

Learning remotely: R&D satellites, intra-firm linkages, and knowledge sourcing

Joel Blit 

Department of Economics, University of Waterloo, 200 University Ave. W., Waterloo, ON, Canada (Email: jblit@uwaterloo.ca)

Abstract

Firms venture abroad not only to access resources and markets but also to learn. Yet there remains limited empirical evidence that headquarters can access geographically remote knowledge by establishing a presence in the remote location. Using U.S. patent data, I show that firm headquarters disproportionately source knowledge from third parties in remote locations where they have an R&D satellite. This “satellite effect” on knowledge flow is economically significant, representing up to 60% of the knowledge-flow premium associated with collocation. Furthermore, the effect seems to be stronger for recent knowledge, as well as in areas of satellite technological specialization, suggesting that firms can target cutting-edge knowledge in specific sectors. In addition, the results show that firms with stronger internal linkages between headquarters and satellites, and those that staff satellites with inventors that previously patented while at other local firms, experience a larger satellite effect on knowledge acquisition.

1 | INTRODUCTION

A principal way that firms build competitive advantage is through innovation. To be successful, firms must both pursue incremental improvements that build upon their existing stock of knowledge (exploitation) and explore broader avenues in search of breakthrough advances (March, 1991). Yet, the latter can be particularly problematic for firms. As Singh and Fleming (2010) show, breakthrough innovation requires access to diverse sources of knowledge. Because these sources of knowledge are often external and can reside in geographically distant locations, firms may need to put in place mechanisms that facilitate access.¹ One such mechanism to access distant external knowledge is the establishment of a satellite R&D unit in the remote location (knowledge-seeking foreign direct investment [FDI] in the international context). But how large is the impact of the R&D satellite on the headquarters' acquisition of remote third-party knowledge? Is the knowledge acquired primarily cutting edge, or mature and, hence, of potentially lower value? How does the satellite facilitate knowledge acquisition? This paper sets out to empirically address these questions.

The idea that firms can tap into remote knowledge through the strategic deployment of subsidiaries is not new (Alcacer & Chung, 2002; Cantwell, 1993; Feinberg & Gupta, 2004; Kuemmerle, 1999). Indeed, firms like Hitachi and BMW have well-established knowledge-seeking R&D units in Cambridge, U.K., and Silicon Valley, respectively (Boutellier, Gassmann, & Von

I am grateful to Ajay Agrawal, Ignatius Horstmann, Joanne Oxley, and Daniel Treffer for their indispensable guidance. The seed idea for this paper originates from a discussion with Gilles Duranton. Alex Oettl generously provided some of my data. I also would like to thank Wilbur Chung, Steve Klepper, Carlos Serrano, Peter Thompson, and others who pushed me to produce a better paper through their insightful critiques. This paper also benefited greatly from the insightful comments of two anonymous referees. I gratefully acknowledge funding from the NBER Productivity, Innovation, and Entrepreneurship Program, the Social Sciences and Humanities Research Council of Canada, and from the Martin Prosperity Institute's Program on Innovation and Creative Industries. Any errors and omissions are my own.

Zedtwitz, 2008). Because knowledge diffusion is partially localized (Audretsch & Feldman, 1996; Branstetter, 2001; Jaffe, Trajtenberg, & Henderson, 1993; Thompson, 2006), collocating a satellite R&D unit with an important source of knowledge can be an effective way to access that remote knowledge. This is particularly true for knowledge that is tacit in nature and hence disseminated through worker mobility or social networks (Almeida & Kogut, 1999; Breschi & Lissoni, 2009). Indeed, strong evidence suggests that remote subsidiaries effectively embed themselves in local networks and make extensive use of local knowledge (Almeida, 1996; Frost, 2001; Singh, 2007). What is less clear, however, is whether the headquarters also taps into this remote knowledge through the medium of their subsidiary.

This paper addresses this question using patent citation data to track the flow of knowledge. Unlike the previous literature that has focused on international knowledge flows, this paper focuses on knowledge flows across different regions within the U.S., and it does so for two reasons. First, the U.S. offers a geographically large but relatively culturally and institutionally homogeneous setting that allows us to focus on the satellite effect on knowledge flows in isolation from the potential effects of cross-country developmental, cultural, and institutional differences. Second, as it is not clear that establishing a subsidiary in Boston would facilitate access to knowledge from San Francisco, the relevant level of analysis needs to be at the level of the city/region.² We would expect, however, that any findings of a satellite effect on knowledge diffusion within the U.S. would also apply to international knowledge flows, though the magnitudes might perhaps depend in part on the characteristics of country pairs.

The analysis focuses on U.S. firms with R&D centers in two different U.S. metropolitan regions. I quantify the satellite effect on knowledge acquisition using a matching methodology in the spirit of Jaffe et al. (1993). To partially address endogeneity in the choice of satellite location and omitted variable bias, I employ a difference-in-differences approach. I show that prior to the establishment of the satellite, the headquarters do not disproportionately source third-party knowledge from the satellite location. However, it does disproportionately source knowledge from the remote location after the satellite has been established. Moreover, this finding is not driven by firms that establish satellites to shift into new sectors. Nonetheless, it should be noted that the analysis is limited to exploring the effect of satellites for firms that have chosen to establish a satellite and the results are therefore only suggestive of the potential benefits for firms that have not.

The principal result is that when a firm has a satellite research center in a remote location, its R&D headquarter patents cite third-party patents from the remote location 40%–60% more often than would be expected given the geographic distribution of innovative activity (accounted for by a set of control patents). To offer a benchmark, a separate analysis is conducted to determine the extent to which patents are more likely to cite patents in their same location (the collocation effect). Patents cite third-party patents in their same region 66%–187% more often than would be expected given the geographic distribution of innovative activity. Clearly then, knowledge sourcing through a satellite cannot fully substitute for being collocated with an important source of knowledge, but it does account for between 32% and 60% of the knowledge-flow premium associated with collocation and an even higher fraction in those technological areas where the R&D satellite is active.

Collocation with a source of knowledge also confers the benefit of timely access to knowledge. We would similarly expect the role of satellites to be most significant in facilitating access to recently developed knowledge. Indeed, this paper presents some evidence that the knowledge received from the satellite location is disproportionately recent, though these results are inconclusive.

A final set of results begins to address the mechanism underlying the satellite effect by examining heterogeneous firm responses. Because effective transmission of knowledge from the satellite to the headquarters is crucial, firms with stronger linkages between the two locations should experience a larger satellite effect on knowledge acquisition. Similarly, satellites with stronger links to other firms in their location will gather more knowledge and hence be stronger channels for knowledge diffusion. I confirm these hypotheses by finding that the magnitude of the satellite effect is positively related to the intensity of cross-location patent coauthorship, the intensity of remote self-citations, and the fraction of satellite inventors that previously patented at other firms in the satellite location.

Taken together, the findings suggest that we need to temper the view that high-tech firms should locate within technology clusters to be successful innovators. From the point of view of knowledge sourcing, the establishment of a satellite is an effective (if imperfect) substitute for collocation with an important source of knowledge.

The paper is organized as follows. Section 2 reviews the literature and discusses the theoretical underpinnings behind R&D satellites as channels for knowledge sourcing. Sections 3 and 4 describe the data and methodology. Section 5 presents the empirical results related to the existence of a satellite effect, studies whether knowledge sourced through the satellite is primarily recent, and examines firm characteristics that influence the magnitude of the satellite effect. Section 6 concludes.

2 | THEORETICAL FRAMEWORK

Knowledge spillovers play an important role in a number of phenomena. The distinctive features of knowledge, namely, that it is a nonrival and only partially excludable good, imply that it can support long-run economic growth (Romer, 1986, 1990). The localized nature of these knowledge diffusion externalities leads to one of three main sources of agglomeration economies identified by Marshall (1890). At a less aggregate level, the localization of knowledge diffusion also has significant implications for firm strategy, both for the acquisition of external knowledge and the protection of one's own intellectual property.

Scholars have proposed several explanations for why knowledge primarily diffuses locally. Von Hippel (1988) cites interfirm linkages in the form of supply relationships as a reason why knowledge diffusion is localized, although Rogers and Larsen (1984) argue that regional social networks facilitate knowledge flow. Agrawal, Cockburn, and McHale (2006) use patent citation data to empirically establish this direct relationship between social relationships and knowledge flow. Almeida and Kogut (1999) find that interfirm mobility of engineers mediates the local transfer of knowledge. Breschi and Lissoni (2009) go further in showing that this inventor mobility explains the majority of knowledge flow localization. Song, Almeida, and Wu (2003) and Singh and Agrawal (2010) also find evidence that suggests a relationship between labor mobility and knowledge flow. Overall, geographic proximity increases the frequency and lowers the cost of person-to-person contact, which facilitates knowledge transfer.

Firms that are collocated with a source of knowledge therefore benefit from increased access, both in terms of the quantity of knowledge received and the speed with which they receive newly generated knowledge (Jaffe et al., 1993; Thompson, 2006; Thompson & Fox-Kean, 2005). It should therefore not be surprising that the localization of knowledge diffusion impacts firm strategy. For instance, a Korean semiconductor firm might establish a subsidiary in Silicon Valley both to access the large market and to learn. Because knowledge markets are notoriously incomplete (Arrow, 1962) and diffusion is partially localized, FDI can be the only way to acquire remote specialized know-how.³ A growing empirical literature supports this view that firms use FDI to seek capabilities abroad. Cantwell (1993) finds that foreign subsidiaries in Britain were primarily in sectors of British technological strength. Almeida (1996) presents more direct evidence in the form of patent citations, finding that Korean and European semiconductor firms offset home country technological weaknesses by setting up subsidiaries in U.S. regions with technological strength. Furthermore, these subsidiaries use local knowledge more than similar domestic firms. Singh (2007) also examines patent citations and concludes that MNC subsidiaries learn from their local peers. Moreover, in technologically advanced countries, these knowledge outflows from the host country to the subsidiary are greater than knowledge inflows from the foreign subsidiary to domestic firms. Branstetter (2006) examines a group of Japanese firms and finds an increase in their number of citations to U.S. patents following FDI to the U.S., particularly when the FDI is in the form of R&D facilities.

Although it is well established that a subsidiary gains access to the knowledge in its region (as we might expect given that knowledge diffusion is partially localized), it has not yet been established whether this knowledge also reaches the headquarters. This is the primary objective of this paper: to determine whether a firm's primary R&D center gains increased access to remote knowledge through the presence of an R&D satellite in the remote location, and whether this R&D satellite effect is economically significant.

For remote knowledge to reach a firm's headquarters, two transitions must take place. First, the knowledge must reach the satellite and it must be absorbed. Second, this knowledge must be transferred by the satellite to the headquarters across geographical space and it must be absorbed by the headquarters. Clearly not all third-party knowledge in the satellite location will successfully make both transitions, but so long as some knowledge does, a "satellite effect" on knowledge diffusion will be observable. Further, by virtue of being collocated with knowledge sources, the satellite may access surrounding new knowledge more quickly, and to the extent that it can promptly transfer it, the primary R&D center will also have access to recent knowledge from the satellite's location.

The mental map of two distinct transitions taking place can also help us understand the mechanisms underlying the satellite effect and what factors mediate its magnitude. We discuss, in turn, each of the two transitions. In order for the satellite to eventually transfer knowledge it must first access and internalize it. One of the mechanisms through which it can access knowledge generated by other local firms is by recruiting their researchers. Rosenkopf and Almeida (2003) find that hiring firms benefit not only from the knowledge of the inventors they hire, but also from increased access to the knowledge of the inventor's previous employer. Thus, satellites that hire engineers and scientists that have previously worked at other firms in the region access more local knowledge and hence potentially generate a bigger satellite effect.

In addition to labor mobility, localized social networks and supply linkages can also facilitate access to local knowledge. We might expect all three of these mechanisms to be stronger within the technological sectors where the satellite is active. Labor generally moves across firms within the same sector, engineers are more likely to have social ties with other individuals working in similar areas, and supply linkages are often stronger within a sector. Moreover, knowledge accessed by the satellite must be absorbed, and this is most easily achieved in technological areas where the satellite has expertise (Cohen & Levinthal, 1990).

Thus, we would expect the satellite to be particularly effective at gathering third-party knowledge in the technological areas where the satellite specializes, and hence for the satellite effect to be largest in these sectors.

Once the satellite has internalized knowledge, it must transfer it to the primary R&D center. Implicit in the concept of a satellite effect is the idea that knowledge can leap across large geographic distances more easily within the firm. This argument that firms are better than markets at sharing and transferring knowledge has been made by Kogut and Zander (1992, 1993) and is central to the knowledge-based theory of the firm. One of the findings of Oettl and Agrawal (2008) is that knowledge indeed flows more easily across remote locations within the firm than outside it.

There are a number of reasons why knowledge might flow more easily within the firm. The transaction-cost literature argues that agents within firms exhibit decreased opportunism, which facilitates knowledge transmission. Other mechanisms mirror those used to explain why knowledge diffusion is localized: input linkages between remote units, movement of employees between units, more frequent interactions, and personal relationships. Thus, we would expect the satellite effect to be largest in firms that exhibit more frequent cross-location patent coauthorship and a higher intensity of cross-location self-citation, as working together and building on each other's work results in more frequent interactions and stronger personal relationships.

3 | DATA

One difficulty in the study of knowledge flows is that they are notoriously hard to measure. As Krugman (1991, p. 53) points out: "knowledge flows... are invisible; they leave no paper trail by which they may be measured and tracked." A crucial breakthrough, then, was provided by Jaffe et al. (1993) when they identified a setting in which these invisible knowledge flows leave a physical paper trail: patent citations. Although now commonly employed in the literature, the use of patent citations to measure knowledge flows is not without its critics. For one, patent citations do not capture all knowledge flows, but rather only those that result in a patented innovation. Further, many citations are added by the patent examiner (Alcacer & Gittelman, 2006). Although on the one hand this makes citations more objective, on the other it results in some citations being to knowledge that the inventors may not have known about.⁴ Notwithstanding, although citations are a noisy measure of knowledge flows, studies comparing citation data with inventor surveys have shown a high enough correlation between patent citations and actual knowledge flows to justify their use in large samples (Duguet & MacGarvie, 2005; Jaffe & Trajtenberg, 2002, Chapter 12).

The principal data set used in this study is the NBER Patent Citations Data File developed by Hall, Jaffe, and Trajtenberg (2001).⁵ It lists detailed information on almost three million patents granted by the U.S. Patents and Trademarks Office, including the location of inventors, the name of the assignee,⁶ and citations made and received. The base sample used consists of all utility patents granted between 1976 and 2002, generated in the U.S., and having a U.S. nongovernment organization (firms) as the assignee.⁷ The sample is further restricted to the 80% of patents where all inventors are from the same Metropolitan Statistical Area (MSA) so as to unambiguously assign a location to each patent based on the residence of the inventors. This also ensures that any measured satellite effect is not occurring as a result of direct knowledge flow through a coinventor in the satellite and thus the estimated magnitude of the satellite effect is a conservative estimate of the full actual effect. Irrespective, as shown in Appendix A, the computed satellite effect does not change significantly if all patents with at least one inventor in the primary R&D center are included in the set of treated patents. The analysis also focuses on single-location patents for the purpose of determining the R&D locations of firms. This is done to increase the threshold for what is considered a satellite, because a satellite must have generated a patent on its own (without remote coinventors). Appendix A also shows that a looser definition of a satellite as any location where the firm has at least one inventor, does not significantly alter the principal results.

This paper focuses on firms with patenting activity in exactly two different MSAs.⁸ The location where the majority of patents were generated is designated the "primary" R&D center, whereas the other location is the "satellite" R&D center. The initial sample consists of 4,610 two-location firms whose primary and satellite R&D centers are geographically distributed as shown in Figures 1 and 2. Table 1 presents summary statistics for these firms.

4 | METHODOLOGY

The baseline methodology builds on Agrawal et al. (2006). There are two types of citing patents in the analysis: treated and control patents. The set of treated patents is comprised of all patents generated in the primary R&D center of these two location firms (regardless of application date). These patents are "treated" in the sense that the innovation benefitted from the presence of (and potential knowledge sourcing through) a satellite in the remote location.⁹

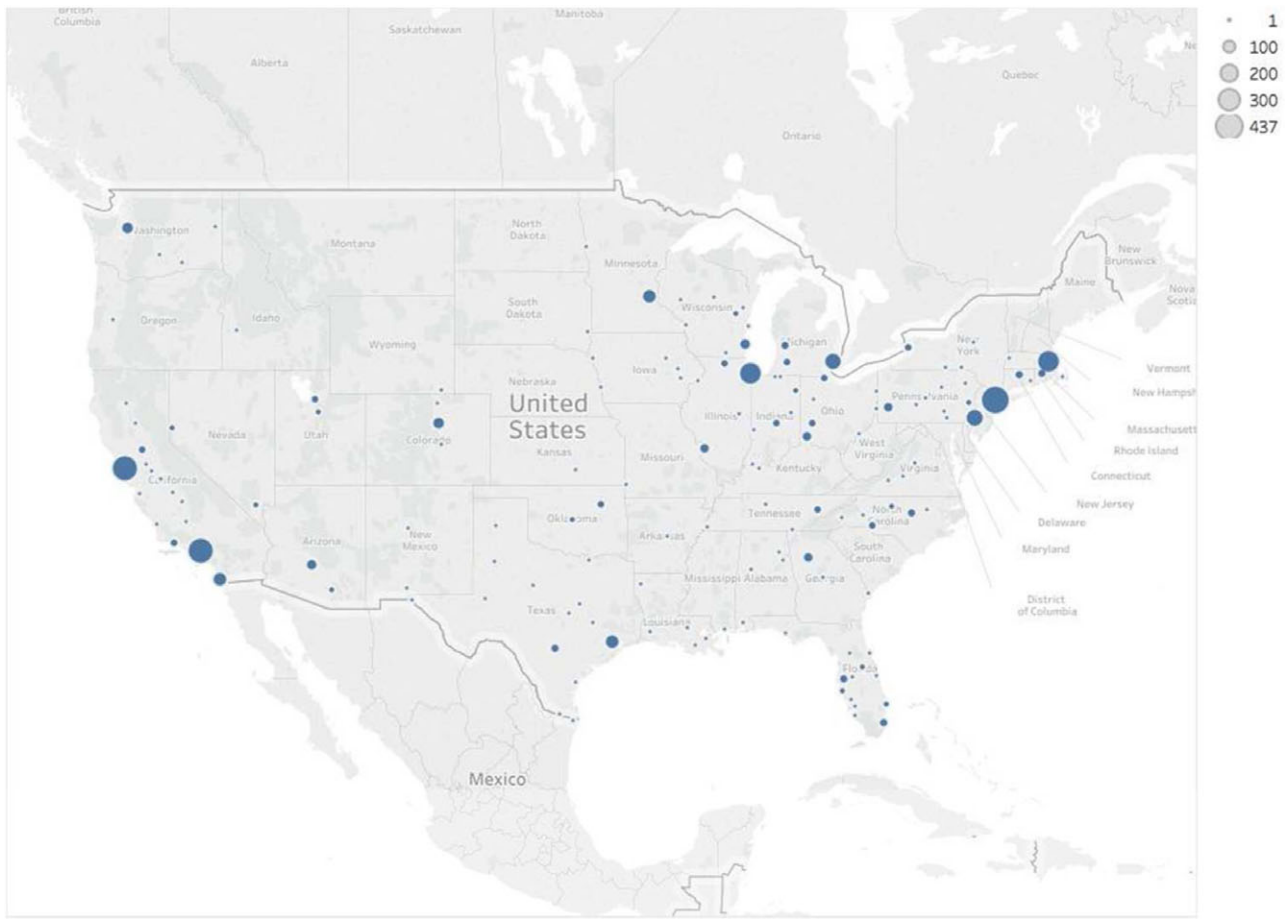


FIGURE 1 Number of primary R&D centers by Metropolitan Statistical Area

[Color figure can be viewed at wileyonlinelibrary.com]

Note: The size of the circle represents the number of primary R&D centers of two-location firms that are located in that MSA. The 4,610 U.S. two-location firms had their primary R&D center in 225 distinct MSAs, with 437 being in New York, 370 in Los Angeles, 359 in San Francisco, 269 in Boston, and 264 in Chicago.

4.1 | Control matching

In order to control for patterns of geographic agglomeration of technological activity that could be related to the choice of location for the satellite, a matched control group is generated from the same subsample of patents.¹⁰ The concern is that patents disproportionately cite other patents in their own technological area. Then, if a firm produces, for example, computer chips, it is more likely both to cite other computer chip patents (many of which are in Silicon Valley) and choose Silicon Valley as the location of its satellite. Not controlling for this would result in a potentially spurious finding of a satellite effect.

To address this, each patent in the treated group is matched to a control patent from the same location, application year, and technology class (this will later be further addressed using time differences). Moreover, the assignee of the control patent cannot have patented in the same location as the primary's satellite (i.e., it cannot have a presence in the same remote location as the satellite). If several potential control patents match the four criteria, the patent belonging to the multilocation firm with the fewest locations is chosen. Patents belonging to single-location firms are chosen as a last resort. If two or more patents are equally suitable controls according to these five criteria, one is chosen at random. Any treated patent with no matching control is dropped from the sample.

An alternative approach is to match treated and control patents as above, but in addition also by technological subclass. Although the finer matches yield results that are more robust to the above criticism, these may suffer from selection bias in that patents in the technology sectors and locations with many peer patents may be systematically different from patents in the technology sectors and locations with few/no peers.¹¹ Finding patents that match by location, application year, class, and subclass is often impossible, which results in having to discard many more potential treated patents due to lack of a suitable control. Although a suitable control was found for 67% of potential treated patents when matching at the class level, a control was

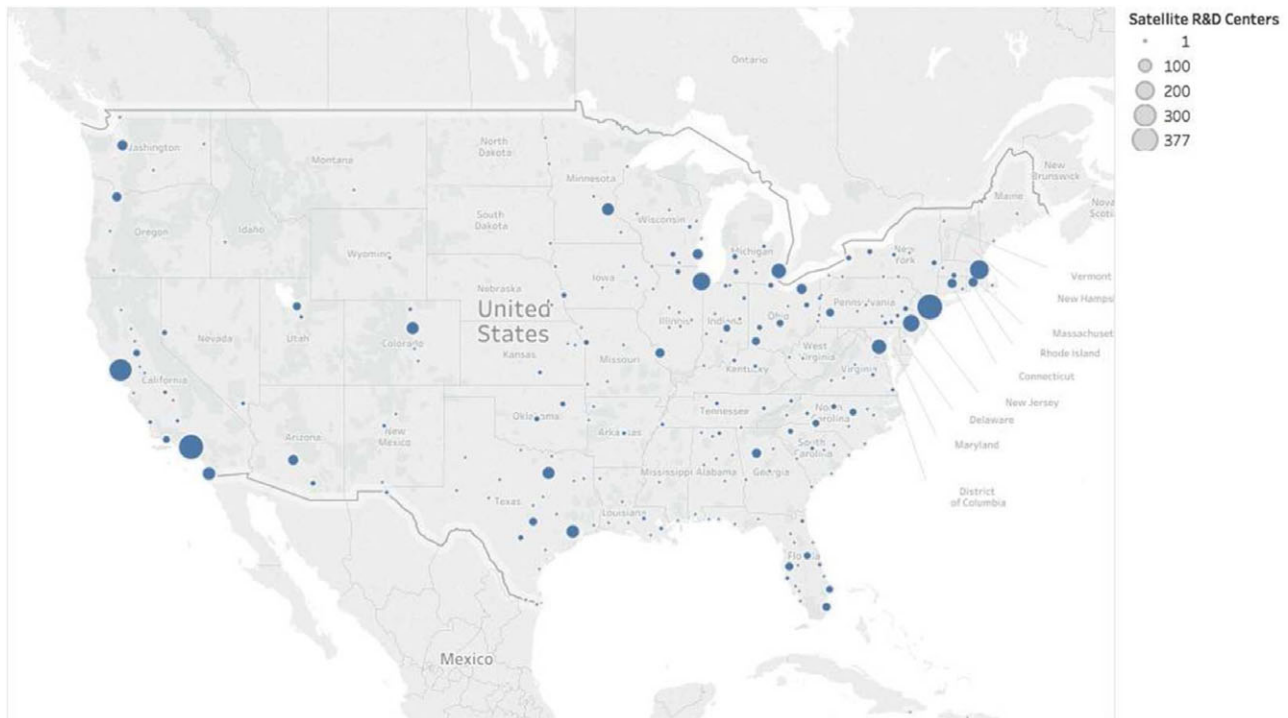


FIGURE 2 Number of satellite R&D centers by Metropolitan Statistical Area

[Color figure can be viewed at wileyonlinelibrary.com]

Note: The size of the circle represents the number of satellite R&D centers of two-location firms that are located in that MSA. The 4,610 U.S. two-location firms had satellites in 239 distinct MSAs, with 377 being in New York, 360 in Los Angeles, 302 in San Francisco, 216 in Boston, and 191 in Chicago.

TABLE 1 Summary statistics

Variable	Mean	Std. Dev.	Min	Max
Number of patents	21.6	58.1	3	2,317
Number of primary center patents	11.5	25.2	2	737
Number of satellite patents	1.7	3.0	1	156
Geographic distance b/w primary and satellite (km)	1,448	1,296	23	7,772
Fraction of satellite inventors with:				
two or more satellite patents	0.19	0.26	0	0.98
previous patent in satellite MSA for other org.	0.30	0.40	0	1
previous patent in primary MSA for primary	0.08	0.25	0	1
previous patent in primary MSA for other org.	0.05	0.21	0	1
previous patent in other MSA for other org.	0.10	0.26	0	1
Cross-location (sat-prim) collaboration intensity	0.12	0.24	0	0.98
Cross-location (sat-prim) self-citation intensity	0.03	0.09	0	1
Technological distance between prim. and sat.	0.90	0.40	0	1.41
Technological distance between prim. and sat. MSA	0.91	0.16	0	1.41

Note: The initial sample contains 4,610 U.S. two-location firms (though the final sample with controls consists of 3,781 firms). On average, firms have a total of 21.6 patents, 11.5 patents where all inventors are from the same primary R&D center MSA, and 1.7 patents where all inventors are from the same satellite MSA. This suggests that the firms have on average 8.4 patents that are either multilocation or have inventors in a foreign country. The summary statistics also show that up to 45% of satellite inventors previously patented at other organizations (30% in the same MSA as the satellite, 5% in the primary MSA, and 10% in some other MSA—though an inventor could fall in more than one of these categories). We also see that on average 12% of satellite patents had a coinventor from the primary R&D center and that 3% of citations by firm patents are citing firm patents from the other location. The technological distance between the primary and satellite is the Euclidean distance between the fraction of patents in the portfolio of primary and satellite patents that are in a particular technology class. The technological distance between the primary and satellite MSA is a similar measure (combining class and subclass) between primary patents and third-party patents from the satellite MSA.

TABLE 2 Control matching frequency

Control Patent Assignee Number of Locations	Matching by Class Number of Matches	Matching by Subclass Number of Matches
1	4,673	1,493
2	17,736	1,957
3	3,650	644
4	1,821	434
5 or more	7,154	2,188
Total:	35,034	6,716
Potential treated patents	52,080	52,080
% Control found	67.3%	12.9%

Note: Although a suitable control matching the same location, application year, and technology class was found for 67.3% of the 52,080 treated patents, a control that also matches by technology subclass was only found for 12.9% of the treated patents. Matching by class results in a sample of 3,781 two-location firms (vs. 1,368 when matching by subclass) with 51% of the control patents belonging to two-location firms (vs. 29% when matching by subclass). As discussed in Section 5.1, matching by subclass seems to generate a downward bias in the estimated satellite effect due to sample selection.

only found for 13% of treated patents when matching at the subclass level (see Table 2). Matching by class results in a sample of 3,781 two-location firms, although only 1,368 such firms remain when matching by subclass. Moreover, when matching by class the matches are closer in the dimension of the assignee's number of locations: 51% of control patents belong to two-location firms when matching at the class level versus 29% when matching at the subclass level. As discussed in Section 5.1, the results indeed suggest that subclass matching introduces a bias due to sample selection. Nonetheless, subclass matching could be preferable in some instances, depending on whether imperfect matching or selection bias is the principal concern, and therefore both sets of principal results are presented.

The above approach for choosing control patents suggests that the primary dimension in which treated and control patents might be different is that control patents will on average belong to firms with a broader geographical presence. Appendix C examines the extent to which the treated and control patents look similar and indeed finds that the main difference is the broader geographical presence of control patent assignees. To the extent that firms with a broader geographical presence are better at gathering remote knowledge, this would attenuate our results. Appendix C shows that this is the case.

An alternative to the control matching methodology presented above uses examiner-added citations as controls (Thompson, 2006) for primary R&D center patents granted in 2001 and 2002. As shown in Appendix E, this approach yields similar results, though the estimated satellite effect is smaller.

4.2 | Computing the satellite effect

The principal hypothesis to be tested is that patents generated in the primary R&D centers disproportionately (relative to the control group) cite third-party patents from the location of their satellite. Pooling across all firms, I obtain a group of treated patents and an associated group of control patents. I compare the proportion of treated group patent citations that are to third-party patents in the satellite's location, to the proportion of control group patent citations that are to third-party patents in the satellite location of their matched treated patent. The ratio of these is the satellite effect (*SE*) on knowledge acquisition, which measures the degree to which having a satellite results in increased knowledge sourcing from the satellite location. Formally, the satellite effect is computed as:

$$SE = \frac{cites_T(loc = sat)}{cites_T} \div \frac{cites_C(loc = sat)}{cites_C},$$

where $cites_T$ is the total number of citations made to third-party single-MSA U.S. patents by treated group patents (aggregating across all firms and treated patents) and $cites_T(loc = sat)$ is the subset of these that cite patents from the location of the satellite. Similarly, $cites_C$ is the total number of citations made to third-party single-MSA U.S. patents by all control patents, and $cites_C(loc = sat)$ is the subset of these that cite patents from the location of their matched treated patent's satellite.

To examine whether the magnitude of *SE* is economically significant, it is useful to compare it with a well-understood benchmark: the knowledge diffusion premium associated with collocation. To enable a fair comparison, the collocation effect is computed using a methodology analogous to the one used to compute the satellite effect. The sample is again restricted to patents where all inventors are from the same MSA. However, to ensure that satellites play no role in knowledge transmission,

the sample is further restricted to patents of assignees having only one location. Then, for each remaining patent (which we label “treated”), a control patent is found with the same application year and technology class but from a different location. The analysis compares the probability that a “treated” patent citation is to a patent from its location, with the probability that a control patent citation is to a patent from the same location as its matched treated patent. Self-citations are discarded. The collocation effect (CE) measures the degree to which being collocated increases knowledge sourcing from that location. It is computed analogously to the satellite effect as:

$$CE = \frac{cites_T(loc = treated)}{cites_T} \div \frac{cites_C(loc = treated)}{cites_C},$$

where $cites_T$ is the total number of citations made to third-party single-MSA patents by treated patents, $cites_T(loc = treated)$ is the subset of these that cite third-party patents from the same location as the treated patent, and $cites_C$ and $cites_C(loc = treated)$ are the counterparts for the group of matched control patents.

To determine whether the satellite effect is larger in the technology sectors where the satellite is active, I determine the proportion of cited patents that are both from the satellite location and in a technology class in which the satellite has patented. In particular, I compute $SE_{satclass}$ as:

$$SE_{satclass} = \frac{cites_T(loc = sat, class = sat)}{cites_T} \div \frac{cites_C(loc = sat, class = sat)}{cites_C},$$

where the variables are defined as before with $cites_T(loc = sat, class = sat)$ being the total number of citations made to third-party patents that are both from the satellite location and whose primary technology class is common to at least one of the satellite’s patents. $cites_C$ and $cites_C(loc = sat, class = sat)$ are defined analogously for control patents.

4.3 | Knowledge recentness

To the extent that the satellite facilitates knowledge transfer, we would expect that new third-party knowledge developed in the satellite region should reach the primary R&D center more quickly. Specifically, we should observe that a disproportionately large share of the recently generated third-party knowledge received by a primary R&D center originates from the location of its R&D satellite (in addition to from its own location). I examine this hypothesis by pooling all citations made by treated and control patents and performing the following citation-level logistical regression:

$$(I_{loc = satmsa})_i = \alpha + \beta_1 age_i + \beta_2 treated_i + \beta_3 age_i \times treated_i + \varepsilon_i.$$

The dependent variable is an indicator equal to 1 when the cited patent is from the satellite location and 0 otherwise. age_i is the age of the cited patent measured by the citation lag (application year of citing patent – application year of cited patent) as in Fabrizio (2007), whereas $treated$ is equal to 1 when the citation is made by a treated patent and 0 if it is made by a control patent. We expect our primary coefficient of interest (β_3) to be negative if younger cited patents are more likely to originate from the satellite location and β_1 to be close to 0 because control patents should not experience a satellite effect.

4.4 | Firm characteristics and satellite effect

Our discussion of the likely mechanisms behind the satellite effect suggests a heterogeneous satellite effect across firms. In particular, satellites staffed with experienced local hires are likely to access and internalize more knowledge from firms in the region. Firms that exhibit more collaboration among inventors from the primary and satellite units and/or more frequently build on previous work performed by the firm at the other remote location are likely to benefit from a more effective transfer of knowledge between the satellite and primary R&D center. *Ceteris paribus*, then, both sets of factors should lead to a larger observed satellite effect.

To determine the impact of these firm characteristics on the satellite effect, I estimate the following OLS regression at the level of the patent:

$$\left(\frac{cites_T(loc = sat)}{cites_T} - \frac{cites_C(loc = sat)}{cites_C} \right)_{ift} = \alpha + \beta invprev_f + \delta_1 crossloccollab_f + \delta_2 crossloccite_f + \varphi X_f + \gamma_t + \varepsilon_{ift},$$

where i indexes the patent, f the firm that owns the patent, and t the patent’s application year. Because SE computed at the patent level would have numerous zeros in the denominator, the dependent variable is instead the difference between the proportion

of a treated patent i 's citations that are to third-party patents in the satellite location and the proportion of its associated control patent's citations that are to third-party patents in the satellite location.

The first set of five explanatory variables (captured collective as *invprev* in the equation above) are measures of the previous patenting experience of satellite inventors. It includes the fraction of satellite inventors that previously patented in the satellite MSA while at a different organization, the fraction that previously patented in the satellite MSA for the satellite/primary firm (i.e., the number of satellite inventors with at least two patents with the satellite), the fraction that previously patented in the primary R&D center MSA for a different organization, the fraction that previously patented in the primary R&D center MSA for the satellite/primary firm, and the fraction that previously patented in an MSA other than the satellite or primary MSA for a different organization. To the extent that learning through hiring is an effective means of accessing knowledge, we expect the satellite effect to be particularly high for satellites with high proportions of inventors with previous local experience at other firms.

The cross-location collaboration variable (*crossloccollab*) measures the intensity of collaboration between firm inventors in the satellite and primary location. It is constructed as the fraction of satellite patents (firm patents having at least one inventor residing in the satellite MSA) that also have at least one inventor in the primary location. The cross-location self-citation measure (*crossloccite*) is constructed by looking across all patents of the firm and counting the proportion of citations to same-firm patents from the other location. We expect to observe a larger satellite effect in firms with higher intensities of cross-location collaboration and citations because working together and building on each other's work results in more frequent interactions and stronger personal relationships that would facilitate the flow of knowledge from the satellite to the primary.

Because the above explanatory variables may in part capture whether the primary and satellite centers work on similar technologies, I construct a control variable that measures technological overlap between primary and satellite patents. In particular, I compute the Euclidean distance (as in Rosenkopf and Almeida, 2003) between the primary and satellite patent portfolios. As a further control, I include the technological distance between the primary patent portfolio and the portfolio of third-party patents from the satellite MSA.¹² I also include a control for geographic distance between the primary and satellite MSAs. Lastly, I include patent application year fixed effects (γ_t).

4.5 | Difference in differences

It could be the case that the finding of a satellite effect is spurious and results from firms choosing to set up satellites in locations from where they already disproportionately source knowledge. For example, firms may establish satellites in locations that specialize in the same technological areas as themselves. Then, if control patents match imperfectly, the fact that patents are more likely to cite within their narrow technology class could lead to finding a satellite effect when none exists.

We can use the time dimension of the data to better address this type of endogeneity. If we assume that potential biases due to imperfect matching are constant across the pre- and postsatellite period, we can compute the average treatment effect on the treated through a difference in differences approach. And in particular, we can show that there was no observed "satellite effect" before the establishment of the satellite, but a significant effect afterwards.

The principal challenge with deploying this approach is that the date of establishment of a satellite is not observed in the patent data. What is observed is a proxy for the date of establishment: the application date of the first patent from the satellite location. Because patents are generally the result of many years of development (and many satellites may act as listening posts prior to innovating themselves), the year of the first application will invariably overshoot the actual year of establishment. Because the years right before the first application date of a satellite patent are the most problematic to accurately define as pre- or postsatellite, they are dropped from the analysis. The presatellite period is therefore defined as ending five years prior to the first satellite patent, though as shown in Appendix F the choice of "gray period" gap size is largely inconsequential to the results. The postsatellite period begins when the first satellite patent application is observed. Clearly, it will still be the case that many satellites are operating and gathering knowledge in what has been defined in this way as the "presatellite" period. But this will bias against finding no effect on knowledge sourcing prior to the establishment of the satellite and an effect thereafter.

To ensure that the pre- and postsatellite patents are similar, the sample is restricted to patents from firms that have primary R&D center patents in each of the pre- and postsatellite periods. As before, for each primary (treated) patent we find a control that is from the same location, application year and technology class. Pooling all patents (both treated and control), we perform the following difference in differences estimation:

$$\left(\frac{\text{cites}(\text{loc} = \text{sat})}{\text{cites}} \right)_{ift} = \alpha + \beta_1 \text{treated}_f + \beta_2 \text{postsat}_{ft} + \beta_3 \text{treated}_f \times \text{postsat}_{ft} + \gamma_f + \gamma_t + \varepsilon_{ift},$$

TABLE 3A Magnitude of collocation and satellite effects—technology class matching

	Satellite Effect All Tech Classes	Satellite Effect Sat Tech Classes	Collocation Effect All Tech Classes
% Treated matching	6.9	2.6	19.6
% Controls matching	4.3	0.9	6.8
<i>t</i> -Statistic	29.95	34.33	87.12
Effect	1.60	2.90	2.87
<i>N</i> - Treated citations	139,010	139,010	103,585
<i>N</i> - Control citations	142,256	142,256	101,765

Note: We analyze two samples: the 35,034 treated patents from a primary R&D center of a two-location firm and the associated 35,034 control patents that match by MSA, application year, and technology class. We consider these patents' citations to third-party single-MSA U.S. patents. Although 6.9% of these 139,010 citations by the treated group are to patents in the MSA of the satellite, only 4.3% of these 142,256 citations by the control group are to patents in the satellite MSA (Column 1). A difference of proportions *t*-test with unequal variances shows this difference to be statistically significant (*t*-statistic of 29.95). The satellite effect is computed as the ratio of these proportions ($SE = 1.60$), suggesting that across all technological classes firms receive on average 60% more knowledge from the location of their satellite than they would if no satellite was present. Column 2 shows that 2.7% of these treated patent citations are to patents in technology classes of the satellite and in the satellite MSA, although this is only 0.9% for control patent citations, yielding a larger satellite effect in technological areas where the satellite has expertise. Bootstrapping shows the difference between the satellite effects computed in Columns 1 and 2 to be significant at better than 0.1%. To compute the collocation effect (Column 3), we begin with all single-MSA U.S. patents of single-location firms (which we call treated) and for each find a control with the same application year and technology class but from a different location. We find that 19.6% of treated patent citations to third-party single-MSA U.S. patents are to patents in their same MSA, although only 6.8% of these control citations are to patents in the same MSA. The estimated collocation effect is the ratio of these (2.87).

where the index *i* is the patent, *f* is the firm that owns the patent, and *t* is the patent's application year. The dependent variable is the fraction of patent *i*'s citations that are to third-party patents in the satellite location. The explanatory variables include a dummy variable (*treated*) that is 1 when the patent is from the treated group and 0 if it is from the control group, a dummy variable (*postsat*) that is 1 when the patent's application year is in its firm's (or its associated treated patent's firm's) postsatellite period and 0 if it is in the presatellite period, and the interaction of these two variables. γ_f and γ_t represent firm and patent application year fixed effects, respectively. Our coefficient of interest, β_3 , provides an estimate of the satellite effect.

4.5.1 | Technological shift

Although the above difference in differences methodology is a more robust approach that addresses a number of potential concerns, it does rule out the possibility that firms might undergo other changes at the same time that they establish a satellite, and that these changes might alter the firm's citation patterns so that they begin to cite the satellite location more frequently. Such a scenario could lead to a spurious finding of a satellite effect. One plausible such scenario is a strategic realignment of the firm toward a new technological sector. To facilitate the headquarters' technological shift the firm might establish a satellite in a region that specializes in the new sector. If the shift is successful, more of the headquarters' patents will be in the new sector in the postsatellite period than in the presatellite period, and as patents more often cite their own technological class, we would observe headquarters' patents citing the satellite's location disproportionately after the establishment of the satellite but not before. To be clear, this problem only arises when technology class matching is inadequate as the statistic of interest is the treated patent's proportion of citations to the satellite location relative to that of a control patent in the same technology class.

To examine the plausibility of this alternative explanation we examine the distribution of a firm's primary patents across sectors in both the pre- and postsatellite period and determine whether the primary R&D center underwent a technological sector shift that brought it closer to the mix of technologies in the satellite's MSA. We compute technological distance using two alternative measures: geometric distance at the class level as in Rosenkopf and Almeida (2003) and geometric technological distance computed at the subclass level and averaged across all technology classes (as described in footnote 12). The difference in differences approach is then applied to the subset of firms that do not exhibit a sectoral shift (approximately half the firms).

5 | RESULTS

5.1 | Satellite effect

Given the ability of firms to transfer knowledge across large distances, we expect to find a knowledge flow premium associated with having a satellite in a remote location. In particular, primary R&D centers should receive a disproportionately large amount of third-party knowledge from the location of their R&D satellite. Table 3A (Column 1) confirms that this is indeed the case.

TABLE 3B Magnitude of collocation and satellite effects—technology subclass matching

	Satellite Effect All Tech Classes	Satellite Effect Sat Tech Classes	Collocation Effect All Tech Classes
% Treated matching	6.4	2.6	18.0
% Controls matching	4.5	1.3	10.8
<i>t</i> -Statistic	10.80	12.52	33.93
Effect	1.40	1.98	1.66
<i>N</i> - Treated citations	34,582	34,582	55,190
<i>N</i> - Control citations	37,329	37,329	53,589

Note: We analyze two samples: the 6,716 treated patents from a primary R&D center of a two-location firm and the associated 6,716 control patents that match by MSA, application year, technology class, and technology subclass. We consider these patents' citations to third-party single-MSA U.S. patents. Although 6.4% of these 34,582 treated group citations are to patents in the MSA of the satellite, only 4.5% of these 37,329 control group citations are to patents in the satellite MSA (Column 1). A difference of proportions *t*-test with unequal variances shows this difference to be statistically significant (*t*-statistic of 10.80). The satellite effect is computed as the ratio of these proportions ($SE = 1.40$), suggesting that across all technological classes firms receive on average 40% more knowledge from the location of their satellite than they would if no satellite was present. Column 2 shows that 2.6% of these treated patent citations are to patents in technology classes of the satellite and in the satellite MSA, although this is only 1.3% for control patent citations, yielding a larger satellite effect in technological areas where the satellite has expertise. Bootstrapping shows the difference between the satellite effects computed in Columns 1 and 2 to be significant at better than 0.1%. To compute the collocation effect (Column 3), we begin with all single-MSA U.S. patents of single-location firms (which we call treated) and for each find a control with the same application year, technology class, and technology subclass but from a different location. We find that 18.0% of treated patent citations to third-party single-MSA U.S. patents are to patents in their same MSA, although only 10.8% of these control citations are to patents in the same MSA. The estimated collocation effect is the ratio of these (1.66).

Although the proportion of treated patent citations that are to a third-party patent in the location of their satellite is $\hat{P}_T = 0.069$, only $\hat{P}_C = 0.043$ of the control patent citations are to a third-party patent in the location of their matched patent's satellite. Performing a difference of proportions *t*-test with unequal variances as in Almeida (1996) yields a *t*-statistic of 29.95.¹³ The satellite effect is computed as $SE = \frac{\hat{P}_T}{\hat{P}_C} = 1.60$, suggesting that firms receive on average 60% more knowledge from the location of their satellite than they would if no satellite was present. Column 2 confirms that the satellite effect is larger in technological classes where the satellite has patented ($SE_{satclass} = 2.90$), presumably because of the satellite's increased ability to identify and absorb knowledge in these areas. Bootstrapping shows the difference between SE and $SE_{satclass}$ to be significant with a *p*-value of 0.03%, suggesting that firms can use satellites to target specific knowledge.

Yet having a satellite in a remote location is likely only an imperfect substitute for full collocation with the knowledge source. Column 3 indeed finds that the knowledge flow premium associated with collocation is larger ($CE = 2.87$). The satellite effect is nonetheless important, accounting for 32% of the collocation premium across all technological classes. As an aside, the magnitude of the collocation premium (187%) is in line with the findings of Jaffe et al. (1993) for firms and collocation at the MSA level.

Table 3B presents the results when matching is performed at the level of the technology subclass. As expected, the magnitude of the effect is significantly smaller ($SE = 1.40$), although it now represents a larger 61% of the collocation premium. This is not surprising as we would expect the upward bias introduced by imperfect matching to be more serious in the computation of the collocation effect.¹⁴ As before, the satellite effect is particularly large in areas of satellite technological specialization ($SE_{satclass} = 1.98$) and the difference with SE is again significant at a better than 0.1% level.

Additional insights are gained by comparing Tables 3A and 3B. For the overall satellite effect (Column 1), the principal reason for a smaller computed satellite effect with subclass matching is not so much an increase in \hat{P}_C (from 0.043 to 0.045) as it is a drop in \hat{P}_T (from 0.069 to 0.064). The increase in \hat{P}_C is expected if indeed subclass matching results in better matches. But \hat{P}_T should not have changed substantially. Its relatively large decline suggests that matching by subclass generates a significant selection bias.

The collocation results are similarly suggestive of a selection bias with the proportion of focal patent citations to collocated patents decreasing from 0.196 to 0.180 with the finer matching. Overall, these findings caution against using subclass matching except when absolutely necessary. For instance, when computing the magnitude of the satellite effect as above, the preferred methodology could be to use subclass matching because imperfect matching might introduce a more serious bias than sample selection. But when considering second-order effects, as in the following sections, the consequences of imperfect matching are less severe (especially if the bias is relatively constant), and so the focus is on results using technology class matching.

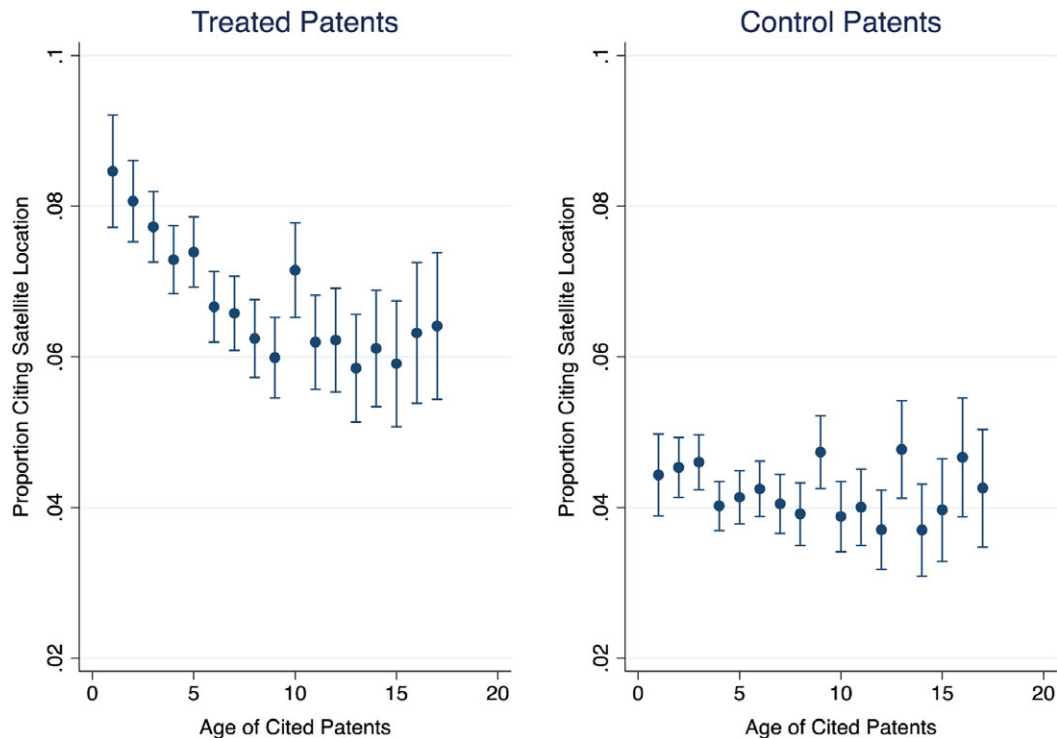


FIGURE 3 Proportion of citations citing satellite location by age of cited patents for both treated (left panel) and control (right panel) patents [Color figure can be viewed at wileyonlinelibrary.com]

Note: The samples consist of the 35,034 treated patents from a primary R&D center of a two-location firm and the associated 35,034 control patents that match by MSA, application year, and technology class. We consider these patents' citations to third-party single-MSA U.S. patents. The left panel shows the proportion of these treated patent citations that are to the satellite location by age of the cited patents. Although approximately 8.5% of one-year-old cited patents are from the satellite location, fewer than 6% of 15-year-old cited patents are from that same satellite location. No such pattern exists for the proportion of control patent citations to patents from the matched treated patent's satellite location (right panel). The error bars represent two standard deviations on either side of the mean (and hence the 95% confidence interval). The figures suggest that satellites might be an effective means of acquiring cutting-edge knowledge.

5.2 | Knowledge recentness

Being collocated with a source of knowledge not only increases the amount of knowledge received, it also ensures that this knowledge is received while it is still recent. To the extent that the satellite can quickly transfer locally gathered knowledge to its primary R&D center, we might expect that a large proportion of recent third-party knowledge received by primary R&D centers originate from the location of their R&D satellite (in addition to from its own location). As shown in Figure 3 (left panel), this indeed seems to be the case. The figure shows the fraction of treated patent citations that are to third-party patents from the satellite location by age of the cited patents. Although approximately 8.5% of one-year-old cited patents are from the satellite location, fewer than 6% of 15-year-old cited patents are from that same satellite location.¹⁵ No such pattern exists for the proportion of control patent citations that are to patents from the matched treated patent's satellite location (right panel). The figures suggest that satellites are indeed an effective means of acquiring cutting-edge knowledge.

Table 4 formalizes the above analysis in a logit regression performed at the level of the patent citation (pooling together citations made by the treated and control patents) where the dependent variable is an index equal to 1 if the cited patent is (a third-party patent) from the location of the satellite. Column 1 presents the coefficients from the logit regression, whereas Column 2 presents the marginal effects. Consistent with the previous finding of a satellite effect, the coefficient on *treated* is positive and significant, indicating that treated patents are more likely to cite (third-party) patents from the satellite MSA. The coefficient on *age* is 0, indicating that for control patents there is no relationship between the age of the cited patent and whether that patent is located in the satellite MSA. On the other hand, for treated citations we find that the more recent the cited patent the more likely it is to be located in the satellite MSA, suggesting that newer knowledge is more likely to originate from the satellite MSA (the coefficient on the interaction between *age* and *treated* is negative). However, while the coefficient is significant at the

TABLE 4 Satellite effect and knowledge recentness

Dependent Variable: $I_{loc = satmsa}$	(1)	(2)
Age of cited patent	0.000	0.000
	(0.008)	(0.000)
Treated	0.634***	0.033***
	(0.203)	(0.011)
Age of cited patent*Treated	-0.017	-0.001
	(0.016)	(0.001)
Observations	277,557	277,557
R^2	0.0080	0.0080

Note: Logit regression at the level of the citation with standard errors clustered by firm (standard errors in parentheses). The starting sample consists of the 139,010 citations to third-party single-MSA U.S. patents made by the 35,034 treated patents from a primary R&D center of a two-location firm and the 142,256 citations to third-party single-MSA U.S. patents made by the associated 35,034 control patents that match by MSA, application year, and technology class. After removing the 1.3% of citations to patents less than a year old the final sample consists of 277,557 citations (the results are unchanged if these citations are kept). Pooling across treated and control citations, we estimate a logit regression at the level of the citation with standard errors clustered by firm and pooling all treated and control. The dependent variable is an indicator for whether the cited patent is located in the satellite MSA, *age* is the citation lag (application year of citing patent—application year of cited patent), and *treated* is 1 when the citation is made by a treated patent and 0 otherwise. Column (1) presents the estimated coefficients of the logit regression. The satellite effect is apparent in that treated citations are more likely to cite the satellite location and this is significant at 1%. For control patents, the age of the cited patent is unrelated to whether the cited patent is from the satellite MSA (coefficient of 0.000). For treated patents, the coefficient is negative (-0.017) suggesting that newer knowledge is more likely to have arrived from the satellite location. However, while this is significant at the 5% level when clustering by patent, it is not significant when clustering by firms (*p*-value of 0.27). Column (2) presents the marginal effects evaluated at the mean of the explanatory variables (8.0 for the age of the cited patent, 0.5 for treated, and 4.0 for their interaction). For treated patents, a one year increase in the age of the citation is associated with a decrease of 0.001 in the probability that the cited patent is in the satellite location (the dependent variable has mean 0.06).

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5% level when clustering standard errors by patent, it becomes insignificant when these are clustered by firms (*p*-value of 0.27), casting doubt on the relationship.

5.3 | Firm characteristics and satellite effect

In Section 5.1 evidence was presented that the satellite effect is greater in the technological sectors where the satellite has expertise. This could be because the satellite has more absorptive capacity in these sectors and is therefore better able to internalize external knowledge, because of supply linkages, because its engineers have social relationships with external inventors working in the same sectors, or because the satellite has recruited scientists and engineers who were previously working on similar technology at other local firms. To gain further insight into the mechanisms behind the satellite effect we examine what satellite and firm characteristics are associated with a larger satellite effect.

Consistent with the learning by hiring hypothesis, satellites with high shares of inventors that have previously patented at other local firms are associated with a larger satellite effect (Table 5, Column 1). This relationship is significant at the 1% level across all specifications. Conversely, the fraction of satellite inventors with previous patenting experience at other firms in different MSAs is inversely related to the magnitude of the satellite effect. This could be either because inventors from other regions do not have the social connections in the satellite MSA that would help access local knowledge. However, it could also be because by leveraging their networks with other MSAs they facilitate knowledge flow from those other MSAs to the primary R&D center (and hence the share of primary citations to the satellite MSA is lower). The other three variables measuring previous experience are not significant. Somewhat surprising is the fact that having a higher share of inventors that have previously patented in the primary R&D center is not associated with a larger satellite effect as that previous experience could facilitate knowledge transfer between the satellite and primary.

Other factors that could facilitate knowledge transfer between the satellite and primary are the intensity of collaboration between inventors from the two locations and the degree to which innovation in each location builds on innovations from the other location (cross-location self-citations). Columns 2 and 3 add these variables to the model. We find that the intensity of cross-location collaboration is positively associated with the magnitude of the satellite effect (this is significant at the 10% level). The intensity of cross-location self-citation is similarly positively correlated with the magnitude of the satellite effect and is significant at the 5% level. Column 4 presents the full model with all the variables. The only significant difference is that the cross-location collaboration measure is no longer significant at the 10% level, likely because the two cross-location measures explain much of the same variation.

TABLE 5 Firm characteristics and satellite effect

Dependent Variable: Difference between Proportion of Citations to Satellite MSA				
	(1)	(2)	(3)	(4)
Fraction of satellite inventors with:				
two or more satellite patents	0.015 (0.011)	0.004 (0.011)	0.008 (0.011)	0.002 (0.011)
previous patent in sat. msa for other org.	0.024*** (0.008)	0.025*** (0.007)	0.024*** (0.008)	0.025*** (0.007)
previous patent in prim. msa for prim.	0.002 (0.011)	0.001 (0.011)	0.001 (0.011)	0.000 (0.011)
previous patent in prim. msa for other org.	-0.002 (0.013)	-0.002 (0.013)	0.002 (0.013)	-0.001 (0.013)
previous patent in other msa for other org.	-0.021*** (0.008)	-0.020*** (0.008)	-0.021*** (0.008)	-0.020*** (0.008)
Cross-location collaboration intensity		0.018* (0.011)		0.017 (0.011)
Cross-location self-citation intensity			0.153** (0.068)	0.145** (0.069)
Geographic distance b-w prim. and sat.	0.002 (0.002)	0.003 (0.002)	0.002 (0.002)	0.003 (0.002)
Tech. distance b/w prim. and sat.	-0.019** (0.009)	-0.018* (0.009)	-0.016* (0.009)	-0.016* (0.009)
Tech. distance b/w prim. and sat. MSA	-0.019 (0.016)	-0.017 (0.015)	-0.021 (0.016)	-0.018 (0.015)
Year fixed effects	Yes	Yes	Yes	Yes
Observations	20,383	20,383	20,383	20,383
R ² -	0.0066	0.0071	0.0072	0.0077

Note: OLS regression at the level of the patent with standard errors clustered by firm (standard errors in parentheses). The starting sample consists of the 35,034 treated patents from a primary R&D center of a two-location firm and the associated control patents that match by MSA, application year, and technology class. The dependent variable is, for a given patent, the difference between the proportion of treated patent citations to third-party single-MSA U.S. patents that are to the satellite MSA and the proportion of control patent citations to third-party single-MSA U.S. patents that are to the same satellite MSA. Therefore, observations where either the treated or associated control patent does not make at least one citation to a third-party single-MSA U.S. patent are dropped. Firms with satellites having a high fraction of inventors that previously patented at other firms in the satellite MSA exhibit a larger satellite effect. Similarly, firms with a higher intensity of collaboration between primary and satellite inventors and with higher rates of cross-location self-citation also exhibit a larger satellite effect.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Among the controls, the only one that is significant is the degree of technological overlap between the patent portfolios of the primary and satellite R&D centers. As we would expect, firms with a smaller technological distance on average experience larger satellite effects. This is consistent with inventors working on similar technologies having more frequent interactions and stronger personal relationships, which facilitate the flow of knowledge across locations.

Taken together, these results suggest that firms can maximize the knowledge flow generated by the satellite by staffing satellites with experienced local hires and by ensuring that inventors in the satellite and primary R&D centers work collaboratively and/or on related technologies.

5.4 | Difference in differences

Given the potential issues associated with imperfect matching, and in the particular the possibility that the primary R&D center might choose to establish the satellite in the location from where it already is disproportionately sourcing knowledge, a more robust way to determine the existence of a satellite effect is to use the time dimension of the data and estimate the satellite effect using a difference-in-differences methodology. We run an OLS patent-level regression with standard errors clustered by firm. The dependent variable is the proportion of the patent's citations that cite the satellite location and the explanatory variables are *treated*, *postsatellite*, and their interaction.

TABLE 6 Difference in differences

Dependent Variable: Proportion of Patent's Citations to Satellite MSA			
	(1)	(2)	(3)
Treated	0.001 (0.005)	0.001 (0.005)	
<i>Postsatellite</i>	0.005 (0.008)	-0.009 (0.007)	-0.012 (0.009)
Treated* <i>Postsatellite</i>	0.026*** (0.008)	0.026*** (0.008)	0.025*** (0.007)
Year fixed effects	No	Yes	Yes
Firm fixed effects	No	No	Yes
Observations	12,649	12,649	12,649
R^2	0.0095	0.0170	0.1817

Note: OLS regression at the level of the patent with standard errors clustered by firm (standard errors in parentheses). Starting with the 35,034 treated patents (and their associated technology class controls), we drop any patent whose application year falls within our defined gap consisting of the five years prior to the satellite's first patent application (because in most instances the satellite probably existed a number of years prior to its first patent application). The variable *Postsatellite* is 0 for patents with application years prior to the gap and 1 for patents with application years after the gap. We keep all firms that have at least one primary patent in each of the pre- and postsatellite periods, resulting in a sample of 18,068 patents and 507 firms with primary R&D centers in 55 different MSAs and satellites in 126 different MSAs. The dependent variable is the proportion of a patent's citations to third-party single-MSA U.S. patents that are to patents in the satellite location. Therefore, observations where the patent does not make at least one citation to a third-party single-MSA U.S. patent are dropped in the regression.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 6 presents the results. The coefficient on *treated* is virtually identical to 0 suggesting that primary patents did not disproportionately cite the satellite location prior to the establishment of the satellite. The positive and significant (at the 1% level) coefficient on the interaction term shows that primary patents do disproportionately cite the satellite location after the establishment of the satellite. In particular, we find that after the establishment of the satellite the proportion of treated patent's citations to the satellite MSA increases by 2.6 percentage points (relative to a mean for the proportion of citations to the satellite MSA of 4.0 percentage points across all patents). This result is identical to our earlier finding (Table 3A) that 6.9% of treated patents cite the satellite MSA versus 4.3% for control patents (the difference between these being 2.6%).

Together, these results suggest that at the very least, the large majority of the satellite effect computed in Section 5.1 is due to knowledge flows through the satellite and not to firms establishing satellites in locations from where they already disproportionately source knowledge.

As an aside, although the results presented are for a gap size of five years between the pre- and postsatellite periods, the actual choice is largely inconsequential. As shown in Appendix F, the interaction term is significant for any choice of gap size. As we would expect given that in most cases, the satellite likely exists for at least a few years before the first patent application, the *treated* variable is positive and significant for gaps of two years or less and virtually identical to 0 and insignificant for any larger choice of gap.

Although the difference-in-differences results are more robust to a number of potential concerns, a spurious satellite effect could be found if at the same time as the firm establishes the satellite it undergoes other changes such as the headquarters' shift into a new sector. Table 7 performs the difference-in-differences analysis for the subset for firms whose primary R&D centers did not exhibit a technological shift toward the satellite MSA between the pre- and postsatellite periods.

The results are virtually unchanged. The coefficient on *treated* remains indistinguishable from 0 and the interaction term is significant and of a similar magnitude to before (larger with technological distance computed at the class level and slightly smaller when it is computed at the combined subclass and class level). The results offer strong evidence that the sectoral shift alternative cannot explain away the finding of an economically significant satellite effect.

Of course, our difference-in-differences results show that firms that choose to set up a satellite experience a satellite effect (the average treatment effect on the treated is significant). But to the extent that unobservables are behind this decision, the results are silent on the potential satellite effect for firms that do not choose to set up a satellite. In particular, unless our controls are truly identical to the treated patents in every way, we cannot interpret our difference-in-differences coefficient as the average treatment effect. This may explain why more firms do not establish satellites, given their apparently large benefits to firms that have them.

TABLE 7 Difference-in-differences—subset of firms with no technological shift

Dependent Variable: Proportion of Patent's Citations to Satellite MSA				
	Class Distance	Class Distance	Subclass Distance	Subclass Distance
Treated	0.000 (0.006)		−0.003 (0.006)	
Postsatellite	−0.012 (0.010)	−0.018 (0.017)	−0.014 (0.010)	0.002 (0.012)
Treated*Postsatellite	0.040*** (0.015)	0.038*** (.014)	0.022** (0.011)	0.020** (0.009)
Year fixed effects	Yes	Yes	Yes	Yes
Firm fixed effects	No	Yes	No	Yes
Observations	4,878	4,878	3,511	3,511
R ²	0.0246	0.1740	0.0160	0.1575

Note: OLS regression at the level of the patent with standard errors clustered by firm (standard errors in parentheses) for the subset of primary R&D center firms that did not exhibit a technological shift toward the portfolio of satellite MSA patents between the pre- and postsatellite periods. In Columns 1 and 2 technological distance is measured as the Euclidean distance computed at the level of the technology class. Of the 507 firms, 257 did not exhibit a shift toward the satellite MSA and these represent the final sample. In Columns 3 and 4, technological distance is measured at the level of the subclass within a patent class and the average over all classes is then computed. Of the 507 firms, 187 did not exhibit a shift towards the satellite MSA and these represent the final sample. Otherwise, the exact same methodology is used as in Table 6. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

6 | CONCLUSION

For all that has been written about our world getting smaller as a result of information technology, knowledge industries remain highly agglomerated and the acquisition of knowledge remains a vital concern. This paper suggests that R&D satellites can substitute for geographic collocation. In particular, the paper shows empirically that the principal research center of firms with remote satellites source a disproportionate amount of knowledge from the remote location. This satellite effect on knowledge acquisition is economically significant, representing between 32% and 62% of the knowledge flow premium associated with collocation. Furthermore, the results suggest that headquarters may be able to target specific knowledge by building satellite expertise in the sought technology areas. Further, while not conclusive, there is some evidence that the knowledge acquired through the satellite tends to be more recent. Yet the simple establishment of a remote R&D satellite may not by itself guarantee a flow of knowledge from the remote location. To reap the full benefits firms must also staff the satellite with scientists and engineers with experience at other local firms and ensure that inventors from the primary and satellite locations of the firm work on related projects and have opportunities to collaborate.

These results are important in a number of contexts. Using a methodology that is more robust to the imperfect matching critique, this paper indirectly provides additional evidence that knowledge flows are partially localized. The results also offer new insights into the process through which knowledge diffuses geographically, suggesting that multilocation firms are an important channel. Perhaps most importantly, the results have implications for firm strategy and in particular for how firms deploy R&D across geographies, how they distribute R&D activities across these locations, how they staff subsidiaries, and how much independence they award remote R&D satellites. Firms can strategically locate headquarters in lower-cost regions and still tap into remote knowledge through the deployment of R&D satellites in key regions. In order to maximize the amount of outside knowledge acquired by the firm, R&D activities should be geographically distributed so that satellites focus on areas of local technological strength. Knowledge acquisition can be targeted by choosing a satellite's area of specialization and by staffing the satellite with local inventors that have experience in those sectors. Independent satellites may more easily embed themselves in local networks and receive local knowledge, but the tradeoff is that this knowledge may not be transferred to other centers within the firm. The results suggest that closer links between primary and satellite research centers are beneficial for remote knowledge sourcing. These links can potentially be fostered by organizing internal conferences, running cross-location training sessions, encouraging job rotation across locations, by facilitating business travel so as to encourage face-to-face interaction, or by investing in communication tools like video conferencing.

Although the current findings have significant implications, several issues remain to be explored. First, we need to develop a better understanding of the mechanisms behind the satellite effect, this paper having taken only initial steps in this direction. Second, we need to better understand how the satellite effect translates to the more complex international setting, paying particular attention to the effect of developmental, cultural, and institutional differences on the satellite effect. As well, we may

ask whether location-specific characteristics such as knowledge stocks impact the satellite effect and whether this effect differs across industries. We also need to determine whether increased knowledge sourcing from a remote location translates into more valuable innovations for the firm? And if so, is the establishment of a satellite the most efficient way to gather remote knowledge? Clearly, further management and policy prescriptions will emerge as we increase our understanding of R&D satellites as knowledge gathering entities.

NOTES

- ¹ Mechanisms for accessing remote knowledge will be particularly important in cases where relevant knowledge is geographically dispersed or where the majority of knowledge resides in a geographical location that is not the location of the firm's headquarters.
- ² Examining the flow of knowledge between countries that are treated as monolithic would significantly bias downward any estimate of the satellite effect.
- ³ Other ways to learn remotely include learning by importing (Coe & Helpman, 1995; Keller, 1998; MacGarvie, 2006) and learning by exporting (Bernard & Jensen, 1999; Clerides, Lach, & Tybout, 1998).
- ⁴ To the extent that the addition of examiner citations does not follow a particular pattern that is correlated to variables of interest, they simply add noise, attenuating the results and biasing against a significant finding.
- ⁵ I extend these data to 2002 using inventor location data obtained from Alex Oettl and Bronwyn Hall's update (<http://www.econ.berkeley.edu/~bhhall/patents.html>)
- ⁶ Assignee names are not standardized, which results in the same firm often appearing under different assignee names. For example, the NBER Patent Data Project reports that "there are over 100 different spellings, misspellings, abbreviations, etc., for the assignees of patents assigned to IBM." Fortunately, James Bessen and the NBER Patent Data Project have gone through the painstaking task of allocating a unique assignee code to each firm and linking this code to all the firm's patents. For a description of the methodology visit <http://www.nber.org/~jbessen/matchdoc.pdf>
- ⁷ Cited patents need not be assigned to U.S. nongovernment organizations.
- ⁸ These represent 59% of all multilocation firms in the data set. As shown in Appendix B, the main results are similar for firms with more locations.
- ⁹ The "amount" of treatment is not explicitly considered in the paper, in part because the size of the satellite (as measured by patent count) does not significantly impact the magnitude of the satellite effect. Regressing the computed satellite effect on satellite size shows the coefficient to be positive but not statistically significant.
- ¹⁰ With the exception that patents belonging to assignees that do not have exactly two locations are no longer discarded.
- ¹¹ For example, the probability of finding a match at the subclass level could be correlated with the firm's ability or motivation to acquire knowledge through the satellite. For example, matches at the subclass level are more likely to be found in MSAs that are important technology clusters. Firms in such clusters may learn less through the satellites since knowledge abounds all around them.
- ¹² This variable builds on Rosenkopf and Almeida's Euclidean distance but computes the distance at the technology subclass level of disaggregation to complement the technology class matching methodology that is also being employed. In particular, geometric distance is computed at the subclass level for each technology class and the final variable is the average across all technology classes. Formally, technological distance for firm i is constructed as: $Tech\ Dist\ Prim\ Satloc_i = \frac{1}{N} \sum_{c=1}^N \sqrt{\sum_{s=1}^M (prim_{i,c} - satmsa_{i,c})^2}$, where $prim_{i,c}$ is the fraction of technology class c primary center patents that are in technology subclass i and $satmsa_{i,c}$ is the fraction of technology class c satellite location patents that are in technology subclass i . N is the total number of technology classes, while M is the total number of technology subclasses in any technology class c . Note: using Rosenkopf and Almeida's Euclidean distance at either the class or subclass level as alternative measures of distance does not affect the results.
- ¹³ Letting $H_0: P_T = P_C$ and $H_1: P_T > P_C$, I calculate the t -statistic as: $t = (\widehat{P}_T - \widehat{P}_C) / \sqrt{\frac{\widehat{P}_T(1-\widehat{P}_T)}{n_T} + \frac{\widehat{P}_C(1-\widehat{P}_C)}{n_C}}$.
- ¹⁴ Imperfect matching biases the collocation results in the presence of any agglomeration. In the case of the satellite effect, we require that, in addition to agglomeration, firms establish subsidiaries in locations that specialize in the same technological subsectors as the primary R&D center.
- ¹⁵ As we might expect, we find a similar downward pattern for collocated knowledge flows: the proportion of treated patent citations that are to third-party patents from the same (primary) location drops from 23% for one-year-old cited patents to 16.5% for 15-year-old patents. The opposite is true for knowledge from locations that are neither the satellite's nor the headquarters'. Note that 68.5% of one-year-old cited patents are from those locations versus 77% for 15-year-old patents (figures are available from the author upon request).

REFERENCES

- Agrawal, A., Cockburn, I., & McHale, J. (2006). Gone but not forgotten: Knowledge flows, labor mobility, and enduring social relationships. *Journal of Economic Geography*, 6(5), 571–591.
- Alcacer, J., & Chung, W. (2002). Knowledge seeking and location choice of foreign direct investment in the United States. *Management Science*, 48(12), 1534–1554.
- Alcacer, J., & Gittelman, M. (2006). Patent citations as a measure of knowledge flows: The influence of examiner citations. *The Review of Economics and Statistics*, 88(4), 774–779.

- Almeida, P. (1996). Knowledge sourcing by foreign multinationals: Patent citation analysis in the U.S. semiconductor industry. *Strategic Management Journal*, 17, 155–165.
- Almeida, P., & Kogut, B. (1999). Localization of knowledge and the mobility of engineers in regional networks. *Management Science*, 45(7), 905–917.
- Arrow, K. (1962). Economic welfare and the allocation of resources for invention. In R. R. Nelson (Ed.), *The rate and direction of inventive activity: Economic and social factors* (Vol. 13, pp. 609–625) of National Bureau of Economic Research Special Conference Series. New Jersey: Princeton University Press.
- Audretsch, D. B., & Feldman, M. P. (1996). R&D spillovers and the geography of innovation and production. *American Economic Review*, 86(3), 630–640.
- Bernard, A., & Jensen, J. (1999). Exceptional exporter performance: Cause, effect, or both? *Journal of International Economics*, 47(1) 1–25.
- Boutellier, R., Gassmann, O., & Von Zedtwitz, M. (2008). *Managing global innovation: Uncovering the secrets of future competitiveness* (3rd ed.). Berlin: Springer.
- Branstetter, L. (2001). Are knowledge spillovers international or intranational in scope? *Journal of International Economics*, 53(1), 53–79.
- Branstetter, L. (2006). Is foreign direct investment a channel of knowledge spillovers? Evidence from Japan's FDI in the United States. *Journal of International Economics*, 68(2), 325–344.
- Breschi, S., & Lissoni, F. (2009). Mobility of skilled workers and co-invention networks: An anatomy of localized knowledge flows. *Journal of Economic Geography*, 9, 439–468.
- Cantwell, J. (1993). The internationalization of technological activity and its implications for competitiveness. In O. Granstrand, H. Hakanson, & S. Sjolander (Eds.), *Technological management and international business* (pp. 137–162). Chichester: John Wiley.
- Clerides, S., Lach, S., & Tybout, J. (1998). Is learning by exporting important? Micro-dynamic evidence from Colombia, Mexico, and Morocco. *Quarterly Journal of Economics*, 113(3), 903–948.
- Coe, D., & Helpman, E. (1995). International R&D spillovers. *European Economic Review*, 39(5), 859–887.
- Cohen, W., & Levinthal, D. (1990). Absorptive capacity: A new perspective on learning and innovation. *Administrative Science Quarterly*, 35(1), 128–152.
- Duguet, E., & MacGarvie, M. (2005). How well do patent citations measure flows of technology? Evidence from French innovation surveys. *Economics of Innovation and New Technology*, 14(5), 375–393.
- Fabrizio, K. (2007). University patenting and the pace of industrial innovation. *Industrial and Corporate Change*, 16(4), 505–534.
- Feinberg, S., & Gupta, A. (2004). Knowledge spillovers and the assignment of R&D responsibilities to foreign subsidiaries. *Strategic Management Journal*, 25(8-9), 823–845.
- Frost, T. S. (2001). The geographic sources of foreign subsidiaries' innovations. *Strategic Management Journal*, 22(2), 101–123.
- Hall, B. H., Jaffe, A. B., & Trajtenberg, M. (2001). *The NBER patent citation data file: Lessons, insights and methodological tools*. NBER Working paper 8498.
- Jaffe, A. B., & Trajtenberg, M. (2002). *Patents, citations, and innovations: A window on the knowledge economy*. Cambridge: MIT Press.
- Jaffe, A. B., Trajtenberg, M., & Henderson, R. (1993). Geographic localization of knowledge spillovers as evidenced by patent citations. *Quarterly Journal of Economics*, 108(3), 577–598.
- Keller, W. (1998). Are international R&D spillovers trade-related? Analyzing spillovers among randomly matched trade partners. *European Economic Review*, 42(8), 1469–1481.
- Kogut, B., & Zander, U. (1992). Knowledge of the firm, combinative capabilities, and the replication of technology. *Organization Science*, 3(3), 383–97.
- Kogut, B., & Zander, U. (1993). Knowledge of the firm and the evolutionary theory of the multinational corporation. *Journal of International Business Studies*, 24(4), 625–645.
- Krugman, P. (1991). *Geography and trade*. Gaston Eyskens Lecture Series. Cambridge: MIT Press.
- Kuemmerle, W. (1999). The drivers of foreign direct investment into research and development: An empirical investigation. *Journal of International Business Studies*, 30(1), 1–24.
- MacGarvie, M. (2006). Do firms learn from international trade? *Review of Economics and Statistics*, 88(1), 46–60.
- March, J. G. (1991). Exploration and exploitation in organizational learning. *Organization Science*, 2(1), 71–87.
- Marshall, A. (1890). *Principles of economics*. London: MacMillan.
- Oettl, A., & Agrawal, A. (2008). Labor mobility and knowledge flow externalities. *Journal of International Business Studies*, 39(8) 1242–1260.
- Rogers, E., & Larsen, J. (1984). *Silicon Valley fever*. New York: Basic Books.
- Romer, P. (1986). Increasing returns and long-run growth. *Journal of Political Economy*, 94(5), 1002–1037.
- Romer, P. (1990). Endogenous technological change. *Journal of Political Economy*, 98(5), S71–S102.
- Rosenkopf, L., & Almeida, P. (2003). Overcoming local search through alliances and mobility. *Management Science*, 49(6), 751–766.

- Singh, J. (2007). Asymmetry of knowledge spillovers between MNCs and host country firms. *Journal of International Business Studies*, 38(5), 764–786.
- Singh, J., & Agrawal, A. (2010). Recruiting for ideas: How firms exploit the prior inventions of new hires. *Management Science*, 57(1), 129–150.
- Singh, J., & Fleming, L. (2010). Lone inventors as sources of breakthroughs: Myth or reality? *Management Science*, 56(1), 41–56.
- Song, J., Almeida, P., & Wu, G. (2003). Learning-by-hiring: When is mobility more likely to facilitate interfirm knowledge transfer? *Management Science*, 49(4), 351–365.
- Thompson, P. (2006). Patent citations and the geography of knowledge spillovers: Evidence from inventor- and examiner-added citations. *Review of Economics and Statistics*, 88(2), 383–389.
- Thompson, P., & Fox-Kean, M. (2005). Patent citations and the geography of knowledge spillovers: A reassessment. *American Economic Review*, 95(1), 450–460.
- Von Hippel, E. (1988). *Sources of innovation*. New York: Oxford University Press.

How to cite this article: Blit J. Learning remotely: R&D satellites, intra-firm linkages, and knowledge sourcing. *J Econ Manage Strat.* 2017;26:757–781. <https://doi.org/10.1111/jems.12213>

APPENDIX A: MULTIPLE MSA PATENTS

The number of utility patents with a U.S. firm as the assignee and at least one inventor residing in a U.S. MSA is 956,807. Of these, 79.89% had all the inventors located in the same MSA, 1.85% had at least one inventor located outside the U.S., 9.34% had at least one inventor living in the U.S. but outside any MSAs, and 11.02% had at least two inventors living in different MSAs (some patents fall into more than one of the categories).

Table A1 below shows separate summary statistics for the 764,348 single-MSA patents and the 192,459 multilocation patents. As expected, multilocation patents have on average almost twice as many inventors because the more inventors on a patent, the more likely that at least two are in different locations. Because patents with more inventors tend to both receive more citations and make more claims, it is not surprising that multilocation patents also exhibit higher means along these dimensions. Lastly, multilocation patents make more citations than single-MSA patents, perhaps because each inventor brings a different piece of prior knowledge. The *t*-test statistic (Column 5) shows that the sample means are statistically different across the two samples.

However, the fact that the two samples are different does not seem to pose a significant problem for the analysis in question. The computed satellite effect is largely unchanged when we include multilocation patents. We can include multi-MSA patents in the analysis in one of three ways: allowing treated (and control) patents to be multi-MSA, allowing citations to be to multi-MSA patents, and using multi-MSA patents to determine the set of two-location firms. I discuss each in turn.

TABLE A1 Comparing single MSA patents with multilocation patents

	Mean	Std. Dev.	Min	Max	<i>t</i> -Stat
Number of Inventors					
Single MSA patents	1.82	1.12	1	23	
Multilocation patents	3.31	1.68	2	41	−370
Citations received					
Single MSA patents	7.33	12.01	0	1,069	
Multilocation patents	6.77	12.34	0	1,046	17.82
Claims					
Single MSA patents	13.82	11.52	1	706	
Multilocation patents	15.35	12.92	1	393	−39.09
Citations made					
Single MSA patents	10.53	14.16	0	745	
Multilocation patents	12.79	18.00	0	770	−51.34

TABLE A2 Satellite effect when including multilocation treated patents—technology class matching

	Satellite Effect All Tech Classes	Satellite Effect Sat Tech Classes
% Treated matching	6.8	2.5
% Controls matching	4.6	1.0
<i>t</i> -Statistic	28.63	34.61
Satellite effect	1.48	2.49
<i>N</i> - Treated citations	182,257	182,257
<i>N</i> - Control citations	186,369	186,369

TABLE A3 Satellite effect when including citations to multi-MSA patents—technology class matching

	Satellite Effect All Tech Classes	Satellite Effect Sat Tech Classes
% Treated matching	4.9	1.8
% Controls matching	3.2	0.7
<i>t</i> -Statistic	31.45	38.63
Effect	1.53	2.81
<i>N</i> - Treated citations	259,299	259,299
<i>N</i> - Control citations	267,048	267,048

(1) Including multi-MSA treated patents

Multilocation treated patents were not included in the analysis to rule out the possibility that remote knowledge might be acquired through having a coinventor located in the remote location (instead of through a satellite effect). If we include as treated patents all patents with at least one inventor in the primary R&D center, our sample size increases substantially but the results are not that different (the computed satellite effect is now 1.48 instead of 1.60). This is shown in Table A2 which reproduces the first two columns of Table 3A.

(2) Including citations to multi-MSA patents

In the main analysis, we only consider citations by the treated and control patents to third-party single-MSA U.S. patents so as to unambiguously assign a location to each patent and so as not to inadvertently introduce a bias. For example, if there are lots of third-party patents with inventors in both the primary and satellite MSAs (perhaps because they are the two locations where a given technology is being developed) the primary might be accessing the knowledge through the collocated (or nearby) inventor and we might attribute it to the satellite effect.

If we consider all citations and treat as a match any third-party cited patent with at least one inventor in the satellite location we obtain the results presented in Table A3.

Although the matching percentage is now lower for both the treated and controls (citations to non-U.S. patents are now included), the magnitude of the satellite effect is similar (1.53 vs. 1.60).

(3) Finding two-location firms with multi-MSA patents

A third impact of restricting the analysis to single location firms is that it affects which firms we define as having two locations. In the main analysis, primary and satellite R&D centers are defined according to the location of patents with inventors that are all from that location. This entails that only satellites that have the ability to innovate separately from the primary R&D centers (no inventors from the primary) are included in the analysis. This constitutes a more robust definition of a satellite where only more “significant” satellites that can innovate in isolation are considered.

Defining satellites as locations with at least one firm inventor significantly increases the number of firms with inventive activities in multiple locations. The number of two-location firms increases from 4,610 to 9,642. Of the 4,610 two-location firms in our main sample, only 2,680 remain two-location firms in the sample with the more permissive definition of a satellite,

TABLE A4 Satellite effect when primary and satellite locations are found using multilocation patents—technology class matching

	Satellite Effect All Tech Classes	Satellite Effect Sat Tech Classes
% Treated matching	4.9	1.2
% Controls matching	3.4	0.3
<i>t</i> -Statistic	18.77	25.21
Satellite effect	1.46	3.81
<i>N</i> - Treated citations	123,751	123,751
<i>N</i> - Control citations	108,182	108,182

the remainder now has more than two locations. It follows that the majority of the 9,642 two-location firms in the new sample were deemed to have only one location in the original sample.

Using the more permissive definition of a satellite yields similar results, though there is an attenuation bias due to increased noise. Only a little over half of the original two-location firms remain and the majority of the sample now consists of “marginal” satellites that under the original definition would not have been considered satellites at all. As such, it is not surprising that the computed satellite effect is smaller as shown in Table A4. Surprisingly though, the satellite effect in technology classes where the satellite is active is now larger.

APPENDIX B: MULTILLOCATION FIRMS

We determine whether the computed satellite effect is similar for firms with multiple locations. For comparability with the results in Table 3A, we conduct the same analysis as in the main paper for each satellite of a multilocation firm and pool the results. For example, in the case of three-location firms, we first count how many of the treated and control patents citations cite the first satellite location, then how many of the same treated and control patents' citations cite the second satellite location, sum the two, and divide by twice the total number of citations. If a satellite effect indeed exists, this will bias downward the computed magnitude of the effect as a treated citation cannot at once be citing both satellite locations. In the limit, a firm with a satellite in each MSA would exhibit no satellite effect.

The results for firms with more than two locations are comparable, though somewhat smaller as expected. These are shown in Table B1 where the first column of Table 3A is reproduced for firms with different numbers of locations.

APPENDIX C: COMPARISON OF TREATED AND CONTROL PATENTS

The sample constructed consists of 35,034 treated patents (patents generated in the primary R&D center of a two-location firm) and an equal number of control patents that are matched by application year, technology class, and MSA. If multiple patents match the treated patent along these dimensions then the patent belonging to the multilocation firm with the fewest locations is chosen. That is, priority is given to patents of firms with two locations, then three locations, four locations, etc. and patents from single-location firms are chosen as a last resort. Given our matching rule, to the extent that larger firms with a broader geographical presence are better at gathering remote knowledge, using firms with more than two locations as controls should attenuate our results. This is indeed what we find.

Table C1 compares some observable characteristics of the 35,034 treated and control patents. The most notable difference is that control patents on average belong to firms that have more locations. However, control patents also tend to have slightly

TABLE B1 Satellite effect across all technology classes by number of locations—technology class matching

	2 Locations	3 Locations	4 Locations	5 Locations	6 Locations
% Treated matching	6.8	4.9	5.3	4.3	3.9
% Controls matching	4.5	3.5	3.5	3.2	2.9
<i>t</i> -Statistic	27.08	16.62	23.51	15.10	15.93
Satellite effect	1.52	1.38	1.50	1.37	1.35
<i>N</i> - Treated citations	141,702	119,500	142,367	111,147	153,644
<i>N</i> - Control citations	142,637	128,420	159,400	127,613	174,657

TABLE C1 Characteristics of treated and control patents ($N = 35,034$)

	Mean	Std. Dev.	Min	Max	<i>t</i> -Stat
Number of Inventors					
Treated patents	1.91	1.19	1	21	
Control patents	1.82	1.16	1	20	9.95
Citations received					
Treated patents	7.36	13.08	0	304	
Control patents	6.97	12.03	0	285	4.09
Claims					
Treated patents	15.11	12.24	1	250	
Control patents	14.76	11.88	1	220	3.15
Citations made					
Treated patents	12.32	18.45	0	430	
Control patents	12.37	17.63	0	430	0.32
Assignee locations					
Treated patents	2	0	2	2	
Control patents	3.44	5.59	1	101	48.20

TABLE C2 Characteristics of treated and control patents for subsample with two-location firm controls ($N = 17,856$)

	Mean	Std. Dev.	Min	Max	<i>t</i> -Stat
Number of Inventors					
Treated patents	2.02	1.27	1	14	
Control patents	1.89	1.21	1	21	9.78
Citations received					
Treated patents	7.16	13.66	0	285	
Control patents	6.99	12.55	0	285	1.19
Claims					
Treated patents	15.39	12.54	1	220	
Control patents	15.27	12.29	1	220	0.68
Citations made					
Treated patents	13.04	20.51	0	430	
Control patents	13.53	20.89	0	430	2.22
Assignee locations					
Treated patents	2	0	2	2	
Control patents	2	0	2	2	-

fewer inventors, receive fewer citations, and make fewer claims. These other differences, though, are largely accounted for by the broader geographical presence of control patent firms. If we restrict the sample to the 17,856 patents for which a control with a two-location assignee was found, our treated and control sample of patents look more similar as shown in Table C2. The only important remaining difference between treated and controls seems to be that treated patents have on average more inventors (2.02) than control patents (1.89). The number of citations made is now also statistically different between the treated and control patents, though it was not for the full sample of patents.

If we compute the satellite effect for this restricted sample of patents with two-location firm controls, we obtain the results shown in Table C3. The fraction of treated citations to the satellite location remains almost the same at 6.8% (vs. 6.9%) although the fraction of control citations to the satellite location has decreased from 4.5 to 4.0. This is what we would expect to observe if firms with a broader geographical presence are better at gathering remote knowledge. The computed satellite effect is slightly larger at 1.71 (from 1.60). Thus, using controls from firms with a broader geographical presence seems to attenuate the results.

TABLE C3 Satellite effect for subsample with two-location firm controls—technology class matching

	Satellite Effect All Tech Classes	Satellite Effect Sat Tech Classes
% Treated matching	6.8	3.0
% Controls matching	4.0	0.9
<i>t</i> -Statistic	25.58	32.36
Satellite effect	1.71	3.56
<i>N</i> - Treated citations	82,080	82,080
<i>N</i> - Control citations	85,956	85,956

APPENDIX D: INVENTOR SELF-CITATIONS TO PREVIOUS SATELLITE MSA PATENTS

One of the roles of the satellite might be to identify talented scientists and engineers in their region so they can be hired by the primary R&D unit. To the extent that this might be a common occurrence, part of the observed satellite effect could be due to these inventors citing their previous work at other firms in the satellite MSA.

Of a total 27,325 distinct inventors in the primary R&D center of our 3,781 two-location firms, there are 723 that were previously patented while working in the satellite MSA for another entity. These inventors were involved in 3,901 of our 35,034 treated patents. To ensure that the computed satellite effect is not only driven by citations from these inventors to their previous work in the satellite MSA, we compute the satellite effect ignoring any citations made by treated patents that involve these inventors to that inventor's previous patents. In total we remove 632 such citations. The results are shown in Table D1. The proportion of treated patent citations that are to third-party patents in the satellite location is now slightly lower at 6.8% (vs. 6.9%) resulting in a slightly smaller satellite effect of 1.56 (vs. 1.60). It would therefore seem that although self-citations to previous work do play a role, it is but a minor one.

APPENDIX E: ANALYZING INVENTOR- VERSUS EXAMINER-ADDED CITATIONS

As a further robustness check, the satellite effect is computed using an identification strategy that follows Thompson (2006). In particular, for the set of patents developed in the primary R&D center of two-location firms (the set of patents that the main analysis labels as "treated"), I compare the fraction of inventor-added citations that cite a third-party patent from the location of the satellite with the fraction of examiner-added citations that cite a third-party patent from the location of the satellite. The key assumption underlying this approach is that inventor-added citations more likely represent true knowledge flows than examiner-added citations.

Data on which citations are added by inventors and examiners is only available for patents granted in 2001 and later. As such, the analysis presented below is for the 7,443 primary R&D center patents granted in 2001 and 2002 (these represent 14.3% of all primary R&D center patents in my full sample). After removing self-citations there remain 43,912 inventor-added citations and 12,638 examiner-added citations on which the analysis is performed. Table E1 presents the results.

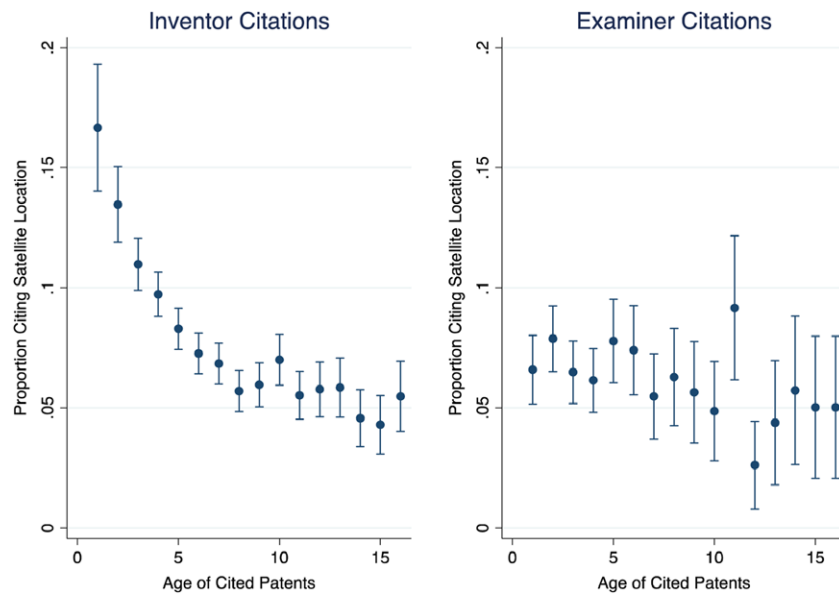
The computed satellite effects are statistically significant, though smaller than those obtained in the main analysis (Tables 3A and 3B). Across all technology classes, the satellite effect is 1.21 (vs. 1.40 with subclass matching) and in technology classes where the satellite has patented the effect is 1.66 (vs. 1.98). Although this is a significant difference, it could be

TABLE D1 Satellite effect after removing previous satellite MSA inventor self-citations—technology class matching

	Satellite Effect All Tech Classes	Satellite Effect Sat Tech Classes
% Treated matching	6.8	2.5
% Controls matching	4.3	0.9
<i>t</i> -Statistic	28.13	32.15
Satellite effect	1.56	2.74
<i>N</i> - Treated citations	138,378	138,378
<i>N</i> - Control citations	142,256	142,256

TABLE E1 Satellite effect—inventor versus examiner-added citations

	H1 Satellite Effect All Tech Classes	H2 Satellite Effect Sat Tech Classes
% Inventor-added matching	7.7	3.4
% Examiner-added matching	6.3	2.1
<i>t</i> -Statistic	5.30	8.86
Satellite effect	1.21	1.66
<i>N</i> – Inventor-added citations	43,912	43,912
<i>N</i> – Examiner-added citations	12,638	12,638

**FIGURE E1** Proportion of citations citing satellite location by age of cited patents for both inventor-added (left panel) and examiner-added (right panel) citations

[Color figure can be viewed at wileyonlinelibrary.com]

Note: The sample consists of 43,912 inventor-added citations and 12,638 examiner-added citations. The left panel shows the proportion of inventor-added citations that are to third-party patents in the satellite location by age of the cited patents. Although approximately 17% of one-year-old cited patents are from the satellite location, fewer than 5% of 15-year-old cited patents are from that same satellite location. No such pattern exists for the proportion of examiner-added citations to patents in the satellite location (right panel). The error bars represent two standard deviations on either side of the mean (and hence the 95% confidence interval). The figures suggest that satellites might be an effective means of acquiring cutting-edge knowledge.

due to having restricted the sample to patents granted in 2001 and 2002. As an aside, if we compute the satellite effect with subclass matching for the same restricted sample the computed satellite effect is nearly identical at 1.17.

I also revisit the result on the satellite effect and knowledge recentness. Consistent with our previous findings, for inventor-added citations a disproportionately high fraction of citations to recent patents are citing the satellite location (Figure E1, Panel A). The same pattern is not present for examiner-added citations (Panel B). The figures are virtually identical to those presented in Figure 3 for treated and control patents, respectively. A logistic regression such as the one whose results are presented in Table 4 finds that the coefficient on the interaction of “inventor added” and “age of cited patent” is negative, as expected, but like in the main results it becomes insignificant when the standard errors are clustered by firm.

APPENDIX F: DIFFERENCE IN DIFFERENCES RESULTS FOR ALTERNATIVE GAPS

Dependent Variable: Proportion of Patent's Citations to Satellite MSA						
Gap size	0	0	2	2	4	4
Treated	0.015*** (0.004)		0.010** (0.005)		0.003 (0.004)	
Postsatellite	0.004 (0.006)	-0.016*** (0.005)	0.004 (0.008)	-0.020** (.009)	-0.012 (0.007)	-0.015* (0.008)
Treated*Post-Sat	0.011* (0.006)	0.025*** (0.005)	0.021*** (0.007)	0.031*** (0.006)	0.028*** (0.008)	0.029*** (0.007)
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Firm fixed effects	No	Yes	No	Yes	No	Yes
Observations	34,187	34,187	21,069	21,069	14,566	14,566
R ²	0.0090	0.2190	0.0159	0.1861	0.0191	0.1974
Gap size	6	6	8	8	10	10
Treated	-0.003 (0.005)		-0.010 (0.008)		-0.006 (0.008)	
Postsatellite	-0.005 (0.008)	-0.006 (0.009)	-0.007 (0.012)	-0.021 (0.015)	-0.020 (0.014)	-0.042** (0.0094)
Treated*Post-Sat	0.025*** (0.008)	0.021*** (0.007)	0.038*** (0.012)	0.027*** (0.009)	0.041*** (0.014)	0.035*** (0.013)
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Firm fixed effects	No	Yes	No	Yes	No	Yes
Observations	10,248	10,248	6,728	6,728	4,024	4,024
R ²	0.0092	0.1484	0.0158	0.1682	0.0264	0.1981

Note: OLS regression at the level of the patent with standard errors clustered by firm (standard errors in parentheses). Starting with the 35,034 treated patents (and their associated technology class controls), we drop any patent whose application year falls within our defined gap consisting of some years prior to the satellite's first patent application (because in most instances the satellite probably existed a number of years prior to its first patent application). The variable *Postsatellite* is 0 for patents with application years prior to the gap and 1 for patents with application years after the gap. We keep all firms that have at least one primary patent in each of the pre- and postsatellite periods. The dependent variable is the proportion of a patent's citations to third-party single-MSA U.S. patents that are to patents in the satellite location. Therefore, observations where the patent does not make at least one citation to a third-party single-MSA U.S. patent are dropped in the regression.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.