Foreign R&D satellites as a medium for the international diffusion of knowledge

Joel Blit*
Department of Economics
University of Waterloo
200 University Ave. West
Waterloo, ON, Canada N2L 3G1
jblit@uwaterloo.ca

October 2017

ABSTRACT: We examine the extent to which foreign R&D satellites of multinational firms act as a medium for the international diffusion of knowledge. Using patents from the United States Patent and Trademark Office, we compare the frequency with which headquarters patents are cited by third-party firms in the satellite's host country relative to a control group of patents, and this both before and after the establishment of the satellite (using a difference in differences approach). The results suggest that the satellite increases the flow of knowledge from the multinational's headquarters to firms in the satellite's host country. This satellite effect on knowledge diffusion is largest in host countries and sectors with strong but not world-class capabilities that have both the motivation and absorptive capacity to learn from foreign parties. The findings also suggest that knowledge diffusion is greatest when satellites are staffed with inventors that have previously either patented with other local firms (thus having stronger local social networks) or with the headquarters (thus having headquarters knowledge).

Keywords: knowledge flows, diffusion, innovation, foreign direct investment, patents
JEL Classification: O3; F2; L2

* I am grateful to Ajay Agrawal, Ignatius Horstmann, Joanne Oxley, and Daniel Trefler for their indispensable guidance, helpful comments, and discussions. David Wavrock provided excellent research assistance gathering data on subsidiaries. Gilles Duranton offered valuable suggestions as a discussant for an early version of this paper. I am also indebted to an editor and two anonymous referees for their contructive comments. I gratefully acknowledge funding from the Social Sciences and Humanities Research Council of Canada, The Institute for New Economic Thinking, and the Centre for International Governance Innovation. Any errors and omissions are my own.

1. Introduction

Multinational enterprises (MNEs) are increasingly conducting R&D abroad, and especially in emerging markets. As of 2010, *Fortune 500* companies had 98 R&D facilities in China and 63 in India.[1] MNEs are establishing themselves in emerging countries not only to access their already large and growing markets but also to tap into the multitude of qualified science, engineering, and computer science graduates that these countries produce. And by doing so, multinationals may be promoting development in their host country. Indeed, many countries, emerging and developed alike, have implemented policies to attract foreign MNEs operating in knowledge-intensive sectors in the belief that the activities of their subsidiaries will generate spillover benefits to local firms.

This paper examines the potential knowledge-bridging effects of multinational firms' R&D satellites. We define R&D satellites as centres outside of the firm's home country where the firm conducts R&D activities (as evidenced by patents). Microsoft, for example, has subsidiaries in more than 100 countries around the world, most of which are sales, marketing and services offices. By contrast, they have relatively few R&D satellites, the largest of which is the Microsoft Research Lab – Asia,[1] which is located in Beijing and employs more than 200 researchers and developers.[2] The establishment of this satellite R&D centre may have facilitated the flow of knowledge from Microsoft's Redmond headquarters, to firms such as Huawei and Alibaba, as evidenced by citations made by these firms' patents to Microsoft Redmond patents. More generally, in the five years prior to the 1998 establishment of the Beijing R&D satellite, Microsoft headquarters patents received but a single citation from patents generated in China, compared to 16 in the five years that followed. This, of course, could be due to the explosion of innovation in China, or even to the rising number of Microsoft headquarters patents. Our more formal analysis, therefore, directly addresses these and other factors, comparing the frequency with which headquarters patents are cited by third-party firms in the satellite's host country relative to a control group of patents, and this both before and after the establishment of the satellite.

---

[1] *The Economist*, "Special report on innovation in emerging markets," April 17, 2010. p. 4.
[2] https://www.microsoft.com/en-us/research/lab/microsoft-research-asia/

Numerous scholars have examined whether FDI confers spillover benefits. Particularly salient is the work of Haddad and Harrison (1993) and Aitken and Harrison (1999), who use micro-level panel data from Morocco and Venezuela to determine whether FDI affected the productivity of domestic firms. More recently, Haskel et al. (2002), Keller and Yeaple (2003), Javorcik (2004), and Kee (2015) have examined the productivity spillover effects of FDI in the U.K., the U.S., Lithuania, and Bangladesh, respectively. Overall, the evidence is mixed, in part because productivity gains as a result of FDI are difficult to measure in the face of numerous possible confounding factors (Gorg and Strobl 2001, Gorg and Greenaway 2004).

Perhaps, then, the way forward is to more closely examine specific mechanisms through which FDI may increase productivity in the host country. For example, it might be the case that productivity increases are achieved through the fostering of indigenous innovation by the presence of multinational subsidiaries. This is consistent with what we observe in China over the period 1985-2001, where patenting activity by foreign multinationals is associated with subsequent indigenous patenting in the same sector (Figure 1). This is also consistent with the findings of Girma, Gong, and Görg (2008), among others. Foreign R&D satellites of multinationals may encourage indigenous innovation in their host country by generating knowledge spillovers that provide indigenous firms with crucial inputs into the innovative process. In particular, satellites may be an important knowledge link between indigenous firms and the outside world.

Knowledge spillovers have received much attention given their central role in economic growth (Romer 1990). Numerous scholars have examined the extent to which these knowledge spillovers are localized; while the results are mixed, the consensus is that there is some degree of localization (Jaffe, Trajtenberg, and Henderson 1993, Audretsch and Feldman 1996, Keller 2002, Thompson and Fox-Kean 2005, Thompson 2006). How then, do firms in emerging countries get access to knowledge generated in the developed world? Coe and Helpman (1995), Keller (1998), and MacGarvie (2006) examine whether the import of manufactured goods can serve as a channel for knowledge flow. More recently, van Biesebroeck (2005), De Loecker (2007), and Lileeva and Trefler (2010) have also examined the possibility of learning by exporting.

A different way that firms in emerging countries may access remotely generated knowledge is through MNE subsidiaries that locate near them. This paper examines this mechanism, presenting evidence that the R&D satellites of MNEs increase the flow of knowledge (as measured by patent citations) from the MNE's headquarters to firms located in the satellite's host country. Due to the localized nature of knowledge, such a medium for the geographic diffusion of knowledge might be particularly crucial for firms in emerging countries that do not yet have the foreign presence or networks to tap into remote sources of knowledge.

Among the first to study the relationship between FDI and knowledge flows was Branstetter (2006), who examines a group of Japanese firms and finds a positive relationship between the firm's level of outward FDI to the U.S. and the number of citations that the firm's patents receive from U.S. patents. As expected, he finds the relationship to be particularly strong for Japanese FDI investments in American R&D and product development facilities. In a more recent working paper, Fons-Rosen (2012) analyzes changes in the degree to which firms in Central and Eastern Europe cite the patents of foreign firms entering the market. He ingeniously employs a difference in difference analysis where the treated group is the foreign firms who won privatization tenders (and hence entered the country) and the control group consists of the losing bidders. He finds that, on average, the winning bidders experienced a 20% increase in citations received relative to the losers.

However, such results do not necessarily imply that subsidiaries facilitate the flow of geographically remote knowledge (from the headquarters) into the host country. It could be that the increased citation rates stem from citations to patents (and knowledge) developed by the subsidiary in the host country since knowledge diffuses locally. In the context of our earlier example, Huawei or Alibaba citing more Microsoft patents does not necessarily imply increased cross-country knowledge flows if the incremental citations are to patents developed in Microsoft's Beijing R&D satellite. To explore the satellite's effect on the international flow of knowledge, this paper examines whether there is an increase in the number of host country firm patent citations to the stock of patents generated in the *headquarters* of the firm doing FDI. The primary finding is that the citations received by headquarters patents are indeed disproportionately (relative to citations received by a set of control patents that are in the same technological class, from the same year, and in the same country as the headquarters) from third-party firms in the satellite's host country. We also employ a difference in differences methodology to show that this disproportionate citation of headquarters patents is not due to firms choosing to establish their satellites in countries where this was already the case.

This paper also differs from previous work on FDI and knowledge diffusion in that, contrary to the large majority of the literature that focuses on one country, or at most one region, the analysis conducted here examines the effect on over 121 host countries. This allows for not only a more general treatment of the topic, but also a cross-country comparison of the country characteristics that facilitate or hinder the flow of remote knowledge through the medium of the satellite. One important finding in this respect is that a host country's level of technological development is a critical factor, with knowledge diffusion through the satellite being greatest in countries and sectors with strong, but not world-class, capabilities. The analysis in this paper also spans a much longer timeframe than other studies (1976-2006).

Finally, this paper also considers whether satellite staffing policy impacts knowledge diffusion though the satellite. While geographic knowledge diffusion through the satellite is beneficial from the perspective of international growth and development, the MNE may be rightly concerned that it weakens its competitive position. The MNE may thus take measures to better protect its knowledge when venturing abroad. We find that their satellite staffing policy is a key determinant of the magnitude of headquarters knowledge diffusion through the satellite. In particular, satellites that have many inventors who previously patented at other firms in the satellite location generate a larger satellite effect, presumably because such inventors have stronger local social networks and because they are more likely to leave the firm with their knowledge. In addition, satellites with more inventors that previously patented at the headquarters are also associated with a larger satellite effect, likely because such expat inventors embody more headquarters knowledge. The analysis, however, does not establish a causal relationship.

The paper is organized as follows. Section 2 briefly discusses why we might expect the satellites of multinational firms to promote the geographic diffusion of knowledge into their host country and why this flow of knowledge might depend on the satellite's staffing policy and the country's existing technical capabilities. Section 3 discusses the data and methodology used to address this question and section 4 presents the results. Finally, section 5 concludes.

2. Satellites and Knowledge Diffusion

For satellites to effectively increase the flow of knowledge from their headquarters to firms in their host country, the MNE headquarters must first share their technological expertise with its satellite. Second, this knowledge must spill out from the satellite to local firms. The fact that, in general, a headquarters shares knowledge with its satellites should not come as a surprise. As Dunning (1977) and Kogut and Zander (1993) argue, the very existence of MNEs may be due to the fact that they are an efficient organizational vehicle through which to transfer and share knowledge across borders. MNEs that fail to share knowledge among their different locations lose many of the advantages of being an MNE. We would therefore expect MNEs to actively share headquarters knowledge with its satellites.

Once a satellite has received knowledge from the parent, this knowledge can spill over to local firms through the usual localized diffusion mechanisms of local input-output linkages, social networks, and labor mobility (Von Hippel 1988, Rogers and Larsen 1984, Almeida and Kogut 1999, Agrawal, Cockburn, and McHale 2006). Whether or not this overall flow of knowledge - from headquarters, to satellite, to other firms in the satellite's location - occurs, is an empirical question and the primary focus of this paper.

We would, of course, expect the magnitude of this overall knowledge flow to be mediated both by characteristics of the MNE and of the local environment in which the satellite resides. MNEs have an incentive to institute policies and mechanisms that curtail outward knowledge flow to potential competitors. One way they might achieve this is for satellites to focus on technologies that are more internally-oriented and hence useless to third parties [Zhao (2006)]. Another way that firms may be able to protect their knowledge is through the strategic staffing of key R&D positions in their foreign units. Hiring local engineers and scientists may seem like the best choice, particularly if the firm is venturing abroad to access talent, but doing so may result in increased knowledge loss to potential competitors. Locals have stronger localized social networks, which, as mentioned before, is one of the key mechanisms behind the localization of knowledge spillovers (Rogers and Larsen, 1984; Agrawal, Cockburn, and McHale, 2006). Another mechanism, is the inter-firm mobility of engineers and scientists (Almeida and Kogut, 1999; Song, Almeida, and Wu, 2003; Breschi and Lissoni, 2009; Singh and Agrawal, 2010) and here again, we would expect locals to leave for another local firm more frequently, taking any acquired knowledge with them. Alternatively, firms may staff the satellite with expatriates. Expatriates are likely to possess more critical technical capabilities and institutional knowledge central to the firm, but they are also less likely to diffuse this knowledge to their local environment since they are largely devoid of local social networks and are less likely to move to another local firm. Overall, the effect of staffing satellites with expatriates on total knowledge diffusion is ambiguous: it increases the stock of headquarters knowledge that potentially can diffuse, but the likelihood of a given piece of knowledge being passed on is probably lower. This is therefore an empirical question.

A second important factor that mediates the overall flow of knowledge from headquarters to firms in the satellite's location is the characteristics of local firms. As Cohen and Levinthal (1990) argue, firms with little expertise in a particular sector may fail to recognize the value of new information or may not be able to assimilate it. In an international context marked by great disparities in technological capabilities across countries and sectors, the absorptive capacity (or lack thereof) of local firms is likely to be a primary driver of whether local firms benefit from the presence of more technologically advanced MNEs. Hence, we would expect knowledge inflows from MNE headquarters to be greater in more technologically advanced countries and sectors.

The caveat is that country/sectors that are already world-class have little to gain from MNEs that are based in less-advanced countries. Instead, they are likely to focus their energy on learning from other firms in their home country. Consistent with this, Singh (2007) finds that knowledge inflows from foreign MNE subsidiaries to local firms are smaller than knowledge outflows only in technologically advanced countries.[3] The combination of these two arguments leads to the hypothesis that knowledge diffusion

[3] Singh's focus is on local knowledge flows between subsidiaries and host country firms. He does not consider remote knowledge flows between the headquarters and firms in the host country.

from the MNE headquarters to local firms is an inverted U function of the host country's capability in the given technological sector.[4] This is indeed confirmed by our empirical findings.

3. Data and Methodology

3.1 Data

To determine whether satellites facilitate the flow of headquarters knowledge to third-party firms in their host country, I use data on patents granted by the United States Patents and Trademarks Office (USPTO). The data comes from two separate sources. The NBER Patent Data Project provides a dataset with the application year, technology class, technology subclass, assignee (owner), and assignee headquarters country of all patents granted between 1976 and 2006, inclusive. This is supplemented with Harvard's Patent Network Dataverse data,[5] generated by Lai, D'Amour, Yu, Sun, and Fleming (2011) to study the co-authorship networks of inventors. It contains inventors' country of residence, a unique ID for each inventor, and citation data, for patents granted up to 2010. The unique inventor ID, allows us to track inventors across different firms and countries of residence, and hence to construct variables measuring inventors' previous experience. The final combined sample consists of all patents granted between 1976 and 2006.

Patents are assigned to an innovating country based on the residence of the inventors. Implicit in this definition is the belief that the inventor's country of residence most closely proxies for the location where the process of innovating occurred. For the analysis, I focus on patents where all inventors reside in the same country so as to be able to unambiguously assign a country of invention.[6] It is also worth noting that the use of USPTO data implies that non-U.S. firms with closer ties to the U.S. are likely to be

[4] More formally, this hypothesis would follow by claiming that both the costs and benefits of knowledge acquisition from foreign satellites are decreasing in their own technological capabilities ($C'<0$, $B'<0$) but that benefits decrease faster than costs ($C''>0$, $B''<0$).
[5] The dataset is available at https://dataverse.harvard.edu/dataset.xhtml?persistentId=hdl:1902.1/15705
[6] For determining the locations of a firm, only considering patents where all inventors reside in the same country is the more conservative approach (a single inventor collaborator in a foreign country does not constitute a satellite). Only 6% of the patents in our sample have inventors from multiple countries.

overrepresented in our sample since, on the margin, they are more likely to apply for USPTO patents.[7]

Knowledge flows are measured using patent citations. Although citations are a noisy measure of knowledge flows, studies comparing citation data with inventor surveys have shown that the correlation between patent citations and actual knowledge flows is high enough to justify their use in large samples (Jaffe and Trajtenberg, 2002, Chapter 12; Duguet and MacGarvie, 2002).

The first step in the construction of the dataset consists of identifying firms with patenting activity in more than one country.[8] While the country of the headquarters is listed directly in the data,[9] the satellite locations are identified as any other countries where the firm has innovative activities (i.e. at least one patent where all inventors reside in that country). It is clear that this definition of satellites captures not only physical R&D centers, but also less formal foreign R&D activities (as long as they generate at least one patent and all the inventors are local). For our purposes, the distinction is not particularly important as either type of foreign R&D activity should facilitate knowledge diffusion. That said, we can expect that a more important R&D presence would yield larger knowledge diffusion. Consistent with this, we find a larger estimated satellite effect when focusing on satellites with larger patent outputs. In Appendix A, we also show that there is a significant overlap between the location of our satellites and that of the firms' subsidiaries as listed in their SEC filings. We also obtain similar estimates of knowledge diffusion when examining instead the effect of these subsidiaries.

In total, there are 8013 firms with foreign innovative activities spread over 14,328 R&D satellites. These satellites have generated on average 7.4 U.S. patents, with 7640 generating only one patent, 2263 two patents, 1135 three patents, and 3290 four or more

---

[7] First, firms with U.S. ties and their satellites will be overrepresented in the sample of treated patents. But because this overrepresentation would apply equally to the treated and control patents it is not clear that it would generate a bias. Second, firms with U.S. ties will be overrepresented in the sample of citing patents. But while such patents are more likely to cite the U.S., this should not generate a bias because if the treated patent is in the U.S. then by construction so is the control.

[8] I ignore patents with multiple assignees since in these cases the extent of a particular assignee's participation in the innovation is not clear (for example, they may have bought the patent).

[9] In more than 95% of firms, the "headquarters" country also coincides with the country with the most patenting activity.

patents. Figures 2 and 3 present the geographical distribution of firm headquarters and satellites, respectively.

3.2 Satellite Effect on Knowledge Diffusion

There are two types of cited patents in the analysis: treated and control. The set of treated patents is comprised of all patents generated in the headquarters of these MNEs. These patents are "treated" in the sense that the knowledge they embody may have diffused to the satellite host country through the medium of the satellite. In order to control for patterns of geographic agglomeration of technological activity that could be related to the choice of location for the satellite, I generate a matched control group.[10] In particular, each patent in the treated group is matched to a control patent from the same country, application year, technology class, and if possible technology subclass. Moreover, the assignee of the control patent cannot have patented in the same country as the treated firm's satellite (i.e., the owner of the control patent cannot have a satellite in the same host country). If several potential control patents match the criteria, I choose the patent belonging to the MNE whose number of locations most closely matches the number of locations of the treated firm. Patents belonging to single-country firms are chosen as a last resort. If two or more patents are equally suitable controls according to these six criteria, one is picked randomly. Any treated patent with no matching control is dropped from the sample. I find controls for 97.2% of the 12,160,638 potential treated patent/satellite combinations, of which 6,609,495 matches are at the subclass level and 5,208,952 at the class level.[11]

Figure 4 provides an illustration of the methodology for Microsoft Corporation's Chinese satellite.[12] The question of interest is whether headquarter patents are disproportionately cited by third parties in the locations where the firm has a satellite. In particular, the analysis compares the proportion of citations received by patents in the treated group that are from third-party patents in the satellite's location ($\hat{P}_T$) to the

[10] The methodology builds on that of Jaffe et al. (1993), Agrawal et al. (2006), and Blit (2017).

[11] There are a total of 428 different three-digit U.S. patent technology classes (with the most active being "drugs" and "semiconductors") and more than 100,000 technology subclasses.

[12] Microsoft Corporation is the U.S. firm with the most patenting activity in China. Research at the Beijing lab focuses on natural user interface, next generation multimedia, data-intensive computing, search and online ads, and computer science fundamentals.

proportion of citations received by patents in the control group that are from third-party patents in the satellite location ($\hat{P}_C$). Pooling all citations to all treated and control patents across all firms and all satellites we compute the satellite effect as the ratio $\hat{P}_T/\hat{P}_C$.

3.3 Difference in Differences

It could be the case that the finding of a satellite effect on knowledge diffusion results from firms choosing to set up satellites in locations that are disproportionately sourcing their headquarters' 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. As a robustness check, then, we conduct a difference in differences analysis to address this and related concerns.

The principal challenge with deploying this approach is that the date of establishment of a satellite is not directly observed. What is observed is a proxy for the date of establishment: the application date of the first patent from the satellite location. Since patents are generally the result of many years of development (and many satellites may initially perform non-innovative tasks), the year of the first application will invariably lag the actual year of establishment. In fact, since most of these satellites have very few patents, we may miss the actual date of establishment by a significant margin. Since the years right before the first application date of a satellite patent are the most problematic to define as pre- or post-satellite, I define a "grey window" as the period starting *w* years before the first satellite patent application and ending the year before the first patent application. This window is dropped from the analysis, and the pre-satellite period is defined as ending *w+1* years prior to the first satellite patent. The post-satellite period begins the year the first satellite patent application is observed. In section 4.2 we show that the results are consistent across all different window lengths.

Nonetheless, even with this grey window, it will still be the case that many satellites are operating and diffusing knowledge in what has been defined in this way as the “pre-satellite” period (especially for smaller windows). However, this will bias against finding a difference in the relative rate of knowledge diffusion between before and after the date of the first application. The imperfectly observed date of establishment does, however, make it difficult to test for the possibility that there might be pre-existing trends in knowledge flows that are correlated with the establishment of a satellite. Any observed pre-trend could be due either to an actual pre-existing trend (which would be problematic) or to the satellite having already been established in what has been defined as the pre-satellite period. Appendix B nonetheless examines the presence of pre-trends and finds one, particularly for the year right before the satellite applies for its first patent.

Pooling all patents (both treated and control), we perform the following difference in differences estimation:

$$\hat{P}_{ps} = \propto + \beta_1 treated_{ps} + \beta_2 postsat_{ps} + \beta_3 treated_{ps} \times postsat_{ps} + \gamma_{f(p)} + \gamma_{t(p)} + \varepsilon_{ps}$$
(1)

Observations are at the level of the patent-satellite, with the index p denoting either a headquarter (treated) patent or a control patent, and the index s denoting one of the headquarter's satellites. The dependent variable is the proportion of citations received by p that are from third-party patents in the country of satellite s. The explanatory variables include a dummy variable (*treated*) that is 1 when the patent is from the treated group and 0 when it is from the control group, a dummy variable (*postsat*) that is 1 when the patent's application year is in the post-satellite period and 0 if it is in the pre-satellite period, and the interaction (*treated_postsat*) of these two variables. $\gamma_{f(p)}$ and $\gamma_{t(p)}$ represent firm and patent application year fixed effects, respectively, where the firm is the assignee of patent p. Our coefficient of interest, β_3, provides an estimate of the satellite effect on knowledge diffusion.

3.4 Heterogeneity in Satellite Effect

3.4.1 Satellite Inventor Experience

To examine the potential impact of satellite staffing policy on outward knowledge diffusion, we construct a variable measuring the satellite's inventors' previous patenting experience using the unique inventor ID in the Harvard patent data. We begin by determining, for each satellite inventor, whether they have previously patented for: other firms in the host country, other firms outside the host country, the headquarters, other satellites of the firm, and the satellite. We then aggregate each of these five different inventor metrics to the satellite-year level as follows:

$$Sat\ Exp_{s,t} = \sum_{y=1960}^{t} \sum_{i=1}^{N_y} Inv\ Exp_{y,i} \quad (2)$$

Exp refers to one of the five previous types of experience described above. The subscript *s* refers to the satellite and *t* to a year. To construct the variables we start by counting the number of satellite inventors that applied for a patent in year *y* that have the type of experience in question (the rightmost sum, where N_y is the total number of satellite inventors that applied for a patent in year *y* and *Inv Exp* is an indicator variable for whether inventor *i* has that previous experience). For any given year, we then compute the total experience of all inventors that have patented at the satellite up to that year (the leftmost sum). The inherent assumption in this approach is that inventors remain at the satellite, even when we don't observe them patenting in a given year (since patenting is a rare event). But inventors that patent frequently will contribute any previous experience multiple times to the satellite level measure. The rationale for this is that such inventors are more central to the R&D operations of the satellite and their previous experience may matter more. Summary statistics for these variables are presented in Table 2.

To examine the relationship between satellite inventor experience and the satellite effect, we estimate the following equation:

$$\left(\hat{P}_T - \hat{P}_C\right)_{ps} = \alpha + \beta_1 locals_{st(p)} + \beta_2 others_{st(p)} + \beta_3 HQexpats_{st(p)} + \beta_4 expats_{st(p)} + \beta_5 satellite_{st(p)} + \delta_1 satpatcount_s + \delta_2 hqpatcount_s + \delta_3 numlocations_s + \gamma_{t(p)} + \gamma_{c(p)} + \gamma_{i(s)} + \gamma_{f(s)} + \varepsilon_{ps}. \quad (3)$$

The unit of observation is the headquarters (treated) patent *p* and satellite *s*. The dependent variable is the proportion of citations to the treated patent that is from the host country minus the proportion of citations to the associated control patent that is from the host country. The five explanatory variables of interest are at the level of the satellite-year. *locals* measures the extent to which satellite inventors have previously patented for another firm in the satellite country, *others* measures previous patenting for another firm in a country other than the satellite's, *HQexpats* measures previous patenting for the firm in the headquarters country, *expats* measures previous patenting for the firm in some country other than the headquarters' or the satellite's, and *satelliters* measures previous patenting for the firm in the satellite country. In addition we include three control variables, the total number of satellite patents, the total number of headquarters patents, and the number of locations (countries) of the firm. $\gamma_{t(p)}$, $\gamma_{c(p)}$, $\gamma_{i(s)}$, and $\gamma_{f(s)}$ are year, technology class, satellite country, and firm fixed effects, respectively.

While $\hat{P}_T/\hat{P}_C$ is the direct measure of the satellite effect, using it as the dependent variable in the above specification would have resulted in a loss of 90% of our sample due to observations with $\hat{P}_C$=0. Since this is likely to bias our results, $\hat{P}_T - \hat{P}_C$ is our preferred dependent variable. Nonetheless, for robustness, Appendix E conducts the same analysis using the reduced sample and the ratio as the dependent variable.

3.4.2 Country Characteristics

We use patent stock per capita at the level of the country-sector-year as a measure of technological capability. Patent stocks are constructed as the discounted number of patents in a given country and technology class, up to a given year. A discount rate of 15% per year is applied to patents from previous years to account for technological obsolescence.[13] These patent stocks are normalized by the host country's population (as obtained from the UN Statistics Division) in the given year to obtain patent stocks per capita. The baseline estimating equation is:

$$\left(\hat{P}_T - \hat{P}_C\right)_{ps} = \alpha + \beta_1 PatStockpc_{i(s)c(p)t(p)} + \beta_2 PatStockpc^2_{i(s)c(p)t(p)} + \beta_3 \boldsymbol{X}_{i(s)t(p)} + \beta_4 \boldsymbol{X}_{i(s)} + \beta_5 \boldsymbol{X}_{i(s)j(p)} + \gamma_{t(p)} + \gamma_{c(p)} + \varepsilon_{ps} \qquad (4)$$

[13] In addition, this is multiplied by the correction term (1-(0.85)$^{2006-1976+1}$)/(1-(0.85)$^{t-1976+1}$), where *t* is the stock year, to adjust for the fact that the data is truncated at 1976.

The unit of observation is the headquarters (treated) patent *p* and satellite *s*. The dependent variable is, as before, the proportion of citations to the treated patent that is from the host country minus the proportion of citations to the associated control patent that is from the host country. While this is our preferred dependent variable because of the numerous zero values of $\hat{P}_C$, we also report the results for $(\hat{P}_T/\hat{P}_C)$ in Appendix F. The variable $PatStockpc_{i(s)c(p)t(p)}$ is the patent stock per capita for the country *i* of satellite *s*, in the technology class of the treated patent, *c(p)*, in the application year of the treated patent, *t(p)*. $\mathbf{X}_{i(s)t(p)}$, $\mathbf{X}_{i(s)}$, and $\mathbf{X}_{i(s)j(p)}$ are vectors of control variables that vary at the level of the satellite country-year, satellite country, and satellite country-headquarters country, respectively. $\gamma_{t(p)}$ and $\gamma_{c(p)}$ are year and technology class fixed effects. Table 1 presents the description and source of these control variables, and their summary statistics are presented in Table 2.

4. Results

4.1 Satellite Effect on Knowledge Diffusion

Table 3 (column 1) suggests that the satellite facilitates the geographic diffusion of knowledge, allowing firms in the host country to receive a disproportionately large amount of knowledge generated in the headquarters. While 3.37% of the citations received by the MNEs' headquarters (treated) patents are from third parties in the country of their satellite, only 2.95% of the citations received by the control group are from third parties in that country ($\hat{P}_T$=.0337 and $\hat{P}_C$=.0295). Performing a difference of proportions t-test with unequal variances as in Almeida (1996) yields a t-statistic of 150.[14] The satellite effect on knowledge diffusion to the host country is $\frac{\hat{P}_T}{\hat{P}_C} = 1.14$, implying that host country firms receive on average 14% more knowledge from the headquarters.

Columns 2-4 present the results when the same analysis is performed separately for developing, middle-income, and developed countries. Countries are categorizes

[14] Letting H_0: $P_T = P_C$ and H_1: $P_T > P_C$, we calculate the t-statistic as: $t = \left(\hat{P}_T - \hat{P}_C\right) / \sqrt{\hat{P}_T\left(1-\hat{P}_T\right)/n_T + \hat{P}_C\left(1-\hat{P}_C\right)/n_C}$, where n_T and n_C are the total number of citations received by the treated and control group, respectively.

according to 1990 per capita GDP measured at current U.S. dollars, as obtained from the United Nations Statistics Division. Countries with a GDP of less than $5000 per capita as classified as developing, those with at least $5000 but less than $15000 as middle-income, and those with at least $15000 as developed. The satellite effect appears to be largest for middle-income countries, consistent with our hypothesis that countries with both the absorptive capacity and the motivation to learn experience the biggest benefit. However, we can expect significant heterogeneity within each of these groups and especially across sectors within a country. This will be analysed more closely in section 4.3.2.

4.2 Difference in Differences

Table 4 presents the results of a difference in differences OLS regression with standard errors clustered by firm, and for different window sizes. The interaction term *Treated*Post Sat* is highly significant in all cases, suggesting that subsequent to the establishment of the satellite there is increased knowledge diffusion from the headquarters to third-party firms in their host country. This effect should not be interpreted as an average treatment effect, but rather as the satellite treatment effect on the treated, since treated firms chose to establish a satellite. Moreover, these firms also chose the satellite's location and the timing of satellite establishment. The positive and significant coefficient on *treated* indeed suggests that firms may have established satellites in countries where they already had significant ties (countries that were already disproportionately citing their patents). However, the positive coefficient on *treated* could also be due to mismeasurement in the date of establishment. As already discussed, the satellite is likely to have been present a number of years before the first observed patent application, resulting in a consistent lag in the measured year of satellite establishment relative to the actual year of establishment. The decrease in the size of the coefficient on *treated* and increase in the coefficient of the interaction term as larger windows are chosen offers some evidence that this mismeasurement bias is indeed present. It should be noted, however, that this mismeasurement is not problematic for determining whether there is a satellite effect because it should bias the coefficient on the interaction term towards zero.

Where the inaccurate measurement of the year of establishment is problematic is in our ability to ensure that there are no pre-trends in our difference in differences estimation. As shown in Appendix B, we find a pre-trend in the interaction coefficient, which could be due either to a real pre-trend before the actual date of establishment or to the consistent lag in the observed year of establishment relative to the actual year. While we cannot rule out a real pre-trend, the fact that the pre-trend is only large and strongly significant for the year directly before the measured year of establishment suggests that this mismeasurement is at play.

It is also worth noting the positive and significant coefficient on *Post Satellite* in Table 4. This suggests that the timing of the establishment may have coincided with firms in the host country increasingly citing the country and/or technology sector of the headquarters. One story consistent with this, is that firms choose to establish satellites in locations with increasing activity in the technological sector of the firm. An alternative explanation is that the timing is consistent with the host country increasing technological (and other) ties with the country of the headquarters. Such scenarios, of course, should not affect our estimated satellite effect since citations to both treated and control patents should be affected equally.

As a robustness check, Appendices C and D estimate a difference in differences using placebo host countries and placebo establishment dates, respectively. In both cases, we find no satellite effect.

Overall, Table 4 suggests that at the very least, a significant portion of the satellite effect computed in Section 4.1 is due to increased knowledge flows through the satellite and not just to firms establishing satellites in locations where third-party firms already disproportionately source headquarters' knowledge.

4.3 Heterogeneity in Satellite Effect

While the focus thus far has been on determining whether there exists a satellite effect and on estimating its magnitude, understanding the effect's heterogeneity across firms and countries is also important as it can provide management and policy relevant insights. We first examine the extent to which firms that venture abroad may be able to limit the outward diffusion of their knowledge through staffing policy and then examine the relationship between host-country technological capability and the satellite effect.

4.3.1 Satellite Inventor Experience

As Table 5 shows, the amount of knowledge diffusion through the satellite is mediated in a significant way by the previous experience of the satellite's inventors. The different columns present the results of estimating equation (3) with different sets of fixed effects. In particular, all four columns include year and technology class fixed effects, but in addition column 2 includes satellite country fixed effects, column 3 includes firm fixed effects, and column 4 includes a satellite fixed effect (interaction of the satellite country and firm) so that identification is off changes in time. We report robust standard errors clustered by firm.

Across all specifications, satellites with numerous inventors that have previous experience at other firms in the same country (*locals*) generate larger knowledge flow from the headquarters to local firms. This is consistent with our hypothesis that *locals* will have stronger local social networks through which firm knowledge can diffuse and that they are also more likely to leave the firm for another local firm, taking their knowledge with them. Consistent with this, satellites with numerous inventors that have previous experience at other firms but in a different country (*others*) do not increase knowledge flow (they do not have a local network). In fact, the estimated coefficient on *other* is negative in three of the four specifications, though it is never significant.

Expatriates have shallower local networks and are probably less likely to leave the satellite for another local firm. On the other hand, expats from the headquarters are likely to possess critical headquarters knowledge, that by carrying it with them to the satellite, are more likely to diffuse to local firms. Empirically, it seems that the latter effect dominates. Across all four specifications, the coefficient on *HQ Expats* is positive and significant. Expatriates from other firm satellites (*Expats*), however, do not have a significant positive effect because they do not possess the same amount of headquarters knowledge.

Satellites with numerous inventors that have previously patented at the satellite (experienced *satelliters*) generate a smaller satellite effect. This could be due to a number of things. First, inventors that have been at the satellite for a long time may have shallower ties to outside firms (especially compared to *locals* who likely are the other big group in most satellites). It could also be that inventors that have amassed a large number of satellite patents have become so specialized that they have less incentive to interact with inventors outside the firm and that their knowledge is of less use to outside firms. An alternative explanation is that strong satellite retention policies could result both in more inventors with multiple satellite patents and fewer inventors leaving the firm with their knowledge.

It is worth noting that adding additional country-level controls does not affect the results. The results are virtually unchanged if we include in the specification all the country variables described in Table 1. As a further robustness check, Appendix E presents the results of the analysis using the alternative dependent variable $(\hat{P}_T/\hat{P}_C)$. The results are similar with the exception of the last specification that includes satellite fixed effects. This may be because there simply is not enough real/accurate time variation in a given satellite's experience measures to obtain credible estimates in the much smaller sample or that the sample selection effect dominates.

It is worth stressing that this section estimates correlations and not causal relationships. For example, the negative coefficient on *satelliters* might be explained by the omitted variable that is satellite retention policies. Notwithstanding, the observed relationships are suggestive of the presence of a satellite effect whose underlying mechanism is consistent with the generally accepted mechanisms that underpin knowledge diffusion.

4.3.2 Technological Capability

We expect an inverse-U relationship between the amount of knowledge diffusion received through satellites and the technological stock of the host country. Table 6 presents the results of estimating equation (4) with different sets of fixed effects. Across all specifications, the coefficient on *patent stock per capita* and its square are highly significant and consistent with an inverted U relationship. The apex of the estimated inverted U is at a *log patent stock per capita* value of between 4.22e-5 and 4.47e-5, depending on the specification, well within the 0 to 7.81e-4 range of the variable, but in the 99.5th value percentile of the variable.[15] This result suggests that the most prominent issue for most countries/sectors is a lack of absorptive capacity that limits their ability to learn from the satellite. However, the most technologically sophisticated countries/sectors also gain less from foreign satellites since these have little to teach them, though this only applies to truly world-class clusters.[16] We note that the inverted-U is present even in the specification with interacted satellite-year fixed effects (Column 4), suggesting that even within country/years, the technology sectors that are advanced but not world-class benefit most from knowledge diffusion through satellites.

None of the estimated coefficients on the control variables are significant with the exception of *GDP per capita*, its square, and *Telephone lines*, and these are only significant in the first specification. Perhaps most surprising is the fact that the coefficient on *largest city population* is not positive and significant. The motivation for including this variable is that, since knowledge diffuses locally, using countries as the geographic level of analysis could bias the results downward.[17] For example, it is doubtful that a satellite in Beijing would help headquarters knowledge diffuse to firms located in Shanghai. Therefore, we expected to observe a larger satellite effect in countries whose population is more concentrated in one city. It is worth noting, that in the set of results for the dependent variable ($\hat{P}_T/\hat{P}_C$), the coefficient on *largest city population* is always positive and significant (see Appendix F).

5. Conclusion

While there is a large established literature examining the impact of foreign direct investment on productivity, the mechanisms through which this may occur are not well understood. This paper examines whether the satellites of multinational firms act as a medium for the international diffusion of knowledge. The results suggest that this is indeed the case, with satellites increasing the flow of knowledge from their headquarters

---

[15] Using instead ($\hat{P}_T/\hat{P}_C$) as the dependent variable, we obtain a similar apex of between 3.08e-5 and 3.62e-5, depending on the specification, which are in between the 99th to 99.3rd percentile of the variable.

[16] It may be the case that this applies more broadly, but the highly aggregated technological sectors are masking the effect. For example, a country may be at the technological frontier (and not learning) in half of the subsectors that make up a sector but some distance behind the frontier (and learning) in the other subsectors.

[17] The geographic unit of analysis is at the country level due to data limitations; reliable city-level inventor locations are only available for the U.S. and Canada.

to third-party firms in their host country by about 14%. Further, this learning through satellites is found to be largest in countries and sectors with strong but not world-class capabilities, presumably because these have both the motivation and absorptive capacity to learn from foreign parties.

We further find that the satellite effect on knowledge diffusion is heterogeneous across satellites. Knowledge diffusion is greater through satellites that have numerous inventors that have either previously patented in other local firms or at the firm's headquarters. A likely explanation for this is that inventors that have previously patented at other local firms are likely locals with strong local social networks that facilitate knowledge transfer. They are also more likely to leave the satellite for a local firm, taking their knowledge with them. Inventors that have previously patented in the headquarters location, on the other hand, are more likely to be expats that have more restricted local social networks, but by virtue of having been in the headquarters, carry more headquarters knowledge with them (some of which will inevitably leak out, either directly or through their local colleagues).

Overall, this paper's finding of a satellite effect on knowledge diffusion lends support to policies designed to attract foreign MNE's operating in technology-intensive sectors. The results further suggest which countries and sectors are most likely to benefit from policies designed to promote FDI and how firms can best protect their knowledge when venturing abroad. The micro-level analysis of the knowledge flows associated with satellites increases our understanding of how FDI may contribute to development and economic growth and, just as importantly, in which instances it is likely to do so.

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: 571-591.

Aitken, B., Harrison, A. 1999. Do Domestic Firms Benefit from Direct Foreign Investment? Evidence from Venezuela. *The American Economic Review* 89:605-618.

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., B. Kogut. 1999. Localization of knowledge and the mobility of engineers in regional networks. *Management Science* 45(7) 905-917.

Audretsch, D. B., Feldman, M. P. 1996. R&D spillovers and the geography of innovation and production. *American Economic Review* 86: 630-640.

Aw, B., Chen, X., Roberts, M. 2001. Firm-level evidence on productivity differentials, turnover, and exports in Taiwanese manufacturing. *Journal of Development Economics* 66:51-86.

Bernard, A., Jensen, J. 1999. Exceptional exporter performance: cause, effect, or both? *Journal of International Economics* 47: 1-25.

Blalock, G., Gertler, P. 2008. Welfare gains from foreign direct investment through technology transfer to local suppliers. *Journal of International Economics* 74: 402-421.

Blit, J. 2017. Learning remotely: R&D satellites, intra-firm linkages, and knowledge sourcing. *Journal of Economics and Management Strategy*, forthcoming.

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: 325-344.

Breschi, S., F. Lissoni. 2009. Mobility of skilled workers and co-invention networks: an anatomy of localized knowledge flows. *Journal of Economic Geography* 9: 439-468.

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: 903–948.

Cohen W., Levinthal, D. 1990. Absorptive capacity: a new perspective on learning and innovation. *Administrative Science Quarterly* 35: 128-152.

De Loecker, J. 2007. Do exports generate higher productivity? Evidence from Slovenia. *Journal of International Economics* 73: 69-98.

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: 375-393.

Dunning, J. 1977. Trade, location of economic activity and the multinational enterprise: A search for an eclectic approach. In B. Ohlin, P.O. Hesselborn & P.M. Wijkman, editors, The international allocation of economic activity. London. Macmillan.

Fons-Rosen, C. 2012. Knowledge flows through FDI: the case of privatisations in central and eastern Europe. Working Paper.

Frost, T.S. 2001. The geographic sources of foreign subsidiaries' innovations. *Strategic Management Journal* 22:101-23.

Ginarte, J.C., Park, W.G. 1997. Determinants of patent rights: A cross-national study. *Research Policy* 26:283-301.

Girma, S., Gong, Y., Görg, H. (2008). Foreign direct investment, access to finance, and innovation activity in Chinese enterprises. *World Bank Economic Review* 22 (2), 367-382.

Gorg, H. and Greenaway, D. 2004. Much ado about nothing? Do domestic firms really benefit from foreign direct investment? *The World Bank Research Observer* 19: 171-197.

Gorg, H. and Strobl, E. 2001. Multinational companies and productivity spillovers: a meta-analysis. *The Economic Journal* 111: F723-F739.

Haddad, M., Harrison, A. 1993. Are there positive spillovers from direct foreign investment? Evidence from panel data from Morocco. *Journal of Development Economics* 42(1):51-74.

Hall, B. H., Jaffe, A. B., Trajtenberg, M. 2001. The NBER Patent Citation Data File: Lessons, Insights and Methodological Tools. NBER Working Paper 8498.

Haskel, J., Pereira, S., Slaughter, M., 2002. Does inward foreign direct investment boost the productivity of domestic firms? NBER Working Paper 8724.

Jaffe, A. B., Trajtenberg, M., Henderson, R. 1993. Geographic localization of knowledge spillovers as evidenced by patent citations. *Quarterly Journal of Economics* 108: 577-598.

Jaffe, A.B., Trajtenberg, M. 2002. Patents, Citations, and Innovations: A Window on the Knowledge Economy, MIT Press: Cambridge and London.

Javorcik, B. 2004. Does foreign direct investment increase the productivity of domestic firms? In search of spillovers through backward linkages. *American Economic Review* 94 (3):605– 627.

Kaufmann, A.K., Mastruzzi, M. 2009. Governance Matters VIII: Aggregate and individual governance indicators, 1996-2008. World Bank Policy Research Working Paper 4978.

Kee, H. L. 2015. Local intermediate inputs and the shared supplier spillovers of foreign direct investment. *Journal of Development Economics* 112: 56-71.

Keller, W. 2002. Geographic localization of international technology diffusion. *American Economic Review* 92: 120-142.

Keller, W., Yeaple, S., 2003. Multinational enterprises, international trade, and productivity growth: firm-level evidence from the United States. NBER Working Paper 9504.

Kogut, B., Zander, U. 1993. Knowledge of the firm and the evolutionary theory of the multinational corporation. *Journal of International Business Studies* 24: 625-645.

La Porta, R., Lopez-de-Silanes, F., Shleifer, A., Vishny, R. 1999. The Quality of Government. *Journal of Law, Economics and Organizations* 15: 222-279.

Lai, R., D'Amour, Yu, A., Sun, Y., Fleming, L. 2011. "Disambiguation and Co-authorship networks of the U.S. Patent Inventor Database (1975-2010)", hdl: 1902.1/15705, Harvard Dataverse, V5, UNF:5:RqsI3LsQEYLHkkg5jG/jRg==

Lileeva, A., Trefler, D. 2010. Improved access to foreign markets raises plant productivity... for some plants. *Quarterly Journal of Economics* 125: 1051-1099.

MacGarvie, M. 2006. Do firms learn from international trade? *Review of Economics and Statistics* 88: 46-60.

Rogers, E., J. Larsen. 1984. *Silicon Valley fever*. Basic Books, New York.

Romer, P. 1990. Endogenous technological change. *Journal of Political Economy* 98: S71-S102.

Singh, J. 2007. Asymmetry of knowledge spillovers between MNCs and host country firms. *Journal of International Business Studies* 38: 764-786

Singh, J., A. Agrawal. 2011. Recruiting for ideas: how firms exploit the prior inventions of new hires. *Management Science* 57(1): 129-150.

Song, J., P. Almeida, G. Wu. 2003. Learning-by-hiring: when is mobility more likely to facilitate interfirm knowledge transfer? *Management Science* **49**(4) 351–365.

Song, J., Shin, J. 2008. The paradox of technological capabilities: a study of knowledge sourcing from host countries of overseas R&D operations. *Journal of International Business Studies* 39: 291-303.

Thompson, P. 2006. Patent citations and the geography of knowledge spillovers: evidence from inventor- and examiner-added citations. *Review of Economics and Statistics* 88: 383-389.

Thompson, P., Fox-Kean, M. 2005. Patent citations and the geography of knowledge spillovers: a reassessment. *American Economic Review* 95: 450-460.

Van Biesebroeck, J. 2005. Exporting raises productivity in sub-Saharan African manufacturing firms, *Journal of International Economics* 67: 373-391.

Von Hippel, E. 1988. *Sources of innovation*. Oxford University Press, New York.

Zhao, M. 2006. Conducting R&D in countries with weak intellectual property rights protection. *Management Science* 56(7): 1185-99.

FIGURES

Figure 1: Chinese Indigenous and Offshored Patent Counts by Application Year

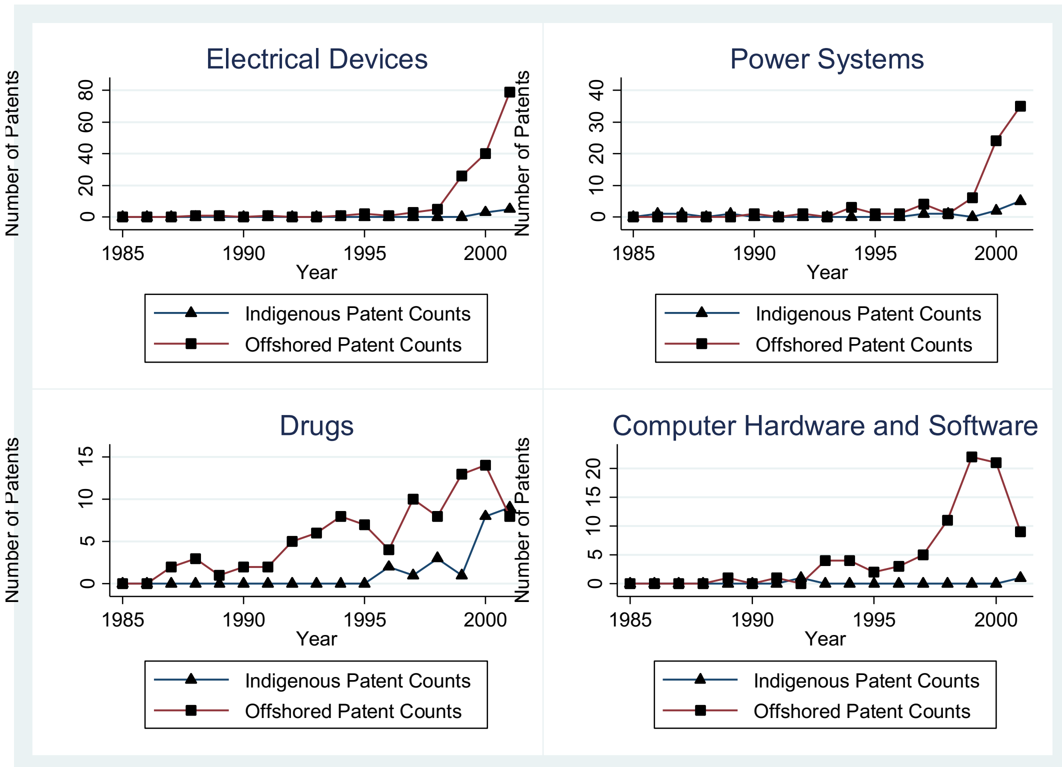

Notes: The figure presents indigenous and offshored USPTO patent counts in the four technology subcategories with the most patenting activity in China. Indigenous patents are patents with all Chinese inventors and a Chinese assignee, while offshored patents are patents with at least one Chinese inventor and a foreign assignee. In each of these sectors, we observe no patenting initially, followed by patenting activity by foreign multinationals with a presence in China, and subsequently by indigenous Chinese patenting.

Figure 2: Number of Headquarters by Country

Notes: The size of the circle represents the number of headquarters that are located in that country. The 8013 firms had their headquarters in 39 distinct countries, with 3868 being in the U.S., 863 in Germany, 468 in the U.K., 464 in Japan, 370 in Switzerland, 351 in France, and 346 in Canada.

Figure 3. Number of Satellites by Country

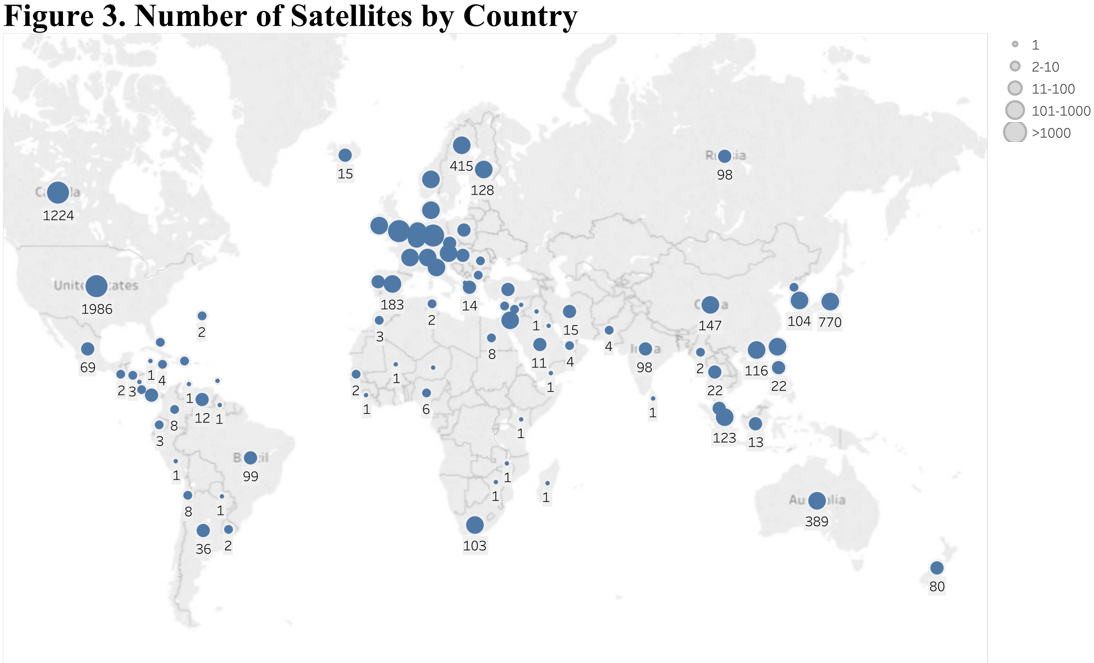

Notes: The size of the circle represents the number of satellite R&D centers that are located in that country. The 8013 firms had a combined total of 14,328 R&D satellites in 93 distinct countries, with 1986 being in the U.S.A., 1682 in the U.K., 1591 in Germany, 1224 in Canada, 950 in France, 770 in Japan, and 677 in Switzerland.

Figure 4: Illustration of Methodology Using Microsoft Corporation's U.S. Headquarters and Chinese Satellite

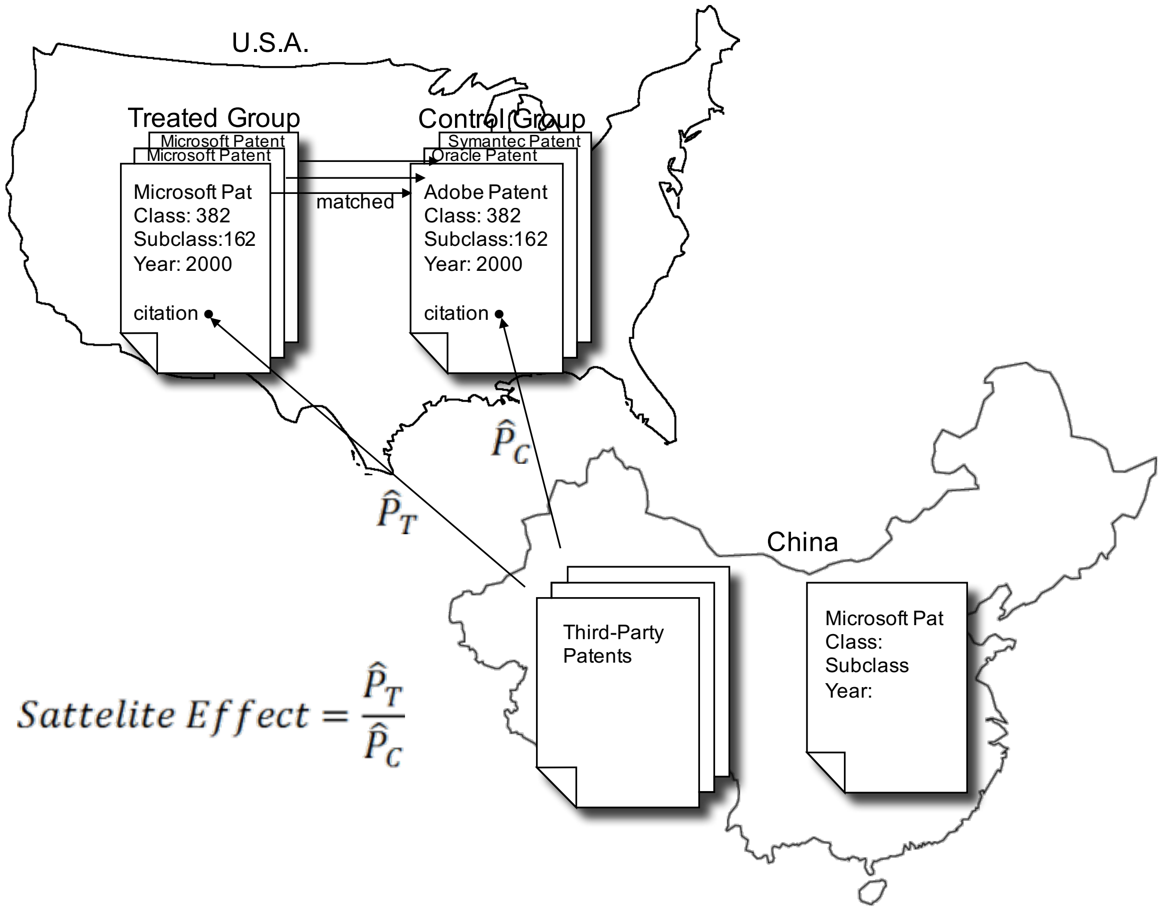

Notes: Microsoft's headquarters are in the U.S. and this satellite is in China (as evidenced by patents with Microsoft as the assignee and inventors residing in China). To determine whether there is a satellite effect on knowledge diffusion, we examine whether the treated patents' proportion of citations that are from third party patents in China ($\hat{P}_T$), is greater than the control patents' proportion of citations received from third party patents in China ($\hat{P}_C$). Treated patents consist of the patents generated in the headquarters of a firm with a satellite in the host country (China in this case), while control patents are patents that are matched one to one to these treated patents so as to be from the same country (U.S. in this case), application year, class, and to the extent possible, subclass. We compute the satellite effect on knowledge diffusion as the ratio $\hat{P}_T/\hat{P}_C$ by pooling all citations across all firms and all satellites.

TABLES

Table 1: Control Variable Descriptions

Variable	Description	Unit of Observation	Source
GDP Per Capita	GDP per capita at current prices in U.S. Dollars	Country/Year (215 countries, 1970-2006)	UN Statistics Division
Geographic Distance	Great Circle distances between capital cities of MNE headquarters and satellite countries in thousands of kilometers	Country pair (9045 country pairs)	Jon Haveman†
Shared Language	Indicator for whether the MNE headquarters and satellite countries share the same primary language	Country pair (14,028 country pairs)	Jon Haveman†
Shared Legal Origin	Indicator for whether the legal origin of the MNE headquarters and satellite countries are the same (among British, French, German, Scandinavian, Socialist)	Country pair (22,366 country pairs)	Laporta et al. (1999)
Largest City Fraction of Total Population	Population of largest city as a fraction of the country's total urban population	Country/Year (118 countries, 1960-2007)	World Bank Urban Development
Telephone Lines	Fixed telephone lines per capita	Country/Year (211 countries, 1975-2008)	UN World Telecom/ICT Indicators Data
Internet Penetration	Estimated fraction of the population that are internet users	Country/Year (206 countries, 1990-2008*)	UN World Telecom/ICT Indicators Data
Patent Rights	Index of patent rights	Country/Year (120 countries, 1960-2010)	Ginarte and Park (1997)
Rule of Law	Rule of law index for 1996	Country (171 countries)	World Bank Governance Matters VIII
Corruption	Index of corruption	Country (126 countries)	Laporta et al. (1999)

*I extend the Internet Penetration variable to earlier years by assuming zero Internet penetration prior to 1990.
†http://www.macalester.edu/research/economics/page/haveman/trade.resources/tradedata.html

Table 2: Summary Statistics

Variable	Mean	Std. Dev.	Min.	Max.
Firm Variables				
Locals	10.96	57.78	0	2461
Others	1.91	9.50	0	334
HQ Expats	4.59	21.51	0	459
Expats	0.63	5.11	0	158
Satellite	35.72	169.95	0	3675
Number of Satellites	15.82	8.90	2	32
Headquarter Patents	12,263	11,457	2	37,327
Satellite Patents	41.50	141.75	1	1952
Country Variables				
Patent Stock Per capita	2.91E-6	6.53E-6	0	0.00078
GDP Per Capita	17,729	11,036	114	143,346
Geographic Distance (Km.)	7296	4049	174	19,007
Shared Language	0.201	0.401	0	1
Shared Legal Origin	0.305	0.460	0	1
Largest City Population %	0.232	0.197	0.026	1.033
Telephone Lines	0.406	0.192	0.001	0.892
Internet Penetration	0.119	0.187	0	0.860
IP Rights	3.687	0.982	0	4.88
Rule of Law	1.41	0.81	-1.57	2.17
Corruption	8.25	1.72	1.67	10

Notes: The statistics above are for the final dataset, where firms with more patents and more satellites are overrepresented (since observations are at the level of the headquarters patent – satellite). The (unweighted) average number of satellites for firms in our sample is 2.78, the average number of headquarter patents 190.61, and the average number of satellite patents is 7.45.

Table 3: Magnitude of Satellite Effect on Knowledge Diffusion to Host Country

	All Countries	**Developing Countries**	**Middle-Inc. Countries**	**Developed Countries**
% Treated Matching ($\hat{P}_T$)	3.37	0.029	0.51	5.18
% Controls Matching ($\hat{P}_C$)	2.95	0.026	0.44	4.55
t-statistic	150.07	4.43	24.83	146.14
Effect ($\hat{P}_T/\hat{P}_C$)	1.14	1.10	1.16	1.14
N - Treated Citations	75,980,240	14,322,071	11,628,819	47,619,688
N - Control Citations	80,535,008	15,445,707	12,206,718	50,452,796

Notes: The first column shows that while 3.37% of the citations received by headquarter (treated) patents are from third parties in their satellite's host country, only 2.95% of the citations received by the associated control patents are from the host country. The ratio of the two is 1.14, suggesting that host country firms receive on average 14% more knowledge from the headquarters as a result of having the satellite. Column 2 presents the results for host countries having a 1990 per capita GDP at current U.S. dollars of less than \$5000 (which includes Argentina, Brazil, China, Hungary, India, Malaysia, Mexico, Poland, South Africa, and Venezuela), column 3 for countries having a 1990 GDP per capita of between \$5000 and \$15000 (which includes Hong Kong, Israel, Singapore, South Korea, and Taiwan), and column 4 for countries having a 1990 GDP per capita of \$15000 or more (which includes Canada, France, Germany, Italy, Japan, Netherlands, Sweden, Switzerland, U.K., and U.S.). Countries for which 1990 GDP per capita was unavailable are omitted from columns 2-4.

Table 4: Difference in Differences Estimation of Satellite Effect

Dependent variable: $\widehat{P}$

	(1) 0 year window	(2) 1 year window	(3) 2 year window	(4) 5 year window	(5) 10 year window	(6) 15 year window
Treated	0.026***	0.025***	0.025***	0.024***	0.022***	0.019***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.004)
Post-Satellite	0.018***	0.019***	0.020***	0.023***	0.025***	0.027***
	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)	(0.002)
Treated*Post-Sat	0.019***	0.019***	0.020***	0.021***	0.023***	0.026***
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
Year fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	17,608,312	16,884,741	16,257,612	14,821,204	13,387,599	12,590,141
R-Squared	.1850	.1861	.1871	.1895	.1930	.1959

Notes: OLS regressions at the level of the patent-satellite, with treated (headquarters) and their associated control patents being separate observations. Standard errors are clustered by firm. The dependent variable, $\widehat{P}$, is the proportion of citations to the (treated or control) patent that are from third-party patents in the satellite location. Treated is a dummy variable equal to 1 if the patent was generated by the headquarters and 0 if the patent is a control. Post satellite is equal to 1 when the patent's application year is after the satellite's first patent application, and equal to 0 otherwise. Since the satellite's first patent application year will invariably be some time after the satellite has been established, I drop from the sample some of the years directly before the first patent application (it is not clear whether these should be considered pre or post). The columns present the results when dropping increasingly large windows from the analysis. As would be expected, the coefficient on treated decreases with the window size and the coefficient on the interaction term increases. Overall, the results suggest that a significant portion of the computed satellite effect is due to knowledge flows through the satellite and not to firms establishing satellites in locations where third-party firms already disproportionately cite them. Standard Errors in Parentheses. *p<0.05, **p<0.01, ***p<0.001

Table 5: Satellite Inventor Experience and Satellite Effect on Knowledge Diffusion

Dependent variable: ($\hat{P}_T$-$\hat{P}_C$) x 100,000

	(1)	(2)	(3)	(4)
Locals (diff. firm in sat. country)	8.040***	7.619***	8.533***	6.027**
	(2.187)	(2.036)	(2.300)	(2.486)
Others (diff. firm not in sat. country)	-8.868	1.791	-10.331	-4.423
	(12.046)	(12.627)	(12.072)	(17.041)
HQ Expats (same firm in HQ country)	8.525*	11.270**	9.765**	17.182***
	(5.030)	(5.291)	(4.976)	(6.023)
Expats (same firm in other country)	1.029	1.099	4.250	8.629
	(6.029)	(5.117)	(4.790)	(5.739)
Satelliters (same firm in sat. country)	-1.955**	-2.129***	-2.182***	-2.221***
	(0.771)	(0.571)	(0.693)	(0.574)
Satellite Patent Count	1.128*	0.544	1.158*	
	(0.637)	(0.397)	(0.622)	
HQ Patent Count	-0.005	0.002		
	(0.004)	(0.004)		
Firm Number of Locations	0.907	-4.663		
	(4.930)	(3.997)		
Year Fixed Effects	Yes	Yes	Yes	Yes
Tech Class Fixed Effects	Yes	Yes	Yes	Yes
Satellite Country Fixed Effects	No	Yes	No	-
Firm Fixed Effects	No	No	Yes	-
Satellite (Firm x Sat Country) FE	No	No	No	Yes
Observations	7,359,175	7,359,175	7,359,175	7,359,175
R-squared	.0008	.0022	.0074	.0164

Notes: OLS regressions at the level of the patent-satellite. Standard errors are clustered by firm. The dependent variable, $\hat{P}_T$-$\hat{P}_C$, is the proportion of citations to the treated patent that are from third-party patents in the satellite location minus the proportion of citations to the associated control patent that are from third-party patents in the satellite location. Our key explanatory variables *Locals*, *Others*, *HQ Expats*, *Expats*, and *Satelliters* are measures of the number of satellite inventors that previously patented for another firm in the satellite country, for another firm but not in the satellite country, for the headquarters, for the same firm but not in the satellite or HQ country, and for the satellite, respectively. All specifications include year and technology class fixed effects. In addition, (2) includes satellite country fixed effects, (3) includes firm fixed effects, and (4) includes satellite (interaction of firm and satellite country) fixed effects. Across all specifications, satellites with more *local* and *HQ expat* inventors generate larger HQ knowledge diffusion to host country firms, presumably because the first group has stronger local social networks and is more likely to leave the firm, and the second group embodies more HQ knowledge. Though the identified relationships are not causal but rather suggestive correlations. Standard Errors in Parentheses. *p<0.10, **p<0.05, ***p<0.01

Table 6: Host Country Characteristics and Magnitude of Satellite Effect on Knowledge Diffusion ($\hat{P}_T$-$\hat{P}_C$)

Dependent variable: ($\hat{P}_T$-$\hat{P}_C$)

	(1)	(2)	(3)	(4)
Ln(1+Patent Stock Per Capita)	1359***	1216***	1126***	1217***
	(491)	(389)	(370)	(431)
(Ln(1+Patent Stock Per Capita))^2	-1.61E7**	-1.43E7***	-1.32E7***	-1.36E7***
	(6.88E6)	(5.22E6)	(4.95E6)	(5.14E6)
Ln(GDP Per Capita)	-0.0133*	0.0026	0.0030	
	(0.0071)	(0.0050)	(0.0042)	
(Ln(GDP Per Capita))^2	-0.0011**	-0.0001	-0.0001	
	(0.0005)	(0.0003)	(0.0003)	
Geographic Distance	0.0001	0.0000		0.0000
	(0.0002)	(0.0001)		(0.0001)
Shared Language	0.0016	0.0022		0.0017
	(0.0014)	(0.0015)		(0.0016)
Shared Legal Origin	-0.0014	-0.0013		-0.0010
	(0.0011)	(0.0013)		(0.0013)
Largest City Fraction Population	0.0061	-0.0337	-0.0292	
	(0.0069)	(0.0255)	(0.0207)	
Telephone Lines	-0.0269**	0.0027	0.0050	
	(0.0113)	(0.0135)	(0.0120)	
Internet Penetration	-0.0031	0.0094	0.0070	
	(0.0093)	(0.0123)	(0.0106)	
Patent Rights	-0.0023	-0.0026	-0.0022	
	(0.0037)	(0.0021)	(0.0017)	
Rule of Law	-0.0023			
	(0.0037)			
Corruption	0.0007			
	(0.0007)			
Tech Class Fixed Effects	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	-
Sat Country Fixed Effects	No	Yes	-	-
Sat-HQ Country Pair Fixed Effects	No	No	Yes	No
Sat-Year Interacted Fixed Effects	No	No	No	Yes
Observations	6,476,276	6,476,276	6,476,276	6,476,276
R-squared	.0014	.0023	.0030	.0036

Notes: OLS regressions at the level of the patent-satellite. Standard errors are clustered by firm. The dependent variable, $\hat{P}_T$-$\hat{P}_C$, is the proportion of citations to the treated patent that are from third-party patents in the satellite location minus the proportion of citations to the associated control patent that are from third-party patents in the satellite location. Our key explanatory variables, the log of patent stock per capita and its square, are significant across all specifications and suggest an inverted-U relationship between the magnitude of the satellite effect and country technological capabilities. Standard Errors in Parentheses. *$p<0.10$, **$p<0.05$, ***$p<0.01$

Appendix A: Subsidiaries and Knowledge Diffusion

While the focus of this paper is to determine the effect of R&D satellites on knowledge diffusion, the effect of subsidiaries more generally is also of interest. Using Exhibit 21 of a firm's SEC filings, we collected data on the subsidiary locations of the 100 U.S. firms in our sample with the most patenting activity. Firms are only required to report "significant"[18] subsidiaries, though in practice they tend to report them more broadly than strictly required. The data was collected for the years 1993 (1994 for some firms) to 2006 since SEC records from prior to 1993 (1994) are not digitized and hence difficult to obtain.

These 100 firms represent 27% of all the treated patents in our sample. Together, they have 4385 foreign subsidiaries, with Procter and Gamble, Cisco Systems, and IBM having subsidiaries in the most countries (101, 99, and 96, respectively). By contrast, in the main analysis of this paper, these same 100 firms have only 987 R&D satellites (as measured by the presence of foreign patenting activity). The overlap between the subsidiary and satellite countries is significant. Of the 987 satellites, 834 (85%) are in countries where the SEC data indicate the presence of a subsidiary for that firm. For the other 153 satellites, it may be that the firm's presence in that country is not deemed to be significant enough to report in Exhibit 21 of the SEC filing. Consistent with this, if we restrict attention to the 297 satellites with at least 10 patents, 95% of these coincide with subsidiary countries in the SEC data.

We estimated the subsidiary effect using the methodology described in sections 3.1 and 3.2, but using the country of subsidiaries as reported in the SEC data instead of the country of satellites. The computed subsidiary effect (column 1 in Table A1 below) is almost identical to that computed in Table 3. For ease of comparison, column 2 below reproduces the main result in Table 3 and column 3 presents the satellite effect (as in Table 3) but computed for the subset of the 100 firms. The only significant differences are that the fraction of citations to treated and controls that match are lower for the subsidiary effect, though these are almost identical to those reported in Table 3 for middle income countries (which are more heavily represented in the subsidiary vs. the satellite data). In addition, we now have many more observations since, for these very large firms which make up a good fraction of our sample of treated patents, there are 4.5 times more subsidiaries than there are satellites.

Table A2, column 1, presents the difference in differences results from estimating equation (1) using the subsidiary data (the sample is restricted to the 100 firms and for the years 1993 to 2006). We continue to find a significant satellite effect, though the magnitude of the coefficient on the interaction term is significantly smaller (for ease of comparison, column 2 reproduces the 0-year window results from Table 4). This could be due to subsidiaries having a smaller effect on knowledge diffusion (once you account for firms establishing subsidiaries in countries that were always disproportionately citing them) or to sample selection. As shown in column (3), the coefficient on the interacted

[18] In general, a subsidiary is deemed to be "significant" if it represents 10% or more of the parent's assets or income.

term is also significantly smaller when we estimate the satellite effect but restrict the sample to the 100 firms.

Table A1: Magnitude of Subsidiary Effect on Knowledge Diffusion to Host Country

	Subsidiary Effect (100 firms)	Satellite Effect (all firms)	Satellite Effect (100 firms)
% Treated Matching ($\hat{P}_T$)	0.51	3.37	1.43
% Controls Matching ($\hat{P}_C$)	0.43	2.95	1.21
t-statistic	99.47	150.07	91.72
Effect ($\hat{P}_T/\hat{P}_C$)	1.18	1.14	1.18
N - Treated Citations	142,750,704	75,980,240	46,406,888
N - Control Citations	153,173,408	80,535,008	49,612,884

Table A2: Difference in Differences Estimation of Subsidiary Effect

Dependent variable: $\hat{P}$

	Subsidiary Effect (100 firms)	Satellite Effect (All firms) (0-year window)	Satellite Effect (100 firms) (0-year window)
Treated	0.00273***	0.026***	0.0114***
	(0.00075)	(0.003)	(0.0015)
Post-Sub/Sat	0.00353***	0.018***	0.0104***
	(0.00060)	(0.001)	(0.0009)
Treated*Post-Sub/Sat	0.00378***	0.019***	0.0073***
	(0.00097)	(0.004)	(0.0016)
Year Fixed Effects	Yes	Yes	Yes
Firm Fixed Effects	Yes	Yes	Yes
Observations	12,760,799	17,608,312	9,297,581
R-Squared	0.0056	0.1850	0.0403

Notes: OLS regressions at the level of the patent-subsidiary or patent-satellite, with treated (headquarters) and their associated control patents being separate observations. Standard errors are clustered by firm. The dependent variable, $\hat{P}$, is the proportion of citations to the (treated or control) patent that are from third-party patents in the subsidiary/satellite location. Standard Errors in Parentheses. *p<0.05, **p<0.01, ***p<0.001

Appendix B: Pre-Trends in Difference in Differences

To examine pre-trends, we include in our difference in differences estimation dummies for one to five years prior to the measured date of establishment, and their interaction with *treated*. As shown in Table B1, there is some evidence of pre-trends. All five pre-establishment interaction terms are positive, though only two are significant. We find a large and strongly significant coefficient for the interaction term in the year before the establishment. In addition, the interaction term for 4 years pre-establishment is also significant at the 5% level, though marginally so.

A real pre-trend would be worrisome for the interpretation of the interaction coefficient as the satellite treatment effect on the treated. But as discussed, it is difficult to determine whether this is due to a real pre-trend or to most satellites having been present a number of years before the first observed patent application. While we cannot rule out a real pre-trend, the fact that the pre-trend is only large and strongly significant for the one year prior to the measured year of establishment suggests that mismeasurement may explain the pre-trend.

Table B1: Difference in Differences Estimation with Pre-Trends

Dependent variable: $\hat{P}$

	(1)
Treated	0.0239***
	(0.0028)
5 Year Pre-Satellite	0.0065***
	(0.0008)
4 Year Pre-Satellite	0.0070***
	(0.0009)
3 Year Pre-Satellite	0.0088***
	(0.0010)
2 Year Pre-Satellite	0.0108***
	(0.0011)
1 Year Pre-Satellite	0.0115***
	(0.0012)
Post-Satellite	0.0238***
	(0.0016)
Treated * 5 Year Pre-Sat	0.0023
	(0.0012)
Treated * 4 Year Pre-Sat	0.0025*
	(0.0012)
Treated * 3 Year Pre-Sat	0.0021
	(0.0013)
Treated * 2 Year Pre-Sat	0.0026
	(0.0015)
Treated * 1 Year Pre-Sat	0.0052**
	(0.0017)
Treated * Post-Satellite	0.0202***
	(0.0038)
Year fixed Effects	Yes
Firm fixed effects	Yes
Observations	17,608,312
R-Squared	.1854

Notes: OLS regressions at the level of the patent-subsidiary or patent-satellite, with treated (headquarters) and their associated control patents being separate observations. Standard errors are clustered by firm. The dependent variable, $\hat{P}$, is the proportion of citations to the (treated or control) patent that are from third-party patents in the subsidiary/satellite location. Standard Errors in Parentheses. *p<0.05, **p<0.01, ***p<0.001

Appendix C: Placebo Satellite Countries

As a robustness check, we perform the analysis to compute the satellite effect but with placebo countries in place of the actual satellite country. In particular, we keep the exact same dataset as before but $\hat{P}_T$ is now computed as the proportion of citations received by treated patents that are from third-party patents in a randomly chosen (from among the 215 countries with USPTO patents) placebo country. Similarly, $\hat{P}_C$ is computed as the proportion of citations received by patents in the control group that are from third-party patents in the placebo country. The results are presented in Table C1.

Table C1: Magnitude of Satellite Effect with Placebo Satellite Countries

	Placebo Effect
% Treated Matching ($\hat{P}_T$)	0.463
% Controls Matching ($\hat{P}_C$)	0.466
t-statistic	-3.41
Effect ($\hat{P}_T/\hat{P}_C$)	0.99
N - Treated Citations	75,980,240
N - Control Citations	80,535,008

As expected, we no longer find a satellite effect. In fact, we now obtain a ($\hat{P}_T/\hat{P}_C$) ratio that is slightly less than unity.

We can similarly use the $\hat{P}_T$ and $\hat{P}_C$ computed using the placebo countries to estimate a difference in differences regression as in equation (1). The results obtained from this exercise are presented in Table C2. All of our coefficients are indistinguishable from zero.

Table C2: Difference in Differences Estimation with Placebo Satellite Countries

Dependent variable: $\hat{P}$

	(1) 0-year window	(2) 1-year window	(3) 2-year window	(4) 5-year window	(5) 10-year window	(6) 15-year window
Treated	-0.0000	0.0000	0.0000	0.0001	-0.0000	-0.0001
	(0.0000)	(0.0000)	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Post-Satellite	-0.0000	-0.0000	-0.0000	0.0000	-0.0001	-0.0000
	(0.0000)	(0.0000)	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Treated*Post-Sat	0.0000	0.0000	0.0000	-0.0000	0.0001	0.0001
	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Year fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	17,608,312	16,884,741	16,257,612	14,821,204	13,387,599	12,590,141
R-Squared	.0040	.0041	.0043	.0046	.0050	.0052

Notes: OLS regressions at the level of the patent-satellite placebo, with treated (headquarters) and their associated control patents being separate observations. Standard errors are clustered by firm. The dependent variable, $\hat{P}$, is the proportion of citations to the (treated or control) patent that are from third-party patents in the placebo country. Standard Errors in Parentheses. *p<0.05, **p<0.01, ***p<0.001

Appendix D: Placebo Date of Satellite Establishments

As a robustness check, we randomly assigned an establishment year (application year of first patent) to each satellite and performed our difference in differences estimation as in (1) using these dates to define the pre- and post-satellite periods. As shown in Table D1, regardless of window size, the coefficient on the interaction term is indistinguishable from zero.

Table D1: Difference in Differences Estimation with Placebo Year of Satellite Establishment

Dependent variable: $\hat{P}$

	(1) 0-year window	(2) 1-year window	(3) 2-year window	(4) 5-year window	(5) 10-year window	(6) 15-year window
Treated	0.0370***	0.0369***	0.0368***	0.0368***	0.0366***	0.0375***
	(0.0033)	(0.0034)	(0.0035)	(0.0038)	(0.0046)	(0.0055)
Post-Satellite	-0.0019**	-0.0021**	-0.0022**	-0.0027**	-0.0052***	-0.0081***
	(0.0007)	(0.0008)	(0.0008)	(0.0010)	(0.0013)	(0.0019)
Treated*Post-Sat	0.0031	0.0032	0.0033	0.0036	0.0043	0.0036
	(0.0043)	(0.0045)	(0.0046)	(0.0049)	(0.0055)	(0.0063)
Year fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	17,608,312	17,168,781	16,730,740	15,403,508	13,559,959	12,333,800
R-Squared	.1785	.1786	.1786	.1801	.1840	.1886

Notes: OLS regressions at the level of the patent-satellite country, with placebo year of satellite establishment. Treated (headquarters) patents and their associated control patents are separate observations. Standard errors are clustered by firm. The dependent variable, $\hat{P}$, is the proportion of citations to the (treated or control) patent that are from third-party patents in the satellite country. Standard Errors in Parentheses. *p<0.05, **p<0.01, ***p<0.001

Appendix E: Satellite Inventor Experience and Satellite Effect ($\hat{P}_T/\hat{P}_C$)

Table E1: Satellite Inventor Experience and Satellite Effect on Knowledge Diffusion

Dependent variable: ($\hat{P}_T/\hat{P}_C$) x 100,000

	(1)	(2)	(3)	(4)
Locals (diff. firm in sat. country)	1.957***	0.335*	2.136***	-0.224
	(0.319)	(0.172)	(0.450)	(0.311)
Others (diff. firm not in sat. country)	-6.277***	0.103	-7.354**	2.844***
	(2.226)	(0.841)	(3.163)	(1.085)
HQ Expats (same firm in HQ country)	3.878***	0.297	4.285***	-0.353
	(0.632)	(0.344)	(1.013)	(0.449)
Expats (same firm in other country)	-0.975	-1.299	4.995	5.930**
	(4.988)	(2.298)	(4.208)	(2.546)
Satelliters (same firm in sat. country)	0.993***	-0.327***	-1.025***	-0.126
	(0.137)	(0.049)	(0.197)	(0.087)
Satellite Patent Count	0.722***	0.235***	0.868***	
	(0.075)	(0.046)	(0.122)	
HQ Patent Count	-0.005***	-0.001		
	(0.002)	(0.001)		
Firm Number of Locations	-9.767***	0.849		
	(2.462)	(1.575)		
Year Fixed Effects	Yes	Yes	Yes	Yes
Tech Class Fixed Effects	Yes	Yes	Yes	Yes
Satellite Country Fixed Effects	No	Yes	No	-
Firm Fixed Effects	No	No	Yes	-
Satellite (Firm x Sat Country) FE	No	No	No	Yes
Observations	754,525	754,525	754,525	754,525
R-squared	.0128	.0369	.0243	.0518

Notes: OLS regressions at the level of the patent-satellite. Standard errors are clustered by firm. The dependent variable, $\hat{P}_T/\hat{P}_C$, is the proportion of citations to the treated patent that are from third-party patents in the satellite location divided by the proportion of citations to the associated control patent that are from third-party patents in the satellite location. This dependent variable, while being a direct measure of the satellite effect, results in the loss of 90% of the sample due to observations where the control patent has received no citations from the satellite country ($\hat{P}_C$=0). Our key explanatory variables *Locals*, *Others*, *HQ Expats*, *Expats*, and *Satelliters* are measures of the number of satellite inventors that previously patented for another firm in the satellite country, for another firm but not in the satellite country, for the headquarters, for the same firm but not in the satellite or HQ country, and for the satellite, respectively. All specifications include year and technology class fixed effects. In addition, (2) includes satellite country fixed effects, (3) includes firm fixed effects, and (4) includes satellite (interaction of firm and satellite country) fixed effects. The coefficients on *local* and *HQ expat* are consistent with those presented in Table 5 for the first three specifications, but adding satellite fixed effects drastically changes the results, presumably because the time variation in our explanatory variables is not accurate and/or because of sample selection. Standard Errors in Parentheses. *p<0.10, **p<0.05, ***p<0.01

Appendix F: Host Country Characteristics and Satellite Effect ($\hat{P}_T/\hat{P}_C$)

Table F1: Host Country Characteristics and Magnitude of Satellite Effect on Knowledge Diffusion ($\hat{P}_T/\hat{P}_C$)

Dependent variable: ($\hat{P}_T/\hat{P}_C$)

	(1)	(2)	(3)	(4)
Ln(1+Patent Stock Per Capita)	40,842***	23,353***	22,687***	25,074***
	(7942)	(6164)	(6179)	(6070)
$(\text{Ln(1+Patent Stock Per Capita)})^2$	-6.63E8***	-3.35E8***	-3.27E8***	-3.46E8***
	(1.69E8)	(1.13E8)	(1.11E8)	(1.06E8)
Ln(GDP Per Capita)	-1.6700**	1.4045***	1.4094	
	(0.6912)	(0.1923)	(0.1817)	
$(\text{Ln(GDP Per Capita)})^2$	0.1046***	-0.0671***	-0.0675***	
	(0.0362)	(0.0103)	(0.0096)	
Geographic Distance	0.0176***	-0.0010		-0.0020
	(0.0065)	(0.0041)		(0.0042)
Shared Language	0.0969	0.2039***		0.1926***
	(0.0757)	(0.0538)		(0.0510)
Shared Legal Origin	-0.0787	-0.1294***		-0.1243***
	(0.0580)	(0.0362)		(0.0352)
Largest City Fraction Population	0.7895**	1.5858***	1.8500***	
	(0.3638)	(0.5623)	(0.5169)	
Telephone Lines	0.2170	0.6530	0.6990*	
	(0.5395)	(0.4189)	(0.3896)	
Internet Penetration	0.9659	0.5457**	0.5199**	
	(0.5988)	(0.2531)	(0.2562)	
Patent Rights	0.3797***	0.0503	0.0502	
	(0.1207)	(0.0481)	(0.0478)	
Rule of Law	-0.1462			
	(0.3209)			
Corruption	-0.1050			
	(0.0856)			
Tech Class Fixed Effects	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	No
Sat Country Fixed Effects	No	Yes	No	No
Sat-HQ Country Pair Fixed Effects	No	No	Yes	No
Sat-Year Interacted Fixed Effects	No	No	No	Yes
Observations	705,968	705,968	705,968	705,968
R-squared	.0305	.0378	.0386	.0036

Notes: OLS regressions at the level of the patent-satellite. Standard errors are clustered by firm. The dependent variable, $\hat{P}_T/\hat{P}_C$, is the proportion of citations to the treated patent that are from third-party patents in the satellite location divided by the proportion of citations to the associated control patent that are from third-party patents in the satellite location. This dependent variable, while being a direct measure of the satellite effect, results in the loss of 90% of the sample due to observations where the control patent has received no citations from the satellite country ($\hat{P}_C$=0). Our key explanatory variables, the log of patent stock per capita and its square, are significant across all specifications and suggest an inverted-U relationship between the magnitude of the satellite effect and country technological capabilities. Standard Errors in Parentheses. *p<0.10, **p<0.05, ***p<0.01

The Canadian Journal of Economics
La Revue canadienne d'économique

Disclosure Statement

Under our *Conflicts of Interest and Ethical Guidelines for Authors*, this form should be appended to the submitted paper so as to be available to the editor and referees. For submissions with more than one author, a form must be completed by and attached for each of the authors. A short statement summarizing the information in this form will be included in the acknowledgment footnote of the published version of the paper. If one of the situations does not apply to you, please write "None".

Paper title:
Foreign R&D satellites as a medium for the international diffusion of knowledge

Sources of financial support:
Social Sciences and Humanities Research Council (SSHRC) PGS A Scholarship
Institute for New Economic Thinking/Centre for International Governance Innovation Research Grant

Interested parties that provided financial or in-kind support:
None

Paid or unpaid positions in organizations with a financial or policy interest in this paper:
Note: I am a Senior Research Fellow at the Centre for International Governance Innovation (CIGI). However, I do not believe that CIGI has a direct policy interest in this paper.

If the paper was partly or wholly written under contract, please write the name of the organization:N/A

Does the paper satisfy all CJE guidelines on originality, authorship, and conflicts of interest as set out in the CJE's Conflict of Interest and Ethical Guidelines for Authors?
Yes

Name, affiliation, and contact information:
Joel Blit
Department of Economics
University of Waterloo
200 University Ave. West
Waterloo, ON
(226)808-3673