

NBER WORKING PAPER SERIES

REVERSED CITATIONS AND THE LOCALIZATION OF KNOWLEDGE SPILLOVERS

Ashish Arora
Sharon Belenzon
Honggi Lee

Working Paper 23036
<http://www.nber.org/papers/w23036>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2017

We thank Wes Cohen for helpful discussions. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2017 by Ashish Arora, Sharon Belenzon, and Honggi Lee. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Reversed Citations and the Localization of Knowledge Spillovers

Ashish Arora, Sharon Belenzon, and Honggi Lee

NBER Working Paper No. 23036

January 2017

JEL No. O32,O34

ABSTRACT

Spillover of knowledge is considered to be an important cause of agglomeration of inventive activity. Many studies argue that knowledge spillovers are localized based on the observation that patents tend to cite nearby patents disproportionately. Specifically, patent citations are interpreted as mapping the transmission of knowledge from the cited invention to the citing invention. The localization of patent citations is therefore taken as evidence that such knowledge transmission is also localized. Localization of knowledge transmission, however, may not be the only reason for why patent citations are localized. Using a set of citations that are unlikely to be associated with knowledge transmission from the cited to the citing invention, we present evidence that challenges the view that localization of citations is driven by localized knowledge transmission. Though localized knowledge transmission may well exist, it is unlikely to be captured by patent citations

Ashish Arora
Fuqua School of Business
Duke University
Box 90120
Durham, NC 27708-0120
and NBER
ashish.arora@duke.edu

Honggi Lee
Duke University
Fuqua School of Business
honggi.lee@duke.edu

Sharon Belenzon
Duke University
Fuqua School of Business
1 Towerview Drive, Durham, NC
United States
sharon.belenzon@duke.edu

Reversed Citations and the Localization of Knowledge Spillovers*

Ashish Arora[†]

Sharon Belenzon[‡]

Honggi Lee[§]

Abstract

Spillover of knowledge is considered to be an important cause of agglomeration of inventive activity. Many studies argue that knowledge spillovers are localized based on the observation that patents tend to cite nearby patents disproportionately. Specifically, patent citations are interpreted as mapping the transmission of knowledge from the cited invention to the citing invention. The localization of patent citations is therefore taken as evidence that such knowledge transmission is also localized. Localization of knowledge transmission, however, may not be the only reason for why patent citations are localized. Using a set of citations that are unlikely to be associated with knowledge transmission from the cited to the citing invention, we present evidence that challenges the view that localization of citations is driven by localized knowledge transmission. Though localized knowledge transmission may well exist, it is unlikely to be captured by patent citations.

Keywords: Patent citations, localization, distance, knowledge spillovers

JEL Classification: O12, O32, O34

1 Introduction

Regional clusters of economic activity are important sources of innovation and economic growth (Porter, 1990). In the United States, there are over 40 clusters, which account for more than 50 percent of traded employment in U.S. Economic Areas (Porter, 2007). Clusters vary in their economic and inventive output (Agrawal et al., 2014; Delgado, Porter, and Stern, 2008; Porter, 2003), and these differences have led scholars to explore potential mechanisms that contribute to their success, including labor market pooling (Audretsch and Feldman, 1996; Delgado, Porter, and Stern, 2014; Krugman, 1991; Rosenthal and Strange, 2001), specialized resources (Porter, 1990, 2003; Swann, 1998), and knowledge spillovers (Audretsch, 1998; Jaffe, Trajtenberg, and Henderson, 1993). The growing importance of innovation and the growing availability of patent data have led scholars to use patent citations to study the extent to which knowledge spillovers are localized.

*We would like to thank Wes Cohen for helpful discussions. All remaining errors are ours.

[†]Duke University, Fuqua School of Business and NBER (ashish.arora@duke.edu)

[‡]Duke University, Fuqua School of Business (sharon.belenzon@duke.edu)

[§]Duke University, Fuqua School of Business (honggi.lee@duke.edu)

In their seminal paper, Jaffe et al. (1993) found that a citing patent is five to ten times more likely to be in the same SMSA as the cited patent than a matched, non-citing, patent. The authors and most scholars in subsequent studies interpret this finding as evidence of localized knowledge transmission from the cited to the citing patent (Alcácer and Gittelman, 2006; Almeida and Kogut, 1999; Belenzon and Schankerman, 2013; Murata et al., 2014). The current paper revisits the finding that patent citations are localized but presents evidence inconsistent with the interpretation that localization of citations is driven by localized transmission of knowledge. While our paper is silent on the question of whether knowledge transmission is localized, to the extent that such localization exists, we argue that it is not captured by patent citations.

There are at least two interpretations for the relationship between patent citations and knowledge spillovers. The first is that a citation from a citing patent A to a cited patent B indicates that the inventor of invention A learns from, or builds on, invention B. Patent citations are localized because inventors are more likely to learn from the inventions of nearby inventors. By definition, the citing invention must come after the cited invention. The second interpretation builds on the view that knowledge is “in the air” and is more likely to be absorbed locally. That is, a citation from patent A to patent B does not indicate that invention A builds on invention B, but rather that some background knowledge embodied in invention B is relevant for invention A. If such citations are localized, it suggests that inventors are drawing from a pool of local knowledge, but does not imply that knowledge spillovers among inventions are necessarily localized. In some cases, the temporal sequence of patent filings may also be reversed.¹ Our paper tests whether patent citations are localized because inventors learn from local inventions. While the pool of knowledge that inventors draw upon may be localized, our results indicate that learning from patented inventions is not, at least as measured by patent citations.

The distinction between drawing upon shared background knowledge and learning from specific inventions is important for both theory and policy. An inventor whose invention is built upon by subsequent inventions can potentially extract licensing revenues from the latter. Indeed, the key question analyzed in models of cumulative innovation is how to divide rents between sequential inventors (Green and Scotch-

¹Patent citations can arise in other situations as well. For instance, a patented invention may accomplish the same task as the cited invention, but do so differently. In this case, there is no presumption that the citation reflects a knowledge flow. However, if related inventive activity is localized, so patent citations will be.

mer, 1996), by choosing patentability criteria or by patent length and breadth (Bessen and Maskin, 2009; O’Donoghue, Scotchmer, and Thisse, 1998). If, instead, the invention contributes to a pool of knowledge from which others can draw, although innovation is cumulative, it would be very difficult for the inventor to identify subsequent inventions that build upon his invention (Laitner and Stolyarov, 2013).

The distinction is also relevant for the study of entrepreneurial spinoffs and regional clusters. Inventors leaving existing employers to start their own firms could build upon either specific inventions (Anton and Yao, 1995) or more general background knowledge (Chatterji, 2009). If spinoffs, which are often collocated with parents, capitalize on discoveries employees make in the course of employment, firms have the ability to design contracts to protect themselves. Such contrivances are less useful if spinoffs exploit more general knowledge learned in the course of their employment. Moreover, firms concerned about protecting inventions from spillover to outsiders may try to disperse inventors (Zhao, 2006) or locate them away from competitors (Alcácer and Chung, 2007; Alcácer and Zhao, 2012; Shaver and Flyer, 2000). This is likely to be more effective in preventing competitors from quickly building upon specific inventions, but less so against the diffusion of the firm’s knowledge more generally. Our findings raise doubts about the usefulness of patent citations to track knowledge transmission. More precisely, they challenge the use of patent citations to measure the extent to which knowledge transmission is localized.

Our empirical methodology is as follows. We identify a set of citations that are unlikely to reflect knowledge transmission, “citation reversals,” and use this set as a benchmark to compare the localization of citations for which knowledge transmission is possible, “citation non-reversals.” If localization of citations is driven by localized knowledge transmission, then we expect citation non-reversals to be more localized than citation reversals.

Citation reversals typically occur during the patent application process when an inventor (or an examiner) cites a relevant patent or application that he becomes aware of. Because ascertaining priority dates is not always straightforward, a patent representing a later invention may be cited as prior art. A patent’s priority date is usually when the patent is filed. However, correctly identifying the priority date becomes more difficult when the patent is an extension of an earlier patent (such as continuation or foreign), in which case the patent takes its priority date from the earlier patent. We refer to this type of citation reversal as “priority reversal.” Because the citing invention is created before the cited invention,

knowledge transmission from the cited to the citing invention is unlikely. The second type of citation reversal is “disclosure reversal,” and it arises when the priority date of the citing patent comes after the priority date of the cited patent but before the cited patent is disclosed to the public. Because the citing invention is created prior to the public disclosure of the cited invention, knowledge transmission is also unlikely for this type of citation.²

Our sample is from the 2014 version of EPO Worldwide Patent Statistical Database (PatStat). The final estimation sample includes 2,090,376 USPTO citations and an equal number of control citations, which are constructed by matching each of the citing patents with a randomly selected non-citing patent on the four-digit international patent classification (IPC) code and publication year.³ We test whether localization is more extensive for citations where knowledge transmission is possible (citation non-reversals) than for citations where knowledge transmission is unlikely (citation reversals). If localization of citations reflects localized knowledge transmission, we expect citation probability to decrease faster in distance for non-reversals than for reversals.

Our findings confirm the stylized fact documented in previous studies, that patent citations are localized. Based on our estimates, a 50-mile increase in distance between inventors lowers the probability of citation by about 20 percentage points (or 40% of the sample’s average citation probability). However, when we apply our test to compare the localization of non-reversals to that of reversals, we find that the negative effect of distance on citation probability is effectively the same between the two types of citations (-0.078 for non-reversals and -0.091 for reversals), a finding that is inconsistent with the interpretation that patent citations are localized because knowledge transmission across inventions is localized.

An important concern about our analysis is that reversals, particularly disclosure reversals, might arise from interactions among co-located inventors. Based on this interpretation of (disclosure) reversals, comparing the localization effect of non-reversals with that of reversals would not allow us to rule out localized knowledge transmission, even if non-reversals were less localized than reversals. We test this alternative interpretation by singling out two main mechanisms by which inventors can interact locally to

²For disclosure reversals, knowledge of the cited invention could have been transmitted by the cited inventor to the citing inventor. For priority reversal, the invention could have taken place up to a year before filing, and the inventor could have transmitted knowledge to others. In subsequent analysis, we examine these issues in greater detail.

³We follow the literature (Jaffe et al., 1993; Belenzon and Schankerman, 2013) by drawing control patents randomly from non-citing patents. The results do not change at all if one does not restrict to non-citing patents, because the likelihood that a given patent cites another is very close to zero.

learn about one another’s inventions. The first mechanism is knowledge sharing within firm boundaries. If local inventors share knowledge about their inventions before they are disclosed to the public, such interaction is more likely to occur within than across firm boundaries. Thus, we examine whether self-citations are more prevalent within reversals than within non-reversals, and also confirm that self-citations are more localized than external citations. A similar concern arises with patent attorneys, who may share information (or add citations) about inventions whose patent applications they handle. Thus, we examine whether citations whose patents are managed by the same attorneys (“same-attorney citations”) are more common within reversals than within non-reversals and are more localized than different-attorney citations. However, excluding these potential sources of localized knowledge transmission from our sample does not change our main finding that citation reversals are at least as localized as non-reversals. Our main results are further supported by the analyses using examiner and priority reversals, which are less prone to the aforementioned concern.

Our study provides evidence in support of previous studies questioning the interpretation that localization of citations is driven by localized knowledge transmission. For instance, building on the notion that technically related inventions might be collocated geographically, Thompson and Fox-Kean (2005) argue that the matching approach in Jaffe et al. (1993) is too coarse and thus may not adequately control for regional specialization. Imposing stricter matching criteria, they find that the localization effect almost disappears at the SMSA and state levels.⁴ Additionally, patent citations can be localized if they are added by intermediaries specializing in patent applications. Wagner, Hoisl, and Thoma (2014) argue that professional service firms build their knowledge repositories while interacting with their clients and reference patents from those repositories when they prepare patent applications. Therefore, if firms engage with patent attorneys that are geographically proximate to them, patent citations would likely be localized simply because they are added by attorneys who draw local knowledge from their repositories.⁵

Our study also contributes to the broader debate on what patent citations mean and to what extent citations can be interpreted as knowledge transmission (Duguet and MacGarvie, 2005; Jaffe, Trajtenberg, and Henderson, 2000; Roach and Cohen, 2013; Callaert, Pellens, and Van Looy, 2014). The current

⁴Adding to the debate, Murata et al. (2014) use a distance-based test to show that citations are localized in more than 30 percent of technology areas even when citing and control patents are matched on six-digit technology classification code.

⁵Wagner et al. (2014) show that, in their sample consisting of European patents, citations remain localized even after controlling for patents represented by the same patent attorneys.

consensus is that, although a significant proportion of citations is not driven by knowledge transmission, patent citation is an adequate proxy for knowledge transmission.⁶ Our results provide evidence that is inconsistent with this consensus.

The rest of the paper is organized as follows. Section 2 describes the data, and Section 3 discusses why the distance effect on citation probability might be stronger for non-reversals than for reversals, if localization of citations is driven by localized knowledge transmission. Section 4 presents the main empirical specification, and Section 5 presents non-parametric results of our main test. Section 6 presents both the main findings from our econometric analysis and the tests conducted to address the concern that citation reversals might be driven by highly localized knowledge transmission. Section 7 presents additional robustness tests, and Section 8 concludes and points to potential avenues for future work.

2 Data

2.1 Sample

Our main sample is from the 2014 version of EPO Worldwide Patent Statistical Database (PatStat), with inventor distances extracted from Google Maps. We follow previous literature and limit our dataset to patents with inventors residing in the contiguous United States. Thus, we exclude Hawaii, Alaska, and offshore U.S. territories, such as Puerto Rico and Guam. The publication years of the citing patents range from 2001 to 2014. (Sample years are based on the availability of examiner citations, which as we later explain are an important part of our analysis.) Finally, we retain only the citations whose priority lag, i.e., difference in years between the priority dates of citing and cited patents, is within five years. This restriction is important because citation reversals have short priority lags, and thus it ensures that reversals and non-reversals remain comparable. This procedure yields 2,090,376 actual citations. For each pair of patents connected by a citation, we match the citing patent with another patent, to create a control pair. We construct the control group by matching each citing patent in the actual citations with a randomly selected, non-citing patent on four-digit IPC and publication year (Belenzon and Schankerman, 2013). These control patents are paired with the cited patents from the actual citations to form control

⁶Roach and Cohen (2013) point out that patent citations are likely to underestimate spillovers from certain types of inventive activities such as basic research and overestimate spillovers when citations are driven by firms' appropriability and citing strategies. Jaffe et al. (2000) show that one third of the respondents did not find out about a cited invention until a patent application process began.

”citations”. The sample includes 4,180,752 observations consisting of both USPTO and control citations.

2.2 Variable definitions

We proceed to describe the main variables used in the analysis. Table 1 also summarizes their definitions.

Priority date. A central piece in our analysis is identifying priority dates, which are used to determine whether a citation is a reversal. Priority date is the earliest application date of the patents that relate to the same underlying invention. (A group of patents with the same underlying invention and priority date is a patent family.) It is the date on which the invention was first recognized by a patent-issuing authority to be in existence. A patent can claim as its priority date the application filing date of an earlier patent if the two patents describe the same invention and share at least one inventor. Depending on the type of the application being filed, application timing requirements also must be met.

Priority dates can be claimed via different types of patent applications. The most common are an application taking advantage of a provisional application filed for the same invention; a continuing application that extends a prior patent application;⁷ and a patent application filed in the United States within twelve months of filing a foreign application for the same invention. In all of these cases, to claim an application filing date of an earlier patent as the priority date, the specifications of the underlying invention described in the child and parent patents must be the same.

For example, US 7333597 is a patent on technology that enables telephone synchronization with software applications and documents. The inventors on the patent are Edward M. Silver, Linda A. Roberts, and Hong T. Nguyen. The application for this patent was filed on November 2, 2004. However, the priority date of the patent is March 29, 2002, based on an earlier patent, US 6873692, which relates to the same invention and has the same inventors. Because US 7333597 is a continuation of US 6873692 and because the underlying invention and at least one inventor are the same between the two patents, the former patent can claim as priority date the application filing date of the latter patent.

We find priority dates by linking child and parent patents via the priority and continuation patent

⁷Continuing applications can be further broken down into continuation, divisional, and continuation-in-part applications. Continuation applications make additional claims based on an existing invention specified in an earlier patent application while divisional applications are filed to separate out distinct inventions from an earlier application usually because the earlier one fails to meet the “unity of invention” requirement. Continuation-in-part application can add extensions to an earlier invention; however, the claims based on new subject matter are assigned, as their priority date, the application filing date of the continuation-in-part application.

datasets in PatStat. In our sample, 56.5 percent of the citing patents and 63 percent of the cited patents claim priority dates based on a provisional application, and 42 percent of the citing patents and 35.2 percent of the cited patents claim priority dates based on continuing applications. The rest are based on foreign applications. (Figure A1 presents the types of applications based on which patents claim priorities.)

Earliest publication date. The earliest publication date is the date on which the first patent from a patent family is publicly disclosed, indicating the first time an invention becomes known to the public. Since the enactment of American Inventors Protection Act of 1999, a patent application or a patent is published usually eighteen months from its earliest filing date, i.e., priority date, unless the applicant requests to have it disclosed earlier or to delay its disclosure until patent issuance. (Figure A2 presents the distribution of patents over number of months it takes for them to be published.)

Citation reversal. We construct two types of citation reversals for which transmission of knowledge is unlikely. The first type occurs when the priority year of the citing patent is earlier than the priority year of the cited patent. We refer to this type as priority reversal. Priority reversals are not likely to reflect knowledge transmission because the cited invention does not exist at least until the citing invention is created.

According to our conversations with a patent examiner and a patent attorney, priority reversals can arise when a patent applicant includes citations on the information disclosure statement (IDS) or when a patent examiner adds additional citations during the patent prosecution process, which starts after the patent application is filed. Some patents have a priority date earlier than their filing date, for instance because they are related to a foreign patent or a continuation of an earlier patent. When these patent applications are examined, the applicant or examiner may cite another patent with a later priority. This type of citation (priority reversal) indicates that the cited patent is related to the citing patent, but, by construction, cannot represent knowledge transmission from the cited to the citing patent.⁸

An example of a priority reversal is patent US 8291504 citing patent US 7962635. (A graphical illustration and transaction highlights are also shown in Figures A3.) US 8291504 was filed on December

⁸Priority reversals happen because priority dates are not always transparent to applicants without an extensive search. Although examiners are supposed to review applicant citations for accuracy and appropriateness, the examination process is not fully automated, and a related patent may be incorrectly cited.

31, 2009, and US 7962635 was filed on December 7, 2009. However, the invention described in US 8291504 was created on December 31, 2008, while the invention described in patent US 7962635 was created on February 24, 2009. Thus, the priority year of the citing patent is prior to the priority year of the cited patent. Priority reversals occur in about 4 percent of all actual citations in our sample.

The second type of citation reversal occurs when the priority year of the citing patent comes after the priority year but before the earliest publication year of the cited patent. We refer to this type as disclosure reversal. Disclosure reversals are not mistakes and thus are more prevalent. They arise often because of overlapping patent prosecution time periods where a prior art patent is published during the citing patent's prosecution process.

An example of a disclosure reversal is patent US 8636913 with the priority date of December 21, 2011, citing patent US 8262919 with the priority date of June 25, 2010, and the earliest publication date of September 11, 2012. (A graphical illustration and transaction highlights are also shown in Figure A4.) We classify this citation as disclosure reversal because the cited invention was not disclosed to the public when the citing invention was created. In this case, it is possible that the citing inventor became aware of the cited invention and built upon it. However, since the cited invention was not publicly disclosed until the citing invention was completed, knowledge transmission would have required means other than learning from the cited patent. About 25 percent of the actual citations in our sample are disclosure reversals.

We illustrate priority and disclosure reversals using the timeline in Figure 1. The horizontal line going from left to right represents the progression of time, and the two vertical lines extending below the timeline represent the priority date and the earliest publication date of a cited patent. The three vertical lines extending above the timeline represent possible instances of the citing patent's priority date. A priority reversal occurs when an invention references another invention that has not been created, as demonstrated by the first instance of the citing patent's priority date coming before the priority date of the cited patent. A disclosure reversal occurs when an invention references another invention that has been created but has not been disclosed to the public, as demonstrated by the second instance of the citing patent's priority date falling in between the priority date and the earliest publication date of the cited patent. In both of these cases, because the cited invention did not exist or was not known to the

public at least until the citing invention was completed, knowledge transmission from the cited to the citing invention is unlikely. A non-reversal occurs when an invention references another invention that has been created and disclosed to the public, as demonstrated by the third instance of the citing patent’s priority date falling after both the priority date and the earliest publication date of the cited patent. In this case, knowledge is more likely to be transmitted from the cited to the citing invention.

Geographical distance. We calculate distances between citation pairs using PatStat’s inventor addresses dataset and Google Maps API. We first extract inventor city and state information from PatStat’s inventor addresses dataset and then use a custom software application that communicates with Google Maps Geocoding API to obtain geographical coordinates and straight line distances between inventors of citing and cited patents.⁹ In addition to continuous distance, we construct dummy variables for distance ranges to examine the nonlinear effect of distance on citation probability. The reference range is 0–25 miles, and the rest of the distance ranges are as follows: 25–50, 50–100, 100–150, 150–250, 250–500, 500–1000, 1000–1500, 1500–2500, and greater than 2500.

Table 2 presents summary statistics for the main variables used in the analysis. The average distance between patents linked by a citation is 919 miles with a standard deviation of 890 miles. 29 percent of the citations are reversals and 47 percent of them are added by examiners. On average, a citing patent receives 6.5 forward citations, and cited patents receive 46 forward citations. Of the citations, 16.4 percent are coast-to-coast citations, and thus we add dummies to control for citations between research clusters that are located in the opposite coasts.

Insert Tables 1–2 and Figure 1 here

3 The Localization Effect of Reversed Citations

We propose the following motivation for why localized knowledge transmission might lead to a stronger negative effect of distance on citation probability for citation non-reversals than for reversals.

We distinguish between citations that reflect knowledge transmission, i.e., where one invention builds upon another, improving or extending it, and citations that reflect other relationships, where no such

⁹In cases of multiple inventor locations for a single patent, we use the city-state combination that occurs most frequently. If there is an equal number of different city-state combinations, we randomly choose a location among them.

knowledge transmission takes place. A patent may cite another patent because both draw upon a common pool of knowledge. Patents may be related in other ways as well. For instance, the citing patent may accomplish the same outcome as the cited patent using a different method, or use similar methods to accomplish different goals.

The distinction bears on whether the citing invention benefited from the knowledge created in the form of the cited invention. The citing inventions that are *related* would be unaffected. The citing invention that builds upon the cited invention would have had to acquire the knowledge created by the cited invention. The extant literature has implicitly or explicitly assumed that, when patent class is controlled for, citations reflect knowledge transmission. We argue that citation reversals, by construction, cannot reflect knowledge transmission from the cited invention to the citing invention.¹⁰ It is conceivable that the inventor may have learned useful knowledge from the inventor of the cited invention, but it is far more likely that a citation reversal represents the recognition of a related invention, albeit one that should not have been cited because it is not prior art.¹¹

We do not observe whether a citation represents transmission of knowledge or merely some type of relatedness. Instead, we propose a simple structure that clarifies how one can infer the relative importance of the two types of citations in generating the observed pattern of localization of patent citations. If citation reversals reflect relatedness, and if knowledge transmission falls with distance, then citation reversals should be less localized than non-reversals.

Consider a pair of patents i and j , where i is the focal patent and j is invented after i . Patent j can build upon patent i with probability π_n if it is near and π_f if it is far. With probability θ_n , patent j can be related to patent i if they are near each other and θ_f if they are far apart. For simplicity, assume a patent that builds upon another will cite it with probability α . Similarly, a patent that is related to another will cite with probability β . The probability that j cites i if they are near is $\alpha\pi_n + \beta\theta_n$.¹²

¹⁰A citation implies that the inventor, the patent agent, or the examiner became aware of the cited invention. However, in citation reversals, the citing invention could not have benefited from this knowledge although the citing patent application may have been modified.

¹¹The issue is subtle. The inventor of a citing patent may acquire the required knowledge from the inventor of the cited patent. Alternatively, the inventor may learn from the cited patent, or the inventor of the cited patent, knowledge that is not unique to the cited invention. We address many of these issues in greater detail in Section 6.2, where we analyze the localization patterns in citations inserted by examiners, self-citations, and citations made to unpublished patents.

¹²We are assuming, for simplicity, that these are mutually exclusive outcomes. However a citation may reflect both building upon and relatedness. In that case, the probability that j cites i if they are near is $\alpha\pi_n + \beta\theta_n - \alpha\beta\pi_n\theta_n$. Both $\alpha\pi_n$ and $\beta\theta_n$ are very small in magnitude, so that the product term $\alpha\beta\pi_n\theta_n$ can be neglected.

Similarly, the probability of a citation if they are far apart is $\alpha\pi_f + \beta\theta_f$. The difference in probability of citation between patents located near each other and patents located far is $\alpha(\pi_n - \pi_f) + \beta(\theta_n - \theta_f)$.¹³ This difference has two components: $\pi_n - \pi_f$ representing the extent to which knowledge transmission decreases with distance, and $\theta_n - \theta_f$ representing the extent to which patents near each other are related relative to those far apart. Each component is weighted by the relevant citation propensity. The consensus is that both components are positive. That is, nearby inventions are more likely to be related, and knowledge transmission is more likely among nearby inventions.

Now consider the probability that i cites j , i.e., a citation reversal. Since i cannot have built upon j , the probability of the patents being related is simply θ_n if they are near each other and θ_f if they are far. If the propensity to cite is β as before, difference in probability of citation between near and far is simply $\beta(\theta_n - \theta_f)$. The effect of proximity on the probability of citation for non-reversals is $\alpha(\pi_n - \pi_f) + \beta(\theta_n - \theta_f)$ so that the difference in the effect of proximity with respect to reversals is $\alpha(\pi_n - \pi_f)$. The extant literature has asserted that knowledge transmission increases with proximity, i.e., $\alpha(\pi_n - \pi_f) > 0$. This implies that the effect of proximity on the probability of a citation reversal should be smaller than the effect of proximity on the probability of a non-reversal citation. In the empirical analysis we use distance, so that we expect that distance should have a smaller absolute effect on reverse citations than on normal or non-reversal citations.

The validity of our test hinges the assumption the underlying knowledge that is described in a patent cannot be changed as the patent moves forward in the examination process. This assumption assures that the priority date is the date when the invention is created and allows us to determine whether knowledge could have been transmitted from one inventor to another. This assumption is validated by the two statutory provisions, 35 U.S.C. 132 and 35 U.S.C. 251, which prohibit introduction of new matter in amendments and the application for reissuance. Examiners are obligated to reject new matters introduced into the abstract, specification, or drawings of a patent application once it is filed. The only changes permitted are rephrasing of a passage without changes in meaning and fixing obvious errors whose correction can be foreseen by a person skilled in the given art.¹⁴ If so, the transmission of knowledge

¹³Existing studies have tried to use patent classes to control for relatedness: It is implicitly or explicitly assumed that two patents in the same class are equally likely to be related independently of whether they are far or near. Formally, for patents in the same patent class, $\theta_n - \theta_f = 0$.

¹⁴The literature on submarine patents discusses how inventors can keep their inventions secret for an extended period of

cannot take place in citation reversals. While this inference is reasonable for priority reversals, disclosure reversals might be associated with knowledge transmission. We take several steps to explore and mitigate this concern in Section 6.2 using information on self-citations, patent attorneys, and examiner-added citations.

4 Econometric Specification

Our empirical analysis tests the implication of the model that we described in the previous section. That is, we examine whether the effect of distance on citation probability is smaller in magnitude for reversals than for non-reversals. As shown in the model, this relationship will hold if localization of citations is driven by localized knowledge transmission.

We follow Jaffe et al. (1993) and match each citing patent with a control, non-citing, patent with the same four-digit IPC code and publication year. Our results are robust to matching at the six-digit IPC code and publication year, though a finer matching naturally reduces the sample size (cf. Appendix table A3). We use a linear probability model to estimate the effect of distance on citation probability for non-reversals and reversals. Our main empirical specification is as follows:

$$Pr(C_{ij} = 1) = \beta_1 \ln D_{ij} + \beta_2 \ln D_{ij} \times reversal_{ij} + \beta_3 reversal_{ij} + Z'\gamma + \eta_j + \epsilon_{ij}$$

where i and j denote citing and cited patents, respectively, C_{ij} is a dummy variable that receives the value of one for an actual citation and zero for a control (non-) citation, D_{ij} is the distance in miles between the location of citing and cited inventors, and $Reversal_{ij}$ is a dummy variable that receives the value of one for a citation reversal (and for the matched control non-citing patent). Z is a vector of dyadic dummies indicating citations between leading research clusters (i.e., Austin, TX; Route 128, MA; Raleigh-Durham, NC; San Diego, CA; and Silicon Valley, CA). These dyadic research cluster dummies are important because patents produced in clusters specializing in similar inventions are likely to cite one another and the clusters are often located on opposite coasts. The stochastic components are represented by η_j , a cited patent fixed effect, and an *iid* error term ϵ_{ij} . Standard errors are always clustered at the cited patent level.

time and change claims using continuation applications (Graham and Mowery, 2004; Reizig, Henkel, and Heath, 2007). The change described by the literature pertains to claims rather than inventions and is thus consistent with our assumption that the underlying invention does not change.

If localization of citations is driven by localized transmission of knowledge, we expect the effect of distance on citation probability to be larger in magnitude for non-reversals than for reversals. Thus, we expect $\hat{\beta}_1 < 0$ to confirm previous evidence on localized citations, and $\hat{\beta}_2 > 0$ to support the view that citations with potential knowledge transmission (non-reversal) are more localized than citations where knowledge transmission is unlikely (reversal).

5 Non-parametric Evidence

Column 1 of Table 3 presents mean comparisons of main variables between citation non-reversals and reversals. The comparison of geographical distances reveals that reversals are at least as localized as non-reversals. The average distance between citation pairs is 956 miles for non-reversals and 826 miles for reversals. Also, the share of citations with citing patents that are within 50 miles of the cited patent is 39 percent for reversals and 26 percent for non-reversals. These findings are inconsistent with the view that localization of citations is driven by localized knowledge transmission.

Figure 2 presents comparisons of average distances across various types of citations and the share of citations with citing patents within 50 miles of cited patents. The figure shows that the average distance is greater for non-reversals than for reversals. The fraction of citing patents within 50 miles of the cited patents shows a similar pattern. Thus, reversals are more localized than non-reversals, and the results hold even when we compare only the applicant citations or only the examiner citations. These findings are inconsistent with the view that citations are driven by localized knowledge transmission.

Insert Table 3 and Figure 2 here

6 Estimation Results

6.1 Reversal *vs.* non-reversal localization effect

Table 4 presents the results from our main test for localized transmission of knowledge in patent citations. Column 1 presents the estimation results for the effect of distance on the probability of citation. Consistent with previous findings in the literature, the results show that patent citations are localized. The coefficient estimate on the distance between citing and cited patents is negative and statistically significant, indicating that two inventors who are geographically close to each other are more likely to cite than inventors who

are far away from each other. Column 2 explores the effect of distance on citation probability using distance dummies. The reference distance is 0–25 miles. Based on our estimates, moving from 0 to 50 miles between inventors lowers the probability of citation by 20 percentage points, or 40 percent of the sample average.

Columns 3–5 present our key findings from comparing the estimated effect of distance on citation probability between non-reversals and reversals. If localization of citations reflects localized transmission of knowledge, then we would observe a significantly larger effect of distance on citation probability for non-reversals than for reversals. Column 3 reveals that the effect of distance on citation probability is -0.075 for non-reversals and -0.098 for reversals, indicating that citation non-reversals are not more localized than reversals.

Columns 4–5 also show no difference in the effect of distance on citation probability for subsamples consisting separately of non-reversals and reversals. Columns 6–7 explore the robustness of the results by allowing for non-linear distance effects. The same pattern of results emerges. For example, moving from 0 to 50 miles reduces the citation probability by 17.8 percentage points for non-reversals and 23.8 percentage points for reversals. These findings are inconsistent with the interpretation that localization of citations reflects localized knowledge transmission since citations do not become more localized even when potential transmission of knowledge is allowed for.

Insert Table 4 here

6.2 Are reversals driven by local inventor interactions?

An important concern in our study is that citation reversals might reflect highly localized knowledge transmission among local inventors who share knowledge about new inventions that have not been disclosed to the public. If this were true, then our test comparing localization of non-reversals with that of reversals would not be able to rule out localized knowledge transmission across inventions even if non-reversals were less localized than reversals.

Table 5 summarizes different mechanisms through which both knowledge about inventions and background knowledge could be transmitted, including those that might lead to citation reversals. As discussed in introduction, our empirical methodology tests whether learning from inventions is localized. There are

three main learning mechanisms indicated by columns 1–3. The first is learning from patent documents—learning about inventions by reading published patent documents. There is no reason to expect this type of learning to be localized. Thus, we are not concerned that learning from patent documents would generate localized citation reversals. The second mechanism is learning about inventions from inventors. In this case, inventors might share among themselves knowledge about new inventions. Such learning might be localized if inventors tend to discuss their inventions more frequently with nearby inventors than faraway inventors. Learning could also occur through intermediaries, such as patent attorneys. By using the same patent attorneys, inventors might learn about new inventions before they are disclosed to the public. Because inventors are more likely to use the same patent attorneys if they are near one another, engaging with patent attorneys might generate localized citation reversals.

We proceed to discussing how we mitigate concerns that local inventor and patent attorney interactions might generate a pattern of highly localized citation reversals.

Insert Table 5 here

6.2.1 Self-citations

If self-citations are more prevalent within reversals and are more localized than external citations, then self-citations might inflate the localization effect of reversals. This bias will lead us not to reject the null hypothesis that the localization effect of non-reversals is the same as that of reversals.

Our data reveal that the share of self-citations is 33 percent for reversals and 19 percent for non-reversals. Self-citations are also more localized than external citations. The share of citations whose citing patent is within 50 miles of the cited patent is 78.4 percent for self-citations and 14 percent for external citations. (These differences are statistically significant at the 1% level.) Within citation reversals, self-citations are also more localized than external citations (84% of citing patents being within 50 miles of cited patents for self-citations relative to 15.6% for external citation reversals). These findings are consistent with the concern that reversals might be driven by highly localized knowledge transmission.

To mitigate the potential bias caused by self-citations, we exclude them from our sample. Columns 1 and 2 of Table 6 present the estimation results. The results show that even after excluding self-citations, reversals are at least as localized as non-reversals. For instance, moving from 0 to 50 miles lowers the

probability of citation by 11.6 percent for reversals and by 9.4 percent for non-reversals. These results show that our main finding continues to hold even after removing the potential source of highly localized knowledge transmission.

6.2.2 Patent attorneys

Patent attorneys are another potential source of localized knowledge transmission that might generate citation reversals. Wagner et al. (2014) argue that patent attorneys develop knowledge repositories through their interactions with clients and reference known patents from the repositories when they prepare patent applications. Thus, if inventors engage with local attorneys and attorneys share knowledge about local inventions before they are publicly disclosed, then interactions between inventors and attorneys might result in citation reversals that reflect highly localized knowledge transmission. To address this concern, we examine whether citations whose patents are managed by the same attorneys (“same-attorney citations”) are more common within citation reversals than within non-reversals and whether they are more localized than different-attorney citations.

To perform this analysis, we extracted attorney information from the weekly compilations of patent publications released by the USPTO for years 2001 through 2014. We first standardized attorney names and then removed any corporate legal offices. We then merged the attorney information with our main dataset. In our sample, the share of same-attorney citations (where citing and cited patents have the same patent attorney) is 13 percent for reversals and 5 percent for non-reversals. Same-attorney citations are also more localized than different-attorney citations. The share of citations that take place within 50 miles of the cited patent is 83 percent for same-attorney citations and 25 percent for different-attorney citations (statistically significant at the 1% level).

To mitigate this concern we exclude from our sample citations where the patents have the same attorney. Columns 1 and 3 in Table 6 show that citation reversals are at least as localized as non-reversals at every distance range even if we exclude citations involving patents with the same attorney. For instance, moving from 0 to 50 miles reduces the citation probability by 10.2 percent for reversals and by 9.4 percent for non-reversals. The results show that our main findings continue to hold even after we remove the potential source of highly localized knowledge transmission and mitigate the concern that our

results are driven by interactions between inventors and attorneys.

6.2.3 Examiner reversals

Citation reversals added by examiners are less prone to biases that could arise from local inventor interactions. If reversals arise due to local interaction among inventors, we expect inventors to generate proportionally more reversals than examiners and that inventor reversals would be more localized than examiner reversals. Inconsistent with this expectation, our data show that the share of reversals is 36 percent for examiner citations and 24 percent for inventor citations. On the other hand, the fraction of citations whose citing patent is within 50 miles of the cited patent is 39.8 percent for inventor reversals and 36.5 percent examiner reversals, indicating that inventor-added reversals are somewhat more localized than examiner-added reversals. While the difference is fairly small, this observation is consistent with the possibility that reversals are driven by inventor interactions.

We also compared the localization of non-reversals with that of examiner-added reversals. Comparing columns 1 and 6 of Table 6 shows that examiner-added reversals are at least as localized as non-reversals. For instance, moving from 0 to 50 miles reduces the effect of distance on citation probability by 10.8 percent for examiner-added reversals and 9.4 percent for non-reversals. Overall, the results in Table 6 indicate that although reversals are more likely for patents filed by the same organization and for patents with a common patent attorney, excluding them from the sample does not change the broad conclusion that localization of patent citations is unlikely to reflect transmission of knowledge from the cited invention to the citing invention. Further, focusing only on citation reversals added by examiners yields a similar effect of distance on citation propensity as non-reversals.

6.2.4 Priority *vs.* disclosure reversals

Priority reversals are also more immune than disclosure reversals to biases that could arise from local inventor interactions because the cited invention in a priority reversal does not exist when the citing invention is created. Thus, we alleviate the concern that reversals might capture highly localized knowledge transmission by comparing localization of non-reversals with that of priority reversals.

If knowledge transmission through local inventor interactions is a valid concern, then we expect both

disclosure reversals and non-reversals to be more localized than priority reversals.¹⁵ Our data show that the average distance is 822 miles for disclosure reversals and 856 miles for priority reversals. The fraction of citations within 50 miles from the cited patent is 33.2 percent for priority reversals and 39 percent for disclosure reversals (statistically significant at the 1% level). While the differences are relatively small, they are consistent with the concern that disclosure reversals might be driven by localized inventor interactions. However, comparing non-reversals and priority reversals in columns 1 and 4 in Table 6 shows that the distance effect on citation probability is very similar between non-reversals and priority reversals within the 50-mile range. For instance, moving from 0 to 50 miles reduces the citation probability by 8.9 percent for priority reversals and 9.4 percent for non-reversals, with the difference being only 0.5 percent. Although the difference widens as the distance range becomes larger, the finding suggests that priority reversals are likely to be as localized as non-reversals within a reasonable distance range for testing localization effect. Thus, the results provide further evidence that localization of citations is not likely to be driven by localized knowledge transmission.

Insert Table 6 here

7 Other Robustness Checks

7.1 Alternative specification

Table A1 in the appendix presents results from an additional test using an alternative specification designed to mitigate the concern that non-reversals and reversals are different along various dimensions that might be correlated with distance. For instance, if cited patents in non-reversals are more likely to be associated with high-quality patents, which are often less localized, then our main test would produce spurious results. We mitigate this concern by comparing localization of non-reversals and that of reversals that cite the same patent while controlling as much as possible for potential differences that may exist across the citing patents.

¹⁵Suppose that patent citations only reflect instances of the citing invention building upon the cited invention. In the sub-sample where, by construction, the citing invention could not have built upon the cited invention or learned from it, patent citations should be less localized than the general population. Furthermore, citations to unpublished patents should be more localized than citations to published patents. In other words, priority reversals should be the least localized, followed by normal citations, and followed by citations to patents unpublished at the time of citation. If, instead, all three samples have the same level of localization, then it follows that the localization of citations reflects some combination of localization of related inventive activity and localization of relevant technical knowledge.

For each citation reversal, we find a non-reversal with the same cited patent and with citing patents that are in the same IPC and publication cohort as those in the reversal. To control for potential differences across citing patents, we add citing patent technology area and publication year fixed effects. Column 1 shows that citing patents in reversals added by examiners are at least as likely to be within a 50-mile range of the cited patent as those in non-reversals. For instance, being a non-reversal is associated with 2.4 percent decrease in the probability that the citing and cited patents are within 50 miles of each other. This result holds even when we separate out the examiner reversals into priority and disclosure reversals (columns 2–3) and when we use inventor reversals instead of examiner reversals (columns 4–6). Thus, our main finding, that localization of citations is not likely to be driven by localized knowledge transmission, holds even when we change the specification to compare reversals and non-reversals within same cited patents.

7.2 Sample period and citation lag

For our main results, we used a sample that contains citations with citing patents published in years 2001 through 2014 because examiner citations became identifiable only in 2001. Table A2 in the appendix presents the results from a test that uses a sample whose citing patents cover years 1977 through 2014 to make sure that our results are not biased by factors inherent to more recent citations.

The results from this larger sample are consistent with the results from the main sample and provide additional support that localization of citations is not likely to reflect localized knowledge transmission. For instance, columns 2 and 3 show that going from 0 to 50 miles reduces the citation probability by 19.1 percentage points for non-reversals and 24.2 percentage points for reversals. This finding continues to hold even if we do not limit our sample to five-year priority lag as in columns 4 and 5.

7.3 Three- and six-digit IPCs

We further examine robustness of our findings by replicating our test with samples whose controls are matched on the three- and six-digit technology classification codes as demonstrated in Murata et al. (2014). Murata et al. (2014) replicate the samples used by both Jaffe et al. (1993) and Thompson and Fox-Kean (2005) to show that using a continuous distance measure allows them to observe localization of citations regardless of whether the controls are constructed by matching on three- or six-digit technology

classification codes. They find that citations are localized in 70 percent of technology classes when controls are matched on three-digit classification code and in 30 percent of technology classes when controls are matched on six-digit classification code. They interpret their findings as evidence that knowledge spillovers are localized within the majority of technology classes.

We test whether our results hold regardless of the level of technology classification code used to construct control citations. We obtained the dataset with controls matched on three-digit classification code used by Murata et al. (2014) from Harvard Dataverse and augmented it with priority dates and earliest publication dates to identify citation reversals. We also constructed a sample with controls matched on six-digit IPC code using PatStat.¹⁶

Table A3 in the appendix presents the results from these two samples. Citations are localized as found in Murata et al. (2014) regardless of the level of technology classification code used to construct control citations. However, columns 2–3 and columns 5–6 show that citation reversals are at least as localized as non-reversals in both cases. For instance, columns 2–3 show that the effect of distance on citation probability is larger in magnitude for reversals (-0.039) than for non-reversals (-0.033).¹⁷ ¹⁸ Also, columns 5–6 show that going from 0 to 50 miles reduces citation probability by 24.7 percentage points for reversals and 16.2 percentage points for non-reversals for the sample matched at six-digit technology classification code.¹⁹ These results are consistent with our main finding that, while citations are localized as indicated by previous studies, their localization is not likely to be driven by localized knowledge transmission.

¹⁶We do not use the sample matched on six-digit classification code from Murata et al. (2014) because the proportion of reversals in the dataset is too small to replicate our test.

¹⁷For the sample matched on three-digit technology classification code, we do not perform analysis with distance dummies because there are not enough reversals at short distance ranges.

¹⁸For the sample matched on three-digit technology classification code, the average distance between citing and cited patents is 873 miles for citation reversals and 899 miles for non-reversals, and the difference is statistically significant. Furthermore, the share of citing patents within 50 miles of cited patents is higher for citation reversals at 11 percent than for non-reversals at 9 percent.

¹⁹For the sample matched on six-digit technology classification code, the average distance between citing and cited patents is 884 miles for citation reversals and 979 miles for non-reversals, and the difference is statistically significant. Furthermore, the share of citing patents within 50 miles of cited patents is higher for citation reversals at 35 percent than for non-reversals at 24 percent.

7.4 Grace period

Under the Patent Act of 1952, inventors are given one-year grace period to file for a patent once they disclose their inventions to the public. If inventors take advantage of this law, then it is possible that our sample over-classifies citations as reversals because the date of invention might be one year prior to the patent application filing date. We address this concern by testing our results after re-classifying reversals with up to one-year priority lag as non-reversals.

Column 1 of Table A4 in the appendix shows that the effect of distance on citation probability is -0.085 for non-reversals and -0.097 for reversals with a priority lag between one and two years, indicating that non-reversals are not more localized than reversals even after the reclassification. A similar pattern holds for reversals with other citation lags in columns 2-5. The results show that the reversals with priority lags greater than one year, i.e., those not affected by grace period, are at least as localized as non-reversals. These results mitigate the concern that our main results might be driven by citations that are incorrectly categorized as reversals.

7.5 Technology areas

The importance of inventor interactions for learning from invention is likely to vary across technology areas. For example, inventions in complex technology areas such as telecommunications are likely to require more tacit knowledge than those in discrete technology areas such as chemicals. Such tacit knowledge might be more localized. To test whether reversals are driven by localized inventor interactions, we first compare the share of reversals across six technology areas: Chemistry, Pharmaceutical, Biotechnology, Medical Technology, Computer Technology, and Telecommunications. If technology areas characterized by tacit knowledge are more prone to inventor interactions, then we expect to see higher shares of citation reversals for complex technology areas than for discrete technology areas.

Inconsistent with the view that reversals are driven by transmission of tacit knowledge, we find that the share of reversals are fairly consistent across the six technology areas, ranging from 24 percent to 38 percent. For instance, the share of reversals is 24 percent in Biotechnology citations and 25 percent in Telecommunications. The share of reversals for the rest of the technology areas range from 34 percent to 38 percent.

We also examine whether the difference between the localization effect for non-reversals and reversals varies by technology area. If reversals in complex technology areas are more localized because inventor interactions are more important for learning, we expect our test to bias against finding a difference in the localization effect between non-reversals and reversals in complex areas, but not in discrete areas.

Tables A5a and A5b in the appendix present the results from comparing localization of citation non-reversals and that of reversals. Inconsistent with the expectation, the overall pattern does not show any systematic variation according to our expectation. For instance, going from 0 to 50 miles reduces the citation probability by 20.8 percentage points for non-reversals (column 5 of Table A5a) and 32 percentage points for reversals (column 6 of Table A5a) in Biotechnology and by 19.5 percentage points for non-reversals (column 1 of Table A5b) and 24.4 percentage points for reversals (column 2 of Table A5b) in Medical Technology.²⁰ These results provide evidence that localization of citations is not likely to be driven by localized knowledge transmission even in complex technology areas.

8 Concluding Remarks

This study examines whether localization of patent citations is driven by localized knowledge transmission. Our main results show that the effect of distance on citation probability is similar for citation non-reversals and reversals, providing evidence that localized knowledge transmission is not likely to be a major driver of localization of citations. The implication is that either patent citations are not driven primarily by knowledge transmission or knowledge transmission is not localized. If the latter, it questions an important, hypothesized reason for the agglomeration of inventive activity. If the former, it brings into question a number of results which have used citations to measure knowledge transmission.

The concern that citation reversals might reflect highly localized knowledge transmission between inventors, either through direct interactions among inventors or through intermediaries, is mitigated by comparing localization of non-reversals with various subsets of reversals that are even less likely to be driven by localized knowledge transmission. The results are consistent with our main finding, that localization of citations does not reflect localization of knowledge transmission.

²⁰The effect of distance on citation probability when moving from 0 to 50 miles is larger in magnitude for non-reversals than for reversals within Chemical and Pharmaceutical. However, the differences in magnitude are very small, at around 6%.

Methodologically, this study provides a way to isolate citations that are unlikely to be associated with knowledge transmission. We identify these special citations by checking whether the citing patent's priority date comes before either the priority date or the earliest publication date of the cited patent. We use these citations to benchmark localization of knowledge transmission as reflected in patent citations.

Our study has important implications for both theory and policy. The distinction between background knowledge and knowledge about specific inventions helps our understanding on how licensing revenues might be divided between sequential inventors. An inventor whose invention is built upon by another inventor might extract licensing revenues from the latter. However, if the subsequent inventor draws upon the background knowledge that the earlier invention contributes to, then it would be difficult to even identify inventions that build on such knowledge.

The distinction also provides insights into entrepreneurial spin-offs and regional clusters. If founders of spin-offs draw knowledge from specific discoveries during their employment at the parent company, then employers can develop contracts to prevent loss of rents. However, if knowledge drawn by the former employees is more general and cannot specifically be identified, then it would be more difficult for employers to protect themselves from former employees utilizing the knowledge. Furthermore, it would be difficult for firms to prevent background knowledge from spilling over to competitors even if they try to disperse inventors or R&D operations across different geographies.

References

- [1] Agrawal, A., Cockburn, I., Galasso, A., Oettl, A., 2014. Why are some regions more innovative than others? The role of small firms in the presence of large labs. *Journal of Urban Economics* 81: 149–165.
- [2] Alcácer, J., Chung, W., 2007. Location strategies and knowledge spillovers. *Management Science* 53(5): 760–776.
- [3] Alcácer, J., Gittelman, M., 2006. Patent citations as a measure of knowledge flows: The influence of examiner citations. *The Review of Economics and Statistics* 88(4): 774–779.
- [4] Alcácer, J., Zhao, M., 2012. Local R&D strategies and multilocation firms: The role of internal linkages. *Management Science* 58(4): 734–753.
- [5] Almeida, P., Kogut, B., 1999. Localization of knowledge and the mobility of engineers in regional networks. *Management Science* 45(7): 905–917.

- [6] Anton, J.J., Yao, D.A., 1995. Start-ups, spin-offs, and internal projects. *Journal of Law, Economics, & Organization* 11(2): 362–378.
- [7] Audretsch, D.B., 1998. Agglomeration and the location of innovative activity. *Oxford Review of Economic Policy* 14(2): 18–29.
- [8] Audretsch, D.B., Feldman, M., 1996. R&D Spillovers and the geography of innovation and production. *The American Economic Review* 86(3): 630–640.
- [9] Belenzon, S., Schankerman, M., 2013. Spreading the word: Geography, policy, and knowledge spillovers. *The Review of Economics and Statistics* 95(3): 884–903.
- [10] Bessen, J., Maskin E., 2009. Sequential innovation, patents, and imitation. *The RAND Journal of Economics* 40(4): 611–635.
- [11] Callaert, J., Pellens, M., Van Looy, B., 2014. Sources of inspiration? Making sense of scientific references in patents. *Scientometrics* 98(3): 1617–1629.
- [12] Chatterji, A.K., 2009. Spawned with a silver spoon? Entrepreneurial performance and innovation in the medical device industry. *Strategic Management Journal* 30(2): 185–206.
- [13] Delgado, M., Porter, M.E., Stern, S., 2008. Convergence, clusters, and economic performance. Institute for Strategy and Competitiveness.
- [14] Delgado, M., Porter, M.E., Stern, S., 2014. Clusters, convergence, and economic performance. *Research Policy* 43(10): 1785–1799.
- [15] Duguet, E., MacGarvie, M., 2005. How well do patent citations measure flow of technology? Evidence from French innovation surveys. *Economics of Innovation and New Technology* 14(5): 375–393.
- [16] Graham, S.J.H., Mowery, D.C., 2004. Submarines in software? Continuations in U.S. software patenting in the 1980s and 1990s. *Economics of Innovation and New Technology* 13(5): 443–456.
- [17] Green, J.R., Scotchmer, S., 1995. On the division of profit in sequential innovation. *The RAND Journal of Economics* 26(1): 20–33.
- [18] Jaffe, A.B., Trajtenberg, M., Fogarty, M.S., 2000. The meaning of patent citations: Report on the NBER/Case-Western Reserve Survey of Patentees (No. w7631). National Bureau of Economic Research.
- [19] Jaffe, A.B., Trajtenberg, M., Henderson, R., 1993. Geographic localization of knowledge spillovers as evidenced by patent citations. *The Quarterly Journal of Economics* 108: 577–598.
- [20] Krugman, P., 1991. *Geography and Trade*. Cambridge: MIT Press.
- [21] Laitner, J., Stolyarov, D., 2013. Derivative ideas and the value of intangible assets. *International Economic Review* 54(1): 59–95.

- [22] Murata, Y., Nakajima, R., Okamoto, R., Tamura, R., 2014. Localized knowledge spillovers and patent citations: A distance-based approach. *The Review of Economics and Statistics* 96(5): 967–985.
- [23] O’Donoghue, T., Scotchmer, S., Thisse, J., 1998. Patent breadth, patent life, and the pace of technological progress. *Journal of Economics & Management Strategy* 40(4): 611–635.
- [24] Porter, M.E., 1990. *The Competitive Advantage of Nations*. New York: Free Press.
- [25] Porter, M.E., 2003. The economic performance of regions. *Regional Studies* 37: 549–578.
- [26] Porter, M.E., 2007. Clusters and economic policy: Aligning public policy with the new economics of competition. White Paper. Institute for Strategy and Competitiveness, Harvard Business School.
- [27] Reitzig, M., Henkel, J., Heath, C., 2007. On sharks, trolls, and their patent prey: Unrealistic damage awards and firms’ strategies of “being infringed.” *Research Policy* 36(1): 134–154.
- [28] Roach, M., Cohen, W.M., 2013. Lens or prism? Patent citations as a measure of knowledge flows from public research. *Management Science* 59(2): 504–525.
- [29] Rosenthal, S.S., Strange, W.C., 2001. The determinants of agglomeration. *Journal of Urban Economics* 50(2): 191–229.
- [30] Shaver, J. M., Flyer, F., 2000. Agglomeration economies, firm heterogeneity, and foreign direct investment in the United States. *Strategic Management Journal* 21(12): 1175–1193.
- [31] Swann, P., 1998. Clusters in the US computing industry. In: Swann, P., Prevezer, M., Stout, D. (Eds.), *The Dynamics of Industrial Clustering: International Comparisons in Computing and Biotechnology*. Oxford University Press, Oxford, pp. 76–105.
- [32] Thompson, P., Fox-Kean, M., 2005. Patent citations and the geography of knowledge spillovers: A reassessment. *American Economic Review*: 450–460.
- [33] Wagner S., Hoisl, K., Thoma G., 2014. Overcoming localization of knowledge: The role of professional service firms. *Strategic Management Journal* 35(11): 1671–1688.
- [34] Zhao, M., 2006. Conducting R&D in countries with weak intellectual property rights protection. *Management Science* 52(8): 1185–1199.

9 Data Appendix

9.1 Priority year

Priority reversals are determined by examining the priority years of the citing and cited patents, and disclosure reversals are determined by examining the priority year of the citing patent and the earliest publication year of the cited patent. We use the priority and continuation dataset in PatStat to determine priority years. We take each of the citing and cited patents and find other patents linked to it via priority established by the Paris Convention of 1883 (PatStat TLS 204), provisional applications (PatStat TLS 204), and continuing applications (PatStat TLS 216). From this patent set, we keep the patent with the earliest application filing year as the priority year.

9.2 Earliest publication year

The earliest publication year of an invention is determined by identifying the earliest publication year from all of the patents within a patent family. The European Patent Office recognizes two types of patent families, simple and extended. We use the simple patent family notion, which groups patents claiming exactly the same set of priorities into a family. We construct this variable by linking the patent family dataset (TLS 218) with application (TLS 201), publication (TLS 211), priority (TLS 204), and continuation (TLS 216) datasets in PatStat.

9.3 Application year

Application year is the year in which an application for a patent is filed. Application year is different from the priority year as it may not indicate the year in which the invention is created, but merely the year in which an application for a patent is filed. We use the applications dataset (TLS 201) in PatStat to find application years of all citing, cited, and control patents.

9.4 Publication year

Publication year is the year in which an application or a patent is disclosed to the public. Publication year is different from the earliest publication year as it may not indicate the year in which the invention is first disclosed to the public, but merely the year in which a patent application or patent is published. Since November 2000, the USPTO has started publishing patent applications 18 months after the priority date. This means that a patent application claiming an earlier priority date may be published less than 18 months from its application date. On the other hand, an applicant may request publication of its application to be delayed until patent issuance if the applicant does not intend to file for a patent in a foreign country. We use the publication dataset (TLS 211) in PatStat to find publication years of all citing, cited, and control patents.

9.5 Inventor distance

We extract inventor cities and states for the citing and cited patents in our sample by linking several datasets in PatStat (TLS 206, TLS 226, TLS 227). We first identify inventors of each patent in our

sample from TLS 227, which is merged with TLS 226 to determine their cities and states. (Inventor addresses for patents whose applications were filed prior to 2005 are extracted from Thomson Innovation because PatStat provides addresses only from 2005.) We use a software application developed in-house, which communicates with Google Maps Geocode API, to find geographic coordinates of each city and state pair and to find a straight line distances between inventors linked by each citation. Whenever we see multiple inventors for a single patent, we take the city-state combination that occurs most frequently in our distance calculation. If there is an equal number of different city-state combinations, then we randomly choose a location from them. An advantage of using Google Maps Geocoding API to obtain geographic coordinates is that the Google API can intelligently handle several misspellings of city or state names. This is beneficial because spelling errors are frequent in the inventor city and state fields.

Table 1. Description of Variables and Concepts

<i>VARIBALE</i>	<i>DESCRIPTION</i>	<i>DATA SOURCE</i>
Actual citation	A reference recorded on the front page of a patent that indicates its use of prior knowledge contained in another patent	EPO Worldwide Patent Statistical Database
Control citation	A non-citing patent pair constructed by matching citing patents on four-digit IPC and publication year	EPO Worldwide Patent Statistical Database
Examiner citation	A citation added by a patent examiner during the patent examination process	EPO Worldwide Patent Statistical Database
Self-citation	A citation where the citing and cited patents have the same original assignee	Thompson Patent Database
In-text citation	A reference to an earlier patent extracted from the specifications of the invention described on a patent	EPO Worldwide Patent Statistical Database
Patent family	A set of patents that relate to the same underlying invention and thus have the same priority date	EPO Worldwide Patent Statistical Database
Priority year	Earliest year in which an invention's novelty, usefulness, and non-obviousness is recognized by a patent-issuing authority (i.e. invention year)	EPO Worldwide Patent Statistical Database
Publication year	Year in which an application or a patent is disclosed to the public	EPO Worldwide Patent Statistical Database
Earliest publication year	Year in which an invention underlying a patent family is disclosed to the public for the first time	EPO Worldwide Patent Statistical Database
Priority reversals	A citation where the priority year of the citing patent is earlier than the priority year of the cited patent	EPO Worldwide Patent Statistical Database
Disclosure reversals	A citation where the priority year of the citing patent is later than the priority year but earlier than the earliest publication year of the cited patent	EPO Worldwide Patent Statistical Database

Table 2. Summary Statistics for Main Variables

VARIBALES	No. Obs.	Mean	Std. Dev.	Distribution		
				10 th	50 th	90 th
<i>Distance (mi)</i>	2,090,376	919.2	899.9	0	674	2,413
<i>Citing priority year</i>	2,090,376	2,003.1	3.3	1999	2003	2008
<i>Cited priority year</i>	2,090,376	2,000.6	3.2	1997	2000	2005
<i>Citations per citing patent</i>	2,090,376	6.5	14.5	0	2	16
<i>Citations per cited patent</i>	2,090,376	45.8	76.9	4	21	110
<i>Priority lag in years</i>	2,090,376	2.6	1.8	0	3	5
<i>Dummy for reversal</i>	2,090,376	0.286	0.45	0	0	1
<i>Dummy for priority reversal</i>	2,090,376	0.036	0.19	0	0	0
<i>Dummy for disclosure reversal</i>	2,090,376	0.251	0.43	0	0	1
<i>Dummy for examiner citation</i>	2,090,376	0.371	0.48	0	0	1
<i>Dummy for self-citation</i>	2,090,376	0.233	0.42	0	0	1
<i>Dummy for $0 \leq \text{Distance} < 25$</i>	2,090,376	7.0	7.5	0	5	19
<i>Dummy for $25 \leq \text{Distance} < 50$</i>	2,090,376	34.0	6.6	26	33	44
<i>Dummy for $50 \leq \text{Distance} < 100$</i>	2,090,376	73.0	14.4	54	72	93
<i>Dummy for $100 \leq \text{Distance} < 150$</i>	2,090,376	125.2	15.4	104	125	146
<i>Dummy for $150 \leq \text{Distance} < 250$</i>	2,090,376	201.4	28.8	160	202	240
<i>Dummy for $250 \leq \text{Distance} < 500$</i>	2,090,376	370.4	65.6	281	362	464
<i>Dummy for $500 \leq \text{Distance} < 1000$</i>	2,090,376	737.7	143.7	548	714	943
<i>Dummy for $1000 \leq \text{Distance} < 1500$</i>	2,090,376	1275.3	165.7	1,044	1,282	1,474
<i>Dummy for $1500 \leq \text{Distance} < 2500$</i>	2,090,376	2032.1	327.1	1,573	2,072	2,431
<i>Dummy for $\text{Distance} \geq 2500$</i>	2,090,376	2594.2	60.6	2,528	2,569	2,682

Notes: This table provides summary statistics for the main variables used in the econometric analysis of the effect of distance on citation probability for the main sample. The sample consists only of actual patent citations. The publication years of citing patents in the sample covers years 2001 through 2014. *Priority lag* is the difference between the priority years of the citing and cited patents. *Dummy for reversal* is a variable that takes 1 if the priority date of the citing patent is earlier than the priority date or the earliest publication date of the cited patent and indicates that knowledge transmission is unlikely. *Dummy for examiner citation* is a variable that takes 1 if a citation was added by a patent examiner. *Dummy for self-citation* is a variable that takes 1 if the citing and cited patents are assigned to the same assignee.

Table 3. Comparisons of Main Citation Characteristics: Non-Reversals vs . Reversals

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Difference in means	Non-Reversals			Reversals		
<i>VARIBALES</i>	(3) minus (6)	No. Obs.	Mean	Std. Dev.	Obs.	Mean	Std. Dev.
<i>Distance (mi)</i>	129.97**	1,492,051	956	891.64	598,325	826	913.52
<i>% of citations within 50 miles</i>	-0.13**	1,492,051	0.26	0.44	598,325	0.39	0.49
<i>Citing priority year</i>	2.028**	1,492,051	2004	3.17	598,325	2002	3.31
<i>Cited priority year</i>	-0.73**	1,492,051	2000	3.20	598,325	2001	3.27
<i>Citations per citing patent</i>	-0.32**	1,492,051	6.43	14.24	598,325	6.75	15.00
<i>Citations per cited patent</i>	7.67**	1,492,051	48.00	80.38	598,325	40.33	66.93
<i>Citation lag in years</i>	2.76**	1,492,051	3.39	1.23	598,325	0.63	1.50
<i>Dummy for examiner citation</i>	-0.14**	1,492,051	0.33	0.47	598,325	0.47	0.50
<i>Fraction of Citations in the Same IPC</i>	-0.08**	1,492,051	0.20	0.40	598,325	0.28	0.45
<i>Fraction of Self Citations</i>	-0.14**	1,492,051	0.19	0.4	598,325	0.33	0.5

Notes: This table presents mean comparisons of main variables between non-reversals and reversals. The sample consists of actual patent citations. The publication years of citing patents covers years 2001 through 2014. ** p<0.01, * p<0.05

Table 4. The Effect of Distance on the Probability of Citation for Citation Non-Reversals vs . Reversals

Dependent Variable:	<i>Dummy for an actual citation</i>						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
<i>VARIBALES</i>	All	All	Distance effect for non-reversals and reversals			Nonlinear distance effect	
			All	Non-reversal	Reversal	Non-reversal	Reversal
<i>Log(Distance)</i>	-0.086** (0.001)		-0.079** (0.001)	-0.078** (0.001)	-0.091** (0.001)		
<i>Log(Distance) x Dummy for Reversal</i>			-0.019** (0.001)				
<i>Dummy for Reversal</i>			0.082** (0.002)				
<i>Dummy for 25 ≤ Distance < 50 miles</i>		-0.210** (0.003)				-0.178** (0.003)	-0.238** (0.005)
<i>Dummy for 50 ≤ Distance < 100 miles</i>		-0.421** (0.004)				-0.376** (0.004)	-0.455** (0.008)
<i>Dummy for 100 ≤ Distance < 150 miles</i>		-0.465** (0.004)				-0.408** (0.004)	-0.521** (0.008)
<i>Dummy for 150 ≤ Distance < 250 miles</i>		-0.484** (0.003)				-0.429** (0.003)	-0.542** (0.006)
<i>Dummy for 250 ≤ Distance < 500 miles</i>		-0.514** (0.002)				-0.454** (0.002)	-0.576** (0.004)
<i>Dummy for 500 ≤ Distance < 1000 miles</i>		-0.500** (0.002)				-0.448** (0.002)	-0.535** (0.003)
<i>Dummy for 1000 ≤ Distance < 1500 miles</i>		-0.547** (0.002)				-0.490** (0.002)	-0.595** (0.004)
<i>Dummy for 1500 ≤ Distance < 2500 miles</i>		-0.542** (0.002)				-0.484** (0.002)	-0.584** (0.003)
<i>Dummy for Distance ≥ 2500 miles</i>		-0.490** (0.002)				-0.439** (0.003)	-0.521** (0.004)
Tech Cluster Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Cited Patent Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Number of Reversals	1,194,951	1,194,951	1,194,951	-	1,194,951	-	1,194,951
Observations	4,180,752	4,180,752	4,180,752	2,985,801	1,194,951	2,985,801	1,194,951
R-squared	0.101	0.104	0.102	0.144	0.389	0.147	0.392

Notes: This table presents the effect of distance on citation probability for non-reversals and reversals. The sample consists of actual USPTO citations and non-citing, control citations that are randomly matched to citing patents on publication year and 4-digit IPC code. Publication years of the citing patents range from 2001 and 2014. Distance dummies are included to show non-linear effect of distance for different distance ranges. (The reference category is 0-25 miles.) The time lag between priority years of the citing and cited year is limited to +-5 years. Cited patent fixed effects are included to control for invariant patent-level characteristics that may influence citation probability. Dyadic dummies indicating citations between leading tech / research clusters (i.e. Austin, MA Route 128, Raleigh-Durham, San Diego, and Silicon Valley) are included. Standard errors are robust to heteroskedasticity and clustered at the cited patent application level to allow for correlation among patents citing the same patent. ** p<0.01, * p<0.05

Table 5. Types and Sources of Knowledge

Source Type	1 Patent	2 Inventor	3 Intermediary (e.g. Patent Attorney)
Invention	Not Localized	Localized	Localized
Background Knowledge	Not Localized	Localized	Localized

Notes: This figure shows the types and sources of knowledge distinguished in the paper. Knowledge about invention is knowledge directly related to the invention while background knowledge is knowledge or methods that might underlie the invention but not directly related to the invention. Knowledge can also have multiple sources that can influence temporal and geographical characteristics of knowledge transmission. The six inner cells indicate whether knowledge tends to be localized given its type and source.

Table 6. Citation Reversals and Local Inventor Interactions

Dependent Variable:	<i>Dummy for an actual citation</i>					
	(1)	(2)	(3)	(4)	(5)	(6)
	Non-reversals	Reversals				Examiner Reversals
<i>VARIBALES</i>	Excl. self cites	Excl. self cites	Excl. self and PA cites	Priority excl. self cites	Disclosure excl. self cites	Excl. self cites
<i>Dummy for 25 ≤ Distance < 50 miles</i>	-0.094** (0.004)	-0.116** (0.007)	-0.102** (0.007)	-0.089** (0.021)	-0.102** (0.008)	-0.108** (0.012)
<i>Dummy for 50 ≤ Distance < 100 miles</i>	-0.211** (0.005)	-0.232** (0.009)	-0.210** (0.009)	-0.181** (0.026)	-0.222** (0.011)	-0.206** (0.016)
<i>Dummy for 100 ≤ Distance < 150 miles</i>	-0.220** (0.005)	-0.252** (0.009)	-0.227** (0.009)	-0.183** (0.026)	-0.246** (0.011)	-0.211** (0.015)
<i>Dummy for 150 ≤ Distance < 250 miles</i>	-0.242** (0.004)	-0.276** (0.007)	-0.251** (0.007)	-0.200** (0.020)	-0.264** (0.008)	-0.259** (0.011)
<i>Dummy for 250 ≤ Distance < 500 miles</i>	-0.256** (0.003)	-0.296** (0.005)	-0.271** (0.005)	-0.198** (0.015)	-0.291** (0.006)	-0.284** (0.009)
<i>Dummy for 500 ≤ Distance < 1000 miles</i>	-0.245** (0.003)	-0.250** (0.005)	-0.225** (0.005)	-0.180** (0.014)	-0.242** (0.005)	-0.248** (0.007)
<i>Dummy for 1000 ≤ Distance < 1500 miles</i>	-0.281** (0.003)	-0.302** (0.005)	-0.276** (0.005)	-0.207** (0.015)	-0.297** (0.006)	-0.294** (0.008)
<i>Dummy for 1500 ≤ Distance < 2500 miles</i>	-0.278** (0.003)	-0.297** (0.004)	-0.272** (0.005)	-0.207** (0.013)	-0.291** (0.005)	-0.291** (0.007)
<i>Dummy for Distance ≥ 2500 miles</i>	-0.241** (0.003)	-0.250** (0.005)	-0.226** (0.005)	-0.185** (0.015)	-0.235** (0.006)	-0.241** (0.008)
Tech Cluster Controls	Yes	Yes	Yes	Yes	Yes	Yes
Cited Patent Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2,661,472	985,296	977,560	249,777	735,519	451,344
R-squared	0.144	0.362	0.362	0.712	0.392	0.378

Notes: This table presents results from comparing localization of citation non-reversals with different subsets of citation reversals. The sample consists of actual USPTO citations and non-citing, control citations that are randomly matched to citing patents on publication year and 4-digit IPC code. Publication years of the citing patents range from 2001 and 2014. Distance dummies are used to show non-linear effect of distance for different distance ranges. (The reference category is 0-25 miles.) The time lag between priority years of the citing and cited year is limited to +5 years. Standard errors are clustered at the cited patent level to allow for correlation among patents citing the same patent. Cited patent fixed effects are included to control for invariant patent-level characteristics that may influence citation probability. Dyadic dummies indicating citations between leading tech / research clusters (i.e. Austin, MA Route 128, Raleigh-Durham, San Diego, and Silicon Valley) are included. Standard errors, in paranthesis, are robust to heteroskedasticity. ** p<0.01, * p<0.05

Table A1. Alternative Specification: Within-patent Comparison of Citation Reversals and Non-reversals

Dependent Variable:	<i>Dummy for a citation with distance less than 50 miles</i>					
	(1)	(2)	(3)	(4)	(5)	(6)
	Examiner Citaion Reversals excl. self-citations			Inventor Citaion Reversals excl. self-citations		
<i>VARIBALES</i>	All Examiner Reversals	Priority Reversals	Disclosure Reversals	All Inventor Reversals	Priority Reversals	Disclosure Reversals
<i>Dummy for non-reversal</i>	-0.024** (0.005)	-0.016 (0.015)	-0.025** (0.005)	-0.012** (0.003)	-0.002 (0.008)	-0.015** (0.004)
Citing Patent IPC Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Citing Publication Year	Yes	Yes	Yes	Yes	Yes	Yes
Cited Patent Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	26,610	2,748	24,130	55,298	10,328	46,008
Sample mean	0.143	0.127	0.146	0.174	0.156	0.178
R-squared	0.543	0.631	0.540	0.645	0.695	0.643

Notes: This table presents a comparison of localization of non-reversals and reversals consisting of the same cited patent and whose citing patents are in the same publication year and IPC. The publication years of citing patents in the sample cover from 2001 to 2014. Cited patent fixed effects are included to control for invariant patent-level characteristics of both citing and cited patents that may influence the outcome. Citing patent technology area (IPC) fixed effects and citing patent publication year fixed effects are included to control for different factors inherent to technology fields and publication years that may influence the outcome. Standard errors are robust to heteroskedasticity and clustered at the cited patent level to allow for correlation among patents citing the same patent. ** p<0.01, * p<0.05.

Table A2. Replicating the Results for Period 1977-2014

Dependent Variable:	(1)	(2)	(3)	(4)	(5)
<i>VARIBALES</i>	All	Nonlinear distance effect		No Priority Lag Restriction	
		Non-reversal	Reversal	Non-reversal	Reversal
Log(<i>Distance</i>)				-0.062** (0.001)	-0.082** (0.001)
<i>Dummy for 25 ≤ Distance < 50 miles</i>	-0.219** (0.002)	-0.191** (0.003)	-0.242** (0.005)		
<i>Dummy for 50 ≤ Distance < 100 miles</i>	-0.430** (0.003)	-0.392** (0.004)	-0.457** (0.007)		
<i>Dummy for 100 ≤ Distance < 150 miles</i>	-0.469** (0.003)	-0.420** (0.004)	-0.521** (0.008)		
<i>Dummy for 150 ≤ Distance < 250 miles</i>	-0.494** (0.002)	-0.447** (0.003)	-0.545** (0.005)		
<i>Dummy for 250 ≤ Distance < 500 miles</i>	-0.523** (0.002)	-0.472** (0.002)	-0.578** (0.004)		
<i>Dummy for 500 ≤ Distance < 1000 miles</i>	-0.512** (0.002)	-0.467** (0.002)	-0.541** (0.003)		
<i>Dummy for 1000 ≤ Distance < 1500 miles</i>	-0.553** (0.002)	-0.503** (0.002)	-0.596** (0.003)		
<i>Dummy for 1500 ≤ Distance < 2500 miles</i>	-0.550** (0.001)	-0.500** (0.002)	-0.588** (0.003)		
<i>Dummy for Distance ≥ 2500 miles</i>	-0.498** (0.002)	-0.454** (0.003)	-0.525** (0.004)		
Tech Cluster Controls	Yes	Yes	Yes	Yes	Yes
Cited Patent Fixed Effects	Yes	Yes	Yes	Yes	Yes
Number of Reversals	1,333,052	-	1,333,052	-	1,957,335
Observations	4,917,874	3,584,822	1,333,052	22,516,227	1,957,335
R-squared	0.106	0.150	0.403	0.051	0.41

Notes: This table presents the effect of distance on citation probability for non-reversals and reversals. The sample consists of actual USPTO citations and non-citing, control citations that are randomly matched to citing patents on publication year and 4-digit IPC code. Publication years of the citing patents range from 1977 and 2014. Distance dummies are included to show non-linear effect of distance for different distance ranges. (The reference category is 0-25 miles.) Cited patent fixed effects are included to control for invariant patent-level characteristics that may influence citation probability. Dyadic dummies indicating citations between leading tech / research clusters (i.e. Austin, MA Route 128, Raleigh-Durham, San Diego, and Silicon Valley) are included. Standard errors are robust to heteroskedasticity and clustered at the cited patent application level to allow for correlation among patents citing the same patent. ** p<0.01, * p<0.05

Table A3. The Effect of Distance on the Probability of Citation using Samples Matched on Three- and Six-Digit Classification Code

Dependent Variable:	<i>Dummy for an actual citation</i>					
	(1)	(2)	(3)	(4)	(5)	(6)
	Controls matched on 3-digit class code (Murata et al., 2014)			Controls matched on 6-digit IPC		
<i>VARIBALES</i>	All	Non-reversal	Reversal	All	Non-reversal	Reversal
<i>Log(Distance)</i>	-0.033** (0.001)	-0.033** (0.001)	-0.039** (0.001)	-0.067** (0.001)		
<i>Log(Distance) x Dummy for Reversal</i>	0.003 (0.003)			-0.020** (0.001)		
<i>Dummy for Reversal</i>	-0.106** (0.020)			0.054** (0.003)		
<i>Dummy for 25 ≤ Distance < 50 miles</i>					-0.162** (0.004)	-0.247** (0.006)
<i>Dummy for 50 ≤ Distance < 100 miles</i>					-0.307** (0.006)	-0.388** (0.010)
<i>Dummy for 100 ≤ Distance < 150 miles</i>					-0.352** (0.006)	-0.454** (0.010)
<i>Dummy for 150 ≤ Distance < 250 miles</i>					-0.357** (0.004)	-0.462** (0.007)
<i>Dummy for 250 ≤ Distance < 500 miles</i>					-0.360** (0.003)	-0.461** (0.005)
<i>Dummy for 500 ≤ Distance < 1000 miles</i>					-0.367** (0.003)	-0.460** (0.004)
<i>Dummy for 1000 ≤ Distance < 1500 miles</i>					-0.406** (0.003)	-0.510** (0.005)
<i>Dummy for 1500 ≤ Distance < 2500 miles</i>					-0.397** (0.003)	-0.492** (0.004)
<i>Dummy for Distance ≥ 2500 miles</i>					-0.372** (0.003)	-0.456** (0.005)
Tech Cluster Controls	Yes	Yes	Yes	Yes	Yes	Yes
Cited Patent Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Number of Reversals	17,294	17,294	17,294	793,636	793,636	793,636
Observations	1,045,828	1,028,534	17,294	2,645,064	1,851,428	793,636
R-squared	0.007	0.009	0.506	0.074	0.133	0.357

Notes: This table presents findings on localization of patent citations using datasets matched on three-digit technology classification code, as in Murata et al. (2014), and on six-digit IPC code constructed from PatStat. The publication years of cited patents in the sample matched on three-digit technology classification code range from 1975 to 1979 and those of citing patents range from 1976 to 1999. The publication years of citing patents in the sample matched on six-digit IPC code range from 2001 to 2014. Cited patent fixed effects are included to control for invariant patent-level characteristics that may influence citation probability. Dyadic dummies indicating citations between leading tech / research clusters (i.e. Austin, MA Route 128, Raleigh-Durham, San Diego, and Silicon Valley) are included. Standard errors are robust to heteroskedasticity and clustered at the cited patent level to allow for correlation among patents citing the same patent. ** p<0.01, * p<0.05

Table A4. Testing for Grace Period Effects

Dependent Variable:	Dummy for an actual citation				
	(1)	(2)	(3)	(4)	(5)
	Reversal lag 1-2 year	Reversal lag 2-3 years	Reversal lag 3-4 years	Reversal lag 4-5 years	Reversal lag 5-6 years
<i>VARIBALES</i>					
<i>Log(Distance)</i>	-0.085** (0.001)	-0.085** (0.001)	-0.085** (0.001)	-0.085** (0.001)	-0.085** (0.001)
<i>Log(Distance) x Dummy for Reversal</i>	-0.012** (0.002)	-0.010** (0.001)	-0.003 (0.002)	0.001 (0.003)	0.003 (0.003)
<i>Dummy for Reversal</i>	-0.165** (0.011)	-0.204** (0.010)	-0.285** (0.013)	-0.339** (0.018)	-0.377** (0.021)
Tech Cluster Controls	Yes	Yes	Yes	Yes	Yes
Cited Patent Fixed Effects	Yes	Yes	Yes	Yes	Yes
Number of Reversals	48,891	67,836	42,273	26,634	17,215
Observations	4,026,794	4,045,739	4,020,176	4,004,537	3,995,118
R-squared	0.109	0.11	0.11	0.111	0.11

Notes: This table presents results from testing potential effects of grace period on localization of non-reversals and reversals. Priority reversals are re-categorized to be non-reversals unless the priority date of the citing patent is at least 1 year before the priority date of the cited patent to account for the possibility of inventors taking advantage of grace period (1 year given to inventors to file for a patent once they publicly disclose their inventions). A *Reversal lag* is a negative priority lag. The sample consists of actual USPTO citations and non-citing, control citations that are randomly matched to citing patents on publication year and 4-digit IPC code. Publication years of the citing patents range from 2001 and 2014. Cited patent fixed effects are included to control for invariant patent-level characteristics that may influence citation probability. Dyadic dummies indicating citations between leading tech / research clusters (i.e. Austin, MA Route 128, Raleigh-Durham, San Diego, and Silicon Valley) are included. Standard errors are robust to heteroskedasticity and clustered at the cited patent level to allow for correlation among patents citing the same patent. ** p<0.01, * p<0.05

Table A5a: The Effect of Distance on the Probability of Citation across Technology Areas

Dependent Variable:	<i>Dummy for an actual citation</i>					
	(1)	(2)	(3)	(4)	(5)	(6)
	Chemical		Pharmaceutical		Biotechnology	
<i>VARIBALES</i>	Non-reversal	Reversal	Non-reversal	Reversal	Non-reversal	Reversal
<i>Dummy for 25 ≤ Distance < 50 miles</i>	-0.299** (0.035)	-0.237** (0.085)	-0.343** (0.018)	-0.282** (0.029)	-0.208** (0.031)	-0.320** (0.043)
<i>Dummy for 50 ≤ Distance < 100 miles</i>	-0.572** (0.049)	-0.616** (0.104)	-0.550** (0.023)	-0.551** (0.037)	-0.428** (0.038)	-0.585** (0.052)
<i>Dummy for 100 ≤ Distance < 150 miles</i>	-0.651** (0.052)	-0.511** (0.139)	-0.581** (0.026)	-0.669** (0.039)	-0.486** (0.038)	-0.760** (0.042)
<i>Dummy for 150 ≤ Distance < 250 miles</i>	-0.719** (0.033)	-0.844** (0.072)	-0.647** (0.017)	-0.666** (0.029)	-0.536** (0.029)	-0.752** (0.041)
<i>Dummy for 250 ≤ Distance < 500 miles</i>	-0.689** (0.022)	-0.766** (0.046)	-0.650** (0.013)	-0.739** (0.020)	-0.456** (0.021)	-0.559** (0.037)
<i>Dummy for 500 ≤ Distance < 1000 miles</i>	-0.760** (0.018)	-0.779** (0.029)	-0.683** (0.012)	-0.765** (0.016)	-0.540** (0.018)	-0.695** (0.024)
<i>Dummy for 1000 ≤ Distance < 1500 miles</i>	-0.783** (0.020)	-0.847** (0.035)	-0.707** (0.015)	-0.803** (0.020)	-0.512** (0.022)	-0.755** (0.026)
<i>Dummy for 1500 ≤ Distance < 2500 miles</i>	-0.753** (0.020)	-0.860** (0.034)	-0.698** (0.010)	-0.799** (0.012)	-0.546** (0.016)	-0.733** (0.017)
<i>Dummy for Distance ≥ 2500 miles</i>	-0.667** (0.040)	-0.723** (0.062)	-0.635** (0.016)	-0.628** (0.028)	-0.506** (0.021)	-0.691** (0.026)
Tech Cluster Controls	Yes	Yes	Yes	Yes	Yes	Yes
Cited Patent Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	29,879	10,997	78,055	41,205	45,827	28,839
R-squared	0.321	0.700	0.294	0.593	0.248	0.515

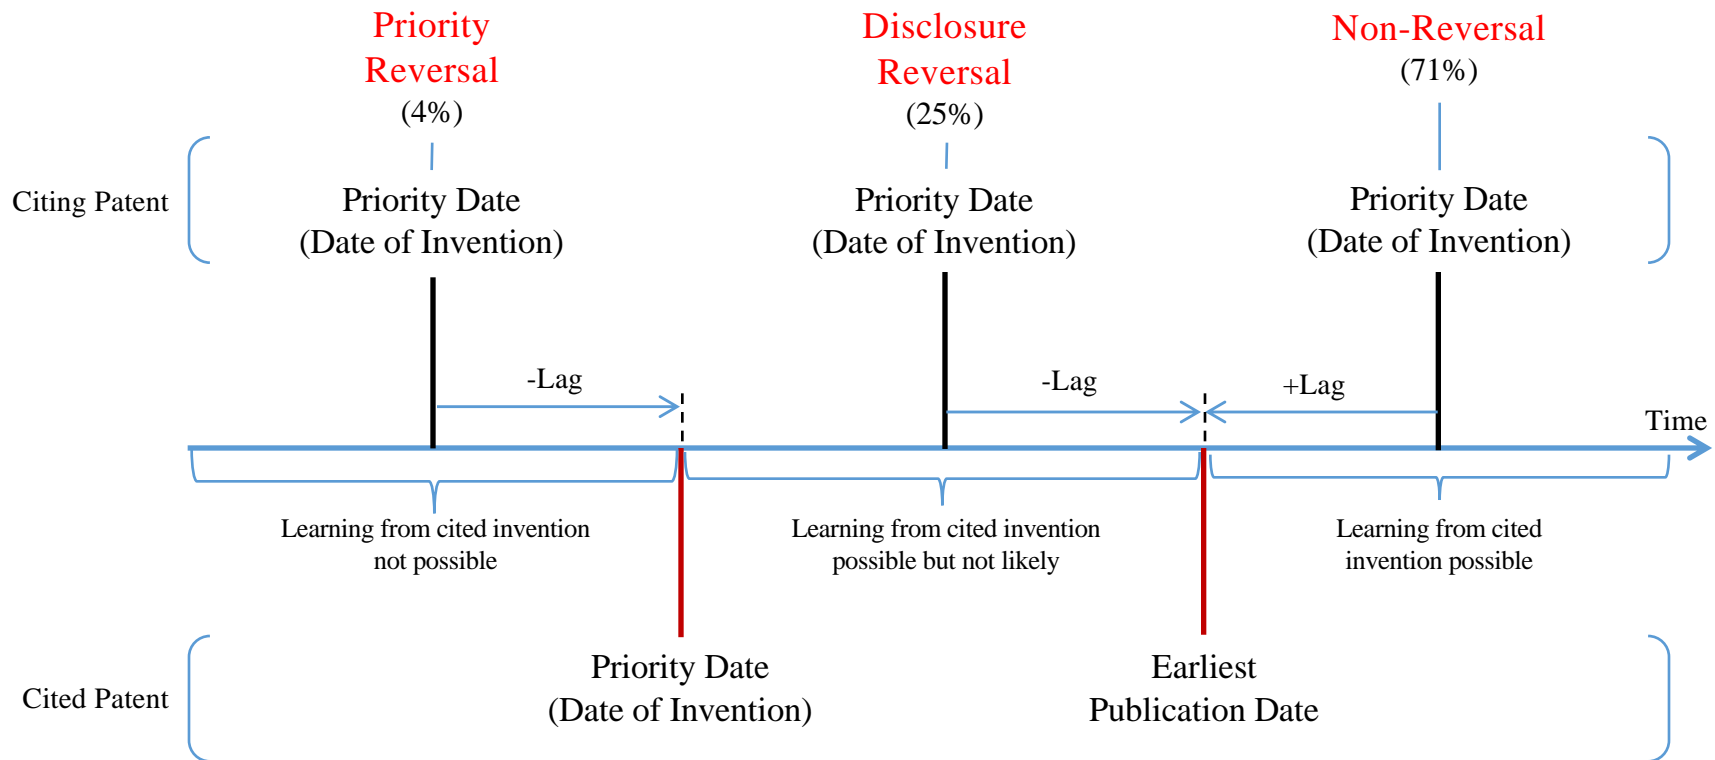
Notes: This table presents the effect of distance on the probability of citations for non-reversals and reversals across technology areas. The sample consists of actual USPTO citations and non-citing, control citations that are randomly matched to citing patents on publication year and 4-digit IPC code. Publication years of the citing patents range from 2001 and 2014. Distance dummies are included to show non-linear effect of distance for different distance ranges. (The reference category is 0-25 miles.) The time lag between priority years of the citing and cited year is limited to ± 5 years. Cited patent fixed effects are included to control for invariant patent-level characteristics that may influence citation probability. Dyadic dummies indicating citations between leading tech / research clusters (i.e. Austin, MA Route 128, Raleigh-Durham, San Diego, and Silicon Valley) are included. Standard errors are robust to heteroskedasticity and clustered at the cited patent level to allow for correlation among patents citing the same patent. ** $p < 0.01$, * $p < 0.05$

Table A5b: The Effect of Distance on the Probability of Citation across Technology Areas

Dependent Variable:	<i>Dummy for an actual citation</i>					
	(1)	(2)	(3)	(4)	(5)	(6)
	Medical Technology		Computer Technology		Telecommunications	
<i>VARIBALES</i>	Non-reversal	Reversal	Non-reversal	Reversal	Non-reversal	Reversal
<i>Dummy for 25 ≤ Distance < 50 miles</i>	-0.195** (0.010)	-0.244** (0.017)	-0.149** (0.007)	-0.191** (0.009)	-0.192** (0.016)	-0.252** (0.024)
<i>Dummy for 50 ≤ Distance < 100 miles</i>	-0.374** (0.014)	-0.470** (0.026)	-0.345** (0.012)	-0.451** (0.015)	-0.353** (0.020)	-0.492** (0.035)
<i>Dummy for 100 ≤ Distance < 150 miles</i>	-0.472** (0.015)	-0.642** (0.030)	-0.330** (0.010)	-0.469** (0.015)	-0.329** (0.023)	-0.544** (0.037)
<i>Dummy for 150 ≤ Distance < 250 miles</i>	-0.485** (0.011)	-0.575** (0.024)	-0.391** (0.008)	-0.536** (0.011)	-0.350** (0.016)	-0.493** (0.024)
<i>Dummy for 250 ≤ Distance < 500 miles</i>	-0.445** (0.007)	-0.586** (0.013)	-0.461** (0.006)	-0.578** (0.007)	-0.348** (0.013)	-0.519** (0.018)
<i>Dummy for 500 ≤ Distance < 1000 miles</i>	-0.441** (0.007)	-0.612** (0.011)	-0.360** (0.005)	-0.453** (0.005)	-0.379** (0.010)	-0.499** (0.014)
<i>Dummy for 1000 ≤ Distance < 1500 miles</i>	-0.442** (0.007)	-0.632** (0.012)	-0.450** (0.005)	-0.562** (0.006)	-0.383** (0.011)	-0.521** (0.015)
<i>Dummy for 1500 ≤ Distance < 2500 miles</i>	-0.457** (0.006)	-0.606** (0.009)	-0.425** (0.005)	-0.537** (0.005)	-0.381** (0.009)	-0.541** (0.013)
<i>Dummy for Distance ≥ 2500 miles</i>	-0.483** (0.008)	-0.640** (0.014)	-0.356** (0.005)	-0.429** (0.006)	-0.411** (0.012)	-0.520** (0.017)
Tech Cluster Controls	Yes	Yes	Yes	Yes	Yes	Yes
Cited Patent Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	270,540	92,930	676,162	363,782	164,771	82,259
R-squared	0.133	0.472	0.150	0.325	0.159	0.385

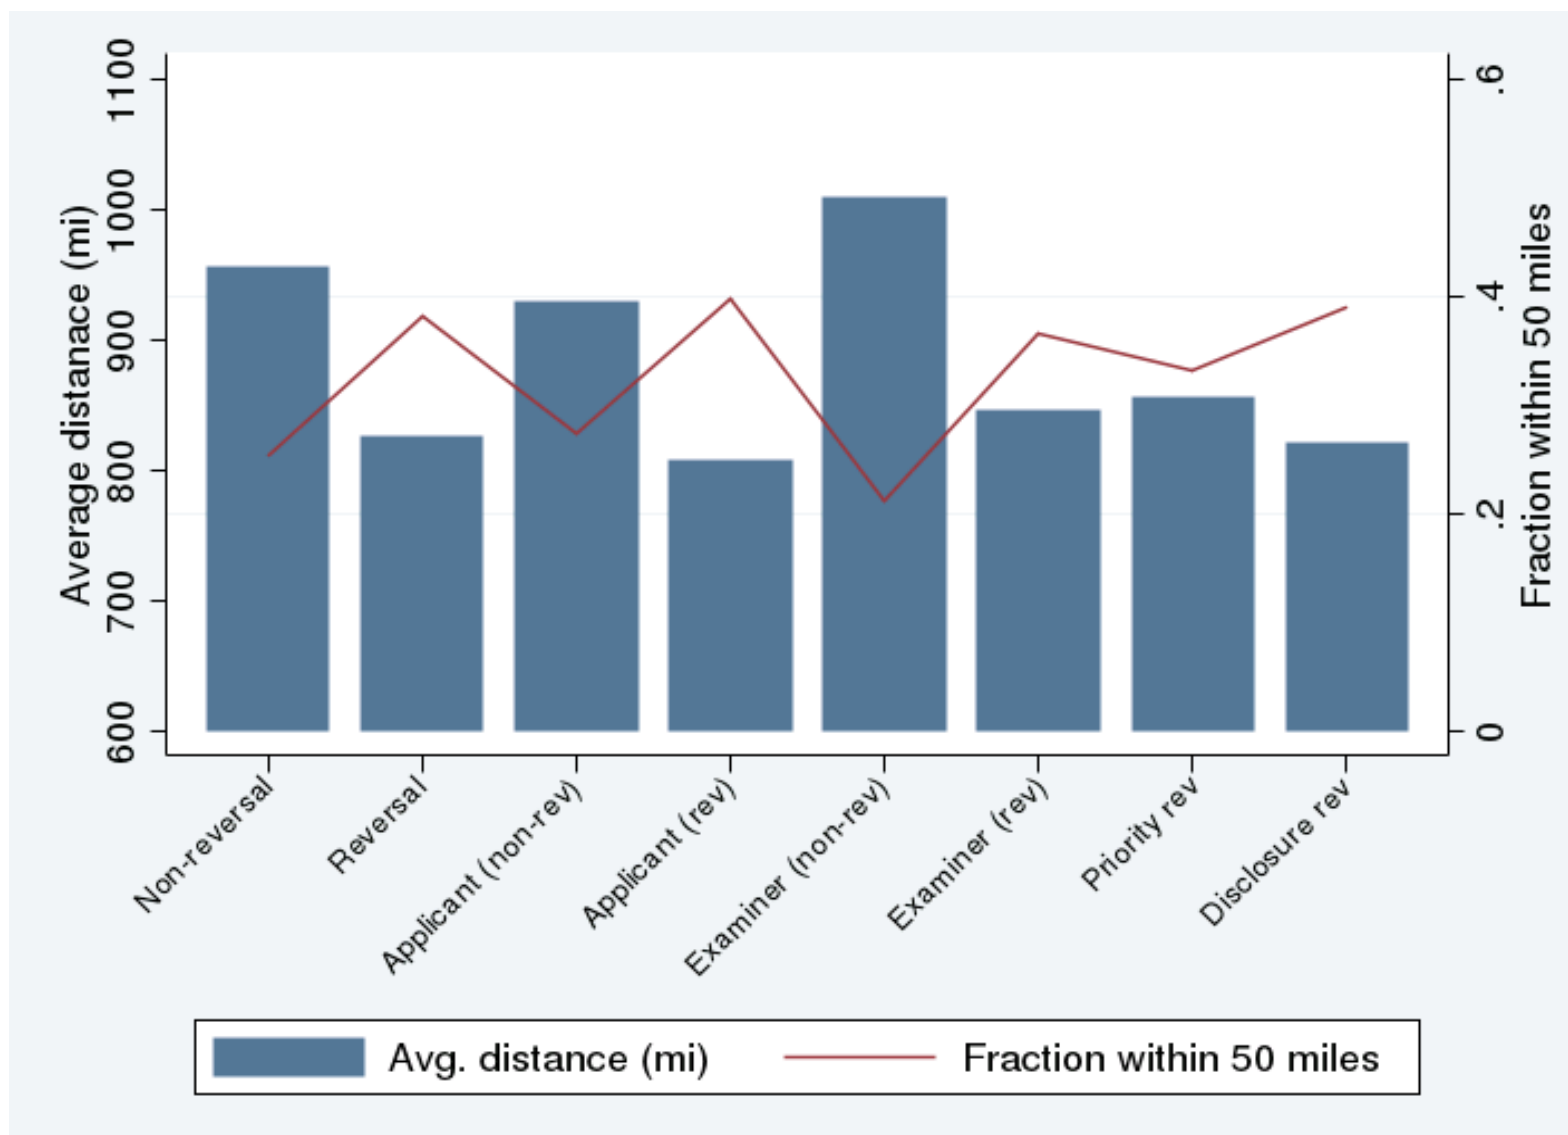
Notes: This table presents the effect of distance on the probability of citations for non-reversals and reversals across technology areas. The sample consists of actual USPTO citations and non-citing, control citations that are randomly matched to citing patents on publication year and 4-digit IPC code. Publication years of the citing patents range from 2001 and 2014. Distance dummies are included to show non-linear effect of distance for different distance ranges. (The reference category is 0-25 miles.) The time lag between priority years of the citing and cited year is limited to +/-5 years. Cited patent fixed effects are included to control for invariant patent-level characteristics that may influence citation probability. Dyadic dummies indicating citations between leading tech / research clusters (i.e. Austin, MA Route 128, Raleigh-Durham, San Diego, and Silicon Valley) are included. Standard errors are robust to heteroskedasticity and clustered at the cited patent level to allow for correlation among patents citing the same patent. ** p<0.01, * p<0.05

Figure 1. Priority and Disclosure Reversals



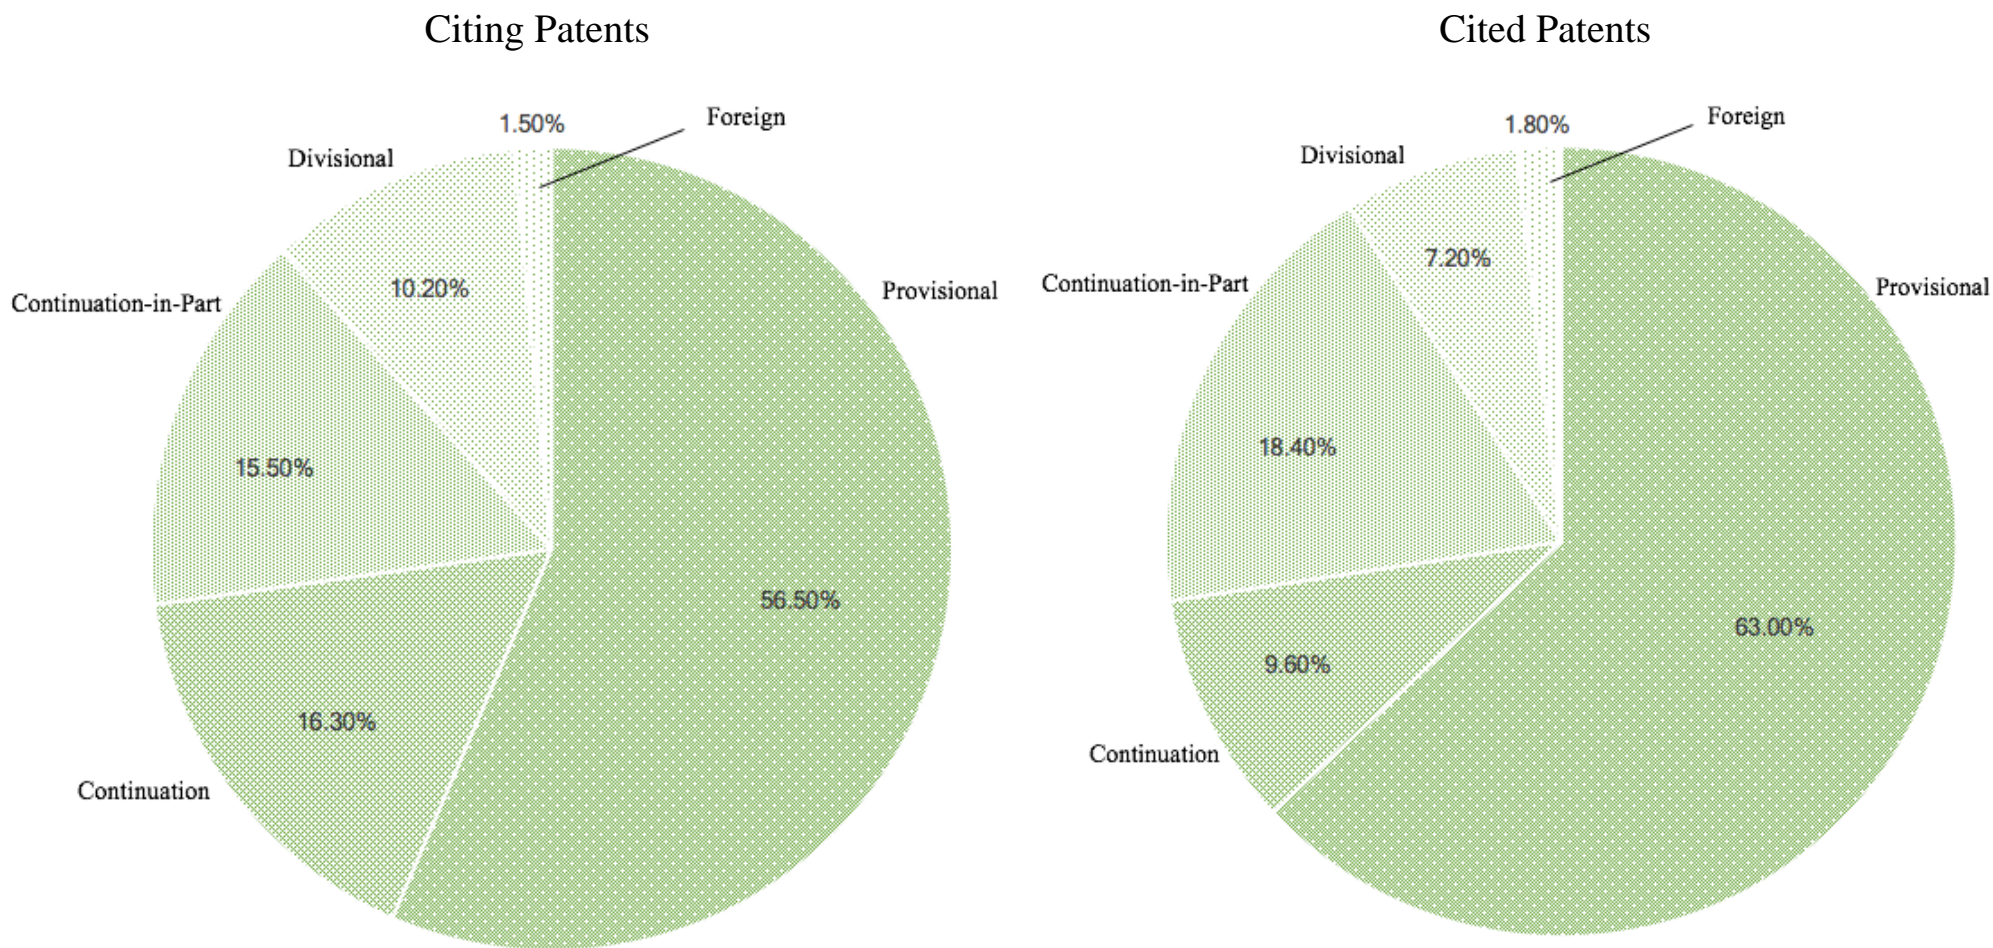
Notes: This figure presents how priority and disclosure reversals occur. Priority reversals are citations where the priority date of the citing patent comes before the priority date of the cited patent. Disclosure reversals are citations where the priority date of the citing patent comes after the priority date of the cited patent but before the earliest publication date of the cited patent. The priority date of the citing patent in non-reversals comes after both the priority date and the earliest publication date of the cited patent.

Figure 2. Average Distance Between Inventors by Citation Type



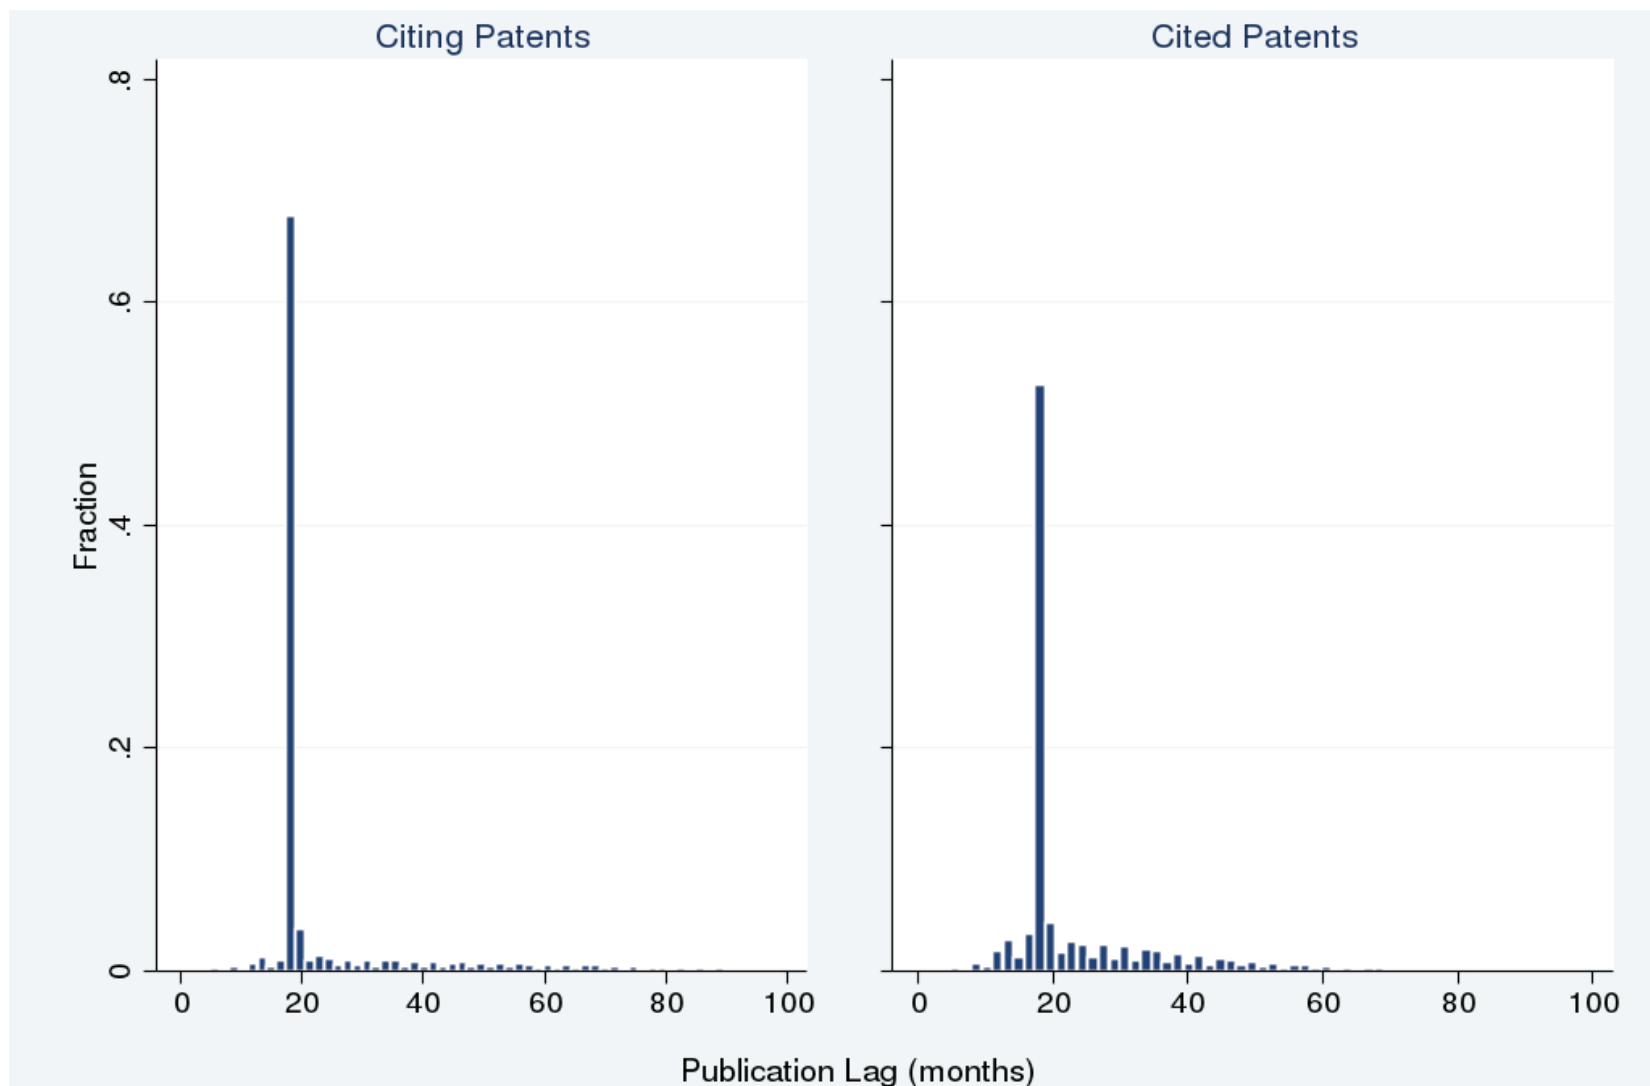
Notes: This figure compares localization of citations across different citation types. “Fraction within 50 miles” is the fraction of citations whose inventors reside within 50 miles of each other. The sample contains actual citations with citing patents covering years 2001 through 2014.

Figure A1. Breakdown of Priority Applications by Type



Notes: This figure presents the breakdown of the types of applications that citing and cited patents claim priorities based on. Provisional applications are those that are filed prior to a non-provisional (regular) application to establish an earlier priority date. (To file a provisional application, the specifications of the invention must be complete, while claims do not have to be.) Continuation applications make additional claims based on an existing invention specified in an earlier patent application. Divisional applications are filed to separate out distinct inventions from an earlier application. Continuation-in-part applications add extensions to an earlier invention. (The claims based on new subject matter are assigned as their priority date the application filing date of the continuation-in-part application.) We identify foreign and provisional applications from the applications table and the rest from the continuing applications table in PatStat.

Figure A2. Distribution of Citing and Cited Patents Over Publication Lag

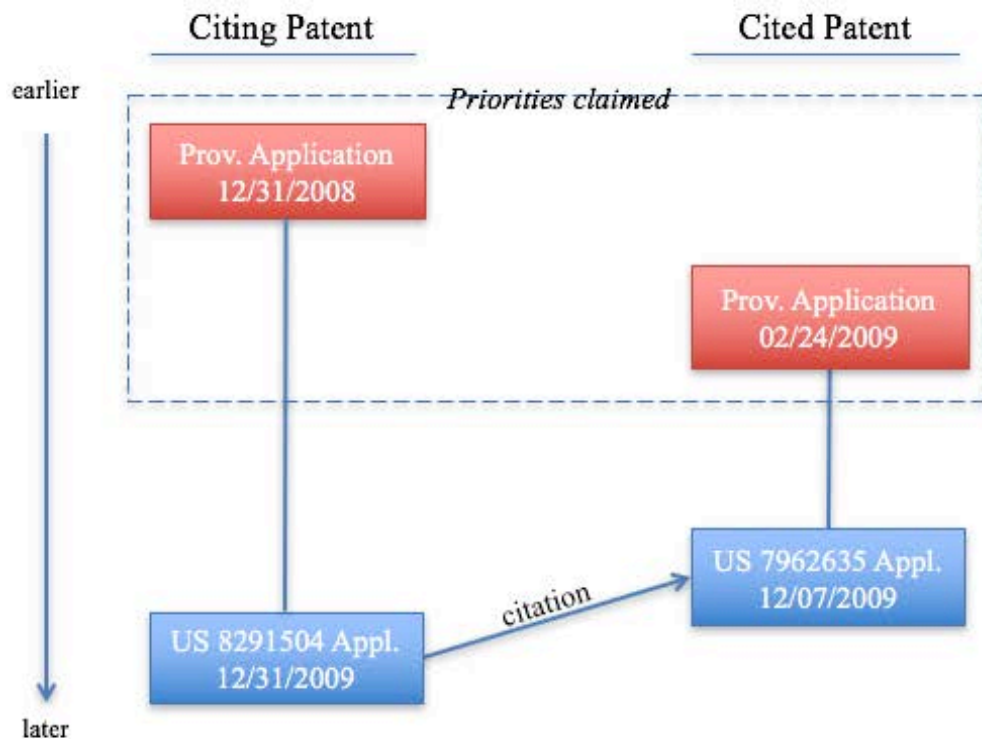


Notes: This figure presents the distribution of citing and cited patents over publication lag. “Publication lag” is time (in months) it takes for a patent or an application to be published and is measured as the earliest publication date minus the priority date.

Figure A3. Priority Reversal

Citing Patent (Akamai Technologies): Enforcing single stream per sign-on from a content delivery network (CDN) media server

Cited Patent (Telcordia Technologies): Systems and methods for single-session management in load-balanced application server clusters



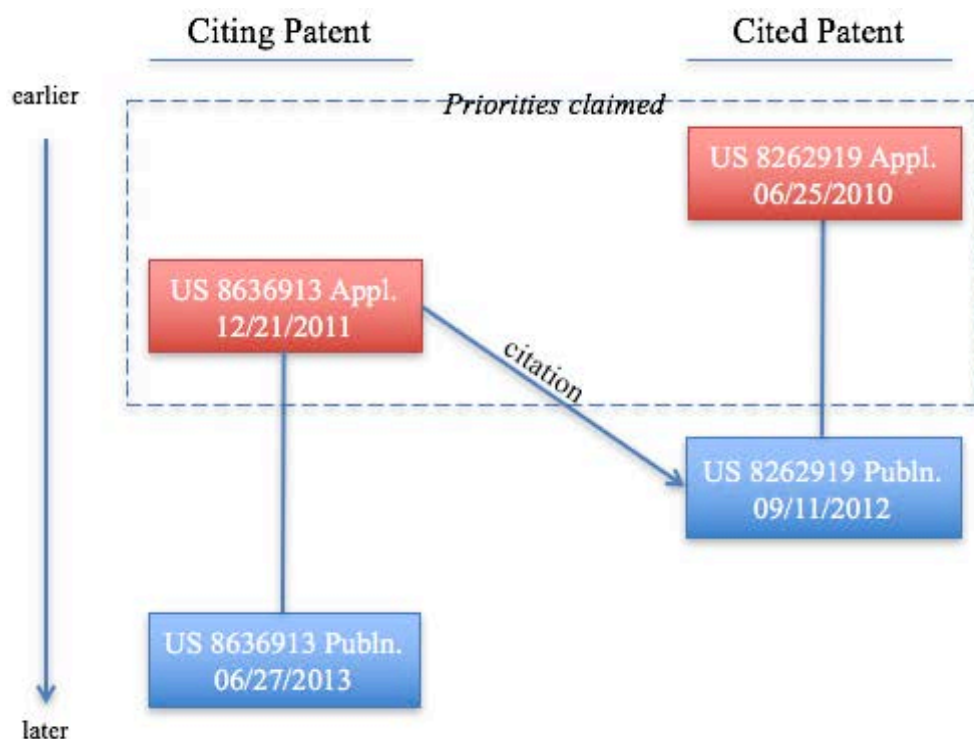
Patent Examination Transaction Highlights

Date	Citing Patent (US 8291504)	Cited Patent (US 7962635)
12/31/2008	Provisional application filed (application No. 61/141,964)	
02/24/2009		Provisional application filed (application No. 61/154,922)
12/07/2009		Application is filed for patent and claims priority based on the provisional application
12/07/2009		Information disclosure statement filed
12/31/2009	Application is filed for patent US 8291504 and claims priority based on the provisional application	
08/26/2010		Pre-grant publication
12/09/2010	Pre-grant publication	
06/14/2011		Patent is published after issuance
06/13/2012	Examiner cites US 7962635	
10/16/2012	Patent is published after issuance	

Notes: The application for the citing patent (12/31/2009) was filed after the application for the cited patent (12/7/2009), but the priority date of the citing patent (12/31/2008) comes before the priority date of the cited patent (2/24/2009), indicating that the citing invention was created before the cited invention was created.

Figure A4. Disclosure Reversal

Citing Patent (HGST Netherlands): Removing residues in magnetic head fabrication
 Cited Patent (Western Digital): Method and system for providing a perpendicular magnetic recording pole using multiple chemical mechanical planarization



Patent Examination Transaction Highlights		
Date	Citing Patent (US 8636913)	Cited Patent (US 8262919)
06/25/2010		Application is filed for patent
06/25/2010		Request for non-publication filed
12/21/2011	Application is filed for patent	Applicant responds after notice of a rejection
12/21/2011	Information disclosure statement filed	
09/11/2012		Patent published after issuance
06/27/2013	Pre-grant publication	
10/23/2013	Examiner cites US 8262919	
01/28/2014	Patent published after issuance	

Notes: The priority date of the citing patent (12/21/2011) comes after the priority date (6/25/2010) but before the earliest publication date (9/11/2012) of the cited patent, indicating that the citing invention was created before the cited invention was disclosed to the public.