier knowledge and scientific production: Evidence from the collapse of international science. *The Quarterly Journal of Economics*.
- Jaffe, A. B., Trajtenberg, M., and Henderson, R. (1993). Geographic localization of knowledge spillovers as evidenced by patent citations. *The Quarterly Journal of Economics*, 108(3):577–98.
- Keller, W. (2001). The Geography and Channels of Diffusion at the World's Technology Frontier. NBER Working Papers 8150, National Bureau of Economic Research, Inc.
- Keller, W. (2002). Geographic localization of international technology diffusion. *American Economic Review*, 92(1):120–142.
- Keller, W. and Yeaple, S. R. (2013). The Gravity of Knowledge. *American Economic Review*, 103(4):1414–44.
- Kerr, W. R. (2008). Ethnic Scientific Communities and International Technology Diffusion. *The Review of Economics and Statistics*, 90(3):518–537.
- Kerr, W. R. and Kominers, S. D. (2015). Agglomerative Forces and Cluster Shapes. *The Review of Economics and Statistics*, 97(4):877–899.
- Kerr, W. R. and Mandorff, M. (2015). Social Networks, Ethnicity, and Entrepreneurship. NBER Working Papers 21597, National Bureau of Economic Research, Inc.

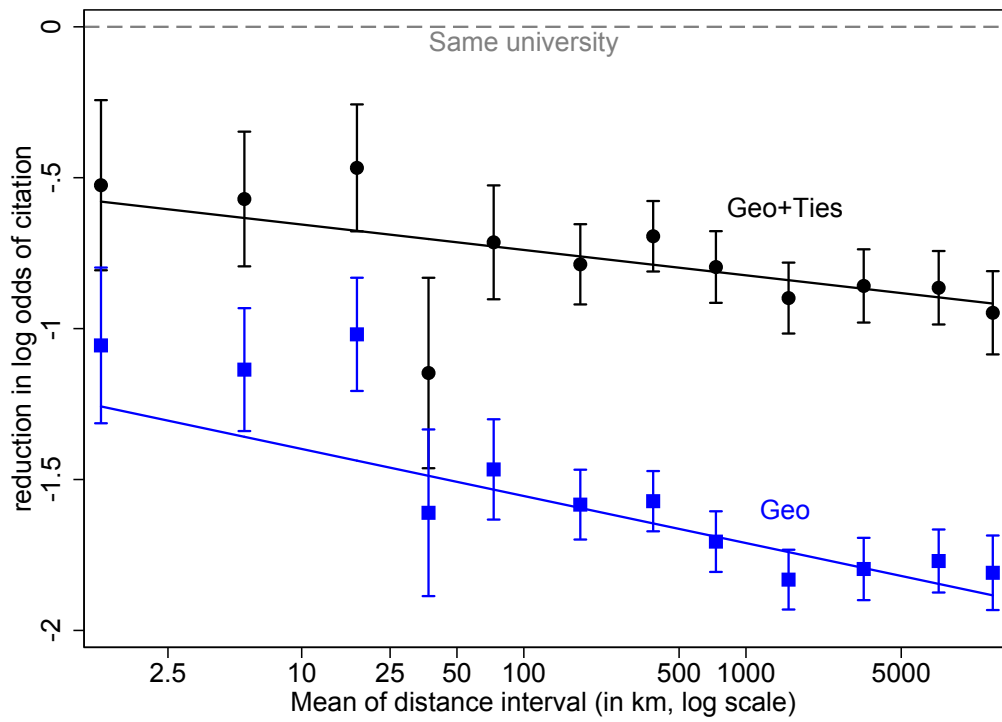
- Li, Y. A. (2014). Borders and distance in knowledge spillovers: Dying over time or dying with age? Evidence from patent citation. *European Economic Review*, 71:152–172.
- Lissoni, F. (2001). Knowledge codification and the geography of innovation: The case of Brescia mechanical cluster. *Research Policy*, 30(9):1479–1500.
- Packalen, M. and Bhattacharya, J. (2015). Cities and ideas. Technical report, National Bureau of Economic Research.
- Peri, G. (2005). Determinants of knowledge flows and their effect on innovation. *The Review of Economics and Statistics*, 87(2):308–322.
- Rauch, J. E. (2001). Business and Social Networks in International Trade. *Journal of Economic Literature*, 39(4):1177–1203.
- Rauch, J. E. and Trindade, V. (2002). Ethnic Chinese Networks In International Trade. *The Review of Economics and Statistics*, 84(1):116–130.
- Singh, J. (2005). Collaborative networks as determinants of knowledge diffusion patterns. *Management science*, 51(5):756–770.
- Singh, J. and Marx, M. (2013). Geographic constraints on knowledge spillovers: Political borders vs. spatial proximity. *Management Science*, 59(9):2056–2078.
- Thompson, P. (2006). Patent citations and the geography of knowledge spillovers: Evidence from inventor- and examiner-added citations. *The Review of Economics and Statistics*, 88(2):383–388.
- Waldinger, F. (2010). Quality matters: The expulsion of professors and the consequences for PhD student outcomes in Nazi Germany. *Journal of Political Economy*, 118(4):787–831.

Figure 1: Distribution of distances for math citations and trade in goods



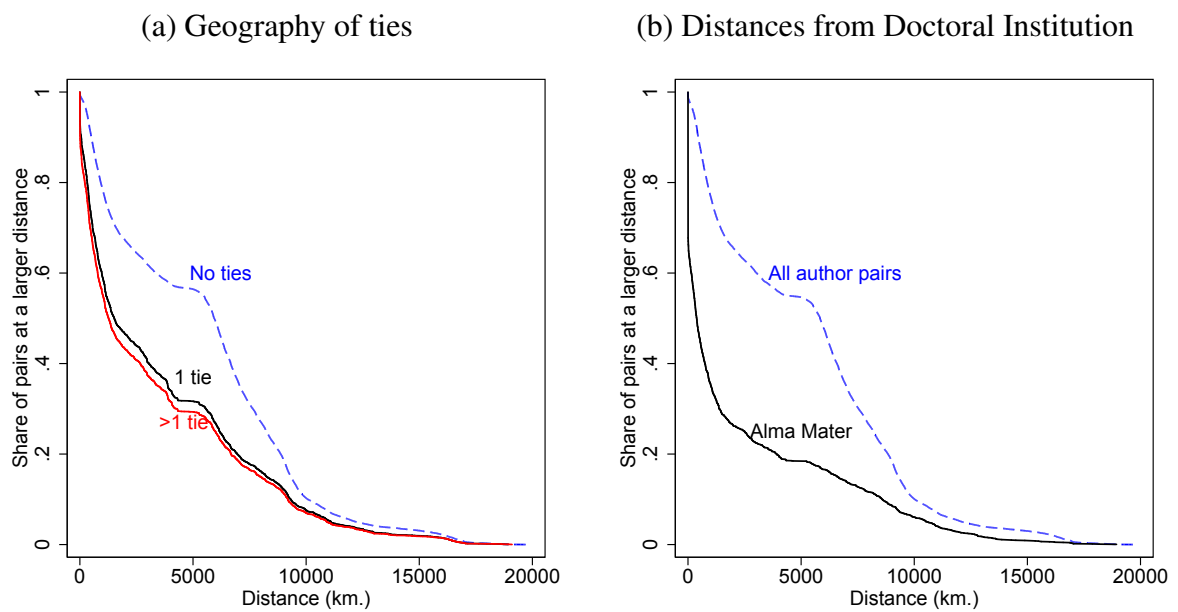
Note: Solid lines represent actual flows of citations or goods between origins and destinations. International flows of value-added in manufacturing (blue lines in left figure) come from the OECD/WTO TiVA database (2009) whereas state-to-state flows come from the FAF database (2007). Dashed lines are “frictionless” benchmarks for citations or trade (calculation detailed in the text).

Figure 2: Non-parametric estimated geography effects



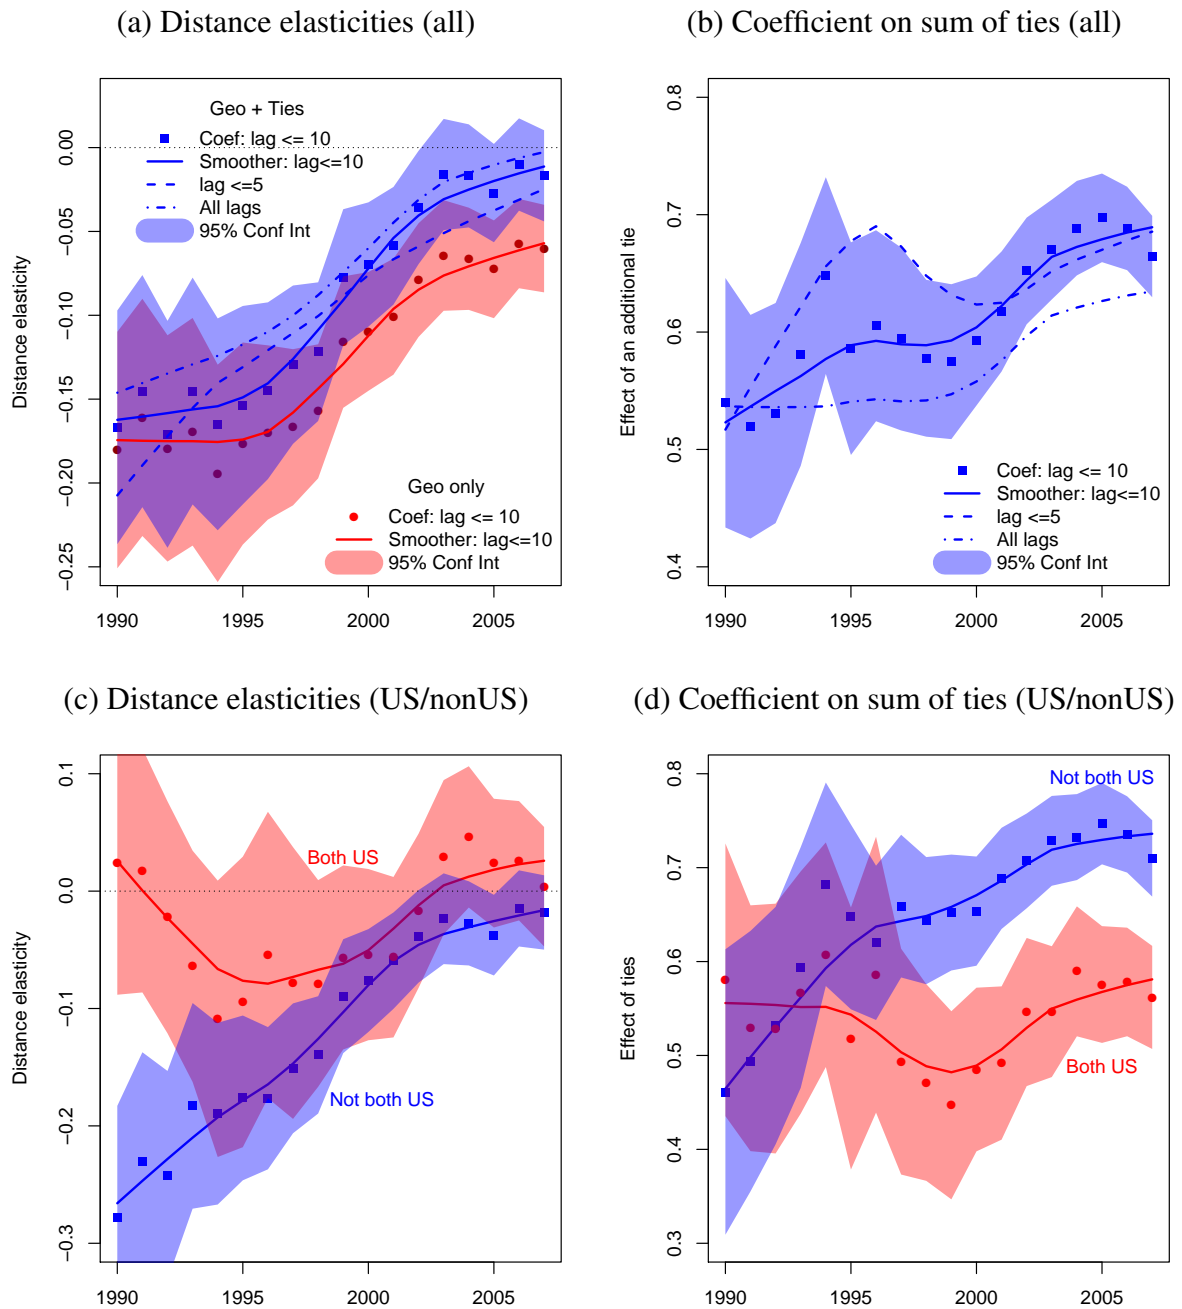
Note: Black circles and whiskers correspond to estimates and confidence intervals for 12 distance bins for specification (6) in Table 1. Blue square estimates are from a specification that omits ties.

Figure 3: The spatial concentration of tied authors



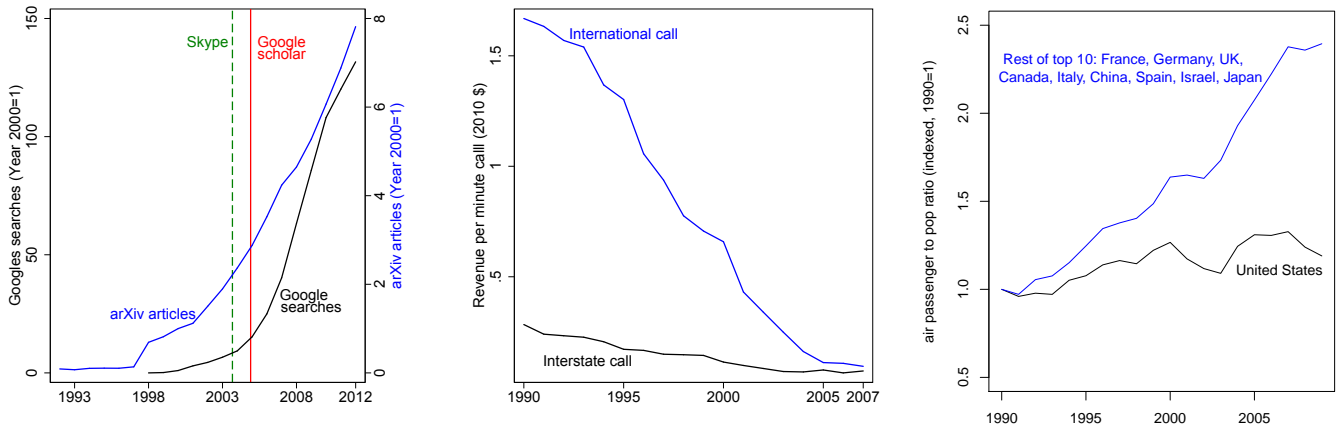
Note: Panel (a) constructs the complementary CDF for distances between authors of 2009 papers, distinguishing author dyads by the presence of ties. The blue line in panel (b) aggregates all the author pairs from panel (a). The black line plots the distribution of distances to the author's Ph.D. granting institution (*alma mater*).

Figure 4: Distance effects shrinking while ties matter more over time



Note: Plotted coefficients are marginal effects on the log odds of citation. Confidence intervals based on standard errors clustered at cited-article level. Estimation window moves by one year for each point, with citing papers published in years $t - 2$ to $t + 2$ and citation lags less than or equal to 5 or 10 years. In Panels (a) and (b) red-shaded interval corresponds to estimates that do not control for ties. Blue-shaded estimates control for the *sum* of 14 ties. In Panels (c) and (d) red-shaded interval corresponds to estimates where citing and cited authors reside in US. Blue-shaded estimates have at least one of citing or cited team from country other than US.

Figure 5: Trends in internet use (arXiv, Google), communication costs, and air transport



(a) Internet use

(b) Call rates

(c) Air passengers

Note: (a) From [arXiv.org](http://arxiv.org) we take annual counts of mathematics articles. Google searches from 2000 to 2012 come from <http://www.internetlivestats.com/google-search-statistics/>. The 1998 and 1999 numbers are extrapolated from reports from September of those years. Series indexed relative to 2000. (b) Average revenue per minute. Data from Federal Communications Commission (2010) Table 13.4, deflated by the CPI. (c) Ratio of air passengers carried over population, indexed relative to 1990 (WDI), Top 10 countries ranked by citing authors in 2009.

Table 1: Baseline estimation (triadic fixed effects logit)

	(1)	(2)	(3)	(4)	(5)	(6)
Sample	WOS	WOS	MGP	MGP	MGP	MGP
Distance > 0	-1.008*	-0.936*	-1.243*		-0.571*	
In Dist Dist > 0	-0.073*	-0.052*	-0.068*	Figure 2	-0.037*	Figure 2
Different country	-0.198*	-0.140*	-0.232*	-0.270*	-0.090*	-0.103*
Different language	-0.104*	-0.066*	-0.082*	-0.079*	-0.025	-0.025
Co-authors		1.672*			1.572*	1.581*
Coincided past		0.712*			0.378*	0.378*
Worked same place		0.478*			0.342*	0.339*
Share Ph.D. (5 years)					0.463*	0.457*
PhD siblings					0.663*	0.666*
PhD cousins					0.365*	0.362*
Advisor citing					1.090*	1.079*
Advisor cited					1.377*	1.375*
Academic grandparent citing					-0.284	-0.254
Academic grandparent cited					1.028*	1.023*
Academic uncle citing					0.227~	0.236†
Academic uncle cited					0.616*	0.619*
Alma Mater citing					0.239*	0.233*
Alma Mater cited					0.120†	0.119†
Observations	537054	537054	441792	441792	441792	441792
<i>pseudo-R</i> ²	0.044	0.085	0.033	0.034	0.091	0.091

Significance levels (based on standard errors clustered by cited article): *, †: 5%, ~: 10%.

Table 2: Obscure, Recent, and Different-field papers are more impacted by ties and geography

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Specification:	Obscure		Recent		Different field		
	base	interact	base	interact	base	interact	
<i>Geography:</i>							
Distance > 0	-0.571*	-0.541*	-0.011	-0.321*	-0.339 [†]	-0.804*	0.147
	(0.073)	(0.083)	(0.166)	(0.107)	(0.137)	(0.136)	(0.198)
ln Dist Dist > 0	-0.037*	-0.031*	-0.037 [~]	-0.021 [~]	-0.028 [~]	-0.026 [~]	-0.004
	(0.008)	(0.009)	(0.019)	(0.011)	(0.015)	(0.014)	(0.021)
Different country	-0.090*	-0.082 [†]	-0.061	-0.095 [†]	0.002	-0.072	0.022
	(0.031)	(0.034)	(0.079)	(0.044)	(0.060)	(0.056)	(0.088)
Different language	-0.025	-0.028	0.014	-0.017	-0.015	-0.037	-0.075
	(0.026)	(0.029)	(0.064)	(0.037)	(0.048)	(0.047)	(0.072)
<i>Ties:</i>							
Average effect of ties	0.652*	0.619*	0.135*	0.543*	0.176*	0.572*	0.223*
	(0.018)	(0.020)	(0.059)	(0.027)	(0.036)	(0.036)	(0.057)
Observations	441792	441792		441792		225768	
<i>pseudo-R</i> ²	0.091	0.092		0.093		0.100	

Notes: 1. Robust standard errors clustered by cited article in parentheses. Significance: *: 1%, [†]: 5%, [~]: 10%. 2. Average effect of ties is the mean of the base and interaction coefficients of 13 ties (3 WOS and 10 MGP). “Obscure” indicates that total citations received for this article are less than or equal to the median number of citations received among all articles, “recent” corresponds to citation lags less than or equal to the median, and “different field” equals 1 if citing article and cited article belong to different 2-digit MSCs.