创业与管理学院
School of Entrepreneurship and Management

SHANGHAITECH SEM WORKING PAPER SERIES
No. 2018-008

The Role of R&D Offshoring in Knowledge Diffusion: New Evidence from China

Guangwei Li
ShanghaiTech University

October 2017
https://ssrn.com/abstract=3160519

School of Entrepreneurship and Management
ShanghaiTech University
http://sem.shanghaitech.edu.cn

The Role of R&D Offshoring in Knowledge Diffusion: New Evidence from China

GUANGWEI LI*

ABSTRACT: While R&D offshoring to emerging economies creates new opportunities for knowledge creation and acquisition, it may also increase the risk of losing critical technological knowledge to indigenous competitors in the host country. Nearly every prior empirical study has shown that the establishment of foreign R&D centers increases knowledge flows to local firms. Business leaders and policymakers have expressed particular concern about this outcome in China, where multinational corporations (MNCs) are rapidly increasing their R&D offshoring, and legal protection for intellectual property rights remains a work in progress. This paper uses new empirical analyses and field interviews to demonstrate that multinational R&D offshoring in China appears to reduce knowledge flows to local firms, not increase them. First, following the literature, I use a matching procedure to construct a sample consisting of focal and control patents, and a t-test to compare citation rates received by focal and control patents from Chinese organizations. Next, I adapt a Difference-in-Differences (DD) approach to address the endogeneity problems of the t-test. Both the t-test and DD estimation suggest R&D offshoring is associated with a lower probability of knowledge flow from the foreign multinational to Chinese organizations. This is surprising because it contradicts not only the existing findings in the literature, but also the perceptions widely held by business leaders and policymakers. Using insights obtained from field interviews, I argue that as multinationals expand R&D in China they also enhance their ability or potential to monitor the learning and invention of their indigenous competitors. By combining this ability or potential with other strategic tools, multinationals can deter indigenous Chinese firms and other organizations to build on MNC R&D. Taking advantage of social network analysis tools, I empirically test this hypothesis and find that a higher monitoring capacity due to offshoring is associated with a lower probability of knowledge diffusion to Chinese organizations. This study calls for a reconsideration of the implications, both for corporate strategy and public policy, of the current trend of multinational R&D offshoring in emerging markets.

1 Introduction

Does the offshoring of multinational corporations' (MNCs') research and development (R&D) to emerging markets inadvertently increase the diffusion of valuable knowledge to domestic organizations? While the answer to this question is crucial, scholars have provided

* I gratefully acknowledge the financial support of the Carnegie Bosch Institute. The views expressed in this paper are those of the author, and I retain sole responsibility for any errors and omissions.

surprisingly little empirical evidence about it.[1] This question is not just for academics, however, but one for businesses and policymakers. From the business perspective, the world's largest MNCs are facing a dilemma. On one hand, they are increasingly interested in conducting R&D in emerging markets (Barrett, van Biljon, and Musso 2011). For many MNCs, emerging markets are just too large to ignore. In order for MNCs to obtain market access, to meet local customer tastes, or to take advantage of local scientific and engineering talent pools, some R&D needs to be conducted locally (UNCTAD 2005; Basant and Mani 2012; Branstetter, Li, and Veloso 2015). On the other hand, MNCs are concerned that conducting R&D in emerging markets may increase the risk of losing critical know-how to domestic competitors especially given the fact that property rights in these markets are often underdeveloped or under-enforced. The IP issue becomes even more salient as scholars have found that multinational R&D subsidiaries function as a channel through which knowledge created by the multinational outside of a particular host country flows to firms located within that host country (Blit 2015). As a result of this knowledge protection concern, many MNCs simply limit the amount of R&D they perform in emerging countries (Barrett, van Biljon, and Musso 2011). This paper provides new empirical evidence on the degree to which multinational R&D offshoring in China affects indigenous Chinese organizations' knowledge acquisition. It seeks to measure the degree to which local R&D subsidiaries of multinationals allow indigenous Chinese organizations to directly tap knowledge created by the multinationals outside of China.

Knowledge flows are hard to measure. However, in their seminal work, Jaffe et al. (1993, p.578) indicate that "knowledge flows do sometimes leave a paper trail, in the form of citations

[1] The research in this area is underdeveloped for two reasons: First, offshoring R&D to emerging markets is a relatively new phenomenon. Second, collecting worthwhile, high-quality data from emerging markets is a very difficult task.

in patents." By tracking patent citations of previous patents (i.e., backward citation), researchers can gain useful insights into knowledge diffusion patterns (Jaffe and Trajtenberg 2002). Thus, in this paper, I use patent citation as an indicator of knowledge flows. Patent citations have their issues as a means of tracking knowledge flows, and I discuss these in detail in the background section.

In this paper, I want to find the change (if any) in patent citations associated with multinational R&D offshoring. To do this, I need to have a sense of the normal level of citations. To find that benchmark level of citations, I follow the "best practice" in the literature in similar settings by adapting a matching procedure to construct samples consisting of both focal patents generated by multinationals with R&D facilities in China and the corresponding control patents—technologically similar patents generated by organizations without R&D facilities in China. I use the control patents as a baseline to capture the "treatment" effect of establishing an R&D subsidiary in China on the path of citations received by the focal patents.

As I conduct this analysis, I have to contend with the challenge of "selection on observables." Prior research has suggested that MNCs choosing to offshore R&D to China are likely to be firms that have special traits, such as hiring better inventors (Branstetter, Li, and Veloso 2015). As a result, the focal patents are not necessarily comparable with the control patents. There might be preexisting differences between the two sets of patents. To control for this issue and other possible confounders, I adapt a Difference-in-Differences (DD) approach.

The final challenge is that the DD estimate might be influenced by pre-existing trends in citations received by focal patents. To address this question empirically, I conduct an event study of focal patents by estimating a series of time indicator coefficients for focal patents during periods both before and after the R&D offshoring.

The empirical work of this study points to an important conclusion: MNCs that have offshored R&D to China seem to experience a decline in the knowledge flow from their R&D operations outside China to Chinese organizations, not the opposite as suggested by the existing literature. I find that patents created outside of China by MNCs with R&D subsidiaries in China receive, on average, significantly fewer citations from Chinese organizations than randomly picked comparable control patents owned by organizations that have no R&D activities in China. In addition, I examine the other two types of knowledge flows: the one from MNCs Chinese R&D subsidiaries to Chinese organizations and the one from MNCs non-China R&D facilities to their Chinese subsidiaries. I find that (1) China-originating patents owned by MNCs with China-based R&D facilities do not appear to receive more citations from Chinese organizations than comparable patents owned by the same MNC originating outside of China and (2) patents created by these MNCs outside of China receive, on average, significantly more citations from their China subsidiaries than from randomly picked comparable control patents owned by organizations that have no R&D activities in China.

The finding that patents created outside of China by MNCs with China-based R&D facilities receive, on average, significantly fewer citations from Chinese organizations than comparable control patents is particularly surprising because it contradicts not only existing findings in the literature (based on data in advanced countries), but also the perceptions commonly held by scholars and business managers. The finding survives various robustness checks and seems to be well grounded. So, what might be the possible mechanism behind this result?

Insights from field interviews suggest that when MNCs offshore R&D to China, there are two possible countervailing effects on Chinese organizations. On one hand, offshoring creates learning opportunities for Chinese organizations. On the other hand, as MNCs expand R&D in

China, they also increase their ability or potential (I call it potential because MNCs may not necessarily realize that they have this capacity) to monitor what is happening locally in terms of technological activity. By leveraging this ability or potential with other strategic tools, MNCs may deter Chinese organizations from building on multinational invention. This finding suggests an unusual benefit of moving R&D operations to emerging markets that has not caught the attention of international business scholar. This paper thus provides a suggestive guidance for future research.

I organize the remainder of the paper as follows: Section 2 provides background information on my rationale for choosing China, some key concepts, and the related literature. Section 3 outlines my empirical framework. Section 4 presents my main empirical results and robustness checks. Section 5 describes the deterrence theory based on insights from field interviews. Finally, Section 6 concludes the paper and discusses the limitations and implications of my findings.

2 Background

2.1 Why China

I focus on China for several reasons. First, China is a leading location of MNC R&D in emerging economies. The joint presence of a large market and a sizeable scientific and engineering talent pool makes China an increasingly important site for foreign MNCs' R&D investment. Over the 1997–2011 period, the total amount of US multinational R&D spending in China increased 33 fold, from 35 million to 1.17 billion US dollars, more than the increase in any other emerging economy.[2] According to the Ministry of Commerce of China, foreign

[2] The US spending increase represents majority-owned affiliates of non-bank US parent companies in China. A majority-owned affiliate is a Chinese affiliate in which the combined direct and indirect ownership interest of all US parents exceeds 50%. Source: US Department of Commerce, Bureau of Economic Analysis, US Direct Investment Abroad: Financial and Operating Data for US Multinational Companies, accessed August 8, 2012, http://www.bea.gov/iTable/iTable.cfm?ReqID=2&step=1.

corporations had set up over 1,600 R&D centers in China by the end of 2011 (Ministry of Commerce of China 2012). Second, China has the largest amount of MNC-sponsored US patent grants among all emerging economies, which allows us to conduct meaningful empirical tests. Third, the leading Chinese domestic enterprises are patenting a considerable number of inventions in the US. At an aggregated level, they own the largest amount of US patents among all other emerging economies like India, Russia, and Brazil. If there is knowledge diffusion from MNCs to domestic firms, it is most likely to happen in China. Fourth, China is regarded by the US government as a market that fails to properly protect intellectual property, and this has been the focus of an intense policy debate. According to the Commission on the Theft of American Intellectual Property, China accounts for 50–80 percent of all theft of US intellectual property, amounting to an estimated $300 billion loss in businesses every year (The Commission on the Theft of American Intellectual Property 2013). Thus, IP-related research based on China is of great relevance to policymakers.

2.2 Knowledge Diffusion, Flow, Spillover, and Leakage

Before turning to the empirical work, I need to define several important and related terms. In this paper, the term "knowledge diffusion" refers to the movement of technological knowledge from one organization to another organization. Obviously, knowledge diffusion consists of both market transactions and externalities (Keller 2004). However, for my purpose, I narrow the definition to include only the movement of technological knowledge due to externalities. By my definition, pure imitation (i.e., direct copying or theft) is also knowledge diffusion. I use the term "knowledge flow" in a similar sense to knowledge diffusion. The difference between knowledge flow and knowledge diffusion is that the former term emphasizes the act of moving itself (i.e., knowledge moves) and the latter term emphasizes the consequence of movement (i.e., knowledge moves from A to B).

The term “knowledge spillover” refers to the process by which one organization learns from the research outcomes of other organizations’ research projects and then uses this knowledge to generate further innovation without fully compensating the other organizations for the values of learning (Branstetter 2006). In contrast to knowledge diffusion, knowledge spillovers is defined as externalities that lead to further innovation. Thus, by my definition, a knowledge spillover is a type of knowledge diffusion. Pure imitation per se is not a knowledge spillover. Yet, knowledge spillovers may contain elements of imitation as long as these elements eventually lead to further innovation.

I use the term “knowledge leakage” mainly for the convenience of my narrative. Throughout the paper, knowledge leakage refers to the knowledge diffusion that is unintended from the knowledge owner’s perspective. This is a relative term since the knowledge owner’s unintended result could be the knowledge seeker’s favored outcome.

2.3 Using Patent Citations to Track Knowledge Flows

The disembodied nature of knowledge poses a challenge to anyone who wants to measure its movement. As Krugman notes, “knowledge flows […] are invisible, they leave no paper trail by which they may be measured and tracked” (Krugman, 1991, p.53). However, in their seminal work, Jaffe et al. (1993, p.578) point out that “knowledge flows do sometimes leave a paper trail, in the form of citations in patents.” By tracking patent citations made to previous patents (i.e., backward citation), researchers can gain useful insights into knowledge diffusion patterns (Jaffe and Trajtenberg 2002). In addition, patent citations also help researchers to distinguish “pure” knowledge spillovers from rent spillovers (Griliches 1979), [3] or the factors that reflect market power (Branstetter and Sakakibara 2002; Bloom, Schankerman, and Van Reenen 2013).

US patent law obligates patent applicants to make "appropriate citations to the prior art" in their patent applications. The primary reason for including citations in patent documents is to disclose the "prior art" upon which inventors build. By explicitly listing the prior art in the patent documents, patent citations serve the important legal function of delimiting the scope of the property rights awarded by the patent. Thus "if patent B cites patent A, it implies that patent A represents a piece of previously existing knowledge upon which B builds, and over which B cannot have a claim" (Hall, Jaffe, and Trajtenberg 2001, p.15).

It is worth noting that patent citations are a noisy measure of knowledge diffusion patterns since many citations are derived from the patent process. By comparing the distribution of examiner-added and non-examiner-added citations, Alcácer and Gittelman (2006) have showed that examiner-added patents indeed add measurement errors when used to measure knowledge spillovers. Fortunately, since 2001, the United States Patent and Trademark Office (USPTO) has included an identifier in the granted patent data to distinguish the examiner-added citations from the non-examiner-added citations. The identifier enables me to remedy some of the measurement errors by using only non-examiner citations. That being said, patent citations could also be added by patent lawyers as well as other parties rather than inventors themselves; patent citations could be added by the inventors after they have completed the invention as well. Similar to examiner-added citations, these citations are probably unrelated to knowledge diffusion, yet there is no way to distinguish these citations from the ones made by inventors themselves before or during

[3] Rent spillovers occur when a product's improved quality is not reflected by a concurrent price increase. Distinguishing the effect of "pure" knowledge spillovers from the effect of rent spillovers is difficult when using the production function approach to infer knowledge spillovers.

the time of invention. Jaffe, Trajtenberg, and Fogarty (2000) provide some insights into the magnitude of the noise caused by these citations. They surveyed inventors to validate patent citations as a measure of knowledge spillovers and found that while perhaps one-half of the responses indicated no spillover, one-quarter of the responses corresponded to fairly clear spillovers, and the remaining one-quarter indicated some possibility of a spillover. The patents they used in the survey did not have identifiers to distinguish examiner-added citations. The examiner-added patents account for about 31 percent of all citations made for patents that were granted during 2001–2012. Thus, it is reasonable to expect that the majority of the non-examiner-added citations that I use in my analysis would reflect at least some possibility of knowledge spillover. The remaining measurement errors are mainly introduced by citations added by firms' patent lawyers.

2.4 Related Literature

Two streams of literature are related to this paper. The first stream appears in the international economics literature, and uses a production function approach to study the impact of foreign direct investment (FDI) on the measured productivity of indigenous firms. Research based on Chinese enterprise- and industry-level data generally finds evidence of a positive impact of FDI on the productivity growth of indigenous Chinese firms (e.g., Du, Harrison, and Jefferson 2012; Lin, Liu, and Zhang 2009; Hu and Jefferson 2002). However, as noted above, the productivity spillovers observed in this stream of studies fail to distinguish "pure" knowledge spillovers from rent spillovers (Griliches 1979), or other factors that reflect market power (Branstetter and Sakakibara 2002; Bloom, Schankerman, and Van Reenen 2013).

The second stream of research to which this paper contributes appears in the economics of innovation literature and in the international business and strategy literature, and it uses patent citation data to study the extent of knowledge flows between MNCs and the host-country organizations. The first sub-stream of this literature focuses on the factual aspects of the issue, with the aim of answering questions such as whether there are knowledge flows from MNCs to host-country organizations and reverse knowledge flows from host-country organizations to MNCs. Data based on organizations in advanced economies, such as the US, Japan, and west European countries, generally point to a consensus that there are knowledge spillovers in both directions (e.g., Almeida 1996; Branstetter 2006; Singh 2007). The gap in this sub-stream of literature, however, is that we do not know to what extent these findings are also applicable to emerging economies, where both the absorptive capacity (Cohen and Levinthal 1990) of local organizations and the legal system protecting IP are quite different from those in advanced countries.

The second sub-stream of this literature focuses on the strategic aspects of the phenomenon, aiming to tackle questions such as what kinds of strategic arrangements can help MNCs prevent knowledge leakage to host-country organizations. Findings in this sub-stream provide interesting insights into how MNCs can leverage internal mechanisms to reduce knowledge leakage to host-country firms (e.g., Alcácer and Zhao 2012). However, this literature generally does not provide much insight into what kinds of strategic arrangements work in emerging economies. One exception is Zhao (2006). She finds that MNCs that are conducting R&D in countries with weak intellectual property rights (IPR) protection show stronger internal linkages in both their technologies developed in countries with weak IPR protection and in their overall technology portfolios. She argues that those firms are using closely knit internal technological structures as an alternative mechanism to protect their intellectual property in countries with weak IPR protection. The goal of this paper is to fill in the gaps in both sub-streams of literature. On one hand, I aim to provide factual evidence about the extent of knowledge diffusion from MNCs to Chinese organizations. On the other hand, I also hope to reveal the implications of my findings for corporate strategy.

3 Empirical Framework

3.1 Data Sources

I use US patent grants originating in China (i.e., with at least one inventor residing in China, as shown in the patent document) as evidence of multinational R&D activity in China. I consider a foreign MNC (that is, an MNC based outside of China) as having offshored R&D to China if it has at least 10 China-originating patents by the end of 2012.[4] My empirical analyses focus on foreign MNCs that meet this criterion.

I collect patent data from several sources. The bibliographic information for patents is drawn from the Disambiguation and Co-authorship Networks of the US Patent Inventor Database (Lai et al. 2011), which contains all granted patents between 1975 and 2010. The citation data are drawn from the Examiner Citation Data (Sampat 2012), which contains examiner-added and non-examiner-added citations between 2001 and 2010. Patents issued in 2011 and 2012 and their citation information are obtained directly from the Google USPTO Bulk Download: Patent Grant Bibliographic Data. I drop withdrawn patents from the datasets[5] and update patent classes to Current Classifications as of the end of 2013 according to the Google USPTO Bulk Download: US Patent Grant Master Classification File (2013 December).

4 This threshold is somewhat arbitrary. Among firms with fewer than 10 China-originating patents, a majority have only one China-originating patent during 1981–2012, according to patent grant years. These firms have hardly engaged in R&D efforts in China. In addition, the threshold reduces noise from firms that have patents with inventors who list China addresses but do not actually conduct R&D in China. For example, field interviews have revealed an instance where a Chinese PhD took a post-doc position in a US lab and got involved in an invention. When the firm filed a patent for the invention, this person had moved back to China and was listed with an address in China. In my research, measurement error related to such circumstances is reduced by choosing 10 China-originating patents as the threshold.

5 USPTO website, http://www.uspto.gov/patents/process/search/withdrawn.jsp.

I regard the first assignee of a patent as the primary owner of the patent. I standardize the assignees according to the USPTO's assignee harmonization system.[6] Based on the harmonized assignee codes, I further code assignees' types, identify the assignees' nationalities, and merge assignee codes to the ultimate owners of the China-originating patents.

3.2 A t-test Approach

As a starting point, I follow previous research in using a two-sample t-test to explore the impact of MNC R&D offshoring on Chinese organizations' knowledge acquisition. To do the test, I first create a data set consisting of all patents generated by inventors residing outside of China and owned by an MNC that has offshored R&D to China by the end of 2012.[7] I refer to these as "focal patents." Next, I match each focal patent with a comparable "control patent" selected from a data set consisting of all patents that are owned by organizations that have no China-originating patents,[8] with the same technological class, subclass,[9] inventor nationality, and application date. In the case in which no exact matches are found, I release the application date criterion to include patents with the closest application date within a one-year range of the application date of the focal patent (either before or after). If several potential control patents are found, one is chosen randomly. Finally, any focal patent with no matched control is dropped from the sample. The purpose of the matching procedure is to reduce concerns about

---

6 The harmonized assignee codes were downloaded from the USPTO website. http://www.uspto.gov/web/offices/ac/ido/oeip/taf/data/misc/data_cd.doc/assignee_harmonization/.

7 I infer the location of invention through the inventors' locations. A patent is considered generated in the US only if all inventors are residing in the US; patents with inventors living in more than one country are dropped.

8 I ignore unassigned patents and patents assigned to individuals because these patents are generally of less technological significance than patents owned by organizations. I matched MNC-owned focal patents with patents owned by all types of organizations rather than firm-owned patents in part because I do not have a reliable method to identify all assignees' type of entity. That being said, the majority of control patents are owned by firms. One can simply think of control patents as randomly picked patents from the universal patent database, which excludes focal patent assignees.

9 This control technique, pioneered by Jaffe et al. (1993), has become the standard methodology in the exploration of knowledge spillovers through patent data. The matching process developed by Jaffe et al. (1993) controls only for patent class; however, Thompson and Fox-Kean (2005) have pointed out that controlling only for three-digit patent classes might induce a systematic bias toward spurious finding of knowledge spillovers. They showed that selecting controls based on patent subclass, a much finer level of disaggregation helps remedy the bias.

endogeneity-related biases. For example, Chinese organizations may disproportionately cite patents owned by MNCs with China-based R&D facilities just because these MNCs are working in technological domains more relevant to Chinese organizations.

Pooling across organizations and time, I obtain a group of focal patents and a matched group of control patents. I then trace out the non-examiner-added citations received by both focal and control patents. The key task here is to compare the citations Chinese organizations make to the focal patents and the control patents, respectively.

Figure 1 illustrates the matching design and the main idea of the t-test for cross-border knowledge flows from MNCs to Chinese organizations, where P_F is the frequency probability that the focal patents are cited by Chinese organizations[10] after the MNCs begin conducting R&D in China, and P_C is the frequency probability that the corresponding control patents are cited by Chinese organizations during the corresponding focal patent's *ex post* period. Formally, I test the following null hypothesis:

$$H_0: P_F = P_c$$

The t-statistic is calculated as: $t = \frac{\hat{P}_T - \hat{P}_C}{\sqrt{\frac{\hat{P}_T(1-\hat{P}_T)}{n_T} + \frac{\hat{P}_C(1-\hat{P}_C)}{n_C}}}$.

[Insert Figure 1 Here]

Several things are notable. First, because I do not have the exact information on MNCs' offshoring time, I simply look at the citations made after the application date of the MNC's first China-originating patent. Second, as mentioned earlier, I am able to distinguish only non-examiner-added citations from examiner-added citations since the grant year 2001, so P_F and P_C

10 I look at not only Chinese firms but also other types of organizations, mainly universities and research institutes, because these organizations are playing important roles in China's technological catch-up. In many cases, research institutes are directly involved in the learning process (including imitation through reverse engineering). Once they master the technology, they may spin off the institute to become a firm; for example, Lenovo was a spin-off of the Chinese Academy of Sciences Computer Technology Research Institute.

are calculated based on the non-examiner-added citations received from the patents granted during the period of 2001–2012. As such, a patent granted before 2001 would have a time window without observable non-examiner-added citations. That being said, both the focal patent and its corresponding control patent are subject to the same truncation. Third, my unit of observation is the citation level. Although the focal patents and the control patents have a one-to-one correspondence, this is not true for citations. There are cases in which a focal patent does not receive any citations while its corresponding control patent does; there are also cases in which a control patent does not receive any citations while its corresponding focal patent does. In both cases, I drop the patents having no citations, resulting in unpaired focal or control patents in the sample.[11]

Since Jaffe, Trajtenberg, and Henderson's (1993) seminal work, this type of t-test is commonly employed by scholars to study knowledge flows using patent citations (e.g., Almeida 1996 and more recently Blit 2014). Yet, such an approach has its limitations. The obvious one is unobserved heterogeneity across firms and patents. For example, larger firms may hire better inventors—their ideas naturally receive more citations from others, including from Chinese organizations—and they might also be more likely to expand R&D to another country. Or, as described previously, focal patents might come from specific technological domains of greater (or less) relevance of Chinese organizations, but the exact degree of technological relevance is too subtle to be captured by the same patent class- and subclass-matching criteria used in my matching procedure. In either case, the unobserved factor (patent quality or patent relevance) confounds the effect of the R&D offshoring. An even more serious problem involves direct selection. Firms might have chosen to offshore their R&D to China for some unobservable

[11] Exclusion of these unpaired patents yields a similar mean difference with a larger t-statistic.

reasons (e.g., better mechanisms to protect their IPRs). All of these cases violate the "selection on observables" assumption, leading to biased results. A Difference-in-Differences (DD) approach therefore helps to address these problems by allowing us to distinguish the confounding effects from the true "treatment" effect.

3.3 A Difference-in-Differences Approach

I now turn to more rigorous DD subsample construction and estimation of the impact of MNC R&D offshoring on knowledge diffusion across borders.[12] To construct a dataset for DD estimation, I start from the t-test sample described in the notes for Figure 1, but only keep observations that meet the following criteria: 1) citing patents have an application date between 2001–2010;[13] 2) cited patents are assigned to a foreign MNC that has its first China-originating patent (according to the application date) between 2004–2007;[14] and 3) cited patents have an application date at least two years before the offshoring date.[15] Figure 2(A) illustrates the timeline of the DD sample. Figure 2(B) illustrates the intuition behind the DD framework. The offshoring date is defined as one year before the application date of the focal patent owner's first

[12] See Singh and Agrawal (2011) for an example of using a DD estimation approach to explore knowledge diffusion via new hires.

[13] I drop the citing patents with an application date before 2001 for two reasons: First, pre-2001 patents owned by Chinese domestic organizations are very rare. Second, there is a truncation issue. Non-examiner-added citations are only observable since the 2001 grant year. Among all citing patents with an application date before 2001, only those with a grant date after January 1, 2001, can be tracked for non-examiner-added citations. Dropping the patents before the 2001 application year fully rules out this type of truncation. I drop the citing patents with an application date after 2010 because of another truncation issue: some patents filed during the recent years are missing because they haven't yet been granted, which means they are still confidential to the USPTO and therefore the citations made by them cannot be tracked. This truncation issue is partially remedied by excluding patents with a post-2010 application date.

[14] This ensures that each cited patent has been exposed to both the pre- and post-offshoring period. The drawback is that this criterion excludes many MNCs. During the period of 1996–2007, 86 offshore MNCs with matched control patents had the first China-originating patents. Among them, 42 had the first China-originating patents during the period of 2004–2007, one in 2008, and 43 during 1996–2003. The 2004–2007 criterion excludes more than half of the offshore MNCs from the sample, including many companies that are engaging in serious R&D efforts in China. Alternatively, shifting the low-bound of this criterion to include 2002 and 2003 would require moving the pre-shoring periods to 1999 and 2000, which causes additional concerns: the non-examiner-added citations observed for these two years will be truncated, and, even more seriously, the inclusion of these years might bias the effects for the pre-offshoring period since few patents were owned by Chinese organizations in 1999–2000. Despite these challenges, I still vary the length of sample in several ways, and DD analyses based on these samples are broadly consistent with the results reported in the paper.

[15] Again, this is to ensure that each cited patent has been exposed to the pre-offshoring period for at least two years.

China-originating patent.[16] As such, the pre-offshoring period is from January 1, 2001, to one year before the application date of the focal patent owner's first China-originating patent, and the post-offshoring period is from one year before the application date of the focal patent owner's first China-originating patent to December 31, 2010. As mentioned earlier, not every patent receives citations. As a result, there are focal patents with citations in cases in which the corresponding control patents have none, or vice versa. To make the comparison simple, I drop these unpaired patents. Finally, I end up with a sample with at least two years of pre-offshoring and a post-offshoring period of at least four years.

[Insert Figure 2 Here]

Formally, I estimate the following equation:

$$\mathrm{cn\_cit_i} = \mathrm{f}(\delta_1 \mathrm{post_i} + \delta_2 \mathrm{post\_focal_i} + \beta_X \mathrm{X_i} + \gamma_i + \psi_t + \varepsilon_i) \quad (1)$$

where cn_cit_i is an indicator variable that equals 1 if the citation is made by a Chinese organization and 0 otherwise. The variable $post_i$ is an indicator variable that equals 1 for citations made in the (focal) post-offshoring period for a matched pair (i.e., focal patent and its corresponding control). $post\_focal_i$ is an indicator that equals 1 for citations received by a focal patent in the (focal's) post-offshoring period. The vector X_i is a set of control variables that may affect citation rates. γ_i is the patent fixed-effect, which I use to control for time-invariant differences across patents (e.g., a patent receives more citations because it conveys more valuable knowledge) and firms. ψ_t is the citing application year fixed effect. The coefficient of interest is δ_i, which is positive and significant if R&D offshoring facilitates cross-border

[16] Patent information is used to infer the offshoring time. The one-year threshold is somewhat arbitrary; however, it is reasonable to assume that there is a gap between the time of inventing and the time of applying for a patent. If the offshoring date were defined to be the first China-originating patent application date, it would yield an obvious downward trend in citations received by focal patents from Chinese organizations starting one year before the offshoring date. This suggests that one year before is a better division of the data.

knowledge flows from offshore MNCs to Chinese organizations and negative and significant if R&D offshoring helps to reduce knowledge flows.

3.4 The Regression Models

Because my dependent variable is a binary response, I start with logit models. The logit regression coefficients have a beneficial feature: they are log odds and can be easily converted to odds ratios. Thus, the effects of R&D offshoring are easy to interpret. However, the drawback of using the logit model is that when we include patent conditional fixed effects, a major fraction of observations that have no within-patent variation in the dependent variable (i.e., the patent receives zero citations from Chinese organizations) are dropped, possibly causing selection bias. As an alternative, I run analogous regressions using linear probability models, which do not drop observations.[17]

4 Empirical Results

4.1 t-test Results

The results of the t-test are provided in Table 1. The "number of citations" corresponds to the total number of citations received by focal patents and control patents, respectively. "Cited by CN organizations %" indicates the percentage of citations received from Chinese organizations. While 0.088 percent of the control citations are received from Chinese organizations, only 0.076 percent of the focal citations are received from Chinese organizations. The t-test for the difference of two groups yields a t-statistic of -3.23 (p-value=0.0012). The effect is computed as P_F/P_C=0.86, suggesting that patents created by the offshore MNCs outside of China receive, on average, 14% fewer citations from China organizations than control patents.

[17] See Angrist and Pischke (2008) for a discussion of the trade-offs between using linear models and using nonlinear models. See Singh and Agrawal (2011) for an example of using linear models to compare patent citation rates.

For comparison purposes, I also explore two additional types of knowledge diffusion: the within-China knowledge diffusion from the offshore MNCs' China subsidiaries to Chinese organizations, as well as the cross-border within-firm knowledge flows from the offshore MNCs' headquarters or subsidiaries in other countries to the MNCs' China subsidiaries. To construct the sample for the within-China test, I use a matching procedure similar to that described in the notes for Figure 1, but I define focal patents and control patents in different ways. Focal patents are now selected from China-originating patents owned by all foreign offshore MNCs, and control patents are now selected from non-China-originating patents owned by these MNCs; because of this change in the population selection, the same inventor nationality criterion is dropped. The t-test results presented in Table 2 illustrate that while focal patents, on average, appear to receive a higher percentage of citations from Chinese organizations than control patents, the difference is not statistically significant. Combining the results from Table 1, these results suggest there is little evidence of knowledge diffusion from the offshore MNCs' China subsidiaries to Chinese organizations. For the cross-border within-firm test, I use the same sample as that described in the notes for Figure 1, but instead of counting citations received from Chinese organizations, I now look at citations received from the China R&D subsidiaries of the multinationals that have established such entities. The results presented in Table 3 illustrate that focal patents receive disproportionately more citations from the MNCs' China subsidiaries than control patents, suggesting that knowledge diffuses across borders within the MNCs' boundaries. These findings are consistent with the view that knowledge flows more easily within a firm than between firms (Kogut and Zander 1992). It is notable that I cannot apply the DD estimation to test these two types of knowledge diffusion because, by definition, foreign MNCs begin to create China-originating patents only after they have offshored R&D to China. As such, I would not be able to either find focal patents for the pre-offshoring period to conduct the within-China test, or define a dependent variable for the pre-offshoring period to conduct the within-firm cross-border test.

[Insert Table 1, Table 2, and Table 3 Here]

4.2 DD Estimation Results

Table 4 summarizes the definitions and the descriptive statistics for the key variables used in the DD estimation. Pooling across type and time, about 0.08 percent of the citations are made by Chinese organizations, which is slightly higher than the population average of 0.07 percent.[18]

[Insert Table 4 Here]

Table 5 reports the percentage of citations received from Chinese organizations for the pre-offshoring and post-offshoring periods corresponding to the focal and control groups. For the pre-offshoring period, focal patents receive a slightly larger percentage of citations from Chinese domestic organizations, yet the difference is not statistically significant. For the post-shoring period, however, focal patents receive a much smaller percentage of citations from Chinese organizations than the control patents, and the difference is statistically significant. One should be very careful about the before-and-after comparison because such a comparison is hardly meaningful due to the time dimensions of patent citations.

[Insert Table 5 Here]

Table 6 reports the DD estimation results for the impact of R&D offshoring on knowledge diffusion from MNCs to Chinese organizations across borders. For the pooled models, I use *gdelay, claims, total_inv*—all in natural log form—to control for the cited patent level heterogeneity (see Table 4 for the definitions of these variables). In the fixed effect models, these

[18] Because only a small number of US patents are owned by Chinese organizations, these patents' backward citations are uncommon in general. For the period in which we can identify non-examiner-added patents (2001–2012 grant years), 18,963,541 citations are made by all kinds of organizations toward patents grated during 1975–2012. Among them, only 13,228 citations, or 0.07%, are made by Chinese organizations.

variables are absorbed by the patent fixed effects. I use *ln_china_pats* to control for increases in the number of patents owned by Chinese organizations over time, and *ln_cit_gap* and citing application fixed effects to count for the time dimensions of citation (citations received by a patent generally follow a right-skewed distribution over citing time). Column (1) – Column (3) present the results using logit models. All the coefficients are expressed as odds ratios. Column (1) shows the cross-sectional findings for the post-offshoring sample only. The coefficient for *focal* is 0.511 and statistically significant, suggesting that focal patents receive on average 49% (calculated as 1-0.511) fewer citations from Chinese organizations than control patents during the post-offshoring period. Column (2) shows the cross-sectional DD results for the sample including both pre- and post-offshoring periods. The coefficient for *focal* is insignificant, suggesting that there is no measured difference between the citation rates received by focal patents and control patents from Chinese organizations before offshoring. The coefficient for *post* is insignificant, suggesting that there is no measured difference between the citation rates received by control patents during the pre-offshoring period and those received during the post-offshoring period. Finally, the DD estimate for the coefficient for *post_focal* is 0.419 and significant, reflecting that focal patents receive 58% (calculated as 1-0.419) fewer citations relative to control patents after offshoring, when other factors that may affect citations are controlled. Combining the estimates for *post* and *post_focal*, the pooled model suggests a negative effect of offshoring on knowledge flows to Chinese organizations. Column (3) reports the results including conditional patent fixed effects. The DD estimate also suggests that R&D offshoring is associated with a decrease in knowledge flows. However, as noted earlier, the model drops all patents that receive zero citations from Chinese organizations, which may cause sample selection bias. Thus, I postpone the discussion of this result until I come to the linear model with patent fixed effects.

We next replicate the analyses in Column (4) – Column (6) using linear models. Column (4) and Column (5) show the pooled results. The results are consistent with the results from the logit models. Column (6) reports my preferred specification for DD regression with patent fixed effects. The DD estimate for the coefficient of *post-focal* is 0.0007 and statistically significant, reflecting that the offshoring reduces the possibility of citations received from Chinese organizations by 0.07% (in absolute terms).[19] As a benchmark, note that the average number of citations received by control patents in the post-offshoring period is 0.158% (Table 5). The DD estimate on *post-focal* implies a 46% reduction in the citation rate.[20] It is also important to note that across models, the coefficients of *post* are always insignificant, suggesting that the DD estimates are not due to citation rate changes in control patents.[21]

[Insert Table 6 Here]

4.3 Testing Pre-Existing Trends

In the DD analyses, I divide the sample into two periods. The concern, however, is that the DD estimate may simply reflect the pre-existing trends of focal patents. To visualize this, imagine that control patents show no obvious trend of citation changes across time, while focal patents have a downward trend across years. By collapsing data into the pre- and post-offshoring periods, the downward trend of focal patents would be reflected by a smaller difference between

[19] A drawback of the linear probability model is that the estimates can lead to probabilities below 0 or above 1. Checking the values of predicted dependent variables suggests that this is not a serious issue in this analysis: there are no predicted values above 1, and only 2.6% of the predicted values are below zero.

[20] In robustness checks I vary the timeline of sampling in two ways: (1) shifting the offshoring date back one more year without changing the lower or upper bound of citing time; (2) shifting both the offshoring date and the lower bound of citing time back one more year, while the upper bound remains the same. Regression analyses using these two alternative samples and linear models with patent fixed effects both yield a negative DD estimate with a p-value between 5 and 6%.

[21] The negative significance of the DD coefficient can be contributed to two possible factors: a decrease of citations for focal patents, or an increase of citations for control patents in the post-offshoring period. Distinguishing these two is important because an increase of citations for control patents in the post-offshoring period does not necessarily indicate the effect of offshoring. However, the insignificant coefficient on post-offshoring reiterates that the DD effect is not caused by the increase of citations in control patents.

focal patents and control patents for the pre-offshoring period and a greater difference for the post-offshoring period. When estimating the DD, I might obtain a negative and significant DD estimate, but the specification of such a DD model itself is actually problematic.

To address this concern, I conduct a kind of event study for focal patents. Instead of using a single indicator variable to divide time into two periods, I now use multiple indicator variables corresponding to five pre-offshoring periods and seven post-offshoring periods in the linear fixed effects regression for focal patents, with the first period immediately prior to the offshoring (pre-offshoring period 1) being the baseline. For the control patents, the time periods are irrelevant, so I define them as the baseline (pre-offshoring period 1). Figure 3 shows the plotting of regression coefficients on these time indicators, with more detailed results reported in Table 7.[22] The figure contradicts the pre-existing trends hypothesis—focal patents show no obvious trend in citations from Chinese organizations in the years leading up to the offshoring date. Figure 3 also facilitates an additional observation: there is a sharp discontinuity in the likelihood that focal patents are cited by Chinese organizations at the time of offshoring, suggesting a negative association between R&D offshoring and knowledge flows to Chinese organizations.

[Insert Figure 3 and Table 7 Here]

4.4 Robustness

4.4.1. The Intentional Omission Hypothesis

It is possible that when MNCs offshore R&D to China, Chinese organizations still build upon the patents (maybe even more frequently) owned by these MNCs, but intentionally omit citations. As such, the negative relationship between R&D offshoring and patents' citations received from

[22] This approach is inspired by McCrary (2007) and Singh and Agrawal (2011).

Chinese organizations hardly represents any reduction in knowledge spillovers. Rather it captures Chinese organizations' strategic efforts to avoid getting caught by the MNCs. This hypothesis can be tested by using examiner-added citations. The idea is based on the fact that the patent examiner, rather than the patent applicant, is ultimately responsible for identifying the prior art. A patent examiner "is supposed to be an expert in the area and hence to be able to identify relevant prior art that the applicant misses or conceals" (Hall, Jaffe, and Trajtenberg 2001, p.15). If the intentional omission hypothesis is true, I would expect that citations deliberately omitted by Chinese applicants would be added by examiners, and a version of the previous regression that used only examiner-added citations would find higher citation rates received by focal patents from Chinese organizations. To test this intentional omission hypothesis, I replicate the DD analyses, as seen in Table 6, using the same patent samples as before, but now I use only examiner-added citations. Table 8 reports the results. As can be seen, the coefficients for *focal_post* are statistically insignificant, suggesting little evidence of intentional omissions by Chinese organizations.

[Insert Table 8 Here]

4.4.2. The Technological Divergence Hypothesis

It is also possible that Chinese organizations and foreign MNCs that have offshored R&D to China possess technological trajectories that are diverging from each other over time. This change happens to coincide with R&D offshoring and is not captured by the patent class- and subclass- matching criteria. In other words, Chinese organizations cite MNCs' patents less often simply because MNCs' technology becomes less relevant to them. Direct evidence regarding the technological divergence hypothesis may be obtained by using a technological proximity measure, as first proposed by Jaffe (1986). The Jaffe technological proximity between entity *i* and entity *j* can be defined as

$$\mathrm{P_{ij}} = \frac{\mathrm{F_i F_j'}}{(\mathrm{F_i F_i'})^{\frac{1}{2}} (\mathrm{F_j F_j'})^{\frac{1}{2}}}$$

where vector $F_i = (f_{i1}, f_{i2}, \dots f_{ik})$ and where f_{ik} is i's share of patents in technology class k. This index ranges between zero and 1. The higher the value, the closer the degree of overlap in technology between two organizations. Since I am interested in the technological proximity between Chinese organizations and the R&D offshoring MNCs as a group, I simply collapse the patent data into two groups (China-originating patents owned by Chinese organizations versus non-China-originating patents owned by the MNCs) to calculate the proximity score across the application years of the DD sample's citing patents. It is notable that while a cited patent is available to all future patents for reference, a citing patent only makes a one-time backward reference to prior patents. To reflect this difference between cited patents and citing patents, I use the patent stock counts up to the citing application year to calculate the *F* vector for the MNC group (i.e., cited patents) and the patent flow counts during the citing application year to calculate the *F* vector for the Chinese organization group (i.e., citing patents). One drawback of the Jaffe technological proximity measure is that it assumes spillovers occur only within the same technology class. To allow spillovers among technology classes, I also calculate the Mahalanobis distance measure, as seen in Bloom, Schankerman, and Van Reenen (2013) for the same data. Similar to the Jaffe technological proximity measure, the higher the value of the Mahalanobis measure, the closer the technology is between two groups. Figure 4 presents the time series plots of the two proximity measures. The figure contradicts the technological divergence hypothesis—the technology difference between Chinese organizations and MNCs seems to decrease over time.

[Insert Figure 4 Here]

Additional evidence against the technological divergence hypothesis may also be obtained from the regression results using examiner-added citations. The logic is straightforward: if there is a true divergence in technology trajectories between Chinese organizations and MNCs, examiners are also less likely to identify focal patents as the “prior art” for the patent applications by Chinese organizations. As such, I expect to see statistically significant negative coefficients for the DD estimate, similar to the results using non-examiner-added citations. However, the results presented in Table 8 suggest that this is not the case.

4.4.3. The Citation Lag Discrepancy Hypothesis

Different patents have different citation lags. It is possible that my results are driven by the discrepancy between the citation lags of focal patents and those of control patents—one group of patents become obsolete faster (or slower) than the other group. To understand this, simply imagine that the underlying distribution of the citation lags for focal patents follows a normal distribution while that for control patents follows a right-skewed distribution, with the left tails of both distributions being identical. Then, if the right tails happen to be situated in the post-offshoring period, I may obtain a negative coefficient for the DD estimate simply because of the different shapes of right tails. This hypothetical example suggests a direct way to test the citation lag discrepancy hypothesis: by comparing the distributions of the forward citation lags for focal patents and control patents to determine whether there is a discrepancy. Figure 5(A) plots the kernel density estimates of the forward citation lags for focal patents versus control patents using the DD sample. Figure 5(B) plots the quantiles of the citation lags for focal patents against the quantiles of the citation lags for control patents. (If the two distributions are similar, the points in the Quantile–Quantile plot should approximately lie on the line $y = x$.) The figures provide straightforward visual evidence against the citation discrepancy hypothesis: focal patents and control patents show very similar distributions in citation lags.

[Insert Figure 5 Here]

5 The Monitoring Function of R&D Offshoring

5.1 The Monitoring Hypothesis

The t-test and DD estimation both point to a consistent conclusion that MNCs conducting R&D in China seem to reduce knowledge flows to Chinese organizations. So, what might be the mechanism behind this result?

Insights from field interviews suggest that when MNCs offshore R&D to China, there are two likely countervailing effects on Chinese organizations. On the one hand, this creates learning (or imitation) opportunities for Chinese organizations. On the other hand, as MNCs expand R&D in China, they also increase their ability to monitor the learning and invention of their indigenous competitors. By leveraging this monitoring ability with other strategic tools, MNCs can limit the ability of indigenous Chinese firms and other organizations to build on MNC R&D.

To test this hypothesis, I draw upon the existing findings in the literature on alliance and the social network analysis tools, estimating the following regression equation:

$$\mathrm{cn\_cit_i} = \mathrm{f}(\alpha_1 \mathrm{cn_i} + \alpha_2 \mathrm{har_i} + \alpha_3 \mathrm{cn\_har_i} + \beta_Z \mathrm{Z_i} + \gamma_i + \psi_t + \varepsilon_i) \qquad (2)$$

Here, cn_cit_i is the dependent variable measuring citations received from Chinese organizations. I measure it as the yearly number of citations an MNC i received from Chinese organizations, where the citations are located in time by using the application years of the citing patents. The variable cn_i equals 1 if the owner of the cited patent is a foreign MNC that has offshored R&D to China (defined as having at least 10 China-originating patents by the end of 2012) and the application date of the citing patent is after the firm's first China-originating patent. If these two conditions are not met, then cn_i equals zero (i.e., the owner of the cited patent is a foreign MNC that has not offshored R&D to China or the owner of the cited patent is an offshore MNC, but the application date of citing patent is before the owner's first China-originating patent). I employ har_i, the harmonic centrality index (Rochat 2009) for i one year before the citing application year, to measure the reach of a foreign MNC in the China alliance network. For node x_i, the normalized harmonic centrality is defined as the sum of the inverse distances from x_i to all other nodes (i.e., alliance participant) divided by (n-1). Formally,

$$\frac{1}{n-1}\sum_{j\neq i}\frac{1}{dist(x_i, x_j)}$$

For two unconnected nodes, the inverse distance of two is defined as zero. cn_har_i is the interaction term of cn_i and har_i, which I use to proxy an MNC's monitoring ability due to R&D offshoring to China. Z_i are a vector of controlling variables. I include two controls: tec_i and $ln\_total\_cit_i$. The variable tec_i is the Jaffe technological proximity index, as described previously, between a foreign MNC's patent portfolio up to one year before the citing year and the patent portfolio owned by all Chinese organizations as a whole in the citing year, which I use to control for the technological relevance of a foreign MNC's patent portfolio to Chinese organizations. As an alternative, I also calculate the Mahalanobis distance index, $maltec_i$, substituting for tec_i in regressions. The variable $ln\_total\_cit_i$ is the log term of total number of citations a foreign MNC receives in the citing application year.

The idea is simple: the higher the degree of an MNC's reach in the Chinese market, the greater the possibility of leaking knowledge to others.[23] At the same time, the higher the degree

[23] Owen-Smith and Powell (2004) find that a firm's betweeness centrality in a geographically dispersed alliance network positively affects its innovation. Schilling and Phelps (2007) find that firms embedded in alliance networks that exhibit high reach have greater innovative output.

of an MNC's reach in the Chinese market, the higher the firm's ability to collect information from others to monitor the Chinese market, and the greater the degree to which it can use this information to limit the leakage of significant knowledge. As such, I expect to observe a positive association between har_i and the dependent variable cn_cit_i, and a negative association between the interaction term cn_har_i and the dependent variable cn_cit_i.

5.2 Alliance Data

To construct a sample for empirical analysis, I merge the alliance data with the patent data by matching alliance participants' names with patent assignees' names. The patent data are drawn from the same sources as described in Section 3. The alliance data are gathered using the SDC Platinum Database. I use only signed alliances that involve at least one Chinese participant with China as the nation of alliance and an effective date between 1998 and 2008. For non-Chinese participants, I only include those with a home country considered an advanced economy. I aggregate all alliances to the ultimate parent corporation.[24] I include alliance types because any type of alliance may provide a path for both knowledge diffusion and information transmission (monitoring). Since alliance termination dates are rarely reported in the SDC Platinum database, following Schilling and Phelps (2007), I assume that alliances last for three years. I create alliance networks based on three-year windows (i.e., 1999–2001, 2000–2002,… 2006–2008)[25], resulting in eight snapshots of networks. I calculate the harmonic centrality based on these eight snapshots of networks. These networks include both foreign organizations and Chinese organizations, with the majority of them being firms. For my purpose, I am only interested in the

Based on these findings, I infer that firms with a higher reach, defined as a higher harmonic index score, are more likely to leak knowledge to Chinese organizations.

[24] Note that I do not trace the whole family structure of the ultimate parent corporation, and thus I may not account for patents assigned to a child company (of the ultimate patent corporation) that has no alliance in China.

[25] I have also created yearly networks using the harmonic centrality scores based on yearly networks, which yields similar empirical results.

relationship between R&D offshoring and patent citations; therefore, I only include foreign MNCs with at least one matched patent in the regression. The final sample contains two types of MNCs: those that have not offshored R&D to China (without any China-originating patents by the end of grant year 2012) and those that have offshored R&D to China (with at least 10 China-originating patents by the end of grant year 2012). The resulting dataset includes 20 MNCs that have offshored R&D to China and 138 that have not.

5.3 Regression Results

Table 9 summarizes descriptive statistics for the key variables used in the regression. Table 10 reports the correlation matrix between variables.

[Insert Table 9 and Table 10 Here]

Table 11 repots the regression results using the Poisson quasi-maximum likelihood (PQML) estimation.[26] Column (1) and Column (3) show the results using the pooled data. Column (2) and Column (4) report with firm fixed effects. Across all models, the coefficient of *har* is positive and significant, suggesting that an MNC that has greater reach (a higher harmonic centrality index) in the China-based alliance networks experiences more knowledge diffusion to Chinese organizations. On the other hand, the coefficient of *cn_har* is negative and significant, suggesting that offshoring R&D to China helps MNCs with a greater reach to reduce their knowledge diffusion to Chinese organizations. Although it is difficult to establish causality for this relationship at this point of time, the association does provide a suggestive guideline for future research on the monitoring hypothesis.

[Insert Table 11 Here]

[26] For a detailed discussion of the PQML model and the trade-offs involved in choosing between PQML and Negative Binomial models in similar settings, see Branstetter, Li, and Veloso (2015).

6 Conclusion and Discussion

This paper has examined the impact of R&D offshoring to China on knowledge diffusion to Chinese organizations. This has been the subject of considerable controversy in management and policy circles—China is widely regarded as a nation that fails to adequately protect intellectual property rights, and there are widespread concerns that the rapidly increasing amount of R&D conducted by MNCs in China could lead to strategically sensitive leakage of valuable technology. In Washington, DC, there are fears that this could not only drive down shareholder value in the long run, but also weaken America's comparative advantage in innovative sectors. Do the data support these concerns?

Using non-examiner-added patent citations as an indicator of knowledge flows, I find that patents created outside of China by MNCs that have offshored R&D to China receive, on average, significantly *fewer* citations from Chinese organizations during the post-offshoring period. This finding proves robust to a number of alternative approaches and tests.

From the point of view of the R&D offshoring MNC, one might infer from this finding that offshoring R&D to emerging economies actually helps to reduce knowledge leakages. For policymakers in an MNC's home country, it is tempting to conclude that offshoring R&D to emerging economies has little to no negative impact on the home country's competitive advantage in technology. These ideas may well have merit, but one must be cautious in making such strong conclusions. There is a fundamental limitation to a causal interpretation of my findings. Recall that I define focal patents as those owned by an MNC that has offshored R&D to China. I interpret this event of offshoring as the "treatment," yet, as I have noted, firms make deliberate choices about when and where to offshore R&D. As such, the "treatment" is endogenous, not random. What I have estimated is the "treatment effect on the treated," rather than the “treatment effect” per se. Although the use of the DD approach helps ameliorate the endogeneity issue to some extent, it remains an issue in the interpretation of my results. Nevertheless, my results are striking—although there is widespread concern that growing multinational R&D may exacerbate problems of unwanted knowledge spillovers to indigenous Chinese organizations, the citation results are simply not consistent with these concerns.

To shed light on the possible mechanism that drives the negative association between R&D offshoring and knowledge diffusion to Chinese organizations, I have conducted interviews with a wide range of R&D practitioners, including managers, researchers, and lawyers, for both MNCs and indigenous Chinese firms in China. Insights from these interviews suggest that when MNCs offshore R&D to China, there appear to be two countervailing effects on Chinese organizations. On the one hand, offshoring R&D may create learning (or imitation) opportunities for Chinese organizations. On the other hand, as MNCs expand R&D in China, they also increase their ability to monitor what is happening locally in terms of technological activity. By leveraging this monitoring ability with other strategic tools, MNCs may limit the ability of Chinese organizations to build closely or directly on MNC R&D, leading them to undertake inventions that are increasingly independent. Using China-based alliance data and social network measures, I find evidence favoring this hypothesis.

The interviews and empirical results provide evidence supporting the “monitoring” hypothesis, yet one should be very careful about regarding it as the only underlying mechanism driving the negative relationship between R&D offshoring and knowledge diffusion. I prefer to consider these results as suggestive rather than dispositive.

One should also be cautious when generalizing my findings. Despite the measurement issues I have discussed in Section 2, patent citations capture only a small proportion of overall knowledge diffusion from MNCs to a very small fraction of Chinese organizations—only those applying for US patents. Although the knowledge flows reflected by patent citations may also include the elements of imitation that contribute to "further innovation," they do not directly measure pure imitations or IP theft.

Although this research is subject to limitations, as mentioned above, this paper does make significant contributions to the literature because it examines an important topic that is often neglected. To the best of my knowledge, this is the first empirical study that uses patent citation data to directly measure the relationship between multinational R&D offshoring and knowledge diffusion from offshore MNCs to domestic organizations in emerging economies. It is also the first paper to find a negative relationship between R&D offshoring and knowledge diffusion, contradicting the prior literature's general consensus of a positive relationship based on data from only advanced countries. My study suggests that future discussion of strategic dynamics and policy arrangements regarding multinational R&D in emerging markets should consider the unique situations of these countries, as they may differ from those in advanced economies.

REFERENCES

Alcácer, Juan, and Michelle Gittelman. 2006. "Patent Citations as a Measure of Knowledge Flows: The Influence of Examiner Citations." *The Review of Economics and Statistics* 88 (4): 774–79. http://www.mitpressjournals.org/doi/abs/10.1162/rest.88.4.774.

Alcácer, Juan, and Minyuan Zhao. 2012. "Local R&D Strategies and Multilocation Firms: The Role of Internal Linkages." *Management Science* 58 (4). INFORMS: 734–53. doi:10.1287/mnsc.1110.1451.

Almeida, Paul. 1996. "Knowledge Sourcing by Foreign Multinationals: Patent Citation Analysis in the U.S. Semiconductor Industry." *Strategic Management Journal* 17 (S2). John Wiley & Sons, Ltd.: 155–65. doi:10.1002/smj.4250171113.

Angrist, JD, and JS Pischke. 2008. *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press.

Barrett, Christie W, Peet van Biljon, and Christopher Musso. 2011. "Research Report: R&D Strategies In Emerging Economies: Results From The McKinsey Global Survey." *Research-Technology Management* 54 (4). Industrial Research Institute, Inc: 17–22.

Basant, Rakesh, and Sunil Mani. 2012. *Foreign R&D Centres in India: An Analysis of Their Size, Structure, and Implications*. Indian Institute of Management.

Blit, Joel. 2014. "Learning Remotely: R&D Satellites, Intra-Firm Networks, and Knowledge Sourcing." https://uwaterloo.ca/economics/sites/ca.economics/files/uploads/files/satellite_effect_2014_06_10_jems_submission.pdf.

———. 2015. "FDI and Indigenous Innovation : Foreign R & D Subsidiaries as a Medium for the Geographic Diffusion of Knowledge." https://uwaterloo.ca/economics/sites/ca.economics/files/uploads/files/dev_mnes_2015_07_02_working_paper.pdf.

Bloom, Nicholas, Mark Schankerman, and John Van Reenen. 2013. "Identifying Technology Spillovers and Product Market Rivalry." *Econometrica* 81 (4). Blackwell Publishing Ltd: 1347–93. doi:10.3982/ECTA9466.

Branstetter, Lee. 2006. "Is Foreign Direct Investment a Channel of Knowledge Spillovers? Evidence from Japan's FDI in the United States." *Journal of International Economics* 68 (2): 325–44. doi:10.1016/j.jinteco.2005.06.006.

Branstetter, Lee, Guangwei Li, and Francisco Veloso. 2015. "The Rise of International Coinvention." In *The Changing Frontier: Rethinking Science and Innovation Policy*, edited by Adam B. Jaffe and Benjamin F. Jones. Chicago and London: University of Chicago Press.

Branstetter, Lee, and Mariko Sakakibara. 2002. "When Do Research Consortia Work Well and Why? Evidence from Japanese Panel Data." *The American Economic Review* 92 (1): 143–59. http://www.jstor.org/stable/3083325.

Cohen, Wesley M., and Daniel A Levinthal. 1990. “Absorptive Capacity: A New Perspective on Learning and Innovation.” *Administrative Science Quarterly* 35 (1). Sage Publications, Inc. on behalf of the Johnson Graduate School of Management, Cornell University: 128–52. doi:10.2307/2393553.

Du, Luosha, Ann Harrison, and Gary H. Jefferson. 2012. “Testing for Horizontal and Vertical Foreign Investment Spillovers in China, 1998–2007.” *Journal of Asian Economics* 23 (3): 234–43. doi:10.1016/j.asieco.2011.01.001.

Griliches, Zvi. 1979. “Issues in Assessing the Contribution of Research and Development to Productivity Growth.” *The Bell Journal of Economics* 10 (1). RAND Corporation: 92–116. doi:10.2307/3003321.

Hall, Bronwyn H., Adam B Jaffe, and Manuel Trajtenberg. 2001. “The NBER Patent Citation Data File: Lessons, Insights and Methodological Tools.” National Bureau of Economic Research Cambridge, Mass., USA. http://www.nber.org/papers/w8498.

Hu, Albert Guangzhou, and Gary H Jefferson. 2002. “FDI Impact and Spillover: Evidence from China’s Electronic and Textile Industries.” *World Economy* 25 (8). Blackwell Publishers Ltd: 1063–76. doi:10.1111/1467-9701.00481.

Jaffe, Adam B. 1986. “Technological Opportunity and Spillovers of R&D: Evidence from Firms’ Patents, Profits and Market Value.” *The American Economic Review* 76 (5). National Bureau of Economic Research Cambridge, Mass., USA: 984–1001. http://www.nber.org/papers/w1815.

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

Jaffe, Adam B, Manuel Trajtenberg, and M.S. Fogarty. 2000. “Knowledge Spillovers and Patent Citations : Evidence from a Survey of Inventors.” *American Economic Review* 90 (2). American Economic Association: 215–218. http://www.nber.org/sloan/sloan.pdf#page=27.

Jaffe, Adam B, Manuel Trajtenberg, and Rebecca Henderson. 1993. “Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations.” *The Quarterly Journal of Economics* 108 (3). Oxford University Press: 577–98. doi:10.2307/2118401.

Keller, Wolfgang. 2004. “International Technology Diffusion.” *Journal of Economic Literature* XLII (September): 752–82.

Kogut, Bruce, and Udo Zander. 1992. “Knowledge of the Firm, Combinative Capabilities, and the Replication of Technology.” *Organization Science* 3 (3): 383–97.

Krugman, PR. 1991. *Geography and Trade*. Cambridge, MA: MIT Press. https://books.google.com/books?hl=en&lr=&id=AQDodCHOgJYC&oi=fnd&pg=PP7&dq=geography+and+trade&ots=Pzh10jOIam&sig=rN8YMMO1DDZzivF97TXZQUhnRiM.

Lin, Ping, Zhuomin Liu, and Yifan Zhang. 2009. “Do Chinese Domestic Firms Benefit from FDI Inflow?” *China Economic Review* 20 (4): 677–91. doi:10.1016/j.chieco.2009.05.010.

McCrary, Justin. 2007. “The Effect of Court-Ordered Hiring Quotas on the Composition and Quality of Police.” *The American Economic Review* 97 (1). American Economic

Association: 318–53. doi:10.2307/30034397.

Ministry of Commerce of China. 2012. "十六大以来商务成就综述之十：利用外资规模稳步扩大 质量水平显著提高." http://www.mofcom.gov.cn/article/ae/ai/201211/20121108421214.shtml.

Owen-Smith, Jason, and Walter W Powell. 2004. "Knowledge Networks as Channels and Conduits: The Effects of Spillovers in the Boston Biotechnology Community." *Organization Science* 15 (1). INFORMS: 5–21. doi:10.1287/orsc.1030.0054.

Rochat, Yannick. 2009. "Closeness Centrality Extended to Unconnected Graphs: The Harmonic Centrality Index." In *ASNA*. http://infoscience.epfl.ch/record/200525/files/[EN]ASNA09.pdf.

Sampat, Bhaven. 2012. "Examiner Citation Data." Harvard Dataverse. doi:hdl/1902.1/18735.

Schilling, M. a., and C. C. Phelps. 2007. "Interfirm Collaboration Networks: The Impact of Large-Scale Network Structure on Firm Innovation." *Management Science* 53 (7): 1113–26. doi:10.1287/mnsc.1060.0624.

Singh, Jasjit. 2007. "Asymmetry of Knowledge Spillovers between MNCs and Host Country Firms." *Journal of International Business Studies* 38 (5): 764–86. http://www.palgrave-journals.com/jibs/journal/v38/n5/abs/8400289a.html.

Singh, Jasjit, and Ajay Agrawal. 2011. "Recruiting for Ideas: How Firms Exploit the Prior Inventions of New Hires." *Management Science* 57 (1): 129–50. doi:10.1287/mnsc.1100.1253.

The Commission on the Theft of American Intellectual Property. 2013. "The IP Commision Report." http://scholar.google.com/scholar?hl=en&btnG=Search&q=intitle:No+Title#0.

Thompson, Peter, and M Fox-Kean. 2005. "Patent Citations and the Geography of Knowledge Spillovers: A Reassessment." *American Economic Review* 95 (1): 450–60. http://www.jstor.org/stable/10.2307/4132691.

UNCTAD. 2005. "World Investment Report 2005: Transnational Corporations and the Internationalization of R&D." *International Business Review*. Vol. 15. New York and Geneva.

Zhao, Minyuan. 2006. "Conducting R&D in Countries with Weak Intellectual Property Rights Protection." *Management Science* 52 (8): 1185–99. doi:10.1287/mnsc.1060.0516.

Table 1. Frequency difference of citations received from Chinese (CN) organizations after MNCs' first China-originating patent: MNC-owned non-China-originating vs. randomly picked non-China-originating

	Number of citations	Cited by CN organizations %
Focal	1,258,912	0.076
Control	1,098,556	0.088
Difference		-0.012
t-Statistic		-3.23
Effect		0.86

Table 2. Frequency difference of citations received from Chinese (CN) organizations after MNCs' first China-originating patent: MNC-owned China-originating vs. MNC-owned non-China-originating

	Number of citations	Cited by CN organizations %
Focal	7,308	0.274
Control	5,814	0.172
Difference		0.102
t-Statistic		1.24
Effect		1.59

Table 3. Frequency difference of citations received from MNC China subsidiaries after MNCs' first China-originating patent: MNC-owned non-China-originating vs. randomly picked non-China-originating

	Number of citations	Cited by MNC China Subsidiaries %
Focal	2,145,665	0.144
Control	2,094,537	0.073
Difference		0.071
t-Statistic		22.11
Effect		1.96

Table 4. DD variable definitions and summary statistics

Variables	Definitions	Mean	SD	Min	Max
cn_cit	Binary indicator for whether the citation is made by a Chinese organization	0.0008	0.0290	0	1
focal	Binary indicator for whether the given patent is a focal patent owned by an R&D offshoring firm	0.52	0.50	0	1
post_focal	For a focal patent, this is a binary indicator for whether the citation is made in the post-offshoring period. For a control patent, it is always 0.	0.26	0.44	0	1
post	For both the focal and corresponding control patents in a matched pair, this is a binary indicator for whether the citation is made in the post-offshoring period according to the focal patent	0.51	0.50	0	1
gdelay	The delay between application date and grant date for the patent (in years)	2.51	1.23	0.18	22.36
claims	Number of claims on the patent	19.86	17.01	1	335
total_inv	Number of inventors on the patent	2.68	2.01	1	38
china_pats	Total number of patents assigned to Chinese organizations in the one year preceding the patent's application date	511.22	291.89	97	1028
cit_gap	Number of years between the citing patent application date and the cited patent application	8.73	4.72	0.00	36.95

Note: Statistics are based on the DD sample as illustrated in Figure 2. The unit of observation is at the citation level. The sample has 199,537 citations for citing application year of 2001–2010, arising from 10,708 focal patents and 10,708 control patents.

Table 5. Frequency difference of patent citations received from Chinese organizations (DD sample): Pre- vs. post-offshoring

	Cited by CN organizations %	
	Before	After
Focal	0.055	0.078
	(N=51,006)	*(N=52,606)*
Control	0.047	0.158
	(N=47,248)	*(N=48,677)*
Difference	0.008	-0.080
t-Statistic	0.580	-3.692

Table 6. Regression analysis (DD sample)

Model:	(1) *Logit (in odds ratio)*	(2) *Logit (in odds ratio)*	(3) *Logit (in odds ratio)*	(4) *Linear*	(5) *Linear*	(6) *Linear*
Sample:	DD sample (post only)	DD sample	DD sample	DD sample (post only)	DD sample	DD sample
focal	0.511***	1.216		-0.000747***	0.000105	
	(0.100)	(0.345)		(0.000218)	(0.000143)	
post_focal		0.419*	0.264*		-0.000861***	-0.000720*
		(0.145)	(0.152)		(0.000261)	(0.000322)
post		0.838	0.645		0.000261	0.0000713
		(0.325)	(0.405)		(0.000227)	(0.000290)
ln_gdelay	0.977	1.057		-0.0000543	0.0000347	
	(0.139)	(0.130)		(0.000169)	(0.000109)	
ln_claims	0.647***	0.666***		-0.000563***	-0.000374***	
	(0.0531)	(0.0548)		(0.000126)	(0.0000869)	
ln_total_inv	1.037	0.945		0.0000557	-0.0000393	
	(0.150)	(0.115)		(0.000168)	(0.0000997)	
ln_china_pats	1.302	1.254	0.783	0.000238	0.000145	-0.000290
	(1.390)	(1.101)	(1.040)	(0.00109)	(0.000676)	(0.000741)
ln_cit_gap	1.786*	1.407*	9.033*	0.000579*	0.000193+	0.000418
	(0.419)	(0.228)	(8.641)	(0.000248)	(0.000100)	(0.000325)
constant	0.000218	0.000217*		0.00189	0.000391	0.00120
	(0.00158)	(0.000930)		(0.00653)	(0.00327)	(0.00341)
Patent FE	No	No	Yes	No	No	Yes
Citing App Year FE	Yes	Yes	Yes	Yes	Yes	Yes
R^2				0.0009	0.0007	0.0003
pseudo R^2	0.04	0.04	0.09			
Chi-square	74.43	126.67	41.01			
F				9.05	4.20	2.08
Pro>chi-square or F	0.00	0.00	0.00	0.00	0.00	0.01
N	93297	199537	1762	101283	199537	199574

Robust standard errors in parentheses

+ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 7. Estimating temporal trend in citations received by focal patent from Chinese organizations (DD sample)

	DD sample
Pre-offshoring period 5-6	-0.000258
	(0.000522)
Pre-offshoring period 4	-0.000423
	(0.000329)
Pre-offshoring period 3	-0.0000980
	(0.000433)
Pre-offshoring period 2	0.000156
	(0.000452)
Post-offshoring period 1	-0.000598^{+}
	(0.000308)
Post-offshoring period 2	-0.000615^{+}
	(0.000369)
Post-offshoring period 3	-0.000989*
	(0.000437)
Post-offshoring period 4	-0.00125**
	(0.000473)
Post-offshoring period 5	-0.000671
	(0.000685)
Post-offshoring period 6	-0.00129^{+}
	(0.000781)
Post-offshoring period 7-8	-0.00148*
	(0.000605)
Observations	199574

Note: The regression uses the "DD" sample and the linear model as in Table 6 Column (6). The reference period is now pre-offshoring period 1 (the omitted timing category). The period for control patents is defined as in pre-offshoring period 1. Figure 3 graphically illustrates the results from this table. Robust standard errors in parentheses. + $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 8. Robustness check using examiner-added citations (DD sample)

	(1)	(2)	(3)	(4)	(5)	(6)
Model:	*Logit (in odds ratio)*	*Logit (in odds ratio)*	*Logit (in odds ratio)*	*Linear*	*Linear*	*Linear*
Sample:	DD sample (post only)	DD sample	DD sample	DD sample (post only)	DD sample	DD sample
focal	0.931	0.912		-0.000390	-0.000140	
	(0.131)	(0.165)		(0.000763)	(0.000304)	
post_focal		1.028	0.910		-0.000241	-0.000537
		(0.235)	(0.268)		(0.000814)	(0.00107)
post		1.058	0.905		0.000329	-0.000171
		(0.382)	(0.415)		(0.00111)	(0.00137)
ln_gdelay	0.969	1.039		-0.000173	0.000125	
	(0.126)	(0.114)		(0.000712)	(0.000342)	
ln_claims	0.863^{+}	0.898		-0.000822^{+}	-0.000324	
	(0.0740)	(0.0638)		(0.000492)	(0.000219)	
ln_total_inv	0.984	0.921		-0.0000827	-0.000236	
	(0.102)	(0.0761)		(0.000563)	(0.000241)	
ln_china_pats	0.893	0.900	1.941	-0.000538	-0.000244	0.00208
	(0.787)	(0.616)	(1.583)	(0.00407)	(0.00176)	(0.00209)
ln_cit_gap	1.042	0.966	1.335	0.000206	-0.0000852	0.000836
	(0.138)	(0.0809)	(0.348)	(0.000647)	(0.000186)	(0.000599)
constant	0.0286	0.00205^{+}		0.00887	0.00309	-0.00978
	(0.167)	(0.00689)		(0.0226)	(0.00863)	(0.0101)
Patent FE	No	No	Yes	No	No	Yes
Citing App Year FE	Yes	Yes	Yes	Yes	Yes	Yes
R^2				0.0015	0.0021	0.0020
pseudo R^2	0.02	0.04	0.11			
Chi-square	50.20	190.48	109.43			
F				3.36	7.83	6.59
Pro>chi-square or F	0.00	0.00	0.00	0.00	0.00	0.00
N	38028	110764	2349	38028	110764	110781

Robust standard errors in parentheses

$^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 9. Descriptive statistics for the monitoring hypothesis test

Variables	Mean	SD	Min	Max
total_cn_cit	1.213	3.196	0	26
cn	0.200	0.400	0	1
har	0.023	0.036	0.002	0.148
cn_har	0.007	0.023	0	0.148
ln_total_cit	4.306	3.276	0	10.859
tec	0.182	0.218	0	0.878
maltec	1.291	1.648	0	7.397

Table 10. Correlations matrix between variables for monitoring hypothesis test

	total_cn_cit	*cn*	*har*	*cn_har*	*ln_total_cit*	*tec*	*maltec*
total_cn_cit	1						
cn	0.60	1					
har	0.02	0.14	1				
cn_har	0.23	0.58	0.51	1			
ln_total_cit	0.55	0.70	0.16	0.43	1		
tec	0.46	0.45	0.18	0.33	0.53	1	
maltec	0.67	0.68	0.00	0.28	0.80	0.628	1

Table 11. Regression analysis on the monitoring hypothesis

DV: total_cn_cit (yearly)	(1)	(2)	(3)	(4)
cn	0.593^{+}	0.269	0.508	0.271
	(0.342)	(0.315)	(0.339)	(0.298)
har	6.839^{+}	10.68**	10.90*	9.923**
	(3.818)	(3.522)	(4.866)	(3.410)
cn_har	-12.45**	-10.27**	-10.08*	-9.666**
	(3.864)	(3.514)	(4.835)	(3.049)
ln_total_cit	0.688***	0.933*	0.608***	0.858^{+}
	(0.0946)	(0.450)	(0.0855)	(0.458)
tec	1.550*	0.326		
	(0.635)	(1.517)		
maltec			0.180^{+}	0.151
			(0.0997)	(0.225)
_cons	-7.144***		-6.819***	
	(0.725)		(0.647)	
Firm FE	No	Yes	No	Yes
Year FE	Yes	Yes	Yes	Yes
pseudo R^2	0.68		0.67	
Number of firms	158	30	158	30
Log pseudolikelihood	-359	-215	-367	-214
Chi-square	561	179	673	266
Pro>chi-square	0	0	0	0
N	436	145	436	145

Robust standard errors clustered by the firm in parentheses

$^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Figure 1. Matching focal patents with controls

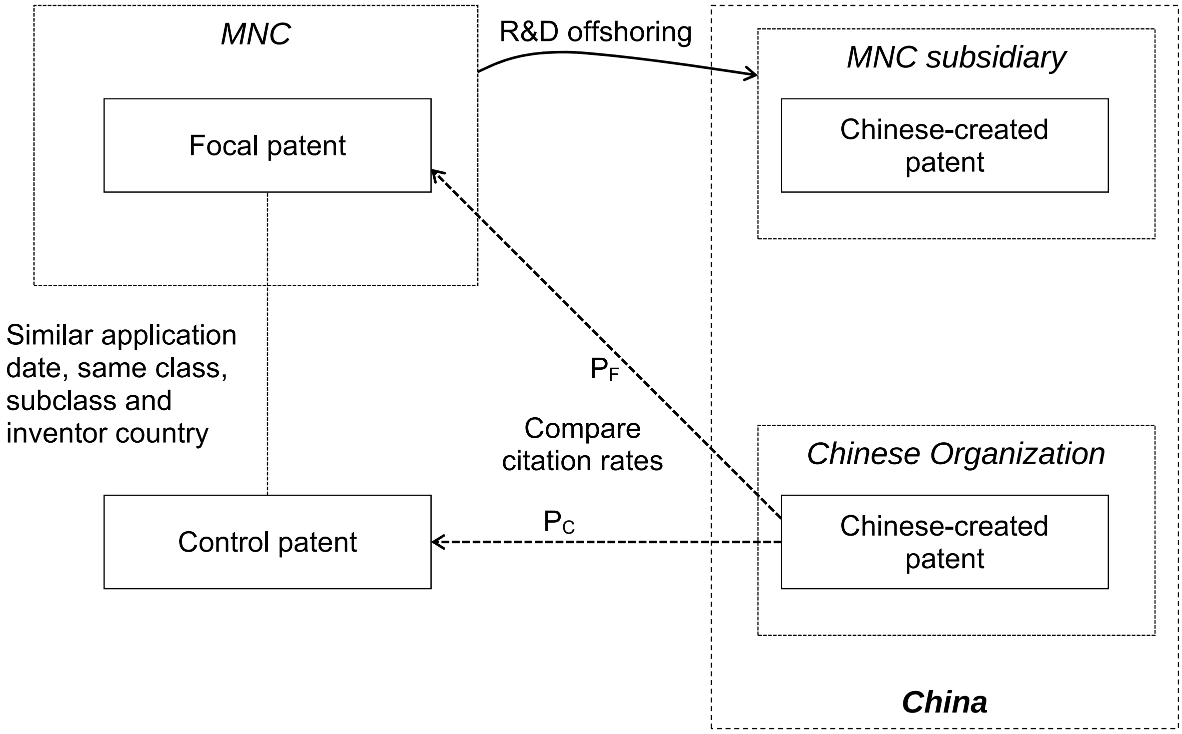

Notes: I begin by identifying a set of focal patents, each created by inventors with a residential address in a country other than China and assigned to a foreign company that has at least one China-originating patent by the end of grant year 2012. I match each focal patent with a control patent (from an organization that has no China-originating patents) such that the control patent has the same technology class, subclass, inventor residential country, and closest application date (within a one-year range). If several potential control patents are found, one is chosen randomly. Using this procedure, 38 percent of the focal patents find a matched control. I determine that an MNC has offshored R&D to China if it has at least 10 China-originating patents; my research has found 123 offshore MNCs. Among then, 113 firms find matched control patents.

Figure 2. DD sample construction and design illustration

A) DD sample timeline

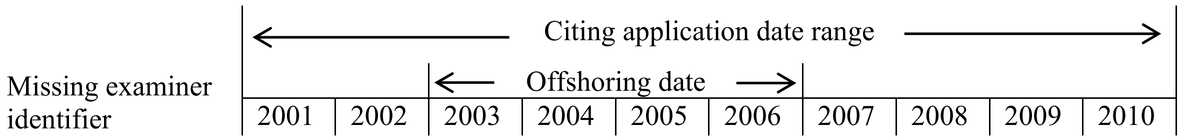

Note: Cited application date is at least two years before the inferred offshoring (treatment date). A cited patent has a 2–5 year pre-offshoring and a 4–7 year post-offshoring period to receive citations.

B) DD design illustration

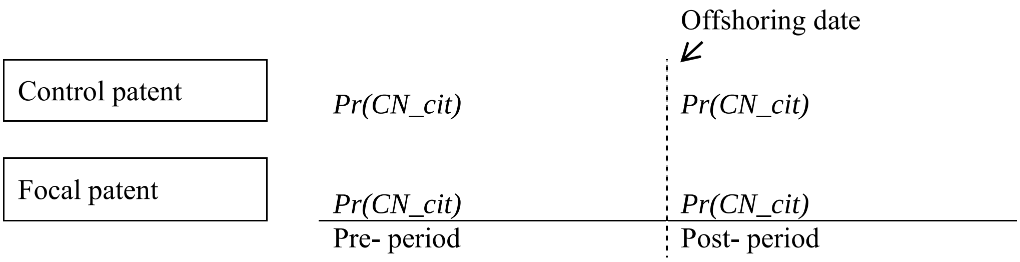

Figure 3. Estimated temporal trend in focal citations received from Chinese organizations

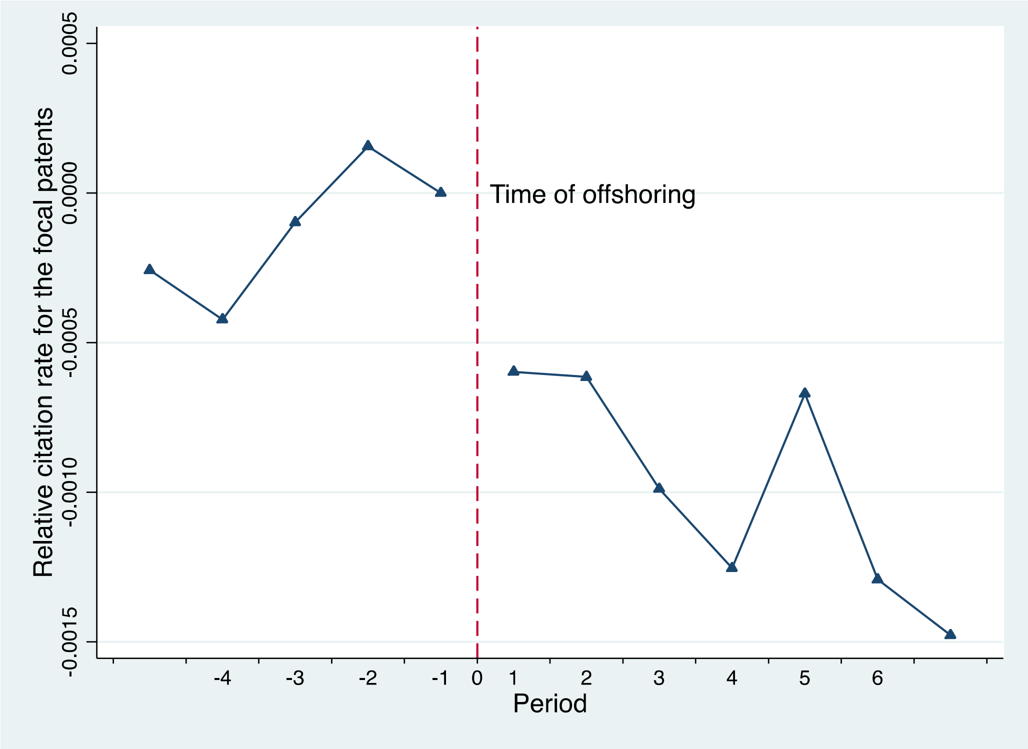

Figure 4. Trends in technological proximity indices between Chinese organizations and offshore MNCs

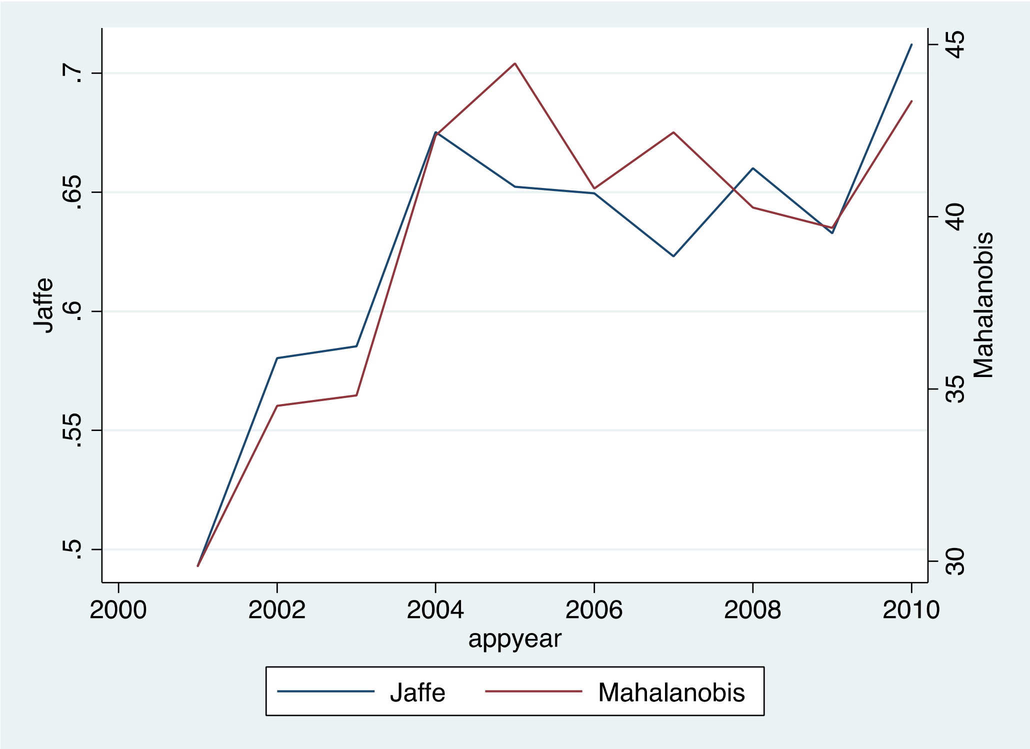

Figure 5. Comparing the distributions of citation lags: Focal vs. control