

INSEAD

The Business School  
for the World®

# Faculty & Research Working Paper

Recruiting for Ideas: A Difference-in-Differences  
Approach for Estimating the  
Effect of Mobility on Access to an  
Inventor's Prior Knowledge

---

Jasjit SINGH  
Ajay K. AGRAWAL  
2009/46/ST

## Recruiting for Ideas:

### A Difference-in-Differences Approach for Estimating the Effect of Mobility on Access to an Inventor's Prior Knowledge

by

Jasjit Singh\*

and

Ajay K. Agrawal\*\*

August 28, 2009

We thank INSEAD, the Social Sciences and Humanities Research Council of Canada, and the Martin Prosperity Institute for funding this research. We are grateful to James Costantini, Pushan Dutt, Lee Fleming, Martin Gargiulo, Javier Gimeno, Henrich Greve, Kwanghui Lim, Matt Marx, Illian Mihov, Johannes Pennings, Andrew Shipilov, Man Zhang, and seminar participants at INSEAD for comments. We also acknowledge the use of the NBER and NUS-MBS patent datasets. Errors remain our own.

\* Assistant Professor of Strategy at INSEAD, 1 Ayer Rajah Avenue, Singapore 138676  
France. Phone: 65 6799 5341, e mail: [jasjit.singh@insead.edu](mailto:jasjit.singh@insead.edu)

\*\* Peter Munk Professor of Entrepreneurship ; Director, Program on Innovation and Creative  
Industries, Martin Prosperity Institute ;Research Associate, National Bureau of Economic  
Research at Rotman School of Management, University of Toronto, 105 St. George Street,  
Toronto, ON M5S 3E6, Canada. Phone: +1 416 946 0203, email:  
[ajay.agrawal@rotman.utoronto.ca](mailto:ajay.agrawal@rotman.utoronto.ca)

A working paper in the INSEAD Working Paper Series is intended as a means whereby a faculty researcher's thoughts and findings may be communicated to interested readers. The paper should be considered preliminary in nature and may require revision.

Printed at INSEAD, Fontainebleau, France. Kindly do not reproduce or circulate without permission.

## **Abstract**

When firms recruit inventors, they may acquire not only the use of their skills but also enhanced access to their stock of prior ideas. In this paper, we examine the extent to which mobility increases a hiring firm's usage of a recruited inventor's prior knowledge, estimating a boost of 197% on average. For this estimation, we employ a difference-in-differences approach novel to this research setting to model the mobility-knowledge flow relationship, comparing both pre-move and post-move citation rates for focal patents versus matched-pair control patents. Our empirical approach has two significant benefits. First, it does not suffer from the severe upward bias we show to be inherent in the commonly employed cross-sectional comparison of just the post-mobility citation rates for focal versus control patents. Second, it enables us to examine the temporal characteristics of the mobility effect, a particularly important issue given the dynamic nature of the knowledge flow process. Furthermore, unpacking the mechanisms behind a firm's increased use of a mobile inventor's stock of prior knowledge reveals that approximately 60% of this increase is actually due to continued self-citation by the mover herself, implying that a new recruit is often responsible for exploiting her prior knowledge without necessarily stimulating broader firm-wide learning. Moreover, such own-knowledge exploitation by the mobile inventor represents an even greater fraction of the firm's increased use of her prior ideas in the years immediately following the mobility event. Finally, consistent with previous research emphasizing the role of mobility for exploration, we find that post-mobility increase in a firm's use of the inventor's prior knowledge is greatest, in percentage terms, for technological domains that are relatively new to the firm.

**Keywords:** Inventor mobility, Knowledge transfer, Knowledge spillovers, Difference-in-differences

## 1. INTRODUCTION

A firm's ability to generate, recombine, and exploit knowledge is critical for competitive advantage (Grant 1996, Kogut and Zander 1996). However, since a firm's internal knowledge development tends to be path-dependent, complementing knowledge developed internally with that from outside is crucial for balancing exploitation and incremental development with distant exploration and recombination (Nelson and Winter 1982, Dosi 1988, March 1991, Weitzman 1998). Hiring experienced personnel has long been recognized as an important means for achieving this balance, and might be particularly useful when the underlying knowledge is tacit and cannot be easily codified (Argote et al. 2003).

Commenting on the potential role of mobility in facilitating inter-firm knowledge spillovers, Arrow (1962: 615) remarks: "Mobility of personnel among firms provides a way of spreading information. Legally imposed property rights can provide only a partial barrier, since there are obviously enormous difficulties in defining in any sharp way an item of information and differentiating it from other similar sounding items." In a similar vein, Stinchcombe (1965: 148) notes that firms lacking certain skills internally can "get by with generalized skills produced outside the organization." Proponents of the institutional theory, such as DiMaggio and Powell (1983), have also suggested inter-firm movement of personnel as a particularly important mechanism through which innovations diffuse among competitors in an industry. Analogously, the resource-based view of the firm also acknowledges that recruitment from outside can enable firms to bypass constraints on growth that relying solely on internally grown resources and capabilities would impose (Penrose 1959, Barney 1991).<sup>1</sup>

Scholars of innovation and technology have particularly emphasized the potential role of mobility in inter-firm knowledge transfer, with the focus naturally being on mobility of technical personnel such as inventors. Mobile inventors bring with them not only technology-specific skills but also tacit knowledge associated with prior inventions they have created, built upon, or been directly

---

<sup>1</sup> Our explicit focus is on an individual moving from one firm to another *existing* firm. It is important to note that there is also an extensive literature emphasizing knowledge transfer through mobility in the context of new firm formation (Bhide 2000, Agarwal et al. 2004, Gompers et al. 2005, Klepper and Sleeper 2005).

exposed to at their prior firm. Such tacit knowledge can enable the carrier of an idea, even a patented idea that is partially codified and disclosed in the form of claims in a patent, to facilitate more effective exploitation of that knowledge. In fact, such transfer of knowledge to the new employer may take place without the old firm being fully compensated for the knowledge the mover carries with her.

Scholars of economic geography as well as economic growth have argued that knowledge flow from inter-firm movement of engineers and researchers has been crucial in the economic success of regions like California's Silicon Valley. As Saxenian (1994: 34-37) notes: "Silicon Valley was quickly distinguished by unusually high levels of job hopping. During the 1970s, average annual employee turnover exceeded 35 percent in local electronics firms and was as high as 59 percent in small firms.... Early efforts to take legal action against departed employees proved inconclusive or protracted, and most firms came to accept high turnover as a cost of business in the region.... This decentralized and fluid environment accelerated the diffusion of technological capabilities and know-how within the region."

Prior empirical research has also reported a positive association between inventor mobility and knowledge flow. In one of the first empirical studies on the subject, Almeida and Kogut (1999) show that locations with greater intra-regional labor mobility between firms also tend to have more localized knowledge flows. In another pioneering study, Song et al. (2003) show that mobile inventors build upon ideas from their previous firm more often than other inventors at the hiring firm who have not previously worked at that firm. In yet another influential article, Rosenkopf and Almeida (2003) examine firm pairs and show that dyads that experience more labor mobility between them also demonstrate greater subsequent knowledge flow. Taken together, these studies have inspired further research employing similar data and methods to sharpen our understanding of various aspects of the mobility-knowledge flow relationship (e.g., Agrawal et al. 2006, Corredoira and Rosenkopf 2008, Agarwal et al. 2009).<sup>2</sup>

---

<sup>2</sup> While we focus on the direct effect of inter-firm mobility on knowledge transfer, it is worth noting that mobility also influences the structure of interpersonal networks within and across firms, which in turn play an important role in further shaping knowledge diffusion patterns (Singh 2005, Fleming et al. 2007a, Breschi and Lissoni 2009).

The above studies have not only advanced the field by establishing an empirical link between mobility and knowledge flow and providing insights into particular nuances of this relationship; they have also raised new questions. Recognizing the inherent limitation of a cross-sectional research design for making causal inferences regarding the observed mobility-knowledge flow correlation, Rosenkopf and Almeida (2003: 764) offer this challenge: “Future research should attempt to utilize fully developed longitudinal databases to explore all possible temporal and causal links.” We take up that challenge here.

To push forward on the causality question in the context of mobility and knowledge flow, we analyze data on year-by-year citations that a patent receives, using within-patent variation in the citation rate over time to estimate the extent to which mobility really facilitates knowledge transfer.<sup>3</sup> We further exploit our fine-grained data to examine who in the recruiting firm is responsible for the increased use of the recruit’s prior knowledge, distinguishing between self-exploitation of the knowledge by the mobile inventor, increased usage by her new collaborators in the recruiting firm and increased usage by others in the firm. We then explore the temporal characteristics of the knowledge transfer process as well as variation across different types of individuals in the timing of their increased use of the recruit’s prior knowledge. Finally, we apply our novel estimation technique to re-examine the prior finding that mobility has a greater impact on knowledge flow when exploring technologies outside the primary domains of activity for the recruiting firm (Song et al. 2003, Rosenkopf and Almeida 2003).

Drawing upon the basic empirical approach from the pioneering works mentioned above, virtually all recent studies investigating various aspects of the inventor mobility-knowledge flow relationship continue to employ a cross-sectional methodology.<sup>4</sup> However, a cross-sectional association between mobility and use of the recruit’s knowledge is in itself not enough to draw conclusions about the extent to which mobility *leads to* an increased use of the recruit’s knowledge.

---

<sup>3</sup> Although they examine research questions different from ours, Furman and Stern (2006) and Murray and Stern (2007) have inspired our basic approach of examining temporal shifts in citation rates.

<sup>4</sup>Oettl and Agrawal (2008) is an exception in that the authors employ longitudinal data, though they do not account for potential shifts in a firm’s technology focus the way we do by applying a difference-in-differences approach.

Establishing causality is actually quite challenging, especially as an inter-firm move is not a random event but the result of deliberate decision-making by the hiring firm as well as the mobile inventor. While our study does not necessarily resolve all aspects of this complex issue, we do address several key causality-related concerns.

Perhaps the most serious concern about the commonly employed cross-sectional approach is regarding unobserved heterogeneity of the underlying knowledge. In particular, inventions associated with future mobility may have an increased likelihood of being used (than, for example, technologically similar “control” inventions used as a benchmark in a cross-sectional comparison) simply because they are of inherently greater quality and/or more relevance for the destination firm. In other words, they would be more intensively used by the firm irrespective of whether the mobility event had taken place or not. Taking this explicitly into account is important to avoid an upward bias otherwise inherent in estimated gains one would attribute to mobility (a bias we demonstrate to be quite severe).

Another issue important to take into account is that the decision to hire itself might just be a manifestation of an overall shift in the firm’s technology strategy. Such a shift in strategy could lead to the decision to hire *and* increased use of the inventor’s prior knowledge, without the former (fully) being a causal precursor of the latter. While the “matched pair” approach in existing studies is indeed intended to deal with this issue to some extent, the problem is inherently unavoidable in any cross-sectional research design as the matching employed can never be perfect. This becomes a particularly serious concern with a majority of existing studies that employ a three-digit technology classification system to construct the matched sample, since such coarse matching would surely fail to capture some critical relevant characteristics of the underlying knowledge. We explain and demonstrate how concerns about inherently imperfect matching become far less critical when employing our empirical framework.

In examining the above issues, we focus our attention specifically on the effect of the mobility event on the hiring firm’s propensity to use the recruit’s own prior inventions, and discuss later how the approach can be generalized to the broader stock of ideas a mobile inventor may bring with her. As such, we construct a longitudinal dataset to exploit within-patent variation over time in

the extent to which a new firm exploits the recruit's prior inventions after the mobility event. We employ a difference-in-differences estimation approach to compare temporal changes in the propensity to build upon the mobile inventor's inventions (before and after the mobility event) with changes in the propensity to build on a matched sample of technologically comparable external ideas created by other inventors. This estimation approach enables us to make both a methodological and conceptual contribution. We come closer to being able to offer a causal interpretation of the mobility-knowledge flow relationship, while providing insights into who in the hiring firm exploits the recruit's prior ideas, when they do so, and what types of ideas they exploit most intensely.

Our research objectives require shifting the focus from the firm as a whole to individuals within a firm. As Simon (1991: 126) notes: "We must be careful about reifying the organization and talking about it as 'knowing' something or 'learning' something. It is usually important to specify where in the organization particular knowledge is stored, or who has learned it.... Since what has been learned is stored in individual heads, its transience or permanence depends on what people leave behind them when they depart from an organization or move from one position to another. Has what they have learned been transmitted to others or stored in ways that will permit it to be recovered when relevant?" For our setting, this translates into distinguishing the extent to which apparent inter-firm knowledge transfer is a result of the mobile inventor leveraging her past knowledge versus broader knowledge exploitation by others in her new firm. In the former case, the recruit rather than the firm might be better positioned to capture gains from mobility, since she can credibly threaten to take the knowledge away by leaving (Becker 1962, Lazear 1986, Coff 1997, Moen 2005, Groysberg et al. 2008). As Peteraf (1993, p. 187) remarks: "A Nobel Prize-winning scientist may be a unique resource, but unless he has firm-specific ties, his perfect mobility makes him an unlikely source of sustainable advantage". However, few existing studies have made such a distinction while examining the use of an inventor's knowledge by the destination firm (One notable exception is a recent working paper by Tzabbar et al. 2009, as described below).

Overall, we do find compelling evidence that mobility facilitates an enhanced use of the mobile inventor's prior ideas: When firms recruit inventors, they increase their utilization of those individuals' prior ideas by 197%, on average, as measured by the hiring firms' citation rates to these

ideas. Importantly, while this is economically quite significant, it is almost an order of magnitude smaller than benefits estimated using only a cross-sectional comparison of post-mobility citation rate to the focal versus control patents. The latter approach (employed in most existing studies) would suggest a much larger (931%) increase in citation rate attributable to mobility, with this upward bias being driven by the fact that patents associated with inventors who subsequently move tend to be more highly cited by the destination firm even before the move occurs (255% more). While our methodology separately accounts for such large systematic differences among patents, a conventional cross-sectional comparison would misinterpret this simply as being a causal result of the move.

After examining the overall effect of mobility, we proceed to disentangle specifically who in the organization is responsible for the increased use of the recruit's prior ideas. Of the 197% boost mentioned above, we find the new recruit's building upon her own prior invention accounts for more than half (118%). This is particularly noteworthy given that much of the prior literature has characterized the knowledge flow associated with mobility as "learning by hiring." To the extent that the new recruit is the only one carrying the knowledge and could leave her new firm, this finding offers an important caveat to a generic firm-learning interpretation: In what sense does the recruit's exploitation of her own prior ideas really reflect "learning" by the hiring firm? While the result is certainly evidence of increased usage or "exploitation-by-hiring," the extent to which it really can be classified as sustainable learning is debatable. Tzabbar et al. (2009) similarly challenge the notion of "learning by hiring" in a study that also finds a significant fraction of the mobility-related knowledge flow to come from the mobile inventor's self-exploitation of knowledge, though they employ a research design closer to traditional cross-sectional studies than our longitudinal difference-in-differences methodology.

Despite the largest component being self-exploitation by the mobile inventor, we find that a significant (81%) boost in the post-move utilization of the recruit's prior ideas derives from others within the firm. Indeed, it seems more reasonable to characterize this component of increased usage of the recruit's prior ideas – by individuals other than the recruit herself - as broader firm-level "learning by hiring". Interestingly, even when we exclude use of the knowledge by individuals in the new firm who become collaborators of the new recruit, we still find a 48% boost in the utilization of

the mover's prior ideas. Finally, when we further exclude inventors even in the recruits' extended network (i.e., her collaborators' collaborators in the destination firm), we still find an increase in the firm's use of the mover's prior ideas by 30% after the mobility event. Overall, the evidence clearly shows that, in making use of the mobile inventor's prior knowledge, the destination firm does not rely just on self-exploitation by the mobile inventor. Instead, other inventors in the destination firm also show an increase in their use of that knowledge. Importantly, however, this increase is more pronounced among direct or indirect collaborators of the mobile inventor, with only limited diffuse and widespread "learning by hiring" among inventors outside the collaborative network of the mobile inventor.

Recognizing that knowledge flow is a dynamic process and that the relevance of prior ideas may change over time, we also examine the detailed temporal patterns of the destination firm's usage of a mobile inventor's prior ideas. We find that, on average, the hiring firm makes its largest boost in utilization of the recruit's prior ideas immediately after the mobility event. Nonetheless, the hiring firm continues to use the inventor's prior ideas at an increased level for many years thereafter. Specifically, we estimate that the average increase in a firm's utilization of a recruit's prior ideas peaks at 235% during the first four years after the move and averages approximately 140% for the next two four-year periods. The mover exploiting her own knowledge appears to largely drive the spike of heavy usage in the years immediately after the mobility event, whereas others in the firm are reasonably constant in their increased usage post-move. As a result, while citations from inventors other than the focal inventor drives only 84% of the 235% boost in citation rate in the four year period right after the move, this increases to 101% of the 144% boost in the period "9-12 years" after the mobility event.

Finally, we apply our difference-in-differences apparatus to one of the most salient prior findings in the mobility-knowledge flows literature: recruiting of new employees as a mechanism to carry out exploration into domains of knowledge new or unfamiliar to the firm (Song et al. 2003, Rosenkopf and Almeida 2003). Consistent with this exploration result, our findings show that firms increase their use of new recruits' prior ideas by only 48% if those ideas are within fields in which the firm is reasonably active, while the use goes up further by a striking 309% if those ideas are outside

the firm's domain of inventive activity. An important caveat in interpreting this result is that although this increase reflects a relatively dramatic jump in percentage terms, it may not indicate a large increase in levels; on average, firms utilize such out-of-domain inventions 77% less to begin with. Furthermore, the increased use of recruits' out-of-domain inventions is not due to the recruit alone: on average, other inventors in the recruiting firm are responsible for a 223% boost in the utilization of such inventions.

Overall, these findings deepen our insight into the mobility-knowledge flows relationship. The mobile inventor herself plays a critical role in recruiting for ideas, especially in the period immediately following the mobility event. However, evidence also clearly shows others in the destination firm, particularly direct and indirect collaborators, increasing their use of the mobile inventor's prior knowledge. Finally, a firm's exploitation of a new recruit's prior ideas is particularly salient when exploring domains that are relatively new to the organization.

In the next section, we describe our overall empirical framework, with emphasis on clarifying how our difference-in-differences methodology differs from conventional methods exploring the link between inventor mobility and knowledge flow. In Section 3, we describe the construction of our dataset and key variables. We present our empirical results in Section 4. Finally, in Section 5, we discuss the implications of our findings as well as the limitations of our study and directions for future research.

## **2. EMPIRICAL FRAMEWORK**

### **2.1. Using patent data to examine the link between mobility and knowledge flow**

Micro-level data suitable for examining the link between mobility and knowledge flow is hard to come by. A notable exception is patent data, which have been commonly employed for this purpose as they include detailed information on each patent's inventor names and location, application and grant dates, assignee organization (if any), technological classification, etc. In particular, chronological tracing of the different organizations where the same inventor appears over her career history allows one to infer inter-organizational movements made by the inventor.

Data on citations a patent receives help the researcher infer how others have subsequently used a piece of knowledge. Admittedly, citations are not perfect in measuring knowledge flow. For example, they are often added for reasons such as avoiding litigation or clarifying claims, and many are in fact added by patent examiners rather than the inventors themselves. Despite this, however, scholars have shown that they correlate quite well with actual knowledge flow, especially when employing large samples (Jaffe and Trajtenberg 2002, Chapter 12; Duguet and MacGarvie 2005).<sup>5</sup> A specific concern might be whether examiner-added citations could bias the results in a particular research setting (Alcacer and Gittelman 2006). However, since inventors may have strategic motives for omitting certain citations, including examiner-added citations might actually be desirable (Lampe 2008). While we would still have liked to do a robustness analysis using just inventor-added citations, we were unable to get these data in machine-readable form for the time period of our study.

## **2.2. The conventional cross-sectional approach for examining the mobility-knowledge flow link**

To compare with previous studies, we start by implementing a method of examining the link between inter-firm mobility and knowledge flow that is similar to research designs employed in existing studies. The “best practice” from existing studies would suggest the following steps (also illustrated in Figure 1). First, we create a sample of “focal patents” representing prior knowledge associated with mobile inventors. Second, we match each focal patent with a technologically comparable “control patent” with the same three-digit technology class and the same application year. Finally, we examine cross-sectional differences between the focal and control patents in terms of citations they receive from the destination firm in the period following the move.

Formally, let us define  $CITES_{i,t}$  as the number of citations patent  $i$  receives from the destination firm in year  $t$  ( $t$  being any of the years following the mobility event). The estimation equation is therefore:

---

<sup>5</sup> It is worth emphasizing that using citations as a measure correlated with actual knowledge flows *does not* assume citations are the mechanism behind these flows. As an analogy, a PhD student’s citation to his advisor’s research paper may tell us that he built upon knowledge that the advisor created, even if most of the actual knowledge transfer happened not through reading the advisor’s papers but by working closely with the advisor.

$$(1) \quad CITES_{i,t} = f(\psi_F FOCAL\_PATENT_{i,t} + \psi_X X_i + \delta_{t-applyear} + \beta_t + \varepsilon_{i,t})$$

Here, FOCAL\_PATENT is a dummy variable equal to one for observations corresponding to an original patent and zero for the corresponding control patent. X is a vector of control variables, which includes the sizes of the overall patent pools of the source and destination firms as well as specific patent characteristics such as number of claims, number of references to other patents, number of non-patent references, inventor age, the inventor's prior patenting experience, U.S. vs. non-U.S. origin, and indicator variables for the two-digit technology categories.

Rather than imposing an arbitrary functional form regarding temporal patterns of citation, the above model non-parametrically accounts for the patent age (using indicator variables  $\delta_{t-applyear}$  for age measured in years) and citing year (using indicator variables  $\beta_t$ ). Since yearly patent citations involve count data skewed to the right (and over-dispersed relative to Poisson), we employ a negative binomial model. Employing robust standard errors with clustering using the cited patent accounts for non-independence of observations pertaining to the same original patent. Our baseline expectation for estimation is that  $\psi_F$  should be positive and significant, representing an association between inter-firm mobility and an increase in subsequent citations to the recruits' prior patents.

Even though commonly employed, a cross-sectional approach like the one above is not satisfactory in dealing with the issue of unobserved heterogeneity across patents. In particular, even control patents drawn from the same three-digit technology class as the focal patent may be systematically different from the focal patents in their overall value and/or relevance for the destination firm, which would affect the number of citations they receive even independent of the move. For example, if patents with better quality are associated with individuals better placed to seek career advancement through mobility, we would observe a positive association between mobility and citations even though greater citation rates are not actually caused by mobility. Similarly, if patents cited more heavily by a destination firm arise from individuals who the destination firm sees as having more relevant skills, a positive association between mobility and citations could arise even if the two are not causally related. The need to disentangle these effects motivates our "difference-in-differences" research design.

### 2.3. A “difference-in-differences” framework for examining the mobility-knowledge flow link

Our “difference-in-differences” approach makes progress on the challenges with the conventional cross-sectional approach by exploiting the fact that we observe citations received by the focal and control patents not just post-move but also in the years preceding the move. While post-move differences confound actual mobility-related knowledge acquisition with just differences in quality and/or relevance, these effects can be separately identified by taking into account differences in citation rates that existed even pre-move. In other words, a pre-move difference in the citation rates for the focal versus control patents can be a benchmark against which to compare the post-move difference, helping us tease out that component of the post-move difference that we can more confidently attribute to the move itself.<sup>6</sup>

We provide Figure 2 to illustrate how our difference-in-differences approach conceptually differs from the cross-sectional approach described above. For ease of exposition, the figure depicts the effect of mobility as an immediate bump in the natural evolution of the citation rate of a patent over time (our actual empirical analysis relaxes this assumption).<sup>7</sup> A cross-sectional approach would take a positive difference in post-mobility citation rates for the focal patent ( $X_a$ ) versus control patent ( $Y_a$ ) as evidence of knowledge flow attributable to the mobility event. However, had the move not taken place, the expected citation rate to the focal patent ( $X_e$ ) would still be greater than that for the control patent ( $Y_e$ ). Comparing citation rate to the focal versus control patent both before and after the move allows us to tease out how the temporal pattern in citation rate changes from the expected path *as the result of mobility*. In other words, what a “difference-in-differences” logic suggests is that the effect attributable to the mobility event is best captured not by the extent to which  $X_a > Y_a$  but by the extent to which  $X_a - Y_a > X_e - Y_e$ .

---

<sup>6</sup> For a general introduction to “difference-in-differences” methodology, see Card and Krueger (1994). See Furman and Stern (2006) and Murray and Stern (2007) for other applications of this in the context of examining citations.

<sup>7</sup> Note that we allow for the possibility of a post-mobility bump even for the control patent. This could happen, for example, if the move coincides with an overall shift in a firm’s technological focus, resulting in a general increase in citation rate for the given technology. In the discussion section, we discuss alternate interpretations.

Implementing the above logic needs a patent-year dataset that includes observations from the years not just after but also before the move. A new dummy *POST\_MOVE* distinguishes between the pre-move and post-move time window, being 0 for the pre-move period and 1 for the post-move period for a given focal patent and the corresponding control. The estimation equation now becomes:

$$(2) \quad CITES_{i,t} = f(\psi_F FOCAL\_PATENT_{i,t} + \psi_{FP} FOCAL\_PATENT_{i,t} * POST\_MOVE_{i,t} + \psi_P POST\_MOVE_{i,t} + \psi_X X_i + \delta_{t-appyear} + \beta_t + \varepsilon_{i,t})$$

Here,  $\psi_F$  captures systematic differences in the focal versus control patents that existed even before the move, and  $\psi_P$  captures systematic shifts in how the destination firm leverages knowledge in this technology irrespective of the move taking place. If mobility leads to an increased use in the inventor's prior knowledge,  $\psi_{FP}$  should be positive and significant in the estimation.

While the above framework allows for a systematic difference between focal and control patents, we can generalize it further to allow even individual patents to be different in unobserved ways (e.g., due to inventor characteristics we do not observe). In particular, patent fixed effects analysis can help detect “abnormal” within-patent changes in the citation rate to a patent after an inventor moves. The estimation relies on deviation of the patent's post-mobility citation rate from an expected rate, where we derive the expected rate by extrapolating from the pre-move citation rate and assuming a temporal trend analogous to other patents from the same technology-year cohort. We express this model as:

$$(3) \quad CITES_{i,t} = f(\gamma_i + \psi_{FP} FOCAL\_PATENT_{i,t} * POST\_MOVE_{i,t} + \psi_P POST\_MOVE_{i,t} + \delta_{t-appyear} + \beta_t + \varepsilon_{i,t})$$

Here,  $\gamma_i$  helps capture fixed effects for the patents. Since the fixed effect absorbs unique patent characteristics, the model no longer includes variables invariant within a patent (*FOCAL\_PATENT<sub>i</sub>* and *X<sub>i</sub>* in the previous model).<sup>8</sup> Given the count data involved, we base our implementation of the

---

<sup>8</sup> Since the lag between patent application and the move varies across original patents, it is in principle possible to carry out the estimation just using the focal patents. However, we prefer to keep the control patents, as doing so reduces concerns regarding endogeneity of a firm's decision to hire in response to a shift in its technology strategy.

above fixed effects model on a Poisson model (Hausman et al. 1984), with the standard errors computed using a Quasi Maximum Likelihood approach (Wooldridge 1999).<sup>9</sup>

### 3. DATASET CONSTRUCTION

#### 3.1. Identifying inventors and detecting inter-firm mobility

We merge patent data obtained directly from the USPTO with patent data made available by the National Bureau of Economic Research (Jaffe and Trajtenberg 2002, Chapter 13) and the National University of Singapore-Melbourne Business School patent dataset. We then enhance these along two dimensions. First, for each assigned patent, we determine the assignee organization by carrying out an assignee name clean-up followed by a parent-subsidary match.<sup>10</sup> Second, we use not just the names of the inventors but also other data fields (i.e., technology classification, inventor address, collaborator names, citation information) to create unique identifiers for each inventor across patents.<sup>11</sup>

In detecting inter-firm mobility, we follow other researchers who have inferred such mobility through observed changes in the assignee firm in successive patents filed by an inventor (Almeida and Kogut 1999, Song et al. 2003, Fleming et al. 2007a). Admittedly, this approach only captures mobility where an inventor successfully files for a patent both pre-move and post-move. Another challenge is that, even when we do observe two successive patents from the same inventor but at different firms, we cannot pinpoint the inventor's exact move date within this window. We thus base all analysis reported in this paper on mobility events for which this window was four

---

<sup>9</sup> We implement this in Stata using the "xtqmlp" procedure (written by Tim Simcoe and available for download at <http://scripts.mit.edu/~pazoulay/Software.html>), which corrects the standard errors from a fixed effects Poisson model for over-dispersion. This procedure addresses the concerns regarding interpreting a conditional fixed effects negative binomial model as a true fixed effects estimator (Wooldridge 1999, Allison and Waterman 2002).

<sup>10</sup> We follow the assignee matching procedure previously used by Singh (2007), which includes using *Compustat* identifiers from NBER, Stopford's *Directory of Multinationals*, *Who Owns Who* directories, and Internet sources.

<sup>11</sup> We base the name matching algorithms used on Singh (2008), which are similar to procedures implemented by Trajtenberg (2006) and Fleming et al. (2007a, 2007b).

years or less, as a detailed temporal examination would not be useful for cases where the move date is too uncertain. Therefore we drop about 30% of all observed inter-firm mobility events.<sup>12</sup>

Lacking further information about the move date, we start by calculating the halfway point between the last observed date at the original firm and the first observed date at the new firm. However, while the move could have taken place anytime after the start of the window, it would almost surely have taken place at least a few months before the end of the window since there would have to be a lag between an inventor joining a new firm and filing for a patent there. Further, we specify the temporal unit of analysis to be the year, rather than a day or month, to avoid any pretence of a precise estimated move date. These two factors lead us to define the beginning of the calendar year of the mid-point calculated above as our actual estimate of when a move happened. Given the high uncertainty in the move date, this somewhat ad-hoc assumption is unlikely to do much worse than more sophisticated heuristics. On the other hand, it facilitates analysis by allowing classification of each calendar year as being either completely before or after a given move, allowing us to work with a patent-year panel dataset.

### **3.2. Constructing the sample of focal patents and corresponding control patents**

In order to ultimately obtain a dataset of large yet manageable size while allowing for a sufficient future time window for observing subsequent inventor mobility as well as annual citation counts, we start with the population of patents with application years 1981-1990 (across all technology classes in the USPTO database). For comparability across patents as well as unambiguous conceptualization of inter-firm mobility (or lack thereof) for any given patent, we restrict ourselves to patents with a single inventor. From this subset, we draw our sample of focal patents for which the inventor exhibits inter-firm mobility anytime between the second and twelfth year following the

---

<sup>12</sup> One might worry about representativeness of the final sample. For example, since longer time windows imply fewer patents per year, the dropped observations could pertain to less productive inventors. To rule out the possibility of any resulting biases, we redo the analyses reported in the paper on a full sample of mobility events as well as on subsamples using different window cut-offs. The main results remain qualitatively unchanged in all cases.

application year (we exclude the first year so that there is at least one year of pre-mobility citation data, which we need for our difference-in-differences calculation).

As shown in Figure 1, we attempt to match each of these “focal” patents with a corresponding “control” patent with the same application year and three-digit technology class. The control patent must originate in another firm and also have a single inventor, but one who did not exhibit any subsequent inter-firm mobility in the 12 years that followed. Having a control patent for each focal patent helps account for general shifts in the technological focus of the firm, as those will get reflected in an increased likelihood of citing not just the original patent but also the control patent. In the relatively infrequent cases where the focal patent cannot be matched with a control patent, we drop the focal patent from the sample. We also examine the robustness of our findings to address the concern that a better “apples to apples” comparison might result from using the more precise nine-digit technology subclasses rather than the broad three-digit technology class (Thompson and Fox-Kean 2005).<sup>13</sup>

The above steps leads to a final sample of 33,886 patents, half of which are focal patents associated with an inventor who subsequently moves, the other half which are corresponding control patents. For each focal or control patent, we calculate the number of annual citations made by the destination firm in the 12 years following the application year. Therefore, we have 12 observations per patent, resulting in a patent-year panel of 406,632 observations.<sup>14</sup>

Our actual implementation of the difference-in-differences research design from the previous section includes two refinements. First, as a way to deal with the uncertainty of the exact move date, we include an additional indicator variable for the two-year “move window” around the calculated move date. Therefore, we actually classify each of the 406,632 observations as being in one of three categories - post-move, in the move window, or pre-move - depending on where it temporally falls

---

<sup>13</sup> Matching on nine-digit subclass has its own drawbacks. First, given that there are more than 150,000 patent subclasses in the database, with a typical subclass having at most a few patents in any given year, it is not possible to find a match for a large majority of the patents. Dropping these not only reduces the sample size drastically, it also raises sample selection concerns (Henderson et al. 2005).

<sup>14</sup> The reason we need such a large sample size is that the citation of a *specific patent* by a *specific firm* is a very rare event, so we need a large sample size to ensure enough non-zero observation of citations in order to carry out a meaningful econometric estimation (especially once we refer to fixed effects models or start to unpack the citations into those coming from the inventor herself, from within her network, or from others).

relative to the calculated mobility date for the inventor of the original patent. We summarize this process in Figure 3.<sup>15</sup>

The second issue we address is that, since matching is not perfect, our model may fail to account for shifts in focus of a firm towards domains related to the focal patent. If this were true, however, we should typically see an increased citation to the focal patent even in the years leading up to the move rather than necessarily right after. Therefore, rather than using the entire pre-time period as the baseline against which we compare post-move citation rate differences, we use the four-year window immediately preceding the move window as the benchmark. Analysis later in the paper allows even the post-move knowledge boost to vary with the time elapsed since the move. This allows for the possibility, for example, that mobility-related use of knowledge happens more in the years immediately following the move than in later years when knowledge might have diffused to the destination firm through alternate channels.

### 3.3. Variables

Table 1a summarizes our dependent variables, as well as the key explanatory and control variables used later in the regression analysis. In line with most existing studies examining the overall use of the inventor's prior knowledge by the destination firm, our first dependent variable, *cites\_all*, includes *all* citations the original patent receives from the hiring firm in the focal year. In particular, it does not distinguish the mobile individual's use of her own prior knowledge from the increased use by others.<sup>16</sup>

The next dependent variable, *cites\_nonself*, excludes self-citations made by an inventor to her own prior patents. This helps distinguish an increase in a firm's ability to leverage the mobile inventor's prior knowledge merely due to her own continued use of her past knowledge from a

---

<sup>15</sup> Recall that our sample excludes moves where uncertainty of the move date exceeds four years. Of the remaining, the uncertainty is zero to two years for 65% of the cases and three to four years for the rest. Thus, an estimate of the move year based on the mid-point is off by not more than one year in two-thirds of the cases, and not more than two years for the rest. The move window in Figure 3 accounts for this, noting that a two-year window typically suffices.

<sup>16</sup> To understand why the mean value for each of the dependent variables is so small, note that we are dealing with highly disaggregated data: these variables capture citations received by a patent from the firm to which the focal patent's inventor subsequently moves. Naturally, the probability of a *specific firm* citing a *specific patent* is extremely low. For example, *cites\_all* takes a non-zero value for only 1.2% of the observations in our sample.

diffusion of the knowledge to other employees in the destination firm. This distinction is important, since self-exploitation by the mobile inventor is more likely to represent a setting where the firm is less able to appropriate the benefits from mobility and the inventor can more easily take the knowledge away by moving again.

Assuming that inter-firm mobility really does contribute to not only *cites\_all* but also *cites\_nonself*, the question remains of exactly which employees benefit in terms of knowledge spillovers from mobility. On the one hand, close interpersonal ties such as those formed by direct or indirect collaboration have been shown to be very helpful for knowledge diffusion, as individuals with whom the mobile inventor interacts closely receive better access to her prior knowledge (Hansen 1999, Reagans and Zuckerman 2001, Singh 2005). On the other hand, other kinds of intra-firm networks or knowledge diffusion mechanisms might enable a wider employee base in the destination firm to obtain access to the knowledge. To see how this empirically plays out, we construct two more dependent variables. The first, *cites\_noncollab*, excludes not just the original inventor's self-citations but also citations made by those who have directly collaborated with the original inventor in the previous five years. The second, *cites\_nonnetwork*, extends this further to also exclude citations made by those at a social distance of two from the original inventor, i.e., collaborators of the recruit's collaborators.

The indicator variable *focal\_patent* distinguishes focal patents from control patents. We classify a given patent-year observation as post-move, in the move window, or pre-move using the indicator variables *post\_move* and *move\_window* (see Figure 3).

We employ several control variables to account for other potential drivers of citations made by the destination firm to the focal patent: *destfirm\_pats* (number of patents assigned to the destination firm in the last five years), *origfirm\_pats* (number of patents assigned to the original firm in the last five years), *claims* (scope of the patent as measured by its number of claims), *patrefs* (number of references made to previous patents), *nonpatrefs* (number of references made to non-patent sources such as trade literature and journals), *inventor\_age* (number of years since the first patent by this inventor), *inventor\_patents* (number of patents the inventor has been involved with in the past), and *US\_inventor* (indicator for whether the inventor has a U.S. address). For variables that

are highly skewed, we use a logarithmic transformation in the actual analysis.<sup>17</sup> As explained earlier, we utilize fixed effects for different citation lags to non-parametrically account for the temporal changes in the citation frequency over time, employ citing year fixed effects to account for any systematic differences in citations from the patent pools from different years, and use technology fixed effects to account for systematic differences in citation rate across different two-digit technology categories.

## 4. RESULTS

We now report our findings on five issues: 1) a comparison of the mobility-knowledge flow as estimated using the traditional cross-sectional method versus our difference-in-differences approach, 2) sensitivity of these two approaches to the matching technique employed in constructing the control sample, 3) pinpointing the individuals in the destination firm most responsible for the increase in knowledge use (the mobile inventor, her direct or indirect collaborators, or other inventors), 4) examining the temporal pattern of different kinds of knowledge flow, 5) comparing benefits from mobility for technology domains that a firm is more versus less familiar with.

### 4.1. Diffusion of the mobile inventor's prior knowledge to the destination firm

Before delving into the regression analysis, we present the basic intuition behind our approach using summary statistics for the subsamples corresponding to the four possible combinations of values the indicator variables *focal\_patent* and *post\_move* can take. We report these statistics in Table 1b (it is also helpful to refer back to Figure 3 in interpreting these). Recall that the indicator variable *focal\_patent* is one for the original patents (involving a move) and zero for the corresponding control patents (not involving a move). We use the indicator variable *post\_move* to distinguish whether a focal observation (for an original or the corresponding control patent) is drawn from the post-move or pre-move period (as defined by the move date for the inventor of the original patent).

---

<sup>17</sup> In doing so, we first add one to any variables that can take a value of zero. The results are robust to changing the size of the offset or using the original variable itself for the analysis instead.

As the column for *post\_move* = 1 indicates, patents from the original sample have a greater average annual citation rate (0.0433 corresponding to *focal\_patent* = 1) than that for the control patents (0.0041 corresponding to *focal\_patent* = 0, giving a difference of 0.0394 between the two). However, it would be misleading to attribute it all to mobility, as the annual citation rate is greater for the original sample even pre-mobility (0.0106 corresponding to *focal\_patent* = 1 vs. 0.0025 corresponding to *focal\_patent* = 0, giving a difference of 0.0081). Therefore, rather than attributing the whole post-mobility difference to the move itself, it seems more justifiable to attribute only the difference between the post-mobility and pre-mobility differences (i.e.,  $0.0394 - 0.0081 = 0.0313$ ) to the move. This analysis is only illustrative, as a regression framework is better for a more thorough analysis.

The baseline analysis reported in Table 2 (like the summary statistics reported in Table 1b) uses *cites\_all* as the dependent variable. Recall that this includes *all* future citations the original patent receives from the destination firm, including self-citations made by an inventor after he or she moves. Column 1 reports findings from a research design comparing the post-move citation frequency of the patents involving an inventor move with comparable patents not involving such a move. The effect of the move appears to be both statistically and economically significant, with the likelihood of a citation implied by the negative binomial regression coefficients being about 10.31 times in the sample where the move does occur versus where it does not.

However, the findings in Column 2 highlight why interpreting the above findings entirely as knowledge acquisition gains from inventor mobility might be inappropriate. Here, we compare the citation frequency of the original and control patents only using the subsample *before* the move date and find that the former systematically receive 3.86 times as many citations even before the move actually takes place. In other words, despite the extraordinary care we take in finding a control sample, the patents in the original sample are still inherently more likely to be cited by the destination firm in part for reasons that cannot be attributed to the move. The subsequent columns take this into account in the analysis by adopting a “difference-in-differences” approach that controls for systematic pre-mobility differences.

Column 3 reports the results from a research design based on a pooled analysis of the entire “original” patent-year panel dataset that includes observations for each of the 12 years following a patent’s application year. Note that this sample includes both the post-move subsample (used in Column 1) and pre-move subsample (used in Column 2), with the indicator variable *post\_move* used to identify whether a focal observation derives from the former. We use another indicator variable *move\_window* to distinguish the two-year window around the estimated move date to account for lack of precision with which the date is actually estimated due to data constraints.

As already discussed, the variable of interest in inferring benefits directly attributable to mobility is the indicator variable *focal\_patentXpost\_move*, which takes a value of one only for post-move observations pertaining to the original patent (but not the control). We find the regression coefficient for this term to be positive and significant, implying that there are significant gains from mobility over and above an increase in the firm’s use of knowledge that happens to coincide technologically with the original patent. However, the implied incidence rate ratio of 2.98 indicates that this effect, while still being large in magnitude, is much smaller than the factor of 10.31 implied by a cross-sectional comparison of post-move citation frequency of original versus control patents as reported in Column 1.

The analysis reported in Column 4 further refines the approach by using the four-year window immediately preceding the move window as a more dependable benchmark to tease out benefits really attributable to the move, ruling out the possibility that the results are driven purely by observations that precede the move date by many years even as the years leading up to the move might demonstrate greater citation rates. We implement this approach by introducing an indicator variable *pre\_move\_yGE6* (defined as one for observations that are ahead of the estimated move date by five years or more). The estimate for *focal\_patentXpost\_move* changes very little, giving confidence in the effect truly being associated with the move itself.

Column 5 helps further address any concern regarding the findings being driven by unobserved heterogeneity across patents, i.e., differences in the intrinsic quality or relevance of patents that the control variables might not fully capture. This analysis is based on a Poisson Quasi-Maximum Likelihood (QML) model with conditional fixed effects for the patents. Given that such

estimation requires within-patent differences in citations received over time, while the number of citations that a most patents receives from a given firm is zero, the usable sample shrinks drastically. Nevertheless, the findings remain very similar to Column 4. This gives more confidence that the findings are not driven by unobserved differences in quality or relevance of the focal patents for the destination firm.<sup>18</sup>

#### **4.2. Using three-digit versus nine-digit technology matching**

To ensure that our results are not too sensitive to the technological classification used for matching (Figure 1), we now compare our findings from the three-digit technology match (reported in Table 2) with what the results would be if we had used a nine-digit technology classification instead. Table 3 reports this analysis. The comparison between the cross-sectional analysis Column 1 of Table 2 and Column 1 of Table 3 demonstrates that using a fine versus relatively coarse technology match really matters when it comes to inferring knowledge spillovers purely through cross-sectional comparison across patents. We find the likelihood of a post-move citation for a patent involving an inventor move to be 4.15 times that for a control patent matched using nine-digit technology match (Table 3), a number much smaller than the factor of 10.31 found earlier in the analysis using the three-digit sample (Table 2).

However, the estimated multiplicative factor of 2.25 from our preferred model (Column 4 of Table 3) is not as far off from the corresponding estimate of 2.97 found earlier with three-digit technology matching (Column 4 of Table 2). This suggests that a traditional cross-sectional empirical design is much more sensitive to the choice of matching technique, while difference-in-differences estimates are more robust to this choice. The intuition for this is that a difference-in-differences design essentially uses the control patent only to anticipate the baseline temporal patterns of citation,

---

<sup>18</sup> To use the full sample, the rest of the paper employs pooled negative binomial analysis, with robust standard errors reported after clustering for non-independence of observations for the same patent. To conserve space, we do not report fixed effects analysis, though our key findings hold using these models as well. Our results are also robust to employing a negative binomial random effects specification, though the strong assumptions behind a random effects specification make it less than ideal: Whether and when an inventor moves might depend on unobserved characteristics of the inventor, implying a correlation between the individual effect and variables like *focal\_patent*.

while allowing the average citation rate (driven by quality or relevance differences) to differ between the focal and control patents.

#### **4.3. Who in the destination firm uses the mobile inventor's prior knowledge?**

Next, we turn to the issue of exactly which individuals in the destination firm end up benefiting from “learning by hiring.” In other words, we are interested in the extent to which increased use of prior knowledge of the incoming inventor is largely restricted to the incoming inventor herself versus spreading to others exposed to her through direct or indirect collaboration and/or more widely to others in the destination firm. The analysis reported in Table 4 examines this by revisiting a pooled analysis analogous to that in Columns 4 and 5 from Table 2 using three additional dependent variables from Table 1a: *cites\_nonsel*, *cites\_noncollab*, and *cites\_nonnetwork*.<sup>19</sup>

For easy comparison, Column 1 reproduces the findings from Column 4 of Table 2 using *cites\_all* as the dependent variable, while Column 2 uses *cites\_nonsel* to exclude the mobile inventor's self-citations to her patents from her previous firm. The estimate for *focal\_patentXpost\_move* becomes weaker, with the magnitude of this effect falling from a multiple of 2.97 in Column 1 to just 1.81 in Column 2. Therefore, although benefits from mobility do appear to extend somewhat beyond the original inventor using her own knowledge, much of the apparent “learning by hiring” reported in Table 2 is actually a result of the original inventor exploiting her knowledge in the new setting rather than the knowledge necessarily spilling over to others.

The results reported in Columns 3 use *cites\_noncollab* as the dependent variable, and reveal that benefits from mobility fall further. The incidence rate ratio of the term of interest is just 1.48 once we restrict the analysis to citations made by individuals with no past collaboration with the mobile inventor. Finally, Column 4 uses *cites\_nonnetwork* as the dependent variable, showing that benefits from mobility are even smaller for individuals with no direct or indirect ties with the mobile inventor; the multiple is now 1.30 and statistically indistinguishable from zero. This suggests that “learning by

---

<sup>19</sup> The findings are qualitatively very similar (typically just marginally greater in magnitude and statistical significance) if we instead employ a fixed effects model analogous to that from Column 5 of Table 2.

hiring” is actually quite localized in the collaborative network that the mobile inventor forms in the destination firm.

#### **4.4. A closer look at the timing of knowledge leveraging**

Table 5 digs deeper into the detailed temporal patterns of patent citations examined above. We define the indicator variables *post\_move\_y1\_4*, *post\_move\_y5\_8* and *post\_move\_yGE9* to be one for the first four-year period after the move window, the second four-year period, and the final four-year period, respectively. Although the coefficient estimate for the interaction of these three variables is positive and significant for all of these time periods in Column 1, the effect is particularly large for the first period. This is in line with the conjecture that knowledge acquisition gains attributable to the move accrue more in the earlier periods rather than later periods post-move. In contrast, the observed benefits from mobility are a little more spread out once we exclude self-citations and in fact exhibit a slightly stronger effect towards the end (though the difference is statistically insignificant). Figure 4 plots these coefficients for easy comparison.<sup>20</sup>

#### **4.5. Knowledge acquisition benefits for unfamiliar versus familiar technologies**

Next, we apply our difference-in-differences framework to re-examine a central finding in the previous literature: Mobility is particularly effective for exploration into technological domains new or unfamiliar to the hiring firm. To investigate this, we define a new indicator variable *unfamiliar\_tech*. If the destination firm has never filed a patent in the focal technology class before, or at least if the fraction of patents from the focal technology class in the firm’s overall patent base in the five years preceding the move does not exceed 2%, the technology is coded as being unfamiliar to the hiring firm.

The findings in Table 6 indicate a clear premium associated with unfamiliar technologies when it comes to benefits from inter-firm mobility: Firms increase their use of new recruits’ prior

---

<sup>20</sup> We find the pattern shown in the figure to be robust to using two-year and three-year time windows instead of four-year windows. However, the large standard errors keep us from being able to rule out statistical equivalence of many of the estimates, which is probably an artifact of a lack of precision of the exact move date in our data.

ideas by only 48% if those ideas are within fields in which the firm is active, while the use goes up further by a striking 309% if those ideas are outside its domain of activity. (It is important to note that although this increase reflects a relatively dramatic jump in percentage terms, it may not indicate a large increase in levels; on average, firms utilize such out-of-domain inventions 77% less to begin with.) Further, the increased use of recruits' out-of-domain inventions is not due to the recruit alone: On average, other inventors in the recruiting firm are responsible for a 223% boost in the utilization of such inventions.<sup>21</sup>

## 5. DISCUSSION AND CONCLUSION

Before turning to an extended discussion, it is useful to recap the key contributions of our study. We have used fine-grained analysis of patent citations to examine and unpack the link between inventor mobility and the destination firm's increased use of the recruit's prior knowledge. The use of longitudinal data and the difference-in-differences approach has enabled us to provide a more causal interpretation to our findings than has been possible in previous cross-sectional research on the subject. In line with previous research, we continue to find robust evidence of mobility leading to a hiring firm's increasing usage of a mobile inventor's prior knowledge. Very importantly, however, we find the extent of this increase to be almost an order of magnitude smaller than what the conventional cross-sectional approach would lead us to believe. We find that the mobile inventor's self-exploitation of knowledge plays a critical role in the increased usage of her ideas by the destination firm post-mobility, especially in the period immediately following the mobility event. While evidence clearly shows others in the destination firm also increasing their use of the mobile inventor's prior knowledge, this effect seems quite localized among the recruit's direct or indirect collaborators in the destination firm; thus, evidence of more diffuse "learning by hiring" is limited. Finally, in line with

---

<sup>21</sup> In additional analysis, we checked whether these findings are just an artifact of benefits from mobility possibly being greater for technologies that are new not just for the *destination firm* but *in general*. To do this, we defined newness of a three-digit technology class based on where the median application year for patents in that class is relatively recent (such classes naturally involved newer fields like computers, semiconductors, biotechnology, etc.). The findings indicate no extra benefits from mobility for the newer technologies. Thus the findings in Table 6 really do reflect an effect of whether a technology is relatively new or unfamiliar specifically for the destination firm.

previous research, knowledge access related benefits from hiring appear to be most salient when exploring domains that are relatively new to the organization.

We now turn to a discussion of six points concerning the limitations, implications, and extensions of our reported results. Specifically, these include: 1) limitations to a causal interpretation of our findings, 2) managerial implications of the mobility effect being restricted largely to the original inventor and her immediate network, 3) suggestive evidence of broader absorptive capacity effects facilitated by recruiting, 4) extending our analysis to a mobile inventor's prior knowledge beyond just her own past inventions, 5) robustness of the difference-in-differences estimates to employing three-digit versus nine-digit technology classification matching, and 6) other identification strategies that could be brought to bear in empirical efforts to further understand the causal effects of mobility.

First, each of our results concerning different aspects of the mobility-knowledge flow relationship – who in the hiring firm increases use of the recruit's prior ideas, when that increase occurs, and what type of ideas experience the greatest boost in use – is predicated on our central methodology for identifying how mobility *facilitates* an increased use of a recruit's prior ideas. Acknowledging limitations to the causal interpretation our framework allows, we purposefully use the term “facilitates” rather than “causes” when referring to the nature of this relationship. Although our use of longitudinal data and estimation techniques enable us to make progress in drawing a causal interpretation from our results, we recognize that there is still a way to go as firms' recruiting decisions as well as inventors' mobility decisions are endogenous in ways our framework does not account for.

An illustrative example clarifies the limits to causal claims we can make. Consider a firm that changes its technological focus to a very specific area, and subsequently decides to hire an inventor with expertise in that area. An increased use of knowledge from this area (including in particular the recruit's prior knowledge) in this case results from the shift in the hiring firm's focus irrespective of the move later taking place. Thus, the increase should not be seen as having been caused by the move. Since our longitudinal approach recognizes the temporal pattern - the firm starts to use the knowledge even before the move - it correctly identifies the relationship as not being causal (noticing

the citation rate is already higher in the years leading up to the move, the period used as the baseline against which we compare citation rates in the post-move period). Furthermore, Figure 4, which plots the estimated citation rate in periods before and after the move, allays concerns that the upwards trend might begin in the years prior to the mobility event: instead we see a clear discontinuity in the period immediately following the move.

Now consider a setting where the firm starts to increase its use of a particular inventor's ideas at exactly the same time that it hires an inventor even though it might have done exactly the same even in the absence of recruiting. In this case, the hiring once again does not cause the firm to increase its use of the inventor's prior ideas (even though the inventor's presence might facilitate that process). However, since the two are now concurrent, our difference-in-differences model is no longer able to distinguish this scenario where the relationship is not causal from a similar one where it is. Future research that identifies a suitable instrument or natural experiment may resolve this issue.

Second, we discuss managerial implications from our results. Given our finding that the majority of a hiring firm's increase in its use of a recruit's prior ideas is due to the inventive activity of the recruit herself, it is tempting to infer that firms could benefit from encouraging more knowledge sharing amongst employees, particularly with respect to new recruits who might bring especially novel knowledge into the firm. However, for this to be warranted, we must assume a divergence between the objective function of the firm and that of its employees. That is because although sharing knowledge creates benefits it also incurs costs. For example, such costs may arise from an increased need for coordination, information overload, and the opportunity cost of the inventors' time.

We assume that researchers engage in knowledge sharing up to the point where their private marginal benefit from doing so equals the private marginal cost. Therefore, the absence of knowledge sharing among particular individuals in a firm reflects that the benefits from doing so did not outweigh the costs from the perspective of at least one of the individuals involved. Firms should only intervene in the knowledge-sharing process if they have reason to believe that employees are under-investing in the sharing of knowledge among each other relative to what is best for the firm. This could be the case if employment contracts or the internal research culture are such that the objectives of the firm and the incentives of its researchers are not well aligned. More broadly, drawing precise

normative implications would also require further research on understanding the extent to which increased knowledge flow attributable to mobility represents true externalities received by the hiring firm versus being effectively getting paid for by the firm directly or indirectly (e.g., as greater wages or licensing fees).

Third, we discuss an intriguing absorptive capacity interpretation suggested by some of the other estimated coefficients. While the focus of existing research (including the present study) has largely been on how mobility contributes to improved access to knowledge from the previous firm, some of the findings we have not emphasized thus far are actually suggestive of a different kind of benefit that has been largely unexplored: hiring as a mechanism for improving a firm's "absorptive capacity" (Cohen and Levinthal 1989). While our difference-in-differences approach has focused on the hiring firm utilizing the recruits' prior ideas much more (255% more, on average; see Table 2, Column 4, top row) than the other technologically similar ideas we use as controls in our matched sample, additional results suggest that mobility might actually help the hiring firm increase its usage even of these other "control" inventions, i.e., of the technological area as a whole.

In particular, the "single difference" (pre-move versus post-move) is striking as the increase in the hiring firm's usage of technologically similar "control" ideas is in itself economically quite significant, 60% on average (as per the coefficient on *post\_move* in column (4) of Table 2). In a conservative interpretation in line with our intention for the original difference-in-differences approach, this might simply be the result of a broad strategic shift (unobservable to the researcher) in the recruiting firm's focus, resulting in a re-allocation of resources (beyond recruiting the mover in question) such that the increased use of ideas technologically similar to the inventor's expertise happens to coincide with the mobility event without being facilitated by the move.

However, it might arguably be too conservative to take the entire effect as being causally unrelated to the move. Instead, at least a part of this increase might reflect real benefits from mobility itself in the form of improved absorptive capacity for the technology in which the incoming inventor specializes. The new hire might bring not just firm-specific but also general expertise in the area to the recruiting firm, thus facilitating the firm's increased utilization of that technology field overall, above and beyond a rise in the exploitation of the mover's own prior ideas. In fact, comparing the

coefficient on *post\_move* in Column 1 with Column 2 in Table 4, we see that the increase in the hiring firm's use of the "control" idea is 60% when we consider all citations but only 33% when we omit citations from the mover, suggesting the mover herself might be responsible for a significant portion of that increase rather than the correlation just being coincidental.

Admittedly, an interpretation of the increase in citation rate to the broader technology (i.e., to both the actual and control patents) post-move is more speculative than our main difference-in-differences findings. For example, cautioning against the absorptive capacity interpretation is an observation that the effect is largely restricted to the first four-year period post-move (coefficient for *post\_move\_y1\_4* in Table 5 Column 1 is positive and significant, but coefficients for *post\_move\_y5\_8* and *post\_move\_yGE9* are statistically indistinguishable from zero), and that the upward trend in this actually starts even a few years before the move (coefficient for *pre\_move\_yGE6* in Table 5 Column 1 is negative and significant). Despite these caveats, the absorptive capacity hypothesis is intriguing and deserves further research.

The fourth main point in the discussion returns us back to our original "difference-in-differences" approach. Our focus in this point is on extending the original framework to examine the recruit's prior ideas more broadly defined. While the analysis reported in the tables uses the inventor's *own* prior inventions as a clean setting for examining inter-firm knowledge flows resulting from mobility, our difference-in-differences methodology is more generally applicable to examine any prior knowledge of which the mobile inventor could be a carrier.

For example, as emphasized in previous research, knowledge acquisition benefits from hiring could also include improved access to other knowledge originating from the mobile inventor's original firm. Although we do not report results here, we apply an extension of our framework to examine how mobility improves access to prior inventions originating from others in the mobile inventor's original firm, with the "focal patents" now being defined not as prior patents by mobile inventors but as patents originating from inventors in the original firm but still belonging to the same three-digit technology class. An analysis similar to Table 2 reveals that inter-firm mobility does help the recruiting firm enhance access even to prior knowledge generated by others in the recruit's source firm, albeit with a much smaller post-mobility boost of 58% (instead of the 197% boost reported in

Table 2). Further, examination of detailed timing of the citations reveal that knowledge from others is more prone to “forgetting”: the benefits from mobility are even more concentrated in the time period right after the move. More detailed examination of “same firm” knowledge as well as investigation of other kinds of prior knowledge that the inventor could possess (e.g., from her prior network ties or geographic context) would be natural extensions of the analysis presented in this paper.

Fifth, we describe a methodological contribution in terms of two useful insights concerning the construction of our matched-pair control sample. Recall that, comparing results between Tables 2 and 3, we show that our difference-in-differences framework is in fact quite robust to the underlying technological classification used for the analysis: there is very little difference in our coefficient estimates of interest between matching on three-digit versus nine-digit USPTO technology classifications. We also show that this is not the case for conventional cross-sectional estimates, which are very sensitive to the level of matching.

Introducing a methodology robust to technological classification is an important contribution, especially given the vociferous debate concerning the strengths and weaknesses of each: Three-digit matching might be too coarse to achieve the goal of precisely controlling for the underlying distribution of technological activity in a given field (Thompson and Fox-Kean 2005), while nine-digit matching might reduce the sample size considerably, thus raising concerns about sample selection bias (Henderson et al. 2005). We show that this issue is of considerable import to our data with respect to the estimates of first differences (e.g., comparing post-mobility citations for focal versus control patents), since the two levels of matching produce significantly different results in such a cross-sectional comparison. However, the difference between the two is mostly “differenced out” in our difference-in-differences model where we compare citations to focal versus control patents both before and after the mobility event, providing greater confidence in our estimates. In fact, recognizing that three-digit results are more similar to nine-digit results in terms of estimated coefficients on the interaction terms (the term of interest) when using the difference-in-differences method may enable research projects to proceed even with limited data, which would be infeasible if requiring nine-digit matching, to proceed using a three-digit matching.

As the sixth and last point in the discussion, we note that we have thus far focused on only one aspect of mobility – how it affects usage of knowledge that already existed pre-move, an issue for which our difference-in-differences research design is relatively well-suited. We have not examined other equally interesting questions, such as measuring changes in productivity of the mobile employee or of other employees benefiting from the resulting knowledge spillovers (Groysberg et al. 2008). Identification is even more accentuated for examining such questions (Lacetera et al. 2004), and our framework, with adjustments, might be suitable in such settings. Alternatively, the researcher could exploit a natural experiment that provides exogenous variation in mobility, such as that arising from closure of establishments (Dahl and Sorenson 2009), changes in non-compete laws (Marx et al. 2009), or death of close colleagues (Azoulay et al. 2009, Oettl 2009). Furthermore, future research on inter-firm mobility would benefit from stronger links to the rich literature on job matching (e.g., Jovanovic 1979, Simon and Warner 1992), explicitly incorporating the dynamic process wherein inter-firm mobility is a result of inventors and firms moving to improve their match, leading to a better understanding of the antecedents as well as consequences of mobility.

More broadly, any line of inquiry related to the antecedents and consequences of mobility will greatly benefit from further work that emphasizes investigating causal relationships and uncovering micro-level mechanisms of knowledge transfer. In this paper, we offer one approach for moving this topic forward. The relationship between inventor mobility and knowledge flow is economically important for both firm strategy and public policy, and we still have much to learn.

## REFERENCES

- Agarwal, R., R. Echambadi, A. Franco, MB Sarkar. 2004. Knowledge transfer through inheritance: spin-out generation, development, and survival. *Academy of Management Journal*. **47**(4) 501–522.
- Agrawal, A., I. Cockburn, J. McHale. 2006. Gone but not forgotten: Labor flows, knowledge spillovers, and enduring social capital. *Journal of Economic Geography*. **6**(5) 571-591.
- Agarwal, R., M. Ganco, R. Ziedonis. 2009. Reputations for toughness in patent enforcement: Implications for knowledge spillovers via inventor mobility. Forthcoming in *Strategic Management Journal*.
- Alcacer, J., M. Gittelman. 2006. Patent citations as a measure of knowledge flows: The influence of examiner citations. *Review of Economics and Statistics*. **88**(4) 774-779.
- Allison, P.D. and R.P. Waterman. 2002. *Sociological Methodology*. **32**: 247-265.
- Almeida, P., B. Kogut. 1999. The localization of knowledge and the mobility of engineers in regional networks. *Management Science*. Vol. **45**(7) 905-917.
- Argote, L., B. McEvily, R. Reagans. 2003. Introduction to the special issue on managing knowledge in organizations: Creating, retaining, and transferring knowledge. *Management Science*. **49**: v-viii.
- Arrow, K. J. 1962. Economic welfare and the allocation of resources for invention. In R. R. Nelson (ed.), *The Rate and Direction of Inventive Activity: Economic and Social Factors*. Vol. 13 of NBER Special Conference Series, Princeton University Press, New Jersey, pp. 609-625.
- Azoulay, P., J. S. Graff Zivin, J. Wang. 2009. Superstar Extinction. *Journal of Industrial Economics*. Forthcoming.
- Barney, J. 1991. Firm resources and sustained competitive advantage. *Journal of Management*. **17**(1) 99-120.
- Becker, G.S. 1962. Investment in human capital: A theoretical analysis. *Journal of Political Economy*. **70**(5) 9-49.
- Bhide, A. 2000. *The Origin and Evolution of New Business*. Oxford University Press: Oxford.
- Breschi, S. and F. Lissoni. 2009. Mobility of skilled workers and co-invention networks: an anatomy of localized knowledge flows. *Journal of Economic Geography* **9**(4) 439-468.
- Card, D., A. Krueger. 1994. Minimum wages and employment: A case study of the fast-food industry in New Jersey and Pennsylvania. *American Economic Review*. **84**(4) 772-793.
- Coff, R.W. 1997. Human assets and management dilemmas: Coping with hazards on the road to resource-based theory. *The Academy of Management Review*. **22**(2) 374-402.
- Cohen, W.M., D.A. Levinthal. 1989. Innovation and learning: The two faces of R&D. *Economic Journal*. **99**(397) 569-596.
- Corredoira, R.A., L. Rosenkopf. 2008. Should auld acquaintance be forgot? The reverse transfer of knowledge through mobility ties. Working Paper.
- Dahl, M.S., O. Sorenson. 2009. The migration of technical workers. *Journal of Urban Economics*. Forthcoming.
- DiMaggio, P., W.W. Powell. 1983. The iron cage revisited: Institutional isomorphism and collective rationality in organizational fields. *American Sociological Review*. **48**(2) 147-160.
- Dosi, G. 1988. Sources, procedures and microeconomic effects of innovation. *Journal of Economic Literature*. **26**(3) 1120-1171.

- Duguet, E., M. MacGarvie. 2005. How well do patent citations measure knowledge spillovers? Evidence from French innovation surveys. *Economics of Innovation and New Technology*. **14**(5) 375-393.
- Fleming, L., C. King, A. Juda, 2007a. Small worlds and innovation. *Organization Science*. **18**(6) 938-954.
- Fleming, L., S. Mingo, D. Chen. 2007b. Collaborative brokerage, generative creativity, and creative success. *Administrative Science Quarterly*. **52**(3) 443-475.
- Furman, J.L., S. Stern. 2006. Climbing atop the shoulders of giants: the impact of institutions on cumulative research. NBER Working Paper 12523.
- Gompers, P., J. Lerner, D. Scharfstein. 2005. Entrepreneurial spawning: public corporations and the genesis of new ventures, 1986 to 1999. *Journal of Finance*. **60**(2) 517-614.
- Grant, R.M. 1996. Toward a knowledge-based theory of the firm. *Strategic Management Journal*. **17**(2) 109-122.
- Groysberg, B., L.E. Lee, A. Nanda. 2008. Can they take it with them? The portability of star knowledge workers' performance. *Management Science*. 1-18.
- Hansen, M.T. 1999. The search-transfer problem: The role of weak ties in sharing knowledge across organization subunits. *Administrative Science Quarterly*. **44**(1) 82-111.
- Hausman, J., B. Hall, Z. Griliches. 1984. Econometric models for count data with an application to the patents-R&D relationship. *Econometrica*. **52**(4) 909-38.
- Henderson, R., A. Jaffe, M. Trajtenberg. 2005. Patent Citations and the Geography of Knowledge Spillovers: A Reassessment: Comment. *American Economic Review* **95**(1) 461-464.
- Jaffe, A.B., M. Trajtenberg, 2002. *Patents, Citations & Innovations: A window on the knowledge economy*. MIT Press, Cambridge.
- Jovanovic, B. 1979. Job matching and the theory of turnover. *Journal of Political Economy*. **87**(5)972-90.
- Klepper, S., S. Sleeper. 2005. Entry by spin-offs. *Management Science*. **51**(8) 1291-1306.
- Kogut, B., U. Zander. 1996. What firms do? Coordination, identity, and learning. *Organization Science*. **7**(5) 502-518.
- Lacetera, N., I.M. Cockburn, R.M. Henderson. 2004. Do firms change capabilities by hiring new people? A study of the adoption of science-based drug discovery. In A.M McGahan & J.A.C. Baum (Ed.), *Advances in Strategic Management*. Vol. 21. Greenwich CT: JAI Press.
- Lampe R. 2008. Strategic citation. SSRN Working Paper #984123.
- Lazear, E.P. 1986. Raids and offer matching. *Research in Labor Economics*. **8** 141-165.
- March, J.G. 1991. Exploration and exploitation in organizational learning. *Organization Science*. **2**(1) 71-87.
- Marx, M., D. Strumsky, L. Fleming. 2009. Mobility, skills, and the Michigan non-compete experiment. *Management Science*. **55**(6) 875-889.
- Moen, J. 2005. Is mobility of technical personnel a source of R&D spillovers? *Journal of Labor Economics* **23**(1) 81-114.
- Murray, F., S. Stern. 2007. Do formal intellectual property rights hinder the free flow of scientific knowledge? An empirical test of the anti-commons hypothesis. *Journal of Economic Behavior and Organization*. **63**(4) 648-687.
- Nelson, R.R., S.G. Winter. 1982. An evolutionary theory of economic change. Cambridge, MA: Belknap.

- Oettl, A., A. Agrawal. 2008. International labor mobility and knowledge flow externalities. *Journal of International Business Studies*.39 1242-1260.
- Oettl, A. 2009. Productivity and Helpfulness: A New Taxonomy for Star Scientists. Mimeo, Georgia Institute of Technology
- Penrose, E.T. 1980. The theory of the growth of the firm. Oxford, England: Blackwell.
- Peteraf, M.A. 1993. The cornerstones of competitive advantage: A resource-based view. *Strategic Management Journal*. **14**(3) 179-191.
- Reagans, R., E. Zuckerman. 2001. Networks, diversity and performance: The social capital of R&D units. *Organization Science*. **12**(4) 502-517.
- Rosenkopf, L., P. Almeida. 2003. Overcoming local search through alliances and mobility. *Management Science*. **49**(6) 0751-0766.
- Saxenian, A.L. 1994. *Regional Advantage: Culture and Competition in Silicon Valley and Route 128*. Harvard University Press, Cambridge, MA.
- Simon, H.A. 1991. Bounded rationality and organizational learning. *Organization Science*. **2**(1) 125-134.
- Simon, C.J., J. Warner. 1992. Matchmaker, matchmaker: The effect of old boy networks on job match quality, earnings, and tenure. *Journal of Labor Economics*. **10**(3) 306-330.
- Singh, J. 2005. Collaborative networks as determinants of knowledge diffusion patterns. *Management Science*. **51**(5) 756-770.
- Singh, J. 2007. Asymmetry of knowledge spillovers between MNCs and host country firms. *Journal of International Business Studies*. **38**(5)764-786.
- Singh, J. 2008. Distributed R&D, cross-regional knowledge integration and quality of innovative output. *Research Policy*. **37**(1) 77-96.
- 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.
- Stinchcombe, A.L. 1965. *Organizations and social structure*. In J.G. March (Ed.), *Handbook of organizations*: 142-193. Chicago: Rand McNally.
- Thompson, P., M. Fox-Kean. 2005. Patent citations and the geography of knowledge spillovers: A reassessment. *American Economic Review*. **95**(1) 450-460.
- Trajtenberg, M. 2006. The "names game": harnessing inventors' patent data for economic research. NBER Working Paper 12479.
- Tzabbar, D., B.S. Silverman, B.S. Aharonson. 2009. Learning by hiring or hiring to avoid learning? Organizational Exploration via Individuals' Exploitation. Mimeo.
- Weitzman, M. 1998. Recombinant growth. *The Quarterly Journal of Economics* CXIII(2) 331-360.
- Wooldridge, J. M. 1999. Distribution-free estimation of some nonlinear panel data models. *Journal of Econometrics*. **90**(1) 77-97.

**Table 1a. Variables and summary statistics for the “original” patent-year sample (N = 406,632)**

		Mean	Std Dev	Min	Max
<b>Dependent variables</b>					
cites_all	All citations to the focal (original or control) patent from the destination firm in the given year	0.0164	0.18	0	18
cites_nonself	Same as <i>cites_all</i> , except that self-cites from an inventor are excluded	0.0117	0.15	0	14
cites_noncollab	Same as <i>cites_nonself</i> , except that citations from the mobile inventor's collaborators in the destination firm are also excluded	0.0099	0.13	0	13
cites_nonnetwork	Same as <i>cites_noncollab</i> , except that citations from the mobile inventor's collaborators' collaborators in the destination firm are also excluded	0.0089	0.12	0	13
<b>Explanatory &amp; control variables</b>					
focal_patent	Indicator that is 1 for original patent and 0 for its control patent	0.50	0.50	0	1
post_move	Indicator that is 1 if and only if citing year is at least one year after the estimated move date for the original patent's inventor	0.57	0.49	0	1
move_window	Indicator that is 1 for the two year window surrounding the estimated move date	0.17	0.37	0	1
destfirm_pats	Number of patents assigned to the destination firm in the last 5 years	676.54	1442.97	0	16983
origfirm_pats	Number of patents assigned to the original firm in the last 5 years	1109.01	1193.24	0	5504
claims	Number of claims made by the patent	11.82	9.59	1	171
patrefs	Number of backward citations that the patent makes to other patents	6.87	6.31	0	114
nonpatrefs	Number of non-patent references made by the patent	0.95	2.68	0	100
inventor_age	Number of years since the first patent by the inventor	4.79	4.56	0	30
inventor_patents	Number of previous patents by the inventor	8.02	15.04	0	256
US_inventor	Indicator that is 1 if and only if inventor has a US address	0.61	0.49	0	1

**Table 1b. Summary statistics for dependent variable *cites\_all* in different subsamples**

	<i>post_move</i>		Total
	0	1	
<b><i>focal_patent</i></b>			
0	mean = 0.0025 sd = 0.0703 N = 87265	mean = 0.0041 sd = 0.0808 N = 116051	mean = 0.0034 sd = 0.0764 N = 203316
1	mean = 0.0106 sd = 0.1407 N = 87265	mean = 0.0433 sd = 0.3012 N = 116051	mean = 0.0293 sd = 0.246 N = 203316
Total	mean = 0.0065 sd = 0.1113 N = 174530	mean = 0.0237 sd = 0.2214 N = 232102	mean = 0.0163 sd = 0.1826 N = 406632

**Table 2. Analysis of citations made by the destination firm to prior patents of the mobile inventor**

This analysis is based on the “original” sample constructed as illustrated in Figures 1 and 3. We start with a sample of “focal” patents whose inventor subsequently moved to another firm. We then match each focal patent with a “control” patent originating in a different assignee but having the same application year and three-digit primary technology class as the focal patent, while also ensuring that it has a different inventor and one who did not subsequently move to the same destination firm. We drop any observations for which a match cannot be found. We then construct our patent-year sample by creating 12 yearly observations for each patent, creating a dataset with the annual citations received from the destination firm in the 12 years directly following its application year.

	(1)	(2)	(3)	(4)	(5)
<i>Dependent Variable</i>	<i>cites_all</i>	<i>cites_all</i>	<i>cites_all</i>	<i>cites_all</i>	<i>cites_all</i>
<i>Regression Model</i>	NBREG	NBREG	NBREG	NBREG	XTQMLP (FE)
<i>Patent-Year Sample</i>	3-digit match (Post move)	3-digit match (Pre move)	3-digit match (All)	3-digit match (All)	3-digit match (FE sample)
focal_patent	2.334*** (0.086) [10.31]	1.351*** (0.16) [3.86]	1.268*** (0.15) [3.55]	1.269*** (0.17) [3.55]	
focal_patentXpost_move			1.093*** (0.17) [2.98]	1.091*** (0.18) [2.97]	1.093*** (0.21) [2.98]
focal_patentXmove_window			0.329 (0.21) [1.38]	0.328 (0.22) [1.38]	0.426* (0.26) [1.53]
focal_patentXpre_move_yGE6				-0.0291 (0.37) [0.97]	-0.0686 (0.55) [0.93]
post_move			0.577*** (0.16) [1.78]	0.473*** (0.17) [1.6]	1.269*** (0.21) [3.55]
move_window			0.471** (0.20) [1.6]	0.352* (0.20) [1.42]	0.667*** (0.24) [1.94]
pre_move_yGE6				-0.716** (0.31) [0.48]	-1.446*** (0.49) [0.23]
ln_destfirm_pats	0.180*** (0.011)	0.440*** (0.022)	0.221*** (0.0095)	0.221*** (0.0095)	0.363*** (0.029)
ln_origfirm_pats	-0.0249 (0.017)	-0.121*** (0.040)	-0.0402** (0.017)	-0.0397** (0.017)	
ln_claims	0.186*** (0.051)	0.158 (0.099)	0.196*** (0.047)	0.196*** (0.047)	
ln_patrefs	0.158*** (0.042)	0.283*** (0.11)	0.191*** (0.039)	0.191*** (0.039)	
ln_nonpatrefs	0.0157 (0.045)	-0.171 (0.11)	-0.0314 (0.042)	-0.0315 (0.042)	
ln_inventor_age	0.193*** (0.047)	-0.116 (0.10)	0.120*** (0.046)	0.120*** (0.046)	
ln_inventor_patents	-0.215*** (0.039)	-0.0448 (0.10)	-0.163*** (0.042)	-0.163*** (0.042)	
US_inventor	0.825*** (0.070)	0.241 (0.16)	0.708*** (0.064)	0.707*** (0.064)	
Observations	232102	106758	406632	406632	34368
Number of unique patents	33886	33886	33886	33886	2864
chi2	1656	2575	3249	3224	1412
df_m	29	29	35	37	24
ll	-21115	-2401	-26935	-26921	-11406

Indicator variables for citation lag, citing year and the 2-digit technology category were included but are not shown here  
Robust and clustered standard errors in parentheses (\*\* p<0.01, \* p<0.05, . p<0.1)  
Incidence rate ratios in square brackets

**Table 3. Re-examining the effect of mobility using nine-digit technology matching instead**

We constructed the patent-year sample employed in this table in a manner analogous to the “original” sample employed in Table 2 (and illustrated in Figures 1 and 3). The only difference was that the matching procedure now required the “focal” and “control” patent to have the same nine-digit technology subclass rather than just the same three-digit technology class. This more stringent matching criterion led to a greater number of focal patents being dropped from the analysis since a nine-digit technology match for these could not be found.

	(1)	(2)	(3)	(4)	(5)
<i>Dependent Variable</i>	<i>cites_all</i>	<i>cites_all</i>	<i>cites_all</i>	<i>cites_all</i>	<i>cites_all</i>
<i>Regression Model</i>	NBREG	NBREG	NBREG	NBREG	XTQMLP (FE)
<i>Patent-Year Sample</i>	Original (Post move)	Original (Pre move)	Original (All)	Original (All)	Original (FE sample)
focal_patent	1.424*** (0.098) [4.15]	0.713*** (0.20) [2.04]	0.618*** (0.20) [1.85]	0.646*** (0.21) [1.9]	
focal_patentXpost_move			0.843*** (0.21) [2.32]	0.815*** (0.22) [2.25]	0.958*** (0.23) [2.6]
focal_patentXmove_window			0.0963 (0.21) [1.1]	0.0682 (0.22) [1.07]	0.236 (0.23) [1.26]
focal_patentXpre_move_yGE6				-0.181 (0.42) [0.83]	-0.569 (0.48) [0.56]
post_move			0.613*** (0.18) [1.84]	0.550*** (0.19) [1.73]	1.328*** (0.20) [3.77]
move_window			0.621*** (0.17) [1.86]	0.545*** (0.17) [1.72]	0.840*** (0.18) [2.31]
pre_move_yGE6				-0.402 (0.30) [0.66]	-0.958*** (0.32) [0.38]
ln_destfirm_pats	0.193*** (0.016)	0.482*** (0.031)	0.248*** (0.014)	0.248*** (0.014)	0.418*** (0.043)
ln_origfirm_pats	-0.00573 (0.026)	-0.114** (0.052)	-0.0138 (0.025)	-0.0136 (0.025)	
ln_claims	0.209*** (0.061)	0.287** (0.13)	0.234*** (0.060)	0.234*** (0.060)	
ln_patrefs	0.247*** (0.063)	0.413*** (0.13)	0.268*** (0.060)	0.269*** (0.060)	
ln_nonpatrefs	0.0973 (0.062)	-0.199 (0.18)	0.0255 (0.065)	0.0249 (0.065)	
ln_inventor_age	0.198*** (0.076)	-0.0425 (0.14)	0.175** (0.070)	0.176** (0.070)	
ln_inventor_patents	-0.230*** (0.065)	-0.0921 (0.13)	-0.207*** (0.060)	-0.207*** (0.060)	
US_inventor	0.787*** (0.10)	0.228 (0.23)	0.646*** (0.098)	0.644*** (0.098)	
Observations	87500	40500	153600	153600	15288
Number of unique patents	12800	12800	12800	12800	1274
chi2	707.4	2068	1360	1355	662.7
df_m	29	29	35	37	24
ll	-8836	-1316	-11951	-11947	-5051

Indicator variables for citation lag, citing year and the 2-digit technology category were included but are not shown here  
Robust and clustered standard errors in parentheses (\*\*\*) p<0.01, \*\* p<0.05, \* p<0.1)  
Incidence rate ratios in square brackets

**Table 4. Unpacking the citations received by the focal patent from the destination firm**

We base this analysis on the original sample already employed in Table 2.

	(1)	(2)	(3)	(4)
<i>Dependent Variable</i>	<i>cites_all</i>	<i>cites_nonself</i>	<i>cites_noncollab</i>	<i>cites_nonnetwork</i>
<i>Regression Model</i>	NBREG	NBREG	NBREG	NBREG
<i>Patent-Year Sample</i>	Original	Original	Original	Original
focal_patent	1.269*** (0.17) [3.55]	1.310*** (0.17) [3.7]	1.285*** (0.17) [3.61]	1.277*** (0.17) [3.58]
focal_patentXpost_move	1.091*** (0.18) [2.97]	0.594*** (0.18) [1.81]	0.393** (0.19) [1.48]	0.264 (0.19) [1.3]
focal_patentXmove_window	0.328 (0.22) [1.38]	0.341 (0.22) [1.4]	0.248 (0.28) [1.28]	0.177 (0.22) [1.19]
focal_patentXpre_move_yGE6	-0.0291 (0.37) [0.97]	-0.0466 (0.37) [0.95]	-0.0162 (0.37) [0.98]	-0.0330 (0.38) [0.96]
post_move	0.473*** (0.17) [1.6]	0.290* (0.17) [1.33]	0.269 (0.17) [1.3]	0.279 (0.17) [1.32]
move_window	0.352* (0.20) [1.42]	0.332* (0.20) [1.39]	0.330 (0.20) [1.39]	0.336* (0.20) [1.39]
pre_move_yGE6	-0.716** (0.31) [0.48]	-0.661** (0.31) [0.51]	-0.645** (0.31) [0.52]	-0.655** (0.31) [0.51]
ln_destfirm_pats	0.221*** (0.0095)	0.312*** (0.011)	0.334*** (0.011)	0.339*** (0.011)
ln_origfirm_pats	-0.0397** (0.017)	-0.0485** (0.019)	-0.0506** (0.020)	-0.0551*** (0.020)
ln_claims	0.196*** (0.047)	0.160*** (0.054)	0.166*** (0.050)	0.177*** (0.052)
ln_patrefs	0.191*** (0.039)	0.257*** (0.045)	0.290*** (0.047)	0.279*** (0.049)
ln_nonpatrefs	-0.0315 (0.042)	-0.143*** (0.050)	-0.139*** (0.049)	-0.150*** (0.051)
ln_inventor_age	0.120*** (0.046)	0.103* (0.055)	0.0887 (0.058)	0.110* (0.061)
ln_inventor_patents	-0.163*** (0.042)	-0.154*** (0.050)	-0.162*** (0.054)	-0.191*** (0.056)
US_inventor	0.707*** (0.064)	0.652*** (0.072)	0.589*** (0.074)	0.540*** (0.076)
Observations	406632	406632	406632	406632
Number of unique patents	33886	33886	33886	33886
chi2	3224	2365	2283	2098
df_m	37	37	37	37
ll	-26921	-20850	-18507	-17210

Indicator variables for citation lag, citing year and the 2-digit technology category were included but are not shown here  
 Robust and clustered standard errors in parentheses (\*\*\*) p<0.01, \*\* p<0.05, \* p<0.1  
 Incidence rate ratios in square brackets

**Table 5. A comparison of temporal patterns of citations**

We base this analysis on the original sample already employed in Tables 2 and 4. Figure 4 graphically illustrates the findings from this table.

	(1)	(2)	(3)	(4)
<i>Dependent Variable</i>	<i>cites_all</i>	<i>cites_nonsel</i>	<i>cites_noncollab</i>	<i>cites_nonnetwork</i>
<i>Regression Model</i>	NBREG	NBREG	NBREG	NBREG
<i>Patent-Year Sample</i>	Original	Original	Original	Original
focal_patent	1.267*** (0.17) [3.55]	1.306*** (0.17) [3.69]	1.282*** (0.17) [3.6]	1.274*** (0.17) [3.57]
focal_patentXpost_move_y1_4	1.209*** (0.18) [3.35]	0.613*** (0.19) [1.84]	0.379** (0.19) [1.46]	0.235 (0.19) [1.26]
focal_patentXpost_move_y5_8	0.861*** (0.22) [2.36]	0.533** (0.23) [1.7]	0.367 (0.23) [1.44]	0.263 (0.23) [1.3]
focal_patentXpost_move_yGE9	0.895*** (0.32) [2.44]	0.701** (0.33) [2.01]	0.683** (0.33) [1.97]	0.599* (0.33) [1.82]
focal_patentXmove_window	0.327 (0.22) [1.38]	0.343 (0.22) [1.4]	0.250 (0.22) [1.28]	0.179 (0.22) [1.19]
focal_patentXpre_move_yGE6	-0.0258 (0.37) [0.97]	-0.0430 (0.37) [0.95]	-0.0131 (0.37) [0.98]	-0.0300 (0.38) [0.97]
post_move_y1_4	0.461*** (0.17) [1.58]	0.364** (0.17) [1.43]	0.360** (0.17) [1.43]	0.376** (0.17) [1.45]
post_move_y5_8	0.0267 (0.21) [1.02]	-0.159 (0.22) [0.85]	-0.144 (0.22) [0.86]	-0.119 (0.22) [0.88]
post_move_yGE9	-0.147 (0.32) [0.86]	-0.408 (0.32) [0.66]	-0.383 (0.33) [0.68]	-0.358 (0.33) [0.69]
move_window	0.310 (0.20) [1.36]	0.295 (0.20) [1.34]	0.299 (0.20) [1.34]	0.307 (0.20) [1.35]
pre_move_yGE6	-0.705** (0.31) [0.49]	-0.655** (0.31) [0.51]	-0.639** (0.31) [0.52]	-0.649** (0.31) [0.52]
ln_destfirm_pats	0.226*** (0.0095)	0.314*** (0.011)	0.337*** (0.011)	0.341*** (0.011)
ln_origfirm_pats	-0.0468*** (0.016)	-0.0526*** (0.019)	-0.0545*** (0.020)	-0.0587*** (0.020)
ln_claims	0.197*** (0.047)	0.161*** (0.054)	0.165*** (0.050)	0.176*** (0.052)
ln_patrefs	0.197*** (0.039)	0.260*** (0.045)	0.291*** (0.047)	0.280*** (0.049)
ln_nonpatrefs	-0.0353 (0.042)	-0.145*** (0.049)	-0.141*** (0.049)	-0.153*** (0.051)
ln_inventor_age	0.130*** (0.047)	0.111** (0.056)	0.0939 (0.058)	0.113* (0.061)
ln_inventor_patents	-0.163*** (0.042)	-0.152*** (0.050)	-0.159*** (0.054)	-0.189*** (0.056)
US_inventor	0.715*** (0.064)	0.658*** (0.072)	0.596*** (0.074)	0.546*** (0.076)
Observations	406632	406632	406632	406632
Number of unique patents	33886	33886	33886	33886
chi2	3971	2555	2397	2188
df_m	41	41	41	41
ll	-26772	-20783	-18465	-17176

Indicator variables for citation lag, citing year and the 2-digit technology category were included but are not shown here  
 Robust and clustered standard errors in parentheses (\*\*\*) p<0.01, \*\* p<0.05, \* p<0.1)  
 Incidence rate ratios in square brackets

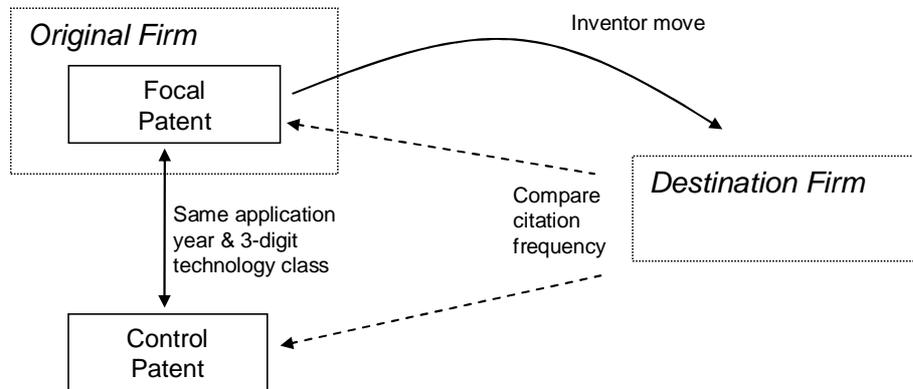
**Table 6. Moderation effect of unfamiliarity of the inventor's technology for the destination firm**

We base this analysis on the original sample already employed in Tables 2, 4, and 5.

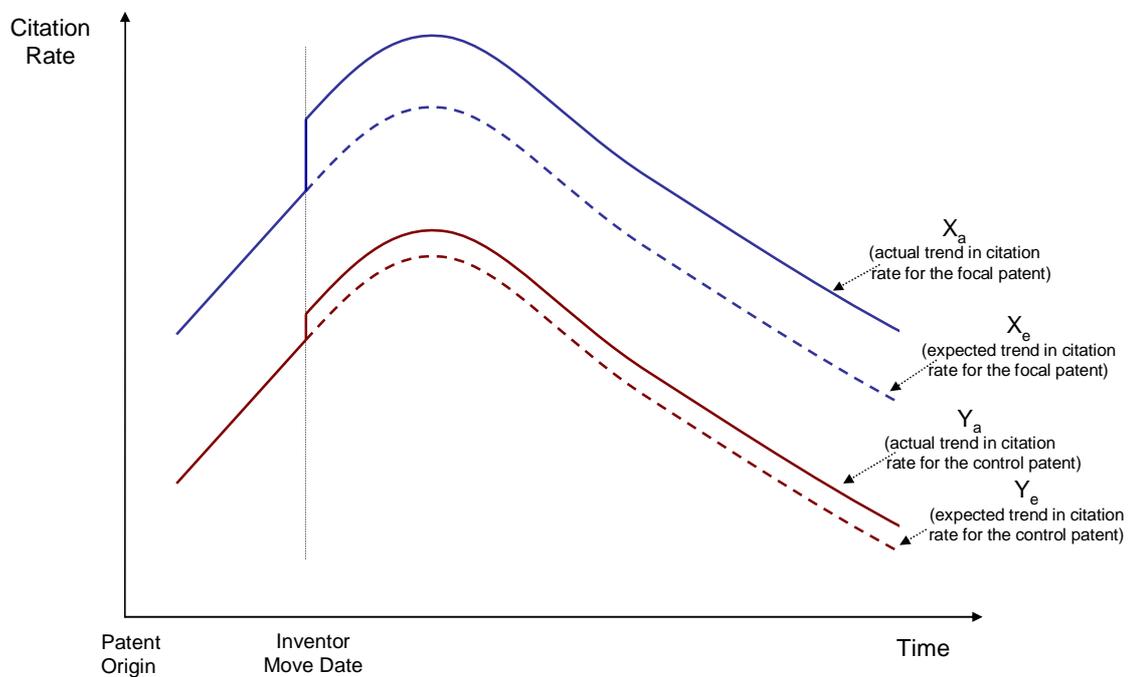
	(1)	(2)	(3)	(4)
<i>Dependent Variable</i>	<i>cites_all</i>	<i>cites_nonself</i>	<i>cites_noncollab</i>	<i>cites_nonnetwork</i>
<i>Regression Model</i>	NBREG	NBREG	NBREG	NBREG
<i>Patent-Year Sample</i>	Original	Original	Original	Original
unfamiliar_techXfocal_patent	-0.267 (0.35) [0.76]	-0.294 (0.35) [0.74]	-0.272 (0.35) [0.76]	-0.268 (0.35) [0.76]
unfamiliar_techXfocal_patentXpost_move	1.409*** (0.39) [4.09]	1.174*** (0.39) [3.23]	1.013*** (0.39) [2.75]	0.910** (0.39) [2.48]
unfamiliar_techXfocal_patentXmove_window	-0.417 (0.45) [0.65]	-0.384 (0.45) [0.68]	-0.517 (0.45) [0.59]	-0.527 (0.45) [0.59]
unfamiliar_techXfocal_patentXpre_move_yGE6	-0.726 (0.75) [0.48]	-0.656 (0.76) [0.51]	-0.673 (0.76) [0.51]	-0.629 (0.76) [0.53]
unfamiliar_techXpost_move	-0.0242 (0.34) [0.97]	0.0150 (0.34) [1.01]	0.0237 (0.34) [1.02]	0.0166 (0.34) [1.01]
unfamiliar_techXmove_window	0.546 (0.41) [1.72]	0.541 (0.41) [1.71]	0.538 (0.41) [1.71]	0.529 (0.41) [1.69]
unfamiliar_techXpre_move_yGE6	0.791 (0.64) [2.2]	0.732 (0.64) [2.07]	0.719 (0.64) [2.05]	0.710 (0.64) [2.03]
unfamiliar_tech	-1.468*** (0.30) [0.23]	-1.474*** (0.30) [0.22]	-1.489*** (0.30) [0.22]	-1.490*** (0.30) [0.22]
focal_patent	1.357*** (0.20) [3.88]	1.408*** (0.21) [4.08]	1.373*** (0.21) [3.94]	1.363*** (0.21) [3.9]
focal_patentXpost_move	0.391* (0.23) [1.47]	0.0557 (0.23) [1.05]	-0.0531 (0.23) [0.94]	-0.126 (0.23) [0.88]
focal_patentXmove_window	0.527** (0.22) [1.69]	0.518** (0.23) [1.67]	0.470** (0.23) [1.59]	0.401* (0.23) [1.49]
focal_patentXpre_move_yGE6	0.340 (0.55) [1.4]	0.276 (0.56) [1.31]	0.311 (0.56) [1.36]	0.275 (0.57) [1.31]
post_move	0.491** (0.21) [1.63]	0.266 (0.21) [1.3]	0.236 (0.21) [1.26]	0.249 (0.21) [1.28]
move_window	0.0974 (0.19) [1.1]	0.0733 (0.20) [1.07]	0.0706 (0.20) [1.07]	0.0810 (0.19) [1.08]
pre_move_yGE6	-1.101** (0.48) [0.33]	-0.997** (0.48) [0.36]	-0.969** (0.49) [0.37]	-0.974** (0.49) [0.37]
ln_destfirm_pats	0.202*** (0.0100)	0.296*** (0.011)	0.318*** (0.012)	0.322*** (0.012)
ln_origfirm_pats	-0.0357** (0.017)	-0.0444** (0.019)	-0.0443** (0.020)	-0.0471** (0.020)
ln_claims	0.184*** (0.047)	0.145*** (0.055)	0.151*** (0.048)	0.162*** (0.050)
ln_patrefs	0.180*** (0.038)	0.242*** (0.044)	0.272*** (0.046)	0.258*** (0.047)
ln_nonpatrefs	-0.0245 (0.042)	-0.138*** (0.051)	-0.132*** (0.049)	-0.147*** (0.051)
ln_inventor_age	0.131*** (0.047)	0.120** (0.056)	0.101* (0.059)	0.126** (0.060)
ln_inventor_patents	-0.172*** (0.042)	-0.168*** (0.050)	-0.175*** (0.055)	-0.210*** (0.056)
US_inventor	0.692*** (0.063)	0.635*** (0.071)	0.564*** (0.072)	0.506*** (0.073)
Observations	406632	406632	406632	406632
Number of unique patents	33886	33886	33886	33886
chi2	3427	2696	2664	2480
df_m	45	45	45	45
ll	-26519	-20449	-18083	-16778

Indicator variables for citation lag, citing year and the 2-digit technology category were included but are not shown here  
 Robust and clustered standard errors in parentheses (\*\*\*) p<0.01, \*\* p<0.05, \* p<0.1  
 Incidence rate ratios in square brackets

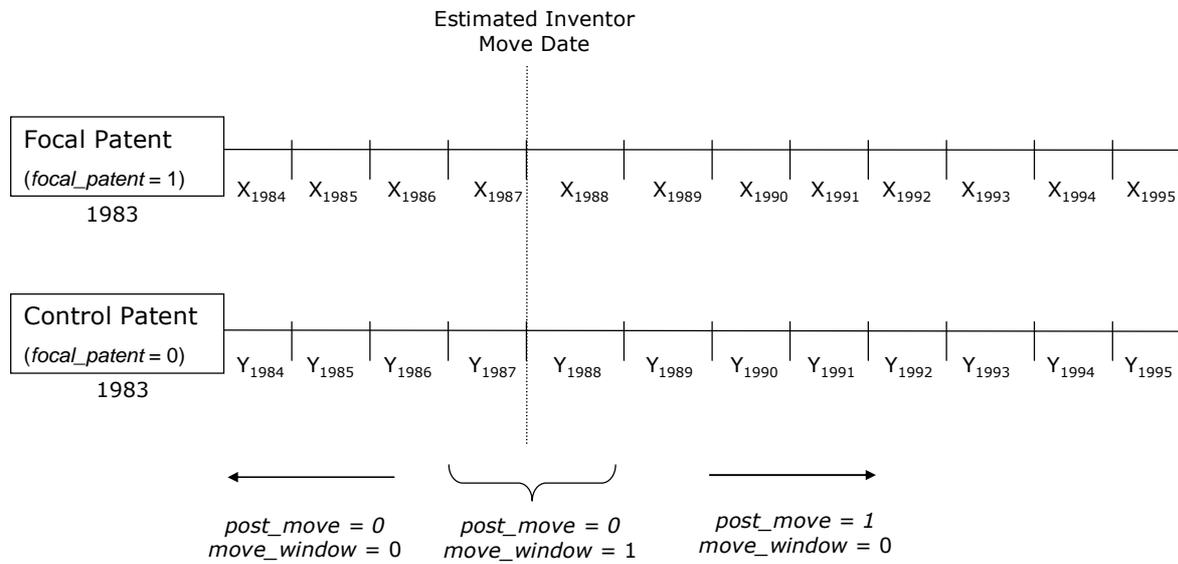
**Figure 1. Constructing a matched sample to examine use of a new recruit’s prior knowledge**



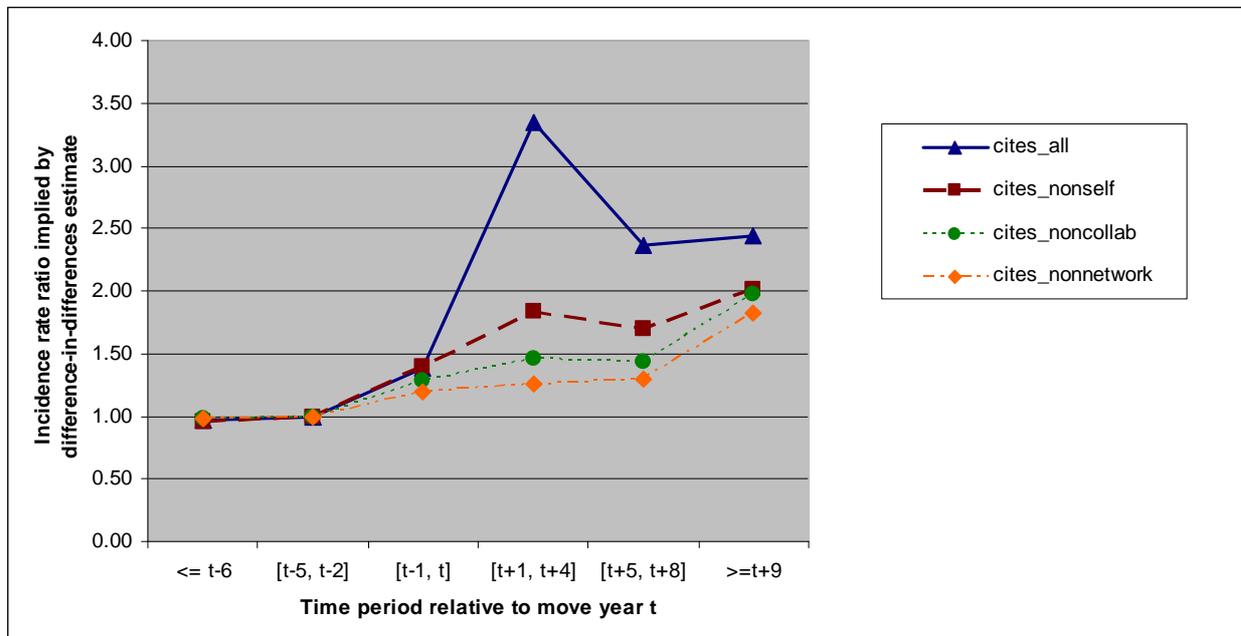
**Figure 2. A “difference-in-differences” approach to examine use of a new recruit’s prior knowledge**



**Figure 3. Constructing a longitudinal dataset of citations received by the focal and control patents**



**Figure 4. Estimated temporal trends in use of a new recruit's prior knowledge**



**Europe Campus**

**Boulevard de Constance**

**77305 Fontainebleau Cedex, France**

**Tel: +33 (0)1 60 72 40 00**

**Fax: +33 (0)1 60 74 55 00/01**

**Asia Campus**

**1 Ayer Rajah Avenue, Singapore 138676**

**Tel: +65 67 99 53 88**

**Fax: +65 67 99 53 99**

**[www.insead.edu](http://www.insead.edu)**

Printed by INSEAD

**INSEAD**



**The Business School  
for the World®**