Federal Reserve Bank of Philadelphia

Working Papers
Research Department

WP 17-44
Revised January 2019
December 2017
https://doi.org/10.21799/frbp.wp.2017.44

The Paper Trail of Knowledge Spillovers: Evidence from Patent Interferences

Ina Ganguli
University of Massachusetts–Amherst

Jeffrey Lin
Federal Reserve Bank of Philadelphia Research Department

Nicholas Reynolds
Brown University

ISSN: 1962-5361

The Paper Trail of Knowledge Spillovers: Evidence from Patent Interferences*

Ina Ganguli† Jeffrey Lin‡ Nicholas Reynolds§

January 2019

Abstract

We show evidence of localized knowledge spillovers using a new database of U.S. patent interferences terminated between 1998 and 2014. Interferences resulted when two or more independent parties submitted identical claims of invention nearly simultaneously. Following the idea that inventors of identical inventions share common knowledge inputs, interferences provide a new method for measuring knowledge spillovers. Interfering inventors are 1.4 to 4 times more likely to live in the same local area than matched control pairs of inventors. They are also more geographically concentrated than citation-linked inventors. Our results emphasize geographic distance as a barrier to tacit knowledge flows.

Keywords: Localized knowledge spillovers, multiple invention, patents, interferences
JEL classification: O30, R12

*We thank Marcus Berliant, Jerry Carlino, Bob Hunt, and John Stevens, and conference and workshop participants at the Urban Economics Association, the Federal Reserve System Committee on Regional Analysis, the NBER Productivity Workshop, EFPL Lausanne, the HBS Entrepreneurial Management Seminar, the CID Growth Lab, and George Washington University for comments and suggestions, and Aaron Rosenbaum for excellent research assistance.

†Department of Economics, University of Massachusetts Amherst, `iganguli@econs.umass.edu`.
‡Research Department, Federal Reserve Bank of Philadelphia, `jeff.lin@phil.frb.org`.
§Department of Economics and Population Studies and Training Center, Brown University, `nicholas_reynolds@brown.edu`.

1 Introduction

Innovative activity happens in cities. Patenting, research and development laboratories, venture capital investments, new product introductions, and new activities are concentrated in large, densely populated places (Carlino and Kerr, 2015; Carlino et al., 2007, Buzard et al., 2017; Chatterji et al., 2014; Feldman and Audretsch, 1999; Lin, 2011). Yet these places also feature higher labor costs and prices for factors such as land (Rosenthal and Strange, 2004). Why then do firms and inventors choose to locate in dense, costly areas? Evidently, dense agglomerations offer some benefits for the creation or application of new knowledge.

One intriguing hypothesis is that inventors benefit from localized knowledge spillovers. Geographic proximity may increase the frequency of interactions. In turn, this may ease the spread of knowledge. Proximity may be especially important for the diffusion of tacit knowledge, or knowledge that is not easily codified or written down (Glaeser, 1999; Feldman, 2000; Ganguli, 2015). This idea dates at least to Marshall (1920): "The mysteries of the trade become no mysteries; but are as it were in the air." Saxenian (1994) described places like the Wagon Wheel Bar near Intel, Raytheon, and Fairchild Semiconductor as "informal recruiting centers as well as listening posts; job information flowed freely along with shop talk." These kinds of knowledge spillovers may be especially important for "high tech hot spots" (Carlino, 2001). Localized knowledge spillovers are important factors in theories of growth and cities. Despite this, the identification of localized knowledge spillovers faces at least two challenges. One, knowledge spillovers are hard to measure: "They leave no paper trail by which they may be measured or tracked" (Krugman, 1991). Two, alternative sources of agglomeration economies yield similar predictions for measured productivity, wages, or other aggregates; i.e., different theories of the spatial concentration of firms and population are "observationally equivalent" (Duranton and Puga, 2004; Audretsch and Feldman, 2004).

To address these two challenges, we construct a novel database of *patent interferences*. Patent interferences measure (nearly) *simultaneous* instances of *identical* invention by two or more *independent* parties. Until 2013, the U.S. had a "first to invent" rule for assigning priority of invention. Two or more independent parties might make identical, nearly simultaneous claims of invention. Then, the U.S. Patent and Trademark Office (USPTO) investigated these competing claims in a patent interference. Under "first to invent," the USPTO awarded patent protection to the inventor who conceived (and reduced to practice) first. (This system contrasts with the "first to file" rule more common in the rest of the world and prevailing in the U.S. today. Under "first to file," the inventor who files first wins the patent.) Famously, simultaneous claims by Bell and Gray in February 1876 provoked an interference over priority for the telephone. More recently, the Broad Institute and the University of California interfered over CRISPR-Cas9, a gene editing technique.

By recording instances of common invention, patent interferences create a unique record of common knowledge inputs. In the spirit of Weitzman (1998), the view that new ideas result from combinations of existing ideas suggests that interfering inventors share similar existing knowledge. For example, interfering inventors may have similar background in chemistry, or they may have similar information about market conditions. Historians of science have noted the frequency of simultaneous, independent discovery, which Merton (1973) called "multiples"—e.g., the independent invention of calculus by Newton and Leibniz, or the independent formulation of the theory of natural selection by Darwin and Wallace. The commonplace nature of multiples has led some to speculate that they must be "in the air, products of the intellectual climate of a specific time and place" (Gladwell, 2008). To give some intuition about how common invention implies common knowledge inputs, consider Gladwell's (2008) description of Bell's and Gray's invention of the telephone:

> They arrived at electric speech by more or less the same pathway. They were trying to find a way to send more than one message at a time along a telegraph wire—which was then one of the central technological problems of the day. They had read the same essential sources—particularly the work of Phillipp Reis, the German physicist who had come startlingly close to building a working telephone back in the early eighteen-sixties.

Following the idea that inventors of identical inventions must share common knowledge inputs, we use interferences to test the hypothesis that geographically proximate inventors are more likely to share common knowledge inputs. To fix ideas, suppose some piece of knowledge K is essential to the invention of I. Further, this piece of knowledge K is discovered in the city where inventors Ann and Bo—who are working on I—reside, while Cam—who is also working on I—lives in a distant city. If knowledge spillovers are geographically localized, Ann and Bo may be more likely than Cam to encounter K via interactions with the discoverer of K (who may or may not be Ann or Bo themselves) or to hear about it through shared social ties. In turn, Ann and Bo may be more likely than Cam to successfully invent I, thus provoking an interference. In contrast, if access to required knowledge input K was unaffected by distance, then all three inventors (holding other factors constant) would be equally likely to invent I and be in interference.

Our approach departs from existing approaches that use the "paper trail" of patent citations to identify localized knowledge spillovers. In principle, a patent citation measures a piece of existing knowledge upon which the new patent builds. First, by focusing on identical inventive *output*, we circumvent the requirement of exactly measuring flows of knowledge across inventors using citations as measures of *inputs*.[1] Thus, our evidence is distinct from the "noisy" signal of knowledge spillovers from citations.[2] A survey of patenting inventors by Jaffe et al. (2000) reported that "one-half of all citations do not correspond to any spillover," and Jaffe et al. (1993) acknowledged that "an enormous number of spillovers [occur] with no citation."[3] Instead, our view is that identical invention likely requires a broad range of shared knowledge inputs. Citations may capture some of these shared knowledge inputs. But many knowledge inputs are not measured by citations—because they are not patentable and therefore not citeable, or perhaps because they are tacit and therefore difficult to codify. Thus, interferences may better capture spillovers of tacit knowledge, or at least knowledge that is not measured by a patent citation.

Second, a common interpretation of a citation is that knowledge has "spilled over" from the cited inventor to the citing inventor. Our setting is more agnostic: it doesn't matter where K originates. It might have been discovered by Ann, and Bo may have accessed it by interacting with Ann directly or with friends and colleagues (or friends of friends) of Ann's. But K may also have originated with another person co-located with Ann and Bo. The key idea is that the accessibility of K varies across geography, with Ann and Bo more likely to acquire it than Cam.[4]

We show that interfering inventors tend to be geographically concentrated, consistent with the hypothesis that geographically proximate inventors are more likely to share common knowledge inputs. Central to identification is whether or not the co-location of interfering inventors can be attributed to localized knowledge spillovers or some other factor. For example, inventors may co-locate to take advantage of thick markets for specialized skills (Bleakley and Lin, 2012), or they may co-locate to share a fixed, non-traded physical input

[1]Further, compared with earlier efforts to measure multiples (Ogburn and Thomas, 1922; Bikard, 2012), our database is large and relies on the real-time declaration of multiples by patent examiners instead of *ex post* measurement by researchers. In a contemporaneous paper, Baruffaldi and Raffo (2017) use citations classified by European Patent Office examiners as showing the claimed invention is not novel to construct a database of duplicated inventions. In contrast, interferences seem likely to measure duplication beyond cited prior art.

[2]Conversely, like citations, interferences are an indirect measure of shared knowledge inputs. Our contribution is that interferences seem likely to be generated in a very different manner compared with citations, providing important complementary evidence on the localization of knowledge spillovers.

[3]Inventors strategically not citing known prior art contribute to this problem.

[4]One could also re-interpret the meaning of patent citations in a similar manner: Instead of a direct spillover from the cited to the citing inventor, a citation is more likely when the citing inventor is more likely to learn about the cited patent from local ties.

(Helmers and Overman, 2017). To deal with the problem of inference under multiple sources of agglomeration economies, we use a matched control approach following Jaffe et al. (1993). For each interfering pair of inventors, we create a control pair by matching an interfering patent or application with a control patent matched on technology class and the date of invention. The control pair represents the expected proximity of inventors working in the same field and time period, except not conditioned on a "knowledge spillover" (i.e., an interference).

We find that interfering inventor pairs are 1.4 to 4.0 times more likely to live in the same local area—a place, city, or region—compared with control inventor pairs. Identification relies on (i) interfering inventors sharing knowledge inputs in common and (ii) matching on observables fully accounting for other factors besides localized knowledge spillovers affecting the geography of invention. We show that our results are robust to conditioning on additional controls. Further, as Jaffe et al. (1993) note, to the extent that control pairs also co-locate to take advantage of knowledge spillovers, we will tend to under-estimate the importance of localized knowledge spillovers. In addition, we avoid scale and border problems by using distance-based tests of localization instead of aggregating inventors to administrative spatial units of arbitrary size (e.g., cities or counties) (Duranton and Overman, 2005). For example, measuring co-location at the county or commuting zone level may understate localization if inventors are clustered on opposite sides of a county boundary. As Murata et al. (2014) show, downwards bias from scale and border problems is large relative to the upwards bias from imperfect matching as emphasized by Thompson (2006) and Thompson and Fox-Kean (2005).

Interfering inventors are more geographically concentrated compared with even *citation-linked* inventors. Pairs of cited-citing inventors may provide an even closer match compared with matching only on other observables. Further, the localization of interfering inventors compared with citation-linked inventors points to the important role of geographic proximity in facilitating tacit knowledge flows. We also analyze the role of previous co-inventor ties in mediating the relationship between localization and interference. Inventors may be linked by a social network defined by inventors who have previously been listed as co-inventors on a patent. Inventor pairs linked by previous co-inventor ties are more likely to interfere with each other compared with inventor pairs not linked by previous co-inventor ties. However, in contrast to recent evidence from Breschi and Lissoni (2005) and Head, Li, and Minondo (2015), we find little evidence that co-inventor ties are an important channel for the localization of knowledge spillovers.

2 Patent interferences: Background

Patent interferences were a unique feature of U.S. patent law. Through March 16, 2013, the U.S. had a "first to invent" rule for assigning priority of invention, versus the "first to file" rule more common in the rest of the world and that prevails in the U.S. today. When the USPTO received patent applications from multiple, independent parties with one or more claims covering "the same, or substantially the same, subject matter" (35 U.S. Code §135) at roughly the same time, it was obliged to investigate the competing claims to determine which party was entitled to patent protection. This investigation was known as a patent interference. It was conducted by the Board of Patent Appeals and Interferences (BPAI, hereafter Board) and determined who was first to invent, meaning both (i) who first had the idea (conception) and (ii) who first put the idea in workable form (reduction to practice). Typically, parties submitted dated laboratory notebooks, testimony by associates, and media reports as evidence of first invention. In this section, we summarize several key institutional features of patent interferences.[5]

An interference could be suggested by a patent examiner during their routine search for prior art, when (i) at least two U.S. patent applications satisfying a timing rule described below or (ii) one U.S. patent application and a recently-issued patent contained the "same patentable invention" (37 Code of Federal Regulations §1.601), or when "the subject matter of a claim of one party would, if prior art, have anticipated or rendered obvious the subject matter of a claim of the opposing party and vice versa" (37 C.F.R. §41.203). The patent examiner would then forward the patent application and a memorandum to the Board, which would declare the patent interference.[6]

The claim(s) of invention must have satisfied standard patentability rules—i.e., the claims must have been in an patent-eligible class, useful, novel, and non-obvious. In addition, the USPTO required that a timing rule be satisfied in order to avoid interferences resulting from the disclosure of patent applications themselves (i.e., publicized patent applications leading to copycat inventions). Thus, in the case of two or more interfering applications, the dates of application must have been no more than 3 months apart. In the case of an interfering issued patent and pending application, (a) the application's date must have been more than

[5]More details about the patent interference proceedings can be found in Calvert (1980), Calvert and Sofocleous (1982), Cohen and Ishii (2006), de Simone, Gambrell and Gareau (1963), and Kingston (2001). Lin (2014) reviews this literature and provides summary statistics for the patent interferences used in this study.

[6]A patent applicant may also in some cases suggest an interference with another application or a patent. Details about the requirements for an applicant to suggest an interference are available in BPAI (2004).

one year before the patent's grant date and (b) the application's date must have been no less than 3 months after the patent's application date.[7]

Upon declaring an interference, the Board defined "counts" corresponding to the interference. Each count was characterized by a distinct invention at issue; each application might claim several distinct inventions so an interference might involve multiple counts. The case was heard before a rotating three-judge panel from the Board. Inventors were assigned a "benefit date"—typically, the date of application, at either the USPTO or a foreign patent office. The inventor with the earlier benefit date was referred to as the "senior party." The burden of proof—i.e., demonstration of an earlier conception and reduction to practice—was on the "junior party." Interfering inventors typically submitted lab notebooks, eyewitness testimony, and other forms of independent corroboration to prove they were the first to invent.

Disposition	Description or example	Patent(s) to
Priority	Y conceived and reduced to practice first	Y
Settlement	X concedes; settlement terms confidential	Y
Abandonment	X concedes; no settlement	Y
Concealment	X kept invention secret, concealed best mode, or delayed filing	Y
Derivation	X stole invention from Y; i.e., prior conception by Y and communication of conception to X	Y
No interference in fact	Claims are actually distinct	X and Y
Common ownership	X and Y work for same conglomerate	X or Y
Unpatentable	Claims are not patentable (e.g., obvious)	No one

Table 1: Common case dispositions for interference between inventors X and Y

Interference cases were terminated by the judges' decision on priority or for some other reason. Table 1 summarizes common interference dispositions. A decision on priority meant a judgment that one party had first conceived of the invention and reduced it to practice. Parties could also settle at any stage. Normally, details of these agreements were kept secret. One party might also concede the case, without a settlement taking place—for example, if they realized their case was weak.

[7]Specifically, interfering claims among pending applications must be made within 1 year of each other (35 U.S.C. 135.b.2). In cases where an application's claims interfere with an already-issued patent, the claims must be made no later than 1 year prior to the patent's issue date (35 U.S.C. 135.b.1), and typically no later than 3 months after the patent's original application date (37 C.F.R. 1.608).

Other potential termination types are useful for distinguishing instances when inventors shared knowledge inputs versus other reasons. For example, according to interference rules, an interfering inventor would lose priority if the inventor had not immediately filed for a patent application following conception and reduction to practice.[8] In particular, the disclosure of the timeline of invention was the primary purpose of the interference proceeding. The alternative to admitting an intentional delay in applying for a patent would be to concede a later date for reduction to practice, weakening the case for priority.[9]

Sometimes, in the course of the proceeding, the Board might decide that the declaration of interference was mistaken, and that there was *no interference in fact*. This reported result is important because it allows us to isolate truly identical inventions, and not just near misses. The USPTO also made sure that interfering inventors did not share other factors. Claims from parties working for the same firm (e.g., different branches of a large corporation) were dismissed. If one party's application was derived—i.e., that one inventor's claims are directly sourced from the competing party through stealing or espionage—that was grounds for an adverse judgment.[10] Finally, in some less common outcomes, the Board ruled that the interference count was anticipated or otherwise unpatentable. As with no interference in fact, these judgments could be interpreted as mistakes by the original patent examiner.

Patent interferences were distinct from patent infringements in several ways. First, interferences were suggested by a patent examiner who specialized in a particular technological area. In some cases, a patent applicant could suggest a possible interference, but an interference is distinct from patent infringement, in which the holder of an existing patent sues an infringing party. In contrast to infringements, private parties could not sue for an

[8]Cohen and Ishii (2006) argue that interferences correspond to an incumbent-entrant game where incumbents decide to keep inventions secret for some period of time before filing a patent application. The requirement to promptly disclose an invention may have reduced the relevance of this margin as decisions went against interfering inventors who chose to keep their inventions secret for some time.

[9]Interference cases appear to vary in terms of whether inventors are aware of each other's efforts. In Lutzker v. Plet (1988), the United States Court of Appeals, Federal Circuit, affirmed that Lutzker was not entitled to a patent for a canape maker, despite having established conception and reduction to practice in early 1976, since he had delayed disclosure and filing for a patent until late 1980. In contrast, Plet received priority by demonstrating her conception and reduction to practice by early 1980, with a filing date of March 3, 1980. The original decision by the Board cited the failure of Lutzker to show renewed activity towards disclosure "until after Plet entered the field" as an important factor in the judgment against him.

[10]Cases involving stealing or espionage could still reflect a spillover if it is a knowledge input rather than the complete invention that was stolen. The "derivation" judgment (see Table 1) occurred when it could be proved that one party completely stole the invention from the other. This judgment is very rare and so including them in the analysis does not affect our results – there is only one case with a "derivation" outcome out of the 1,329 in our database.

interference. Second, interferences must involve parties with nearly simultaneous pending applications for patents. This feature makes interferences distinct from patent infringements, which typically involve leaders and followers.

Several features of interference practice allow us to rule out alternative explanations—outside of shared knowledge inputs—for multiple invention. First, interferences between parties with common ownership interests were not allowed. Thus, they seem unlikely to result from other shared factors or from within-firm spillovers. Second, as mentioned earlier, cases of no interferences in fact help us to distinguish between identical inventions and near misses. Third, the USPTO records the number of application claims involved in an interference, allowing us to determine the degree of overlap between interfering patent applications. Overall, interferences appear to involve very similar inventive claims.[11] Thus, *contra* Schmookler (1966), interferences can identify identical inventions, versus near misses.

Finally, to the extent that interference cases are costly to prosecute, interferences are likely to involve valuable patents, and thus actual inventions.[12]

3 Data and methodology

3.1 Interferences, patents, and applications

We constructed a database of patent interference cases for our analysis. Our database starts with information from 1,329 interference *decisions* issued by the USPTO Board of Interferences between 1998 and 2014.[13] These decisions were downloaded from the Board's "e-FOIA Reading Room." From each decision, we record information about the *case*, the *parties*, the *application(s)* and/or *patent(s)*, the *claims*, and the *inventors*. [14]

[11] See Appendix Table A1.

[12] A recent interference case decided in February 2017 involved patent rights to the CRISPR gene-editing technique. (The decision date puts the case outside our sample.) In that decision, the Board found no interference in fact—the inventions claimed by the competing inventors, assigned to the Broad Institute and the University of California, were separate and did not overlap. The validity of the Broad patents was a "surprise" to researchers in the field; as a sign of the value of the invention, the stock of the licensee to the Broad patents went up sharply following the decision (Pollack, 2017).

[13] There are a few decisions related to interferences declared well before 1998, as far back as the early 1980s, including a famous case over the method of producing the hepatitis B antigen. On average, however, the lag between interference declaration and decision dates is a few years. See Calvert and Sofocleous (1982, 1986, 1989, 1992, and 1995.)

[14] Note the following: (1) each case is argued between two or more parties; (2) each party may have one or more (co-)inventors; (3) each party may also have one or more applications and/or patents in interference; (4) each application or patent makes one or more claims; (5) one or more of these claims are declared by the examiner to be in interference.

The decisions typically report the: (i) names of the interfering inventors; (ii) seniority status of each party; (iii) associated patent and application numbers; (iv) assignees; (v) judges' names; (vi) application claims in interference; (vii) decision on priority at the claim level (if there was one) or other disposition of the case; (viii) legal counsel; and (ix) hearing and decision dates. Sometimes, terse decisions omit some of these details. When available, we collect these details using additional documents found on the USPTO's "eFile" site or the Patent Application Information Retrieval (PAIR) service. The eFile site sometimes lists the notice declaring the interference, from which we can observe (x) inventors' location of residence. For cases with documents available on the eFile site, we also record (xi) notices of settlement agreements. These notices acknowledge the existence of a settlement agreement, as opposed to a decision on priority or some other outcome.[15] The PAIR service provides an alternate source of information on assignees, case disposition, the decision date, and inventors' locations.[16] Note that for inventors never (eventually) issued a patent, information on inventor location is available *only* on the notice of interference on the eFile site or the PAIR service.

Table 2 summarizes case dispositions for our sample. The first two columns display frequencies and the share of cases by disposition for our full sample. Nearly 20 percent of cases resulted in a judgment on priority, while nearly 60 percent of cases were conceded. Concessions occur when one party files a request for adverse judgment. An abandonment occurs when one party fails to file at some stage of the case. We code these outcomes as they are noted in the decisions. Absent detail in the decision, it is difficult to ascertain the motivations for concessions (and we cannot rule out abandonments if failure to file is not mentioned.) However, for the sub-sample of 977 cases that we are able to match to documents on the eFile site, we code cases including an acknowledgement of settlement. In this sub-sample, settlements constitute the majority of concessions and nearly one-third of all cases. We are unable to characterize the remaining conceded cases that have no acknowledgement of settlement or text in the decisions referring to a failure to file.

Most of our analysis focuses on interference cases where the board's decisions report a settlement or judgment on priority. Shared knowledge inputs are more likely in these cases compared with other case dispositions. For example, cases dismissed for no interference in fact seem less likely to involve exactly common knowledge inputs. The frequencies of other

[15] Unfortunately, settlement agreements are sealed. Thus, we can note their existence, but we cannot analyze their contents.

[16] In nearly every case where they overlap, information available on the PAIR record confirms information recorded from the decision.

dispositions are listed in the bottom half of Table 2. About 9 percent of cases were dismissed because the claims were deemed unpatentable. Five percent were dismissed because the interfering parties were discovered to have assigned rights to a common owner, e.g., a multinational firm.[17] Three percent of cases were dismissed after a finding of no interference in fact.

Disposition	*Full sample*		*eFile*
Number of cases	1,329		977
Decision on priority	260	19.6%	19.7%
Conceded, total	781	58.8	58.1
... settled	.	.	32.8
... abandoned	92	6.9	5.5
... all other reasons	.	.	19.8
No interference in fact	46	3.5	3.4
Common ownership	64	4.8	4.7
Unpatentable	122	9.2	9.6
Other	56	4.2	4.5

Table 2: Distribution of interference case dispositions

Interferences appear to involve similar inventive claims, based on the counts declared by the Board and the corresponding claims of invention in interfering applications. In Appendix Table A1, we show that interfering applications overlap substantially in their inventive claims. On average, over three-quarters of a party's application claims are found to be in interference. Further, partial decision are rare. Typically, the interfering claims are awarded entirely to one party or the other.

We use the USPTO technology classifications to construct controls. We obtain a list of all the technology classifications assigned to a patent from the USPTO's Master Classification File. Patents are classified according to a 3-digit technology class or a 6-digit subclass. This information is available for all issued patents, but only for patent applications filed in 2001 and later. For earlier patent applications, we instead obtain classification information from the PAIR service, which records only a single main technology classification, considered the "primary classification" of that invention.

[17]In at least one case, a merger appeared to have been *caused* by the pending interference.

Next, we use information from patents and applications to measure the locations of inventors. For issued patents, we match our interference case database to the inventor disambiguation dataset of Lai et al. (2013). For all patents issued by the USPTO between 1975 and 2010, this includes the names and locations of all inventors, the patent application date, and assignee. This dataset also includes a unique inventor identifier that is consistent across all patents, as a result of a name disambiguation algorithm.

If an interference is decided against a party without an issued patent, then the losing party's patent application is never passed for issue. Since most *patent* databases include data only for issued patents, the decisions, the notices of interference from eFile, and the PAIR database as additional sources of inventor location are essential. The importance of supplementing the Lai et al. (2013) database with information from PAIR, the interference decisions or the notices of interference from eFile is illustrated by the frequency of inventor locations by data source. While 75 percent of inventors are located in the Lai et al. (2013) database, the remaining 25 percent of inventors' locations are recorded *only* in the decisions, eFile, or PAIR. At the case level, we are able to record inventor locations for at least 1 inventor in each party in 88 percent of interference cases, and all inventors in 85 percent of interference cases.

We then compute the distance between pairs of inventors, either those on opposing teams in an interference or control pairs which we construct and are described later. We use the place of residence of inventors — the place named in the bibliographic description of the patent or application. Places can be large cities, such as "San Francisco, CA," but also can be small towns or even unincorporated places. Unfortunately, the data are not detailed enough to describe co-location within a place. We obtain the latitude and longitude of each inventor's place of residence from the database of Lai et al. (2013), or by matching the place of residence of the inventor from either the PAIR service or the notice of interference to the Census Gazetteer file, which maps place names to latitude and longitude. We then compute the minimum geodesic distance across all possible pairings of inventors within an interference case.[18]

We also construct a geographic matching measure of co-location, in order to provide results comparable to earlier work and to highlight proximity at short distances. For example,

[18]For an interference between a party with m co-inventors and another party with n co-inventors, there are mn pairwise combinations of inventors. We report results using the minimum pairwise distance across all of these pairwise combinations. We also experimented with using the median or mean distance for each interfering pair of parties; as well as, following the convention of earlier work, the distance between the first-named inventors of each patent. Results using alternative measures are similar to those reported in the paper.

we define a variable that indicates whether or not an inventor pair shares the same place of residence, or if their places of residence are within within 161 kilometers (100 miles) of each other. By these measures we intend to capture localized interactions via social ties, workplace relationships, or random meetings. The 100-mile cutoff is comparable to a metropolitan area (used in Jaffe et al., 1993) or a commuting zone (as in Autor and Dorn, 2013).[19] An advantage of our distance-based measure compared with commuting zones is they avoid border and scale problems, as explained in more detail in Section 3.3.

3.2 Conceptual framework

We test for tacit knowledge flows by comparing the geographic proximity of interfering inventors to a set of matched control pairs, following the strategy of Jaffe et al. (1993). We construct control pairs which include one interfering application and one issued control patent matched on technology class and application date (within 180 days of the application date of the interfering application.) The idea is that the spatial distribution of control patent pairs is an appropriate counterfactual to the observed spatial distribution of interfering inventors. That is, control inventors likely face a similar location choice problem compared with interfering inventors. By comparing the geographical localization of interfering inventors to this counterfactual, we hope to control for all factors except common knowledge inputs. As Jaffe et al. (1993) note, this may be a conservative estimate in the sense that only the concentration *in excess* of the control pairs is attributed to localized knowledge spillovers. However, if matched control pairs do not adequately control for unobserved factors (other than knowledge spillovers) affecting the geography of invention then our estimate of knowledge spillovers will be biased upwards.

To fix ideas, consider the following reduced-form model describing the relationship between inventors' shared knowledge, geographic location, and the declaration of interference cases. First, assume that the probability $P(int_i)$ that an arbitrary pair of patents i are involved in an interference depends on their shared knowledge inputs A_i, observable factors X_i, and an unobservable factor n_i,

$$P(int_i) = g(A_i, X_i, n_i) \tag{1}$$

Further, assume that the degree of knowledge inputs shared by two inventors is a function

[19] On average, commuting zones are similar in size to a circle with diameter 80 miles; the largest consolidated statistical area (New York) is similar in size to a circle with diameter 130 miles.

of the geographic distance between them, $dist_i$, and an unobservable factor e_i,

$$A_i = f(dist_i, e_i) \tag{2}$$

We propose to test for the existence of localized knowledge spillovers by comparing the geographic distances between inventors involved in interference cases to that of inventors of similar control patent pairs. This will be a valid test under two assumptions. First, interferences must be more likely to be declared between patents whose inventors have more knowledge inputs in common — that is g must be monotonically increasing in A_i. This assumption follows from the logic that new ideas result from combinations of existing ideas, and therefore inventors who make the same discovery must share the same knowledge inputs.

Second, the observable variables used to create matched-control pairs must fully account for other factors besides localized knowledge spillovers, which may drive interferences to be geographically localized. That is, n_i must be independent of $dist_i$.

Under these two assumptions, if the probability that a pair of patents is involved in an interference is decreasing in the distance between the inventors of that patent, conditional on other factors X_i, this is evidence of localized knowledge spillovers. Following the prior literature, our general approach is to account for X_i by selecting control patent pairs matched on technology classification and application date.

First, we operationalize this test non-parametrically using distance-based methods following Duranton and Overman (2005) and Murata et al. (2014). We compute distances between pairs of interfering inventors to overcome scale and border problems in using spatial units (e.g., counties or metropolitan areas) of arbitrary size. Then we compare the kernel density of the distances between interfering patents to a counterfactual density of distances between control patents, which is discussed in more detail in Section 3.5.

Second, we use a linear probability model which compares the probability of interference between patent pairs above and below different co-location thresholds. The linear probability models we estimate can be motivated by assuming particular functional forms for the relationships in equations (1) and (2). Specifically, assume that the impact on the probability of interference depends on shared knowledge, observable characteristics, and unobservables as $P(int_i) = \beta_1 A_i + \mathbf{X}_i \beta_X + \eta_i$. Further, assume shared knowledge inputs depend on distance and unobserved factors as $A_i = f(distance_i) + \epsilon_i$. Combining equations yields

$$P(int_i) = \beta_1 f(distance_i) + \mathbf{X}_i \beta_X + \beta_1 \epsilon_i + \eta_i. \tag{3}$$

Equation 3 illustrates that the effect of distance on interference based on a naïve comparison between interfering pairs and other arbitrarily-chosen patent pairs will be biased if other (omitted) factors η_i that influence the probability of interference are correlated with the distance between the pair of inventors. Ideally, the matched-control estimator can still identify β_1, the effect of distance on interference. To see this, take expectations of equation 3 and note that by assumption $E(\mathbf{X}_i\beta_X|IC) = 0$. That is, conditioned on the sample of interfering and control pairs IC, the expected effect of other factors on interference is zero. However, to the extent that matched controls do *not* satisfy this condition, then estimates of $\hat{\beta}_1$ will still be biased. In other words, perhaps the matched controls are "not similar enough," and therefore, the identifying assumption is not satisfied.

3.3 Matched control pairs

Next, we describe the method we use to select matched control patents and to construct control pairs. We further discuss their validity as counterfactuals.

Since some interference cases involve more than two parties, our database of 1,329 interference cases involves 1,401 interfering pairs of inventing parties.[20] We select a set of control patents that are similar to the invention described by each party's patent(s) or application(s) declared in interference. First, we require a control patent to share at least one 3- or 6-digit technology classification with a party's patent(s) or application(s). Second, control patents must have an application date within 180 days of the application date of an interfering application. We are able to obtain a set of suitable control patents for nearly every interfering inventor pair—only 24 pairs lack suitable 3-digit controls and 32 pairs lack suitable 6-digit controls.

Next, each of these control patents is now eligible to form a *control pair*. A control pair matches a control patent to an interfering patent or application. Control patents are matched to the opposing party's interfering application(s). This matching structure is identical to Jaffe et al. (1993). In their application, *cited* patents are used to identify control patents. Control patents are then matched to *citing* patents. In our sample, the interfering application(s) of the first party are used to identify control patents. Control patents are then matched to the interfering application(s) of the second party. The pool of control patents is large. Conditioned on finding a control, the average interfering inventor pair is associated with 706 control patents matched on the 6-digit technology class and 6,457 controls matched

[20]Twenty-four cases involve three parties, three cases involve four parties, and one case involves five independent parties.

on the 3-digit technology class.

As described above, a concern is that control pairs of patents may imperfectly capture unobservable factors. This is the subject of the analysis by Thompson and Fox-Kean (2005), who show that geographic matching results following Jaffe et al. (1993) are sensitive to matching on technological classification. In particular, conditioning matched control patents on (6-digit) technology *subclass*, the finest detail available in the U.S. patent classification, rather than the 3-digit technology class in Jaffe et al. (1993), shows little localization of citations compared with control pairs. Henderson et al. (2005) note considerable uncertainty in the "right" way to select matched controls, whether at the 3-digit class level (with about 450 classifications) or the 6-digit subclass level (with about 150,000 subclasses). While 6-digit technology classifications are quite detailed (Thompson and Fox-Kean (2005) note that the 3-digit class "231–Whips and whip apparatus" contains 7 distinct subclasses), Henderson et al. (2005) counter that there is little evidence suggesting that selecting on 6-digit subclasses achieves "closer" technologically matched controls.

Whether or not matched control pairs can adequately control for unobserved factors (other than knowledge spillovers) affecting the geography of invention is of course a key question for the credibility of our results. A useful set of results comes from Murata et al. (2014). First, the degree of localization of citations appears to be considerably understated by the geographic matching tests of Jaffe et al. (1993) and Thompson and Fox-Kean (2005). Murata et al. (2014) show that a *distance-based* test shows localization of patent citations, even when selecting matched controls of 6-digit subclasses. The reason that geographic-matching tests understate the localization of citations is that aggregating inventor locations to metropolitan areas and states introduces border and scale problems. Border problems arise because counties, the basis of metropolitan areas, may split clusters of inventors. This tends to bias downwards measures of localization, since inventors split by a county boundary will not be localized by a commuting zone measure. Scale problems arise because aggregating the data to metropolitan areas or commuting zones allows only analysis at one spatial scale. Further, since commuting zones vary from very small to very large in area, analyses based on commuting zones mix different spatial scales. Thus, the sensitivity of the results on the localization of citations appear to be dwarfed by the (opposite) bias introduced by spatial aggregation. More formally, Murata et al. (2014) perform a sensitivity analysis following Rosenbaum (2002) to bound the degree of bias in the presence of unobserved factors affecting localization. They find that biases from such unobserved factors would have to be extremely large before reversing the conclusion that patent citations are geographically localized.

Aside from concerns about imperfect controls, one might also be concerned that as inventions become more technologically similar (in "idea space"), that other sources of agglomeration economies—say, labor pooling—might also become stronger. If this were true, a hypothetically perfect control patent would be unable to distinguish localized knowledge spillovers from other factors.[21] For example, if there are more potential employees who know about the details of reprogramming yeast cell reproduction in San Francisco, then inventors who want to reprogram yeast cells and are concerned about insuring against the idiosyncratic risk of losing an employee will be more likely to locate in San Francisco. On the other hand, localized knowledge spillovers may still be identified using a matched-control approach if the "transportation cost" of ideas exceeds that of other factors.

To see this, consider a firm's location choice. It may benefit from two types of agglomeration economies—knowledge spillovers and labor pooling. Benefits from knowledge spillovers are realized if it co-locates with another firm. The idea here is that knowledge spillovers are more likely within a small geographic area like a neighborhood or office park, because of lower communication costs, a greater probability of chance meetings, and a higher likelihood of social relationships. However, these benefits decline swiftly with distance—chance meetings, or casual conversations with social ties are much less likely even at modest distances (Allen, 1984; Arzaghi and Henderson, 2008). Firms might also benefit from labor pooling, but the geographic scope of these benefits is larger than that of benefits from localized knowledge spillovers. For example, they may be able to realize benefits from labor pooling when they are located in neighboring counties within the same commuting zone. This structure follows from the idea that the "transportation cost" of ideas—particularly tacit knowledge—is higher compared with the cost of transporting workers, i.e., commuting, or the cost of transporting other factors (some evidence in favor of this hypothesis is in Rosenthal and Strange, 2001). Workers may be indifferent between many potential commutes, some quite far apart, within a commuting zone. Thus, firms need not co-locate within the same office park to take advantage of labor pooling benefits.

In this environment, where do firms locate? Two firms will co-locate if, in equilibrium, the productivity gains from one or both sources of agglomeration economies exceed higher production or congestion costs caused by competition for the same fixed factors, e.g., land or space. Two firms that benefit from *both* localized knowledge spillovers and labor pooling will be more likely to co-locate compared with two firms that *only* benefit from labor pooling, but not from localized knowledge spillovers. Thus, even if other sources of agglomeration

[21] We are grateful to an anonymous referee for highlighting this issue.

economies become stronger as inventions become more technologically similar, we may still identify localized knowledge spillovers if the cost of accessing knowledge still exceeds that of other factors.

3.4 Illustrative example

We briefly discuss one interference case to illustrate how interferences reflect shared knowledge inputs and how the matched control strategy works in practice. Interference number 103435 involved competing claims of invention of an intraocular lens—a lens that is implanted in the eye during procedures such as cataract removal. The claims at issue involved the method by which haptics, side struts which hold the implanted lens in place inside the eye, are attached to the optic (lens). Both parties found that exposing the haptics to corona discharge (a plasma curtain that is created when air around a conductor gets ionized) increases the strength of the bond between the haptic and the lens.

The junior party included 3 employees (Richard Christ, David Fencil and Patricia Knight) in the R&D Department at Allergan Inc., a pharmaceutical company located in Irvine, California.[22] The senior party was Larry Blake, inventor and owner of the small company Pharmacia Iovision, Inc., which was also located in Irvine. The Board determined that Blake had both conceived the invention and reduced it to practice sometime in August 1987, while the Christ team had conceived the invention sometime in 1985 and reduced the invention to practice in December 1987.[23]

A few aspects of the case are relevant for our analysis. These details are described in the Board's decision, which was based on evidence presented during the case—primarily lab notebooks and eyewitness testimony.

First, both parties appear to have had common knowledge inputs that led to the common discovery. They were aware of the problem of attaching haptics to lenses and knew that corona discharge had been "used for surface treatment of plastics and rubber to improve adhesion."

Second, the geographic proximity of the inventors facilitated knowledge flows. Most relevantly, Blake had been a short-term consultant at Allergen between June 1986 and April 1987. His consulting work at Allergan related to silicone lenses. Notably, he was *not* working

[22]Allergan initially had an R&D focus on eye care therapies and its flagship product is Botox. It purchased American Medical Optics, located in Irvine in 1986. It was one of the largest intraocular lens producers at the time.

[23]Priority was awarded to Blake as Christ "failed to show diligence between the period just prior to August 1987 and the reduction to practice date established for the Junior party of December 31, 1987."

with the Christ team nor in a similar research area. Nonetheless, the decision cites testimony from Allergen employees stating that Larry Blake would have been aware of the project through his physical presence at Allergen: "...there were no secrets in the R&D Department of Allergan during the mid-1986 time frame and everyone shared resources." One employee testified that the "plasma and corona work was being conducted in the Technology and Ventures laboratory" and Blake was seen at least once in the building and "anyone could go in and out of the building freely."

Both firms were located in Irvine, and the individual inventors lived nearby: Blake, Christ, and Knight lived in surrounding Orange County communities (in Cota de Caza, Laguna Beach, Laguna Niguel, respectively), while Fencil lived in Goleta, in Santa Barbara County. (These places of residence are the locations recorded in their corresponding patents and applications.)

Third, examining the control patents matched to the inventions in interference suggests that the control inventions are quite similar to those in the dispute, but the inventors are geographically distant. For example, one control patent matched to the Christ application was filed by 3 inventors all living in the Minneapolis–St. Paul metropolitan area and working at Minnesota Mining and Manufacturing Company (3M) in St. Paul, Minnesota. The invention appears quite similar to the interfering patents; it is for an intraocular lens and discloses a method for fixing a haptic having an anchoring filament to a lens. Another control patent matched to the Blake application was filed by an inventor living in Arlington, Texas, and also concerned attaching a haptic to a lens for an intraocular lens (titled "Method of Attaching a Haptic to an Optic of an Intraocular Lens").

3.5 Inference and simulation of counterfactual distances

Control patents and control pairs in hand, we turn to estimating the presence of localized knowledge spillovers following the approaches of Duranton and Overman (2005) and Murata et al. (2014). Here, our goal is to compare the distribution of distances between pairs of interfering inventors to a counterfactual distribution of pairwise distances using our sample of control pairs. In this way, we compare the geographic distribution of interferences with the distribution that would be expected to occur randomly conditioned on the geography of invention represented by the controls. As described above, a flexible distance-based method overcomes border and scale problems as it uses the distance between inventors' places of residence rather than arbitrary spatial units (e.g., metropolitan areas or commuting zones).

First, we compute the density of distances between pairs of interfering inventors. (Recall that our main analysis uses the minimum distance between inventing parties in cases with multiple co-inventors per party, and we try the median or first-named inventor distance with similar results.) For each interfering pair indexed by $i = 1 \ldots I$, we compute the geodesic distance d_i. The estimator of the density of pairwise distances (the kernel density) at any distance d is then

$$\hat{K}(d) = \frac{1}{2h} \sum_{i=1}^{I} f\left(\frac{d - d_i}{h}\right),$$

where h is the bandwidth and f is the kernel function. All densities are computed using a Gaussian kernel and with the bandwidth set as in Silverman (1986).[24]

Second, we use Monte Carlo simulations to construct a counterfactual distribution of pairwise distances using our sample of control pairs and test for significant departures from our counterfactual benchmarks and estimate confidence intervals. Our simulation strategy closely follows that of Murata et al. (2014), which is itself is an extension of the approach of Jaffe et al. (1993) to compare pairs of citing patents to counterfactual control patents. Specifically, Murata et al. (2014) expand on Jaffe et al. (1993) in two important ways: (1) they make repeated random draws from a set of potential control patents, generating counterfactual confidence bands, and (2) they compare the entire distributions of citation and counterfactual distances, giving a more complete picture of the pattern of localization. We extend their approach by focusing on pairs of patents in an interference case, rather than cited-citing pairs, and by comparing interference and control pairs across a number of other statistics in addition to geographic distance.

We perform 1,000 Monte Carlo simulations. In each simulation, we randomly sample a matched control pair for each interfering pair from the set of permissible control pairs. With this random draw of control pairs, we estimate the kernel density for the distribution of pairwise distance. Then, we repeat this exercise by drawing a new set of control pairs.

After performing these simulations, we evaluate significant departures from our counterfactual benchmarks by constructing local confidence intervals. We consider all distances between 0 and about 3,500 km, which is the median distance between all pairs of control and interfering inventors.[25] Comparison of the distance densities above 3,500 km is redundant because densities must sum to one over the entire range of distances. In other words, if the density of interfering pairs is *lower* compared with control pairs at distances greater than

[24] To deal with boundary problems at zero due to the non-negative domain of distances, we use the reflection method of Silverman (1986), following Duranton and Overman (2005).

[25] The median distance is 3,367 km for the sample of interfering and 6-digit control pairs and 3,844 km for the sample of interfering and 3-digit control pairs.

3,500 km, it will be *higher* compared with control pairs at distances less than 3,500 km. Thus, we need only examine distances below the median to infer localization.

We rank the simulated kernel density estimates at 100 evenly spaced distances, and select the 50th ranked simulated kernel density at each distance to construct the lower 5% confidence band and the 950th-ranked to construct the upper 5% confidence band. These are the lower 5% and upper 5% confidence levels denoted by $\underline{K}(d)$ and $\overline{K}(d)$, respectively. When $\hat{K}(d) > \overline{K}(d)$, we infer that interferences exhibit localization at distance d at a 5% confidence level. That is, interfering pairs are more likely to be d km apart compared with control pairs. Graphically, if the density of pairwise distances for interfering pairs is above the upper confidence interval (e.g., at 1,000 km), we would say that interfering inventor pairs are localized at 1,000 km at a 5% level. (Similarly, when $\hat{K}(d) < \underline{K}(d)$, we infer that interferences exhibit dispersion at distance d at a 5% confidence level.)

These inferences are *local* in the sense that they only allow us to make local statements (i.e., at a given distance d) about the relationship between the interfering and counterfactual distributions. However, even if interfering pairs were distributed randomly with respect to control pairs, there is a high probability that interfering pairs will exhibit localization at some distance, since by construction there is a 5 percent probability for each particular distance that a random draw of control pairs will exhibit localization.

To detect localized knowledge spillovers, we define global confidence bands. In this way we can make statements about the overall location patterns of interfering inventors. We search for identical upper and lower confidence intervals such that when we consider them for all distances between 0 and 3,500 km, only 5% of our randomly generated simulated kernel densities hit them. That is, we define a global upper confidence band $\overline{\overline{K}}(d)$ as the band that is hit by 5% of our simulations between 0 and 3,500 km. Interfering pairs are considered globally localized (at a 5% confidence level) when $\hat{K}(d) > \overline{\overline{K}}(d)$ for *at least* one $d \in [0, 3500]$. Naturally, the global confidence bands are wider compared with the local confidence intervals. In general, we end up selecting approximately the 10th-ranked simulated kernel density to construct the lower 5% global confidence level and the 990th-ranked simulated kernel density to construct the upper 5% global confidence level.[26]

Graphically, interferences are globally localized if the density of pairwise distances for interfering pairs is above the upper global confidence band for at least one distance d up to the sample median. (Again, by construction, only 5% of simulated densities will have this property.) Conversely, interferences are globally dispersed if the density lies below the lower

[26]See Duranton and Overman (2005) for more discussion.

confidence band and never lies above the upper confidence band.

4 Results

4.1 Shared codified knowledge inputs

Interfering inventors tend to share common knowledge inputs. In this section, we provide evidence of shared knowledge inputs among interfering and control pairs by analyzing two commonly-used bibliometric measures of patents and applications – shared technology classifications and shared backward citations. These results connect interferences to more familiar measures of shared knowledge. They also motivate our later analysis conditioning on additional controls.

Interfering patents and applications have a similar number of total technology classifications and sub-classifications compared with control patents. Recall that patents may be classified with multiple technology classifications and sub-classifications when their subject matter overlaps with multiple technology areas. Table 3 displays summary statistics for interfering applications or patents and associated controls. Panel A compares the total number of citations and technology classifications between interfering and control patents and applications. The first column shows means for interfering patents and applications. The subsequent columns compare this statistic to corresponding statistics for control patents. We report means and lower 5% and upper 5% confidence levels from our Monte Carlo simulations for eligible 3- and 6-digit control patents.

Patents can cite other patents as prior art, so that the number and type of "backwards citations" can then be derived from the citations made by a patent. Interfering patents and applications have a similar number of total backwards citations, or citations to previously issued patents as "prior art," compared with control patents. On average, interfering patents and applications contain 11.1 backwards citations. Control patents have a similar number of citations, 11.8, and the confidence interval for controls includes the mean for interfering patents and applications.

Interfering patents and applications tend to have slightly fewer U.S. patent classifications and subclassifications compared with control patents matched on 3-digit technology classes and 6-digit subclasses. These differences are statistically significant, but small in magnitude, 1.93 patent classifications for interfering versus 2.28 and 2.24 for control patents. For subclassifications, the means are 4.93 for interfering patents and applications and 5.95 and 7.06 for the controls. These small differences are entirely accounted for by the rate of interfering

	Interfering	3-digit controls μ	C.I.	6-digit controls μ	C.I.
A. Means for applications/patents					
Backwards citations	11.1	11.8	(11.1, 12.7)	11.8	(11.0, 12.6)
USPC classes	1.93	2.28	(2.23, 2.34)	2.24	(2.18, 2.29)
USPC subclasses	4.93	5.95	(5.71, 6.19)	7.06	(6.75, 7.41)
B. Means for pairs of applications/patents					
Backwards citations shared	3.53	0.03	(0.01, 0.06)	0.33	(0.20, 0.62)
USPC classes shared	1.25	0.85	(0.82, 0.87)	1.02	(0.99, 1.04)
USPC subclasses shared	1.50	0.10	(0.08, 0.13)	0.68	(0.64, 0.71)

Table 3: Interfering inventors share codified knowledge inputs

This table compares means of interfering patent pairs to simulated means, 5th- and 95th-percentile estimates for control pairs. A control pair includes 1 interfering application and 1 issued control patent that share a technology class and application date. Simulated CIs based on 1,000 random draws from eligible control pairs. The sample is cases with decisions on priority and concessions.

applications—about a quarter—that are never passed for issue. For these applications we observe only one primary class and subclass in the PAIR database.[27]

Overall, interfering patents and control patents are similar in terms of the total number of backwards citations to prior art and the total number of technology classes and subclasses. These results support the validity of the matched control strategy, since these factors were not used in selecting matched controls.

However, interfering pairs differ in terms of the number of *shared* backwards citations and technology classes and subclasses. This is an indicator that interfering pairs overlap more in terms of shared knowledge inputs, as measured by these common bibliographic measures of patents. Panel B reports statistics for the number of backwards citations and technology classifications shared within interfering pairs and control pairs. Despite the overall similarity in the total number of backwards citations and technology classes and subclasses, interfering pairs share substantially more backwards citations and technology classes and subclasses. Interfering pairs share 3.5 backwards citations, compared with 3-digit control pairs that share 0.03 backwards citations and 6-digit control pairs that share 0.33 backwards citations. Interfering pairs also share more U.S. patent classifications. Interfering pairs tend to share

[27] If we compare only primary classes and subclasses, both interfering and control pairs have on average 1 primary class and subclass, by construction. The actual number of primary classes and subclasses for interfering pairs is slightly more than 1 owing to a small number of interfering parties with claims spread across multiple patents or applications.

1.25 classes and 1.50 subclasses. This is more than the number of classes and subclasses shared by both 3- and 6-digit control pairs. As noted earlier, for about a quarter of interfering applications we observe only one primary class and subclass. If we focus on primary classes and subclasses, 62 percent of interfering classes share a primary class and 25 percent share a primary subclass, compared with 48 percent and 7 percent of 6-digit control pairs and 35 and 1 percent of 3-digit control pairs, respectively. Thus, the result that interfering pairs share more classes and subclasses is robust to conditioning on primary classes and subclasses, for which we have complete data. Overall, these results suggest that interfering inventing teams may share more *codified* knowledge inputs compared with control pairs.[28]

4.2 Geographic proximity

Interfering pairs of inventors are more geographically localized compared with control inventor pairs not linked by an interference. Figure 1 compares the density of geographic distance between pairs of interfering inventors with that of similar, non-interfering control pairs, which controls for the overall distribution of invention in the same technological classes and time periods as the interfering inventors. We use the geodesic distance between the places of residence of inventors from opposite teams.[29] The black line shows the estimated kernel density of geographic distance between all interfering pairs for which we were able to find suitable controls. The dotted red lines show the upper 5% and lower 5% local confidence levels. The dashed blue lines show the upper 5% and lower 5% global confidence bands. (The construction of these bands was described in detail in Section 3.5.)

The estimated kernel density of pairwise distances for interfering pairs exceeds that upper 5% local confidence level for a range of distances up to about 1,000 km (620 miles). For example, since $\hat{K}(d) > \overline{K}(d)$ for $d = 1000$, we infer that interfering inventor pairs are more likely (at a 5% confidence level) to be 1,000 km apart compared with control inventor pairs.

Recall that even if interfering pairs were distributed randomly with respect to the set of control pairs, there is a high probability that interfering pairs would exhibit localization at some distance, since by construction there is a 5% probability for each particular distance that a random draw of control pairs exhibit localization. We therefore use global confidence bands to conduct inference on the localization of interfering inventors. Recall that our global

[28]In Sections 4.2 and 4.3, we also check that our results are robust to conditioning on these bibliometric measures of similarity.

[29]For example, for a pair of patents with two and three inventors respectively, we compute distances for all 6 possible inventor pairs and use the minimum distance. We obtain similar results using the median distance or the distance between the first-named inventors.

A. Interfering and 3-digit control pairs

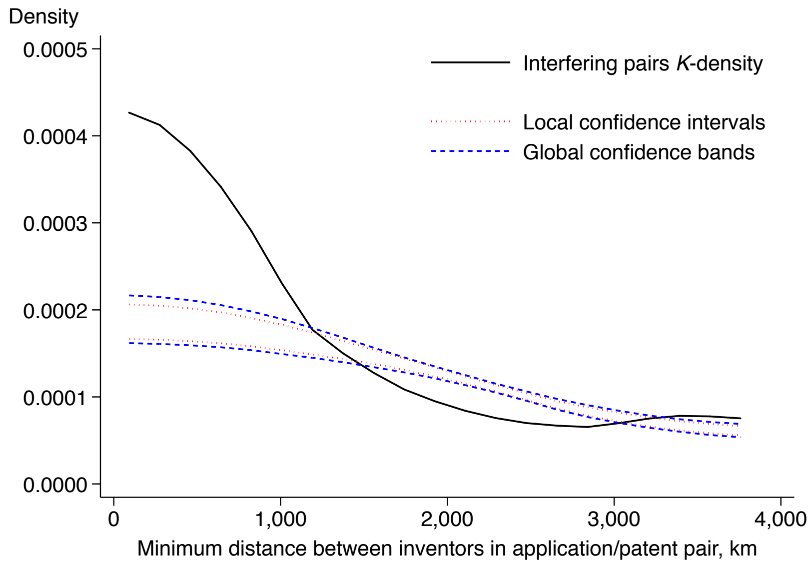

B. Interfering and 6-digit control pairs

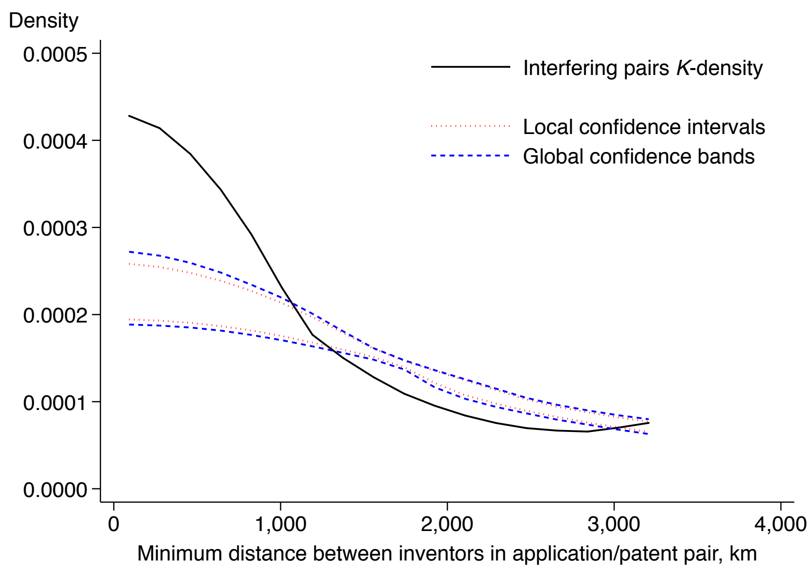

Figure 1: Interfering pairs are more localized compared with control pairs

These graphs compare the estimated kernel densities of geographic distances between pairs of interfering inventors to the distribution for similar non-interfering control pairs. We use the minimum geodesic distance between the places of residence of inventors from opposite teams. The black line shows the estimated kernel density function for all interfering pairs for which we were able to find suitable controls. The dotted red lines show the upper and lower 5% local confidence intervals of the kernel density for non-interfering control pairs. The dotted blue lines show the upper and lower 5% global confidence bands of the kernel density for non-interfering control pairs. The estimated densities end at the median distance between inventing pairs. Interfering pairs are considered localized if the estimated kernel density of interfering inventors is above the upper global confidence band for at least one distance d up to the sample median.

confidence bands are defined as identical upper (or lower) confidence levels such that only 5% of our randomly-generated simulated kernel densities hit them. Interfering pairs are considered localized (at a 5% confidence level) if $\hat{K}(d) > \overline{\overline{K}}(d)$ *at least once* at any distance up to the sample median. Thus, Figure 1 shows that interfering pairs of inventors are localized compared with control pairs of inventors.

These results are robust to matching on 3-digit technology class or 6-digit technology subclass. In Panel A, we show results using 3-digit control pairs; in Panel B, we show results using 6-digit control pairs. Thus, while the observed distribution of pairwise distances for interfering pairs is the same in both panels, the counterfactual distribution is more geographically concentrated for the 6-digit control pairs shown in Panel B. This is consistent with Thompson and Fox-Kean (2005), who show that the localization of patent citations is sensitive to the selection of controls. However, unlike Thompson and Fox-Kean (2005), we find that for either counterfactual, interferences are indeed significantly geographically localized. This echoes Murata et al. (2014), who show that the localization of patent citations is robust to matching on 6-digit subclasses when using a distance-based test, as we are doing here. Overall, this evidence is consistent with geographic proximity facilitating the sharing of common knowledge inputs.

Our results on geographic localization are robust to the choice of proximity measure and conditioning on decision type. Table 4 presents results separately for priority decisions (excluding concessions) and three different measures of inventor proximity. Panel A compares the average distance between inventors in an interfering pair, to that of control pairs, with simulated confidence intervals. On average, interfering pairs in priority decisions and concessions are 3,451 km apart, compared with 4,778 km separating 3-digit control pairs of inventors and 4,425 km separating 6-digit control pairs of inventors. Average distances separating interfering and control pairs of inventors are similar when we focus on priority decisions only.

Though interfering inventors are closer together on average, the average pairwise distances may obscure the relevant range of geographic proximity for localized knowledge spillovers. Panels B and C present results that examine inventor co-location or "geographic matching" as in the main tests performed by Jaffe et al. (1993). Panel B shows the share of inventor pairs that report the same place (e.g., a town or city) of residence. Nearly 3 percent of interfering inventors share a place of residence, compared with 1 and 2 percent of 3- and 6-digit control pairs, respectively. (The difference compared with 3-digit control pairs is statistically significant.)

Table 4: Interfering pairs are co-located compared with control pairs

	Type of case	
	Priority decisions and concessions	Priority decisions only
A. Average distance between pair of inventors in kilometers		
Interfering pairs	3,451	3,603
3-digit control pairs	4,778	4,714
	(4,563, 4,985)	(4,321, 5,122)
6-digit control pairs	4,425	4,282
	(4,219, 4,623)	(3,890, 4,688)
B. Share of inventor pairs with same place, town or city of residence		
Interfering pairs	2.7%	2.8%
3-digit control pairs	0.8%	0.7%
	(0.4%, 1.4%)	(0.0%, 1.8%)
6-digit control pairs	2.0%	2.0%
	(1.3%, 2.9%)	(0.7%, 3.6%)
C. Share of inventor pairs with places of residence within 161km or 100mi		
Interfering pairs	13.8%	11.7%
3-digit control pairs	5.2%	5.2%
	(4.0%, 6.4%)	(3.1%, 7.6%)
6-digit control pairs	8.2%	8.2%
	(6.7%, 9.7%)	(5.4%, 11.2%)

This table reports statistics for interfering and control pairs of inventing teams. Panel A reports the average minimum distance between places of residences of inventing teams. Panel B reports the share of inventor pairs that share a place of residence. Panel C reports the share of inventor pairs where the minimum distance between places of residence is within 161 km. Simulated upper and lower 5% confidence intervals for control pairs shown in parentheses.

Panel C shows the share of inventor pairs that report places of residence within 161 kilometers or 100 miles of each other. By this measure we intend to capture localized interactions via social ties, workplace relationships, or random meetings. The 100-mile cutoff is comparable to a metropolitan area (used in Jaffe et al. 1993) or a commuting zone (as in Autor and Dorn, 2013). As described in Section 3.3, we prefer to use this distance-based cutoff as opposed to explicitly calculating commuting zone matches, because it avoids border and scale problems. Thus, this test is similar to the original matching-rate tests at the metropolitan area level reported by Jaffe et al. (1993), except that we leverage the micro-geography of inventor location compared with co-locating within the same county or set of counties.

Between 12 and 14 percent of interfering inventor pairs are less than 161 km (100 miles) apart, compared with 5 to 8 percent of control inventor pairs. These differences are significant. In sum, interfering inventor pairs are 1.4 to 4.0 times more likely to locate in the same place or region compared with control inventor pairs.

These results are also robust to conditioning on bibliographic measures of pairwise similarity. To show this, we estimate the following analog of equation 3:

$$Pr(int_{i[g]}) = \mu_g + \beta_1 f(distance_i) + \mathbf{X}_i \beta_X + \varepsilon_i. \tag{4}$$

The dependent variable is a dummy indicating whether an inventor pair is involved in an interference. The regressions are run on the entire sample of interfering pairs and matched controls. All specifications include a fixed effect for each "pair group," defined by an interfering inventor pair and all associated matched control pairs. Thus, the effect of co-location on interference is identified by variation within the group of an interfering pair and its associated control pairs. We cluster standard errors at the pair-group level.

We use several measures for measuring the proximity of inventor pairs, following the results reported in Table 4: (i) the logarithm of distance between pairs of inventor teams; (ii) an indicator equal to 1 when the minimum distance between an inventor pair is within 1.6 km (1 mile); and (iii) an indicator equal to 1 when the minimum distance between the locations of residence reported by an interfering or control pair is within 161 km (100 miles).

Estimation results are reported in Table 5. For each combination of sample (interfering pairs plus 3- or 6-digit control pairs) and proximity measure we show two specifications. The first, in columns (1) and (3), reports the univariate relationship between inventor pair distance and interference. The second, in columns (2) and (4), reports conditional estimates, controlling for three measures of bibliometric similarity: (i) the number of technological classes shared by the pair of inventors, (ii) the number of shared subclasses by the pair of inventors, and (iii) the number of backwards citations to prior art shared by the pair of inventors. Recall that we only match controls based on a single shared (3-digit) class or (6-digit) subclass. Including controls for additional shared classes and subclasses conditions our estimate of the effect of distance on even more similar inventions—that is, those pairs that share similar numbers of classifications. Similarly, including controls for the number of backwards citations conditions our estimates on pairs of inventions that probably, by evidence of citing the same prior art, are technologically quite similar.

Closer proximity leads to a higher likelihood of interference, as suggested by the estimates in Panel A. In the sample of 3-digit matched control pairs and interferences, a doubling of the minimum distance between a pair of inventor teams decreases the probability of interference by 0.008 percent (column 1), compared with the mean probability of interference of 0.04 percent. In the sample of 6-digit matched control pairs and interferences, a doubling of the minimum distance between a pair of inventor teams decreases the probability of interference by 0.04 percent (column 3), compared with the mean probability of interference of 0.4 percent. Similarly, in panels B and C, co-location in the same place of residence or with places of residence within 161 km increases the likelihood of interference by 0.06 percent and 0.04 percent, respectively, compared with 3-digit control pairs. These effects are precisely estimated.

Columns (2) and (4) add bibliometric controls measuring patent-pair similarity. We control for the number of shared technology (3-digit) classes, the number of shared (6-digit) subclasses, and the number of shared backwards citations to prior art. Most of the estimates of the effect of proximity on interference are robust to the inclusion of these controls and are precisely estimated. The effect of distance attenuates by about one-quarter when controlling for our measures of patent-pair similarity (panel A). Similarly, the effect of co-location within 161 km attenuates by about one-quarter when including controls (panel C). These results suggest that, at least to the extent that we are able to introduce additional controls for factors contributing to the geography of invention, our results do not seem to be driven by imperfectly matched controls.

The effect of co-location in the same city, town or place (panel B) attenuates more when controlling for shared classes and citations but is estimated imprecisely. One explanation may be that measuring geographic proximity may be imprecisely measured using place of *residence*. One might want to use place of *work* or other places where inventors might be exposed to localized knowledge spillovers. However we observe only place of residence.

Table 5: Robustness of co-location effect on interference

	3-digit controls			6-digit controls		
	μ [σ]	(1)	(2)	μ [σ]	(3)	(4)
A. Distance between pair of inventors						
Log distance	8.386	-0.008^c	-0.006^c	8.171	-0.041^c	-0.033^c
	[1.752]	(0.001)	(0.001)	[1.925]	(0.007)	(0.007)
No. shared classes	0.812		-0.012^c	1.141		0.008
	[0.662]		(0.003)	[0.844]		(0.022)
No. shared subclasses	0.130		0.166^c	0.788		0.286^c
	[0.502]		(0.015)	[1.078]		(0.033)
No. shared citations	0.011		0.099^a	0.056		0.130
	[0.631]		(0.057)	[1.676]		(0.096)
B. Co-location with same place, town, or city of residence						
1(Co-located in same place)	0.009	0.060^c	0.014	0.014	0.115	-0.042
	[0.092]	(0.020)	(0.023)	[0.118]	(0.112)	(0.126)
No. shared classes	0.812		-0.012^c	1.141		0.007
	[0.662]		(0.003)	[0.844]		(0.022)
No. shared subclasses	0.130		0.167^c	0.788		0.289^c
	[0.502]		(0.015)	[1.078]		(0.033)
No. shared citations	0.011		0.100^a	0.056		0.131
	[0.631]		(0.057)	[1.676]		(0.097)
C. Co-location with places of residence within 161km (100mi)						
1(Co-located within 161km)	0.057	0.043^c	0.030^c	0.077	0.206^c	0.156^c
	[0.232]	(0.007)	(0.007)	[0.266]	(0.048)	(0.048)
No. shared classes	0.812		-0.012^c	1.141		0.007
	[0.662]		(0.003)	[0.844]		(0.022)
No. shared subclasses	0.130		0.166^c	0.788		0.287^c
	[0.502]		(0.015)	[1.078]		(0.033)
No. shared citations	0.011		0.099^a	0.056		0.131
	[0.631]		(0.057)	[1.676]		(0.096)
Pairs	5,712,342	5,712,342	5,712,342	604,828	604,828	604,828
Pair-groups		831	831		821	821

This table shows estimates from a regression of an indicator for interference on a measure of proximity and controls. The sample includes interfering pairs of inventing teams and 3- or 6-digit controls, as indicated by the column group heading. The dependent variable is an indicator for interference $\times$ 100, with mean 0.04 percent in the sample with 3-digit controls and mean 0.4 percent in the sample with 6-digit controls. Sample means and standard deviations are reported in the columns labeled $\mu[\sigma]$. Panel A reports estimates where the measure of proximity is the minimum log distance between the pair of inventing teams. Panel B reports estimates where the measure of proximity is an indicator for whether or not the pair of inventing teams is co-located in the same place, town, or city of residence. Panel C reports estimates where the measure of proximity is an indicator for whether or not the pair of inventing teams is co-located with places of residences within 161km or 100mi. Standard errors, clustered on pair groups, are reported in parentheses. ^{a}—$p < 0.10$, ^{b}—$p < 0.05$, ^{c}—$p < 0.01$. The number of observations reported in the bottom panel applies to all regressions in that column.

Relatedly, we also note that place of residence may refer to a very small spatial scale: a suburban town or village, or even an unincorporated community. At this spatial scale, other factors such as local amenities may contribute more to the sorting of inventors to places compared with localized knowledge spillovers. Finally, as already discussed, the co-location results are subject to bias related to spatial aggregation. For these reasons, we place more weight on the distance-based results in panel A or the co-location results at the spatial scale of a commuting zone in panel C.

Finally, we note that geography appears to contribute to interference over and above the number of shared citations. If all knowledge spillovers were recorded in citations, then controlling for citations should remove effect of distance on interferences. Instead, our results suggest that proximity matters for localized spillovers of knowledge not recorded in citations, or tacit knowledge. We further explore this idea in the next section.

4.3 Comparison with citation-linked inventors

Are spillovers of *tacit* knowledge more localized than other forms of knowledge that are more easily codified? Arzaghi and Henderson's (2008) results suggest that the external benefits to advertising agencies in Manhattan attenuate quickly over space—they dissipate in as little as 750 meters. To the extent that interferences can capture spillovers of both tacit and codified knowledge, their localization could provide evidence that tacit knowledge spillovers require even closer proximity. To test this hypothesis, we compare the observed distribution of geographic proximity between interfering inventors to a particularly strong counterfactual: pairs of control patents and interfering patents linked by citation. Thus, our *control pairs* in this exercise are the *treated pairs* in Jaffe et al. (1993). For each interfering pair, we identify potential controls as patents *cited by* one of the interfering parties. Then, we form cited-citing control pairs by matching an interfering application to one of the these cited control patents.

This exercise also suggests that our main result is robust to matching controls on other factors. Here, matched control patents are selected based on citation links, rather than technology class and date of invention as in the main analysis presented in Section 4.2. As citation-linked patents have been used extensively following Jaffe et al. (1993) as evidence themselves of knowledge spillovers, we view these controls as a particularly strong counterfactual. Citation-linked patents seem even more likely to share similar production and demand factors, in addition to shared knowledge as measured by citations.

Figure 2 shows the comparison between interfering inventor pairs and cited-citing control

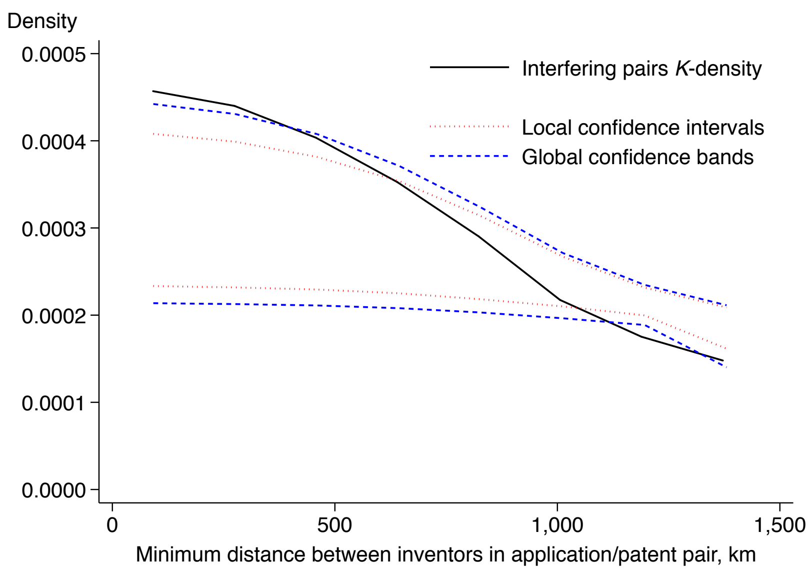

Figure 2: Distribution of geographic distances vs. citation-linked controls

This graph compares the estimated kernel densities of geographic distances between pairs of interfering inventors to the distribution for control pairs that include one interfering patent and one patent cited by the interfering patent. We use the minimum geodesic distance between the places of residence of inventors from opposite teams. The black line shows the estimated kernel density function for all interfering pairs for which we were able to find suitable controls. The dotted red lines show the upper and lower 5% local confidence intervals of the kernel density for non-interfering control pairs. The dotted blue lines show the upper and lower 5% global confidence bands of the kernel density for non-interfering control pairs. Control patents are restricted to backwards citations from one of the interfering patents. Interfering pairs are considered localized if the estimated kernel density of interfering inventors is above the upper global confidence band for at least one distance d up to the sample median.

pairs. The black line again shows the observed density of pairwise distances between interfering inventors. (This is the same density reported in Figure 1, except for the change in sample to interferences with eligible *citation-linked* controls.) The dotted and dashed lines show local confidence intervals and global confidence bands, respectively, for the density of pairwise distances between cited-citing control pairs. Even compared against the geography of citation-linked patents, interfering patents are significantly more localized. A natural interpretation is that forms of input knowledge not captured easily by citations—perhaps tacit knowledge—contribute to the localization of patent interferences over that of patent citations. Thus, this result is consistent with the Jaffe et al. (1993) conjecture that citations are a *lower bound* on the strength of localized knowledge spillovers.

4.4 Localization and co-inventor ties

How is localized knowledge transmitted? Recent work suggests that network ties mediate the relationship between geographic proximity and localized knowledge spillovers. Social and professional ties may be especially important for the transmission of tacit knowledge. For example, Breschi and Lissoni (2009) use past co-inventorship as recorded on U.S. patents as a measure of social ties. They found that controlling for this measure greatly reduced estimates of the effect of geographic localization on citiations. More recently, Head, Li, and Minondo (2015) find that controlling for measures of the network ties among mathematicians halves the estimated effect of geographic distance on citations.

To shed light on the role of social ties in mediating the relationship between localization and interference, we follow Breschi and Lissoni (2009) and consider previous co-inventorship as a proxy for social ties. We use the Lai et al. (2013) database of inventors of U.S. patents to construct measures of ties. We define a network with each inventor represented as a node and connections or edges between any inventors that have been co-inventors on an issued patent. (Unique inventors are identified with name disambiguation, available in the Lai et al. (2013) database. Their algorithm uses not only name similarity but also inventor location, assignee, and technological class information.) The network distance between two inventors is the minimum path distance between them in the network—the number of edges along the shortest path from one inventor node to the other. This network distance is conditioned on co-inventor links up to 5 years *before* the earliest application date in interference. For a pair of patents A and B, we assign the shortest network distance between any inventor on patent A and any inventor on patent B.

An important limitation of this database is that it includes only information about co-inventor ties from issued U.S. patents. For interfering inventors without an issued patent—about one-quarter of inventors—we are unable to include them in this analysis, since they do not appear in the Lai et al. (2013) database. Further, social and professional ties outside of co-inventorship are not observed.

Interfering inventor pairs are more than twice as likely to be connected by previous co-inventor ties. Over two-fifths of interfering inventor pairs are connected by previous co-inventor ties, compared with about one-fifth of control inventor pairs[30]. This result confirms the findings of earlier studies suggesting that social and professional ties are important for facilitating knowledge flows.

However, we find only modest evidence that the localization of interfering inventors is accounted for by previous co-inventor ties. Table 6 shows the result of interference-rate regressions conditioned on co-inventor ties. Compared with Table 5, there are two more controls: (i) an indicator if two inventing teams *are not* connected by a previous co-inventor tie ($1(Network|observed) = 0$), and (ii) an indicator if two inventing teams *are* connected by a previous co-inventor tie ($1(Network|observed) = 1$). Unobserved network ties (for inventors outside the Lai et al. (2013) database) are the omitted category. Column (1) repeats estimates displayed in Table 5, column (3) for comparison. Columns (2)-(3) condition on previous co-inventor ties. Column (3) adds controls for patent-pair similarity as introduced in Table 5, column (4).

A few patterns stand out. One, there appears to be an interference premium for those connected by previous co-inventor ties compared with those that are not. The estimates in columns (2)-(3) of both panels suggest that pairs connected by previous co-inventor ties are about 0.06–0.07 percent more likely to be in interference compared with pairs not connected by previous co-inventor ties. This gap is robust to measures of geographic proximity and additional controls for patent pair similarity.

Two, pairs where network information is not observed are much more likely to be in interference, by about 0.5–0.6 percent. This pattern is accounted for by the fact that we do not observe network information for inventing parties who do not have an issued patent. Because control patents are selected based on having an issued patent, by construction, only the network information for interfering inventors is censored.

Three, the estimated effect of proximity does not change when we condition on observed co-inventor ties. In contrast to Breschi and Lissoni (2005), we find little evidence that the localization of interfering inventors is mediated via social ties, at least as proxied by

[30]See Appendix A.2 for details.

Table 6: Co-location and co-inventor ties

A. Distance between pair of inventors				
	μ [σ]	(1)	(2)	(3)
Log distance	8.171	-0.041^c	-0.040^c	-0.032^c
	[1.925]	(0.007)	(0.007)	(0.007)
No. shared classes	1.141			0.032
	[0.844]			(0.023)
No. shared subclasses	0.788			0.292^c
	[1.078]			(0.034)
No. shared citations	0.056			0.132
	[1.676]			(0.096)
1(Network|observed)=0	0.318		-0.539^c	-0.634^c
	[0.466]		(0.085)	(0.104)
1(Network|observed)=1	0.326		-0.471^c	-0.585^c
	[0.469]		(0.082)	(0.103)
R^2		0.065	0.066	0.070
Pairs	604,828	604,828	604,828	604,828
Pair-groups		821	821	821
B. Co-location with places of residence within 161km (100mi)				
	μ [σ]	(1)	(2)	(3)
1(Co-located within 161km)	0.077	0.206^c	0.195^c	0.146^c
	[0.266]	(0.048)	(0.048)	(0.049)
No. shared classes	1.141			0.031
	[0.844]			(0.023)
No. shared subclasses	0.788			0.293^c
	[1.078]			(0.034)
No. shared citations	0.056			0.132
	[1.676]			(0.096)
1(Network|observed)=0	0.318		-0.543^c	-0.637^c
	[0.466]		(0.085)	(0.104)
1(Network|observed)=1	0.326		-0.466^c	-0.582^c
	[0.469]		(0.082)	(0.104)
R^2		0.065	0.066	0.070
Pairs	604,828	604,828	604,828	604,828
Pair-groups		821	821	821

This table shows estimates from a regression of an indicator for interference on a measure of proximity and controls. The sample includes interfering pairs of inventing teams and 6-digit controls. The dependent variable is an indicator for interference $\times$ 100, with mean 0.4 percent. Sample means and standard deviations are reported in the column labeled $\mu[\sigma]$. Panel A reports estimates where the measure of proximity is the minimum log distance between the pair of of inventing teams. Panel B reports estimates where the measure of proximity is an indicator for whether or not the pair of inventing teams is co-located with places of residences within 161km or 100mi. Standard errors, clustered on pair groups, are reported in parentheses. ^{a}—$p < 0.10$, ^{b}—$p < 0.05$, ^{c}—$p < 0.01$.

previous co-inventor ties. These regression results suggest that this measure of social ties is independent of geographic localization. Of course, it is worth reiterating that an important limitation is that we do not observe other ties not captured by co-inventorship on U.S. patents. For only about two-thirds of the sample do we observe co-inventor ties, so part of the effect of log distance is identified by the one-third of the sample where we do not observe co-inventor ties.

5 Conclusions

We present new evidence of localized knowledge spillovers using a novel database of patent interferences—instances of simultaneous, identical invention by multiple, independent parties. By evidence of common, identical invention, interfering inventors share common knowledge inputs. Interfering inventor pairs show significant geographic localization compared with the counterfactual of inventor pairs sharing similar invention dates and technology classification. Thus, our results provide verification of the existence of localized knowledge spillovers and are distinct from the literature using patent citations. Interfering inventor pairs are even more localized compared with cited-citing inventor pairs, consistent with the Jaffe et al. (1993) conjecture that citations are a lower bound on the strength of localized knowledge spillovers.

Our results suggest that, in contrast to conventional wisdom about "the death of distance," geographic distance continues to matter, especially for flows of tacit, or difficult to codify, forms of knowledge. These are the types of knowledge flows where the lack of a "paper trail" has hampered the availability of evidence. Interferences therefore provide a unique and useful window into localized knowledge spillovers. In future work, it would be useful to leverage the potential of interferences to measure shared knowledge inputs to investigate other features of the invention process.

6 Works cited

Allen, Thomas J. (1984). *Managing the Flow of Technology: Technology Transfer and the Dissemination of Technological Information Within the R&D Organization.* Cambridge, MA: MIT Press.

Arzaghi, Mohammad, and J. Vernon Henderson (2008). "Networking off Madison Avenue," *The Review of Economic Studies*, 75 (4): 1011–1038.

Audretsch, D.B., and M.P. Feldman (2004). "Knowledge Spillovers and the Geography of Innovation," *Handbook of Regional and Urban Economics*, vol. 4, 2713–2739.

Autor, D.H., and D. Dorn (2013). "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market," *American Economic Review* 103 (5): 1553–1597.

Baruffaldi, Stefano, and Julio Raffo (2017). "The Geography of Duplicated Inventions: Evidence from Patent Citations," *Regional Studies*, 51 (8): 1232–1245.

Bikard, Michaël (2012). "Simultaneous Discoveries as a Research Tool: Method and Promise," working paper.

Bleakley, H., and J. Lin (2012). "Thick-Market Effects and Churning in the Labor Market: Evidence from U.S. Cities," *Journal of Urban Economics*, 72 (2-3): 87–103.

Breschi, Stefano, and Francesco Lissoni (2005). "'Cross-Firm' Inventors and Social Networks: Localized Knowledge Spillovers Revisited," *Annals d'Économie et de Statistique* 79–80: 189–209.

Breschi, Stefano, and Francesco Lissoni (2009). "Mobility of Skilled Workers and Co-Invention Networks: An Anatomy of Localized Knowledge Flows," *Journal of Economic Geography* 9 (4): 439–468.

Buzard, Kristy, Gerald A. Carlino, Robert M. Hunt, Jake K. Carr, and Tony E. Smith (2017). "The Agglomeration of American R&D Labs," *Journal of Urban Economics* 101: 14–26.

Calvert, Ian A. (1980). "An Overview of Interference Practice," *Journal of the Patent Office Society* 62: 290–307.

Calvert, Ian A., and Michael Sofocleous (1982). "Three Years of Interference Statistics," *Journal of the Patent Office Society* 64 (12): 699–707.

Calvert, Ian A., and Michael Sofocleous (1986). "Interference Statistics for Fiscal Years 1983 to 1985," *Journal of the Patent and Trademark Office Society* 68 (8): 385–393.

Calvert, Ian A., and Michael Sofocleous (1989). "Interference Statistics for Fiscal Years 1986 to 1988," *Journal of the Patent and Trademark Office Society* 71: 399-410.

Calvert, Ian A., and Michael Sofocleous (1992). "Interference Statistics for Fiscal Years 1989 to 1991," *Journal of the Patent and Trademark Office Society* 74: 822–826.

Calvert, Ian A., and Michael Sofocleous (1995). "Interference Statistics for Fiscal Years 1992 to 1994," *Journal of the Patent and Trademark Office Society* 77: 417–422.

Carlino, Gerald A. (2001). "Knowledge Spillovers: Cities' Role in the New Economy," *Federal Reserve Bank of Philadelphia Business Review* 2001 (Q4): 17–26.

Carlino, Gerald A., Jake K. Carr, Robert M. Hunt, and Tony E. Smith (2017). "The Agglomeration of American R&D Labs," *Journal of Urban Economics* 101: 14–26.

Carlino, Gerald A., Satyajit Chatterjee, and Robert M. Hunt (2007). "Urban Density and the Rate of Invention," *Journal of Urban Economics* 61: 389–419.

Carlino, Gerald A., and William R. Kerr (2015). "Agglomeration and Innovation," in *Handbook of Regional and Urban Economics* vol. 5, Gilles Duranton, J. Vernon Henderson, William C. Strange (Eds.), Elsevier, 349–404.

Chatterji, A., E. Glaeser, and W.R. Kerr (2014). "Clusters of Entrepreneurship and Innovation," in *Innovation Policy and the Economy*, vol. 14, J. Lerner, S. Stern (Eds.), University of Chicago Press, 129–166.

Cohen, Linda R., and Jun Ishii (2006). "Competition, Innovation, and Racing for Priority at the U.S. Patent and Trademark Office," working paper.

De Simone, Daniel V., James B. Gambrell, and Charles F. Gareau (1963). "Characteristics of Interference Practice," *Journal of the Patent Office Society* 45: 503–591.

Duranton, Gilles, and Henry Overman (2005). "Testing for Localization Using Micro-Geographic Data," *Review of Economic Studies* 72: 1077–1106.

Duranton, Gilles, and Diego Puga (2004). "Micro-Foundations of Urban Agglomeration Economies," in J. Vernon Henderson and Jacques-François Thisse (Eds.), *Handbook of Regional and Urban Economics.* Amsterdam: Elsevier, vol. 4, chap. 48, 2063–2117.

Feldman, M. P. (2000). "Location and Innovation: The New Economic Geography of Innovation, Spillovers, and Agglomeration" in *The Oxford Handbook of Economic Geography*, G. Clark, M. Feldman and M. Gertler (Eds.). Oxford: Oxford University Press, 373–394.

Feldman, M.P., and D.B. Audretsch (1999). "Innovation in Cities: Science-Based Diversity, Specialization and Localized Competition," *European Economic Review*, 43 (2): 409–429.

Ganguli, Ina (2015). "Immigration and Ideas: What Did Russian Scientists 'Bring' to the United States?" *Journal of Labor Economics*, 33 (S1), U.S. High-Skilled Immigration in the Global Economy (Part 2): S257–S288.

Gladwell, Malcom (2008). "In the Air: Who Says Big Ideas Are Rare?" *The New Yorker*, May 12, 2008.

Glaeser, Edward (1999). "Learning in Cities," *Journal of Urban Economics* 46 (2); 254–277.

Head, Keith, Yao Amber Li, and Asier Minondo (2015). "Geography, Ties, and Knowledge Flows: Evidence from Citations in Mathematics," working paper.

Helmers, Christian and Henry Overman (2017). "My Precious! The Location and Diffusion of Scientific Research: Evidence from the Synchrotron Diamond Light Source," *Economic Journal* 127 (604): 2006–2040

Henderson, Rebecca, Adam Jaffe, and Manuel Trajtenberg (2005). "Patent Citations and the Geography of Knowledge Spillovers: A Reassessment: Comment." *American Economic Review*, 95 (1): 461–464.

Jaffe, Adam B., Manuel Trajtenberg, and Michael S. Fogarty (2000). "Knowledge Spillovers and Patent Citations: Evidence from a Survey of Inventors," *American Economic Review* 90 (2): 215–218.

Jaffe, Adam B., Manuel Trajtenberg, and Rebecca Henderson (1993). "Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations," *Quarterly Journal of Economics* 108 (3): 577–598.

Kingston, William (2001). "Light on Simultaneous Invention from U.S. Patent Office 'Interference' Records," *World Patent Information* 26: 209–220.

Krugman, Paul (1991). *Geography and Trade.* Cambridge: MIT Press.

Lai, Ronald, Alexander D'Amour, Amy Yu, Ye Sun, and Lee Fleming, 2013,"Disambiguation and Co-Authorship Networks of the U.S. Patent Inventor Database (1975 - 2010)," `http://hdl.handle.net/1902.1/15705`. The Harvard Dataverse Network [Distributor] V5 [Version].

Lin, Jeffrey (2011). “Technological Adaptation, Cities, and New Work,” *Review of Economics and Statistics*, 93 (2): 554–574.

Lin, Jeffrey (2014). “The Paper Trail of Knowledge Transfers,” Federal Reserve Bank of Philadelphia *Business Review* Q2: 1–6.

Lucas, Robert (1988). “On the Mechanics of Economic Development,” *Journal of Monetary Economics* 22: 3–42.

Marshall, Alfred. *Principles of Economics.* London: Macmillan, 1920.

Merton, Robert K. (1973). *The Sociology of Science: Theoretical and Empirical Investigations.* Chicago: University of Chicago Press.

Murata, Yasusada, Ryo Nakajima, Ryosuke Okamoto, and Ryuichi Tamura (2014). “Localized Knowledge Spillovers and Patent Citations: A Distance-Based Approach,” *Review of Economics and Statistics* 96 (5): 967–985.

Ogburn, William F., and Dorothy Thomas (1922). “Are Inventions Inevitable? A Note on Social Evolution,” *Political Science Quarterly* 37 (1): 83–98.

Pollack, Andrew (2017). “Board Rules on Patents for Process to Edit DNA.” *The New York Times*, February 16, page B3.

Rosenbaum, Paul R. (2002). *Observational Studies*, 2nd ed. New York: Springer-Verlag.

Rosenthal, Stuart S., and William C. Strange (2001). “The Determinants of Agglomeration,” *Journal of Urban Economics* 50 (2): 191–229.

Rosenthal, Stuart S., and William C. Strange (2004). “Evidence on the Nature and Sources of Agglomeration Economies,” in J. Vernon Henderson and Jacques-François Thisse (Eds.), *Handbook of Regional and Urban Economics* (Amsterdam: Elsevier), 2119–2171.

Saxenian, AnnaLee (1994). *Regional Advantage: Culture and Competition in Silicon Valley and Route 128*, Harvard University Press.

Schmookler, Jacob (1966). *Invention and Economic Growth*, Harvard University Press.

Silverman, Bernard W. (1986). *Density Estimation for Statistics and Data Analysis.* New York: Chapman and Hall.

Thompson, Peter (2006). “Patent Citations and the Geography of Knowledge Spillovers: Evidence from Inventor- and Examiner-Added Citations,” *Review of Economics and Statistics* 88 (2): 383-388.

Thompson, Peter, and Melanie Fox-Kean (2005). “Patent Citations and the Geography of Knowledge Spillovers: A Reassessment.” *American Economic Review* 95 (1): 450–460.

United States Patent and Trademark Office. “e-FOIA Reading Room,” http://e-foia.uspto.gov/Foia/PTABReadingRoom.jsp, accessed November 3, 2014.

United States Patent and Trademark Office. “Efiling for Patent Trial and Appeal Board,” https://acts.uspto.gov/ifiling/PublicView.jsp, accessed May 6, 2015.

United States Patent and Trademark Office. “Master Classification File,” http://patents.reedtech.com/classdata.php, accessed January 16, 2014.

United States Patent and Trademark Office. “Public Patent Application Information Retrieval,” https://portal.uspto.gov/pair/PublicPair, accessed June 11, 2013.

Weitzman, Martin L. (1998). “Recombinant Growth,” *Quarterly Journal of Economics* 113 (2): 331–360.

A Appendix

A.1 Claims and decisions

Table A1 displays statistics on claims of invention by parties according to their seniority status. (Recall that seniority status is typically determined by first-to-file and that the burden of proof in showing earlier conception and reduction to practice is on the junior party.) Panel A shows statistics for interference cases where we are able to observe both the claims made on all involved applications and the claims in interference. Panel B shows statistics for interference cases where we are able to observe only the claims in interference.

Interference counts correspond to most of application claims. The first row of Panel A shows that junior parties tend to make about 3 more claims on average compared with senior parties. (This difference in means is significant at the 5 percent level.) Of the 26 and 23 claims, respectively, made by junior and senior parties on average, 20 and 19 claims are declared in interference. (For the sample for which we only observe interfering claims, slightly fewer claims, 18 and 17, are declared in interference.) Thus, interferences tend to involve mostly identical competing claims of invention.

Senior parties tend to win two-thirds of the time. For the sample described in Panel A, junior parties lose 13 out of the 20 claims in interference, or 65 percent.

Partial decisions are rare. Out of the 688 junior parties, 465 (68 percent) lost *all* of their claims declared in interference. Out of the 687 senior parties, 221 (32 percent) lost all of their claims declared in interference. (The difference in these rates is significant at the 1 percent level.) These rates are similar to the larger sample described in Panel B (for which we observe only claims in interference). Overall, the patterns are also similar to those seen in the subsamples of cases decided on priority or cases conceded. Thus, winners and losers tend to win or lose all of the claims in interference.

A.2 Co-inventor ties

Interfering inventor pairs are more than twice as likely to be connected by previous co-inventor ties. Over two-fifths of interfering inventor pairs are connected by previous co-inventor ties, compared with about one-fifth of control inventor pairs (see Appendix Figure A1). This result confirms the findings of earlier studies suggesting that social and professional ties are important for facilitating knowledge flows.

	All cases		*Priority decisions*		*Conceded*	
	Jr.	Sr.	Jr.	Sr.	Jr.	Sr.
A. Application claims observed						
Claims in application(s)	26.0	22.8^b	26.7	24.1	24.7	21.1^b
and/or patent(s)	(26.9)	(26.6)	(35.9)	(31.4)	(22.0)	(24.1)
Claims in	20.0	18.8	21.0	19.3	18.1	17.6
interference	(21.5)	(22.5)	(21.5)	(24.4)	(16.8)	(20.8)
Claims lost	13.3	6.6^c	17.9	3.9^c	12.8	5.0^c
in decision	(19.4)	(16.5)	(20.6)	(12.5)	(17.4)	(11.1)
Number of parties	688	687	121	116	358	360
Lost all appl./pat. claims	276	173	73	14	131	89
Lost all interfering claims	465	221	101	20	251	114
B. Only interfering claims observed						
Claims in	17.6	17.2	18.8	17.7	16.4	16.5
interference	(18.9)	(20.0)	(19.8)	(20.5)	(15.3)	(18.8)
Claims lost	12.7	6.5^c	15.9	4.1^c	12.0	6.0^c
in decision	(17.3)	(15.2)	(19.5)	(11.7)	(15.3)	(18.8)
Number of parties	1,236	1,102	257	214	627	563
Lost all interfering claims	910	405	212	47	469	221

Table A1: Claims for senior and junior parties

This table reports means and standard deviations (in parentheses) for senior and junior parties. Seniority is determined before an interference proceeding begins, according to the earliest benefit date on file. (Typically, the benefit date is the earliest date of application to the USPTO or a foreign patent authority.) Number of claims is the sum of claims across all applications filed by each independent interfering party. H_0: Difference in means by seniority is zero. ^{a}—$p < 0.10$; ^{b}—$p < 0.05$; ^{c}—$p < 0.01$.

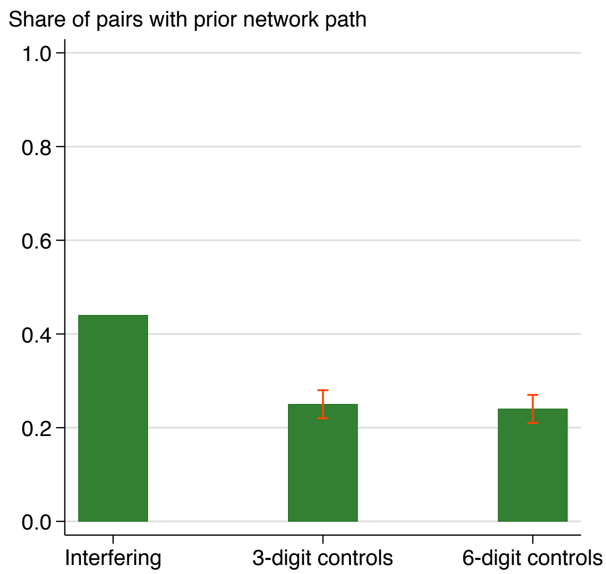

Figure A1: Interfering inventors are more likely to be connected by previous co-inventor ties

Priority decisions only. Simulated 90 percent confidence intervals shown.