The Effect of Inventor Mobility on Network Productivity

Brit Sharoni*
University of Michigan
Ross School of Business

February 11, 2026

Abstract

When inventors move to new locations, they carry knowledge and expertise, which may be a loss to their former collaborators. But they might also serve as a bridge between otherwise disconnected innovation hubs, facilitating information flows. This paper studies the effect of an inventor's relocation on their former collaborators' productivity. A simple patent production model formalizes relocators' dual role as former collaborators and information intermediaries and guides the empirical analysis. Using a novel dataset linking USPTO records to online employment histories, I find sizable positive effects on collaborators' productivity, primarily driven by greater access to novel information and new networks.

*Email: bsharoni@umich.edu. I am deeply grateful to my advisors Zoë Cullen, Claudia Goldin, Oliver Hart and Josh Lerner for invaluable guidance and support. I would also like to thank Ed Glaeser, Larry Katz and Florian Trouvain for their helpful suggestions and Dev Patel, Marta Prato, Jesse Shapiro, Stefanie Stancheva and seminar participants at Harvard University and the NBER SI Innovation 2024 Meeting for valuable comments and discussions. I also thank Ben Zweig from Revelio Labs for help with the data. This research benefited from financial support from LEAP (The Lab for Economic Applications and Policy) at Harvard University.

1 Introduction

Inventors' mobility matters for innovation. When inventors relocate to innovation clusters they experience an increase to their patenting productivity (Moretti, 2021).[1] This effect can be attributed to several factors, including agglomeration externalities, knowledge spillovers, and the collaborative nature of innovation, where patents are often the result of a teamwork the diverse skills and knowledge of its members (Jones, 2009).

While a large literature has focused on the effect of a reallocation of an inventor on their own patenting productivity, surprisingly little is know about the effects on the productivity of the movers' colleagues, who remain in the origin location. Given that most patenting activity involves team efforts, evidenced by more than 70 percent of patents filed with the United States Patents and Trademark Office in 2022 having at least two inventors listed, it seems reasonable to conclude that the reallocation of a colleague has an impact on the productivity of the inventors who stayed in the original location.[2] And this effect should be as important as the productivity effects related to the mover's own patenting activity. In this paper, I focus on the productivity of inventors who remained in the original location after their collaborator relocated and quantify the effect of such relocations on the productivity of the inventor's former co-inventors, both theoretically and empirically.

Ex-ante, it is unclear whether a colleague's departure positively or negatively impacts the productivity of their former collaborators. On the one hand, when inventors relocate they take their skills and expertise with them. So, in a world where the inventors who remained in the origin location lose access to all of their productive colleagues, the effect is expected to be negative, as was shown by Azoulay, Graff Zivin and Wang (2010) and Jaravel, Petkova and Bell (2018) who use death shocks as a source of an exogenous variation. However, on the other hand, in a world where their former collaborators introduce them to new techniques and ideas, the effect becomes less obvious and more complex to assess. In this scenario, it is unclear whether the inventors who remain in the original location are better off, even if they stop co-patenting with the mover entirely, simply because they have access to information from inventors in the destination location. Understanding which of these scenarios applies is crucial for innovation.

I develop this argument formally in a simple model of team production and patent cre-

[1]See also Ellison and Glaeser (1997) and Bloom, Hassan, Kalyani, Lerner and Tahoun (2021) who highlighted the existence of innovation clusters. And Ellison and Glaeser (1999), Ellison, Glaeser and Kerr (2010) and Greenstone, Hornbeck and Moretti (2010) who document the advantageous outcomes innovation clusters offer both inventors and firms deciding to relocate there.

[2]Rising collaboration rates appear in the ascending pattern observed in both the average count of inventors listed on a patent and the proportion of patents generated through collaborative efforts, as demonstrated in Jones (2009). It is part of a longer trend that reflects a shift towards collaborative knowledge creation, facilitated by information sharing (Wuchty, Jones and Uzzi, 2007).

ation. I build on recent advances in network theory to illustrate the distinct roles information sharing and collaboration play after a collaborator's relocation. In the model, the net effect of the move depends on the network characteristics, such as the intensity of the collaboration between inventors and the extent to which a colleague's move offers access to a new network and novel information. While the answer to this question is empirical in nature, the model helps to fix ideas, to discuss potential mechanisms, and to guide the empirical tests.

A significant obstacle in using relocations as a source of variation is the issue of endogenity. Relocations are influenced by various observed and unobserved factors, with both workers and firms determining the timing and destination of the move. This raises the concern of that the inventors who experience the relocation of one of their collaborators are selected. In this case, the inventors who stay in the original location when their collaborator relocates may differ significantly from those who do not experience such relocations, and the estimated results might reflect these differences rather than the true effect.

To overcome this challenge, I build on the methodology introduced in Jaravel et al. (2018). Specifically, I create a control group of inventors who have never experienced a relocation of a collaborator but have collaborated with an inventor who is similar to a mover at the time of the move. I begin by creating similar groups of movers and non-movers, based on various observable characteristics using an exact matching procedure. I then identify all inventors who have collaborated with movers and non-movers in these groups. This approach allows me to compare the treatment group (co-inventors of movers) to the control group (co-inventors of matched non-movers) using a difference-in-difference research design. This procedure helps mitigate endogeneity concern, by testing for the parallel trends assumption, enabling a more robust estimation of the effect of co-inventor relocations on the productivity of those who stay.

An additional hurdle is related to the structure of the patent data. Although the USPTO dataset allows to follow inventors over time, observations are tied to patenting activity, meaning that information on inventors, such as location, is only observed when they apply for patents. Since most inventors do not patent every year, it becomes impossible to locate inventors with certainty during the time intervals between patents. This means that using the USPTO dataset alone does not allow for identifying the exact time of a move or the movers who stop patenting after relocating.

To surmount this challenge, I build a novel dataset by merging information about inventors in the United States with their online professional profiles. This enables me to create a comprehensive longitudinal record of inventors based in the United States and, in turn, allows me to track their locations over time. Consequently, I can identify the movers and the timing of their moves, gaining valuable insights into the characteristics of the move. These specifics, including details about the firms where the movers are employed and the geographical distances involved, enable a comprehensive exploration of the impacts of co-inventor relocations on innovation networks and productivity, as well as the underlying factors shaping these effects.

The procedure I employ results in a longitudinal data on approximately 300,000 inventors based in the United States between 1990 and 2022. I identify 49,902 inventors who relocated during this period. Following a co-inventor's move, I find that the inventors who remain in the original location increase their annual number of patent applications by an average of about 9 percent, when compared to the control group. This overall effect becomes evident over approximately three years and weakly increases over time. Notably, this increase in the quantity of patents does not come at the expense of quality, and I estimate an increase of almost 15 percent in the average number of adjusted forward citations, compared to the control group.[3]

I then explore the specific mechanism that drives the estimated positive effects. I first address a concern related to common team productivity shocks, which could bias the results upwards and overstate the positive effect of the move. If the timing of the relocation coincides with periods of heightened patenting success, it is possible that inventors remaining in the original location, who are also contributors to these successful patents, may benefit from this success in different ways—such as through research grants or improved working conditions—that impact their productivity. I show that the size of the effect changes in ways that do not align with the common shock reasoning, thereby rejecting this explanation.

Next, I present evidence demonstrating that information from the destination location of the mover flows to inventors who remained in the original location. I show that following the relocation, inventors who stayed in the original location begin to cite patents from the mover's destination more frequently than they did before the move. Additionally, I provide evidence showing that, after the relocation, inventors who stayed in the original location expand their co-inventors network into the mover's destination location, relative to the control group. The findings imply that some information flows from the mover's destination to their former collaborators after the move. This can be the results of an exposure to patents produced in that location or through a direct collaboration with inventors who are located there.

Finally, to establish that this mechanism predominantly hinges on the access to new information networks, I examine the effect by contrasting two scenarios. In the first scenario,

[3] Hall, Jaffe and Trajtenberg (2001) and Lerner and Seru (2021) emphasize the importance of weighted number of patents to account for the potential bias generated by patenting trends over time and field. Hall, Jaffe and Trajtenberg (2005) and Lerner, Sorensen and Strömberg (2011) show that the adjusted number of citations not only sheds light on an inventor's patenting activity and the influence of the patent, but is also shown to carry economic value.

I analyze cases where the inventor who stayed had previously collaborated with inventors in the destination location *before* the move, and compare that to a scenario in which such collaborations did not take place. This comparison allows me to assess the difference between scenarios where the move grants the inventor access a new cohort of inventors potentially holding novel information and situations where the information exchange might have occurred prior to the move. I find a positive and statistically significant impact when the network is new, while observing a statistically insignificant effect when the collaboration with inventors in the destination location was already established prior to the move. Specifically, in comparison to the control group, the stayers experience a 10 percent increase in their annual number of patents and a 21 percent increase in their annual number of adjusted citations.

Collectively, these findings suggest that the main driving force behind the observed effects is access to a new network and, consequently, to new information resulting from co-inventor relocations.

This paper is connected to various strands of the literature. Firstly, it relates to studies that examine the effects of migration on productivity in the origin location (Kerr, Kerr, Özden and Parsons, 2016, 2017; Kerr, 2008; Waldinger, 2012) and the horse race between "brain drain" and "brain gain."[4] I, not only focus specifically on domestic migration within the United States and the effects thereof, but also take into account some characteristics of the move and the mover, which explain some of the differential effects I estimate.

Additionally, it relates to papers such as Azoulay et al. (2010) and Jaravel et al. (2018) that study adjacent topics, and in particular the significance of team-specific capital and network structure. Both of them report a negative effect on the productivity and labor market outcome of inventors and scholars experiencing the unexpected death of one of their colleagues. In a sharp contrast, I estimate positive productivity effects that evolve over time, following a colleague's relocation. These results are quite surprising and hinge on the idea that in contrast to a death shock, which implies a complete discontinuation of any relationship, a colleague's relocation introduces opportunities for spillovers. This could occur though continued collaboration or even in the form of network expansion, both of which are impossible in the event the colleague passes away.

This research is also connected to Moretti (2021), who studies how the size of an inno-

[4] In a more recent paper, Prato (2022) examines the impact of migrants from Europe on innovation activity in the United States, as well as the reciprocal effect. Similarly, Bernstein, Diamond, Jiranaphawiboon, McQuade and Pousada (2022) study the influence of high-skill migration on innovation within the United States. Moser, Voena and Waldinger (2014) use the emigration of Jewish Germans from Germany to the United States to address this topic in the setting of Chemical innovation. Other papers highlight the possibility that reallocation of skilled labor into the innovation sector can have negative effects on technology adoption in other parts of the economy, resulting in a form of brain drain (see for instance recent work by Trouvain (2022)).

vation cluster affects the innovative output of inventors within that cluster. Leveraging the variations introduced in this paper, as well as the definition of economic regions, my study focuses on understanding the effects on inventors who stayed in the origin location when their collaborators relocate, rather than concentrating on the relocating individuals themselves. This approach enables me to delve deeper into the effects associated with relocations from a different perspective, and study the implications for innovation in the "losing" location.

My results add to the results by Agrawal, Kapur and McHale (2008), by showing that collaboration with an inventor who moves to a new location, can serve as a channel though which the stayer can get access to a new network and, therefore, to promote knowledge spillovers between inventors who are not co-located. It also corroborates the findings of Zacchia (2020), which demonstrate that collaborative patent activities spanning different firms facilitate the transfer of knowledge between these firms. In contrast, my focus is on the individual productivity of inventors, and I use prior collaboration as a prerequisite for subsequent knowledge spillovers, rather than leveraging current collaboration for immediate spillovers. Moreover, I examine situations where an inventor relocates and analyze how this relocation affects the productivity of inventors who remain in their original location. This analysis includes cases where the relocation leads to the termination of the mover and the stayer collaborative relationships, and I do not assume a static framework. Moreover, some of the effects I capture follow from within firm relocations and spillovers.

My research is also closely related to Zacchia (2018). However, there are notable differences between our studies. First, I do not restrict the relocations to be undertaken by superstar collaborators, thus examining the question in a broader context. Second, my dataset offers more comprehensive information about inventors, including their locations even in years in which they do not patent. This allows me to create a comparable control group through exact matching and use a different identification strategy that leverages recent advancements in the field. Additionally, I provide insights into the mechanisms behind the observed results, a facet not investigated in the referenced paper.

Lastly, this study is connected to other related literature that study team-based innovation and the significance of collaborative efforts in idea generation (Crescenzi, Nathan and Rodríguez-Pose, 2016; Jaffe, Jones et al., 2015; Jones, 2009; Wuchty et al., 2007). It also has relevance to the theoretical literature concerning team-based or network-based human capital (Chillemi and Gui, 1997; Mailath and Postlewaite, 1990) and theories encompassing the transfer of knowledge among inventors (Jarosch, Oberfield and Rossi-Hansberg, 2021; Lucas Jr, 2009; Lucas Jr and Moll, 2014; Stein, 2008). Additionally, while I focus on collaboration between geographical locations within the United States, Kerr and Kerr (2018) study the characteristics of collaborative patents where the inventors listed on them are located both within and outside of the United States.

The reminder of the paper is organized as follows. In Section 2, I outline the model. Section 3 describes the data and sample construction. Section 4 reports the estimates. Section 5 covers the mechanisms and Section 6 concludes.

2 A Model of Collaboration

In this theoretical section, I provide a simple framework that sets the stage for understanding the potential directions of effect and the mechanisms driving it. It is used to form predictions and to guide the empirical tests.

The model incorporates two key drives of inventors' productivity: collaboration, through team-production, and information sharing on a network. While both information acquisition and collaboration are integral components of the patent production function, the channels through which they affect one's productivity are distinct. Therefore, the model allows me to explore scenarios where the benefits of accessing new information outweigh the cost of losing a collaborator due to their relocation.

2.1 Basic Framework

2.1.1 Inventors' Network

A society of n inventors is connected via a directed and weighted network, which has an adjacency matrix $\mathbf{W} \in [0,1]^{n \times n}$. A general element $w_{ij} \in [0,1]$ represents the status and the intensity of the relationship between inventor i and inventor j. One can think about w_{ij} as the share of the patents both inventor i and inventor j are listed on out of the total number of patents inventor i is listed on. Similarly, w_{ii} is the fraction of solo-patents produced by inventor i by themselves. In other words, w_{ij} is a proxy for the fraction of time inventor i spends with inventor j producing patents. Specifically, an entry $w_{ij} = 0$ implies that inventor i and inventor j do not collaborate on patents, and therefore are not connected. A higher w_{ij} corresponds to a stronger relationship. Note that although the matrix $\mathbf{W}$ is not symmetric by assumption, collaboration is a reciprocal relationship, and therefore $w_{ij} > 0$ if and only if $w_{ji} > 0$.[5]

Figure 1 presents an example of a network with three inventors. In this example, inventor 1 collaborates with inventor 2 and with inventor 3, with whom they spend 0.5 and 0.15 of

[5]The reciprocity is only with respect to zero. That is, either both elements w_{ij} and w_{ji} are greater than zero, or both are equal to zero. However, when positive, I impose no assumption on whether they are equal or not.

their time, respectively. Inventor 2 and inventor 3 do not collaborate with one another, and they spend 0.25 and 0.35 of their time patenting with inventor 1, respectively.

Figure 1: Network of Inventors

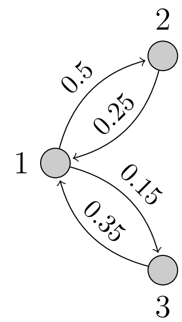

Inventors also have a fixed ability level $\alpha_i \in \mathbb{R}_+$ and an initial knowledge level $k_i \in \mathbb{R}_+$. These concepts play a central role in patent production and information sharing. Specifically, inventors exchange their knowledge through collaboration and apply their skills when creating patents.

2.1.2 Information Acquisition

Inventors acquire information through their network of inventors, not just from their immediate connections, but also from inventors who are located farther away. In particular, inventors can acquire information from inventors they collaborate with directly, as well as their second degree connections.[6]

To formalize this concept, it is helpful to introduce the definition for second degree connections.[7]

Definition 1 (Second-degree Connection)**.** Inventor j is said to be inventor i's second degree connection if

1. Inventor i and inventor j are not directly connected ($w_{ij} = w_{ji} = 0$), and

2. There exists an inventor m such that inventor m is directly connected to both inventor i and inventor j ($w_{mj}, w_{jm} > 0$, $w_{im}, w_{mi} > 0$).

[6]The assumption that information can be acquired only through first- or second-degree inventors simplifies the notations and enhances the comparability with the empirical tests and the estimation conducted in the following sections. It, by no means, restricts the validity of the propositions. In Appendix B, I present an extension where I do not impose this assumption.

[7]This is usually referred to as a distance of two between inventor i and j.

In this case, the intensity of the (indirect) relationship between inventor i and inventor j, through inventor m, is given by multiplying the intensities of the relationships between inventor i and inventor m and inventor m and inventor j ($w_{im} \cdot w_{mj}$ and $w_{jm} \cdot w_{mi}$, respectively).

Building upon Definition 1, I can now generalize notion of collaboration intensity to second-degree connections. Specifically, let $\bar{w}_{ji}$ denote the intensity of the (direct or indirect) relationship between inventor i and inventor j from inventor j's perspective.[8] In the case where inventors i and j are directly connected, this weight equals the intensity of their direct relationship, denoted as w_{ji}. However, if inventors i and j are each other's second-degree connections, these weights are determined by the sum of all intensities across all the inventors who connect between them. Specifically, for all i and j it is given by,

$$\bar{w}_{ji} = \begin{cases} w_{ji} & \text{if } w_{ji} > 0 \\ \sum_{m=1}^{n} w_{jm} \cdot w_{mi} & \text{if } w_{ji} = 0 \end{cases}$$

In Figure 1, inventor 1 and inventor 2 are directly connected. And inventor 3 is inventor 2's second degree inventor, with inventor 1 connecting between them. The intensity of the indirect relationship in this case is equal to $\bar{w}_{23} = 0.35 \cdot 0.5 = 0.175$ and of the direct relationship is $\bar{w}_{21} = 0.5$.

I define the total information held by inventor i as the result of combining between their initial knowledge and the information they acquire through first- and second-degree interactions with other inventors. Formally,

$$I_i = k_i + \sum_{j=1}^{n} \bar{w}_{ji} k_j \quad \forall i \tag{1}$$

where weighting the collaborators' knowledge by $\bar{w}_{ji}$ is meant to capture the idea that the information inventor i can learn from inventor j is proportional to a measure of the dependence between them. Intuitively, if inventor j spends only half of their time innovating with inventor i, it is unlikely that inventor i can learn everything inventor j knows during that time. Moreover, it implies that if inventor i and inventor j are only connected through an intermediate inventor m, then the amount of information inventor i gathers from inventor j is less than the amount inventor m gathers. This is because inventor m has a direct access to inventor j, whereas inventor i must rely on an intermediary to obtain any information from inventor j. Additionally, it implies the implicit assumption that information cannot

[8] Since the adjacency matrix is not necessarily symmetric, the intensity of the relationship depends on which inventor we focus on. For example, inventor i might co-invent 50 percent of their patents with inventor j4 while inventor j only co-invents 10 percent of their patents with inventor i. Therefore, the perspective of the specific inventor matters.

be gathered from both direct and indirect connections simultaneously. When inventors are connected through both direct and indirect links, they choose to use the direct connection for information acquisition. The idea is that a direct relationship offers superior access to information compared to an indirect one. Consequently, inventors are more inclined to act on the direct link rather than the indirect one.

2.1.3 Patent Production Function

Inventors produce patents in teams. Each inventor's total output is a linear combination between their own individual output, which is determined by their ability and the total information they hold, as well as the output contributed by their direct connections.[9] Inventor i's output is, therefore, given by

$$\begin{aligned} y_i &= \alpha_i + I_i && \forall i \\ Y_i &= y_i + \sum_{j \neq i} w_{ji} y_j && \forall i \end{aligned} \tag{2}$$

where y_i is inventor i's individual contribution, and it can be though of as the rate of producing solo-patents. Note the contribution of inventor i's collaborators is weighted by the time inventor j spends collaborating with inventor i, denoted as w_{ji}. It captures the idea that inventor j contributes more to the total output of inventors they spend more working on patents with.

This production function reflects the substitution between different sources that drive patent production. It emphasizes the tradeoff between the information inventors access through the network and the direct benefit the inventor gains from co-patenting, which comes in the form of the output contributed by their direct collaborators. This tradeoff will be the main focus of the next subsection.

2.2 Two Period Model

In general, patents are produced in various geographical locations. As a consequence, inventors may move around. In this part, I study the effect of a relocation on the total output of the mover's former collaborators through the eyes of the model. Specifically, the magnitude and the direction of this effect will be contingent on the specifics of the connections between these inventors, and on how their network changes in response to the relocation.

[9] The linearity assumption suggests that there are no peer effects in production, and inventors' individual outputs are perfect substitutes. Although peer effects are considered important, incorporating them would not provide deeper insights and would only complicate the model.

To start with, consider two time periods and two geographical locations. Let $t = 1$ represent the time before any relocation occurs, and denote by $t = 2$ the time after the relocation has taken place. Assume that the weight matrix $\mathbf{W}$ is fixed across time periods. I impose this assumption for compatibility with the empirical section where the intensity of the relationship between inventors is measured only prior to the relocation and not after, especially since there are limitations on what is observed and measurable.

Next, since the collaboration network can undergo changes between the two time periods, I introduce new notations which capture the state of the network in each one of these periods. Denote by $\mathbf{G}(t) \in \{0,1\}^{n\times n}$ the undirected adjacency matrix at time t. The ij-th entry represents the collaboration status between inventors i and j at time t, with the entry equal to one if they co-patent, and zero otherwise. This relationship is reciprocal. Additionally, let $\mathbf{S}(t) \in \{0,1\}^{n\times n}$ be the undirected information exchange network at time t. This is a symmetric matrix whose entries are equal to one whenever the inventors are engaged in information sharing. In the initial period, information sharing occurs exclusively when inventors co-patent.[10] However, after the relocation, in period $t = 2$, inventors who previously co-patented in $t = 1$ may cease their collaboration in period $t = 2$, and yet still engage in information sharing. Formally,

$$
\begin{aligned}
g_{ij}(1) = 1 \quad &\Longleftrightarrow \quad s_{ij}(1) = 1 \qquad\qquad \forall i,j \\
g_{ij}(2) = 1 \quad &\Longrightarrow \quad s_{ij}(2) = 1 \qquad\qquad \forall i,j
\end{aligned}
$$

Lastly, $I_i(t)$ is the level of information held by inventor i at time t, where the directed and second degree connections are now measured based on the entries of the network $\mathbf{S}(t)$.[11]

2.2.1 Information Acquisition the Two Period Model

To accommodate some degree of continuity across the two periods, I impose the following assumption:

Assumption 1. The information inventors acquire in the first period cannot be forgotten and therefore, it is not subject to relearning.

This implies that inventor i begins period $t = 2$ with information level that is equal to $I_i(1)$, rather than k_i, as it is at the beginning of period $t = 1$. Intuitively, once techniques

[10] This assumption follows by the empirical limitations. In the data, I cannot observe any relationship between inventors that is not tied to co-patenting.

[11] In the first period, the elements of the matrices $\mathbf{S}(1)$ and $\mathbf{G}(1)$ are equivalent to the indicators $\mathbb{1}\{w_{ij} > 0\}$. Therefore, the direct and second degree connections on $\mathbf{W}$, $\mathbf{S}(1)$ and $\mathbf{G}(1)$ are the same. However, in the second period, since information sharing can take place even when the inventors do not collaborate, it does not necessarily hold.

and ideas are acquired, they cannot be unlearned. Once learned, inventors can apply them again without relearning. In particular, equation (1) becomes

$$\begin{aligned} I_i(1) &= k_i + \sum_{j=1}^{n} \bar{w}_{ji} k_j && \forall i \\ I_i(2) &= I_i(1) + \sum_{j=1}^{n} \mathbb{1}\left\{\bar{w}_{ji}(2) - \bar{w}_{ji}(1) > 0\right\} \cdot \left[\bar{w}_{ji}(2) - \bar{w}_{ji}(1)\right] k_j && \forall i \end{aligned} \tag{3}$$

where $\bar{w}_{ji}(t)$ corresponds to the indirect weight and it is based on the links in the matrix $\mathbf{S}(t)$, which are weighted by the matrix $\mathbf{W}$.[12] Note that the multiplication by the element $\mathbb{1}\{\bar{w}_{ji}(2) - \bar{w}_{ji}(1) > 0\}$ enforces the condition that the information acquired in the second period was not previously acquired in the first period, as information from the first period cannot be lost or relearned.

2.2.2 Patent Production in the Two Period Model

Production in both periods follows the same structure as in equation (2), with output depending on both individual and scaled collaborative outputs. That is,

$$\begin{aligned} y_i(t) &= \alpha_i(t) + I_i(t) && \forall i \\ Y_i(t) &= y_i(t) + \sum_{j=1}^{n} g_{ji}(t) w_{ji} y_j && \forall i \end{aligned}$$

Define by $\Delta Y_i = Y_i(2) - Y_i(1)$ the change in inventor i's total output between the two periods.

2.2.3 Predictions

Let inventor i and inventor j be two inventors who are listed on at least one joint patent at time $t = 1$. And assume that inventor j relocates. In order to isolate the effect stemming directly from alternations in the patenting relationship with inventor j, assume that any change to inventor i's network involves adjustments related directly to inventor j.[13] To put it differently, my assumption is that the only connection of inventor i that could potentially change is the one with inventor j, and all other attributes, such as the abilities and information

[12] Although the matrix $\mathbf{W}$ is fixed across the two time periods, the potential differences in the matrices $\mathbf{S}(2)$ and $\mathbf{S}(1)$ can lead to differences in $\bar{w}_{ji}(t)$ across the time periods, if new links are formed and/or old links are severed.

[13] Empirically, I show that the number of inventors who use inventor j's old or new connections is very small, so I can abstract away from that in the model. However, I still address this possibility empirically.

held by inventor i's other connections, remain constant across both time periods.

As a result of inventor j's relocation, the relationship between inventor i and inventor j can evolve in different ways, which will determine the sign of the effect.

Proposition 1. *If, in period $t = 2$, inventor i and inventor j no longer collaborate or engage in information exchange with one another ($g_{ji}(2) = 0$ and $s_{ji}(2) = 0$), then the effect of the relocation on inventor i's output is negative ($\Delta Y_i = 0$). However, if information sharing takes place, the effect can be strictly positive.*

Proof. See Appendix A.

This Proposition states that, ex-ante, the effect on inventor i's productivity is ambiguous. If inventor i and inventor j break their collaboration and cease any communication, then inventor i only bears the output costs to that are associated with the loss of a collaborator. This loss is proportional to inventor j's output in the first period.[14] On the other hand, the effect driven by exposure to a new informational network operates in an opposite way to that of losing a collaborator. Therefore, as long as the information acquired through inventor j in the second period is valuable enough, it can offset the negative effect of losing inventor j as a collaborator and result in a positive effect.[15]

This proposition underscores the significance of information sharing. It asserts that, in contrast to death shocks, the termination of collaboration between inventors does not necessarily imply they stop engaging in information sharing. Hence, the discontinuation of collaboration does not inevitably lead to a negative impact on the inventors who remain in the original location after their collaborator relocates. The sign of effect hinges on the tradeoff between the benefits from shared information and the potential cost of losing a collaborator.

The potential informational benefits depend on the identity of inventor j's new connections after the move, as outlined in the next proposition.

Proposition 2. *Denote by $N_i(t)$ and $N_j(t)$ the set of inventor i's and inventor j's collaborators at time t, respectively. Holding the collaboration and information status between inventors i and j fixed, as well as the level of information inventor j gets access to after the relocation, if*

[14] This alteration to inventor i's network resembles that of a death shock caused by an inventor's death. The result in this proposition is consistent with results in previous studies (see, for example, Azoulay et al. (2010) and Jaravel et al. (2018)) which utilize death shocks in innovative sectors to study the effect on the survivors and find negative effects.

[15] It is also important to note that if the inventors maintain their collaborative links, inventor i benefits from a weakly higher level of information and a weakly stronger collaborator. In that case, both effects operate in the same direction, leading to a positive effect.

1. *The number of inventors j collaborates with after the move is fixed across two scenarios* $\left(\left|\tilde{N}_j(2)\right| = \left|N_j(2)\right|\right)$*, and*

2. *The number of new connections j establishes with individuals who were not collaborators of inventor i differs across the two scenarios* $\left(e.g.\ \tilde{N}_j(2) \smallsetminus N_i(1) \subseteq N_j(2) \smallsetminus N_i(1)\right)$

then, the size of the effect under $N_j(2)$ is greater. In other words,

$$\Delta Y_i\left(N_j(2)\right) \geq \Delta Y_i\left(\tilde{N}_j(2)\right)$$

With strict inequality as long as $s_{ji}(2) = 1$.

Proof. See Appendix A.

This proposition states that, holding everything fixed but the connections between inventor j's new collaborators and inventor i's collaborators, a greater overlap between inventor i's collaborators and inventor j's new collaborators leads to a lower effect on inventor i's output. The intuition relies on Assumption 1, which posits that inventor i cannot relearn information they already possess. Consequently, when inventor j relocates, inventor i and their immediate connections remain connected, meaning that inventor i doesn't acquire new information through inventor j's new connections, provided that these connections are formed with inventor i's existing connections. However, when inventor j forms connections with inventors not directly linked to inventor i, inventor i gathers additional information mediated by inventor j. And this results in a higher output, since the assumptions in the proposition impose that inventor j's individual output and relationship to inventor i is fixed across these two cases. This emphasizes that it is not just the act of sharing information that results in a potentially positive effect, but rather also the opportunity to access new and previously unknown information.

Conversely, if a change in the output is associated with a common shock following, for example, a significant success of a patent on which both inventor j and inventor i are listed, then the effect would be synchronous with the timing of the relocation but would not be triggered by any modifications in the network. This situation is equivalent to a shock affecting the innate abilities of both inventor i and inventor j, denoted as α_i and α_j, respectively. Consequently, it would not hinge on the identity of the inventors in the destination location, as Proposition 2 suggests.

Proposition 3. *Let $\alpha_i(t), \alpha_j(t)$ be the innate abilities of inventors i and j at time t, respectively. And assume that $\alpha_i(2) > \alpha_i(1)$ and $\alpha_j(2) > \alpha_j(1)$. If the conditions in Proposition*

2 hold, and the effect is driven solely by the changes in the inventors innate ability, then

$$\Delta Y_i \left(N_j \left(2\right)\right) - \Delta Y_i \left(\tilde{N}_j \left(2\right)\right) = 0$$

That is, the effect does not depend on the characteristics of the network inventor i gets access to.

Proof. See Appendix A.

The assumption that the effect is solely driven by a shock to inventors' abilities implies that the same amount of information is acquired in both scenarios. In that case, since the abilities are also equal across these cases, the effect should be the same.

Note that this proposition is not limited to that specific scenario, and a similar behavior should take place regardless of the which variation in the network I utilize.[16]

Another important feature of the collaboration network is the strength of the connections between the inventors. Given that both the amount of information acquired and the output produced through the network are scaled by these weights, it plays a crucial role in determining the magnitude of the effect.

Proposition 4. *Let $\tilde{w}_{ji} > w_{ji} > 0$ be two weights. Assuming the status of collaboration and the status of information sharing in the second period is the same under these two weights, the size of the effect under $\tilde{w}_{ji}$ is greater in absolute terms. In other words, $\left|\Delta Y_i \left(\tilde{w}_{ji}\right)\right| > \left|\Delta Y_i \left(w_{ji}\right)\right|$.*

Proof. See Appendix A.

This proposition highlights that a more intense collaborative relationship between inventor j and inventor i leads to a more pronounced effect, measured in absolute terms. This follows since the intensity of the connection between the inventors scales the information gains and also the benefits or costs associated with continued collaboration or severance of the links. Therefore, if the difference between the information gained and the collaboration effect is negative, multiplying this difference by a larger number amplifies the negative effect. Similarly, if the overall effect is positive, scaling it by a larger number magnifies its magnitude.

[16] In the empirical section, I use different variations such as the gender or distance of the move, which I do not model here.

3 Data and Descriptive Statistics

The dataset used for conducting the analysis is the product of a combination of two data sources, which together allow me to follow the innovative activity of inventors, as well as their locations within the United States over time. The first is the patent data from the US Patent and Trademark Office (USPTO). The second is provided by Revelio Labs, and it includes public employment information and other characteristics available on online professional profiles. The merge between these datasets opens up the opportunity to track inventors' employment history, including, but not limited to the employing firm, their office location, and their role in the firm.

3.1 Patent Data

The patent data covers all the U.S. patents granted between 1976 and 2022 and it was downloaded directly from PatentsView.[17] For each patent, this dataset includes information on the dates at which it was applied for and eventually granted, the individuals who worked on the patent, the firm to which the patent was assigned, the CPC classes and subclasses, as well as front-page backward and forward citations.[18,19] Furthermore, the city and the state the inventor resides in at the time of the application are also part of this dataset.[20]

The raw data initially lacks consistent identifiers for patent inventors, making it challenging to accurately track and analyze inventor information. However, PatentsView provides a valuable dataset that has undergone a disambiguation process. This process involves assigning similarity scores to inventors using various algorithms, allowing for a reliable and standardized inventor identification. Additionally, PatentsView follows a similar procedure

[17]According to their website, PatentsView is "a collaboration between the USPTO, American Institutes for Research (AIR), University of Massachusetts Amherst, New York University, University of California, Berkeley, Twin Arch Technologies, and Periscopic which started in 2012." It allows a bulk data download of raw as well as disambiguated, or processed data on patents and applications applied to with the USPTO.

[18]CPC or Cooperative Patent Classification is a patent classification system. It was developed in a collaboration between the USPTO and the EPO (European Patent Office) in the goal of constructing a consistent classification system across these two entities. The CPC is in use since 2013, but older patents were given this classification retro-actively. It has five hierarchies, where each layer in the hierarchy refines the subject to which the invention relates to.

[19]There is a distinguish between front-page citations and in-text citations. While front-page citations most often refer to prior art and can be added by the examiner as well, in-text citations refer to the history of the patent, its development and the usage of the innovation. Backward citations are the patents or papers a patent cites. While forward citations are the citations it receives.

[20]This is the inventor's home address and not the location of the firm the patent is assigned to (although that is provided as well). The reason one would prefer to focus on the inventor's home address is that the listed address for the assignee is usually the location of the headquarters, which might not be the the actual establishment the inventor works at. Since I am interested in the physical location where the inventor works at, using their home address is more appropriate due to the likelihood of proximity between the one's home and office.

for assignees, ensuring that a disambiguated assignee information is consistently represented in the dataset. Using this dataset, I can create a annual panel of inventors over time, where the observations in each year include information on all granted patents the inventor applied for that year.

Since I follow relocations within the United States, I do not include any inventor whose home address was listed outside of the United States for any of the years 1976-2022.[21] There are about 1.2 million unique inventors in the final dataset.

3.2 Online Professional Profiles Data

This is individual level database was built and provided by Revelio Labs. It is comprised of about 1.25B professional profiles and it provides their entire employment and education history as appears online by the end of 2022.[22] The information in these profiles is made available and published by the individuals themselves. It includes, but is not limited to, the firm they are employed in, the role they take in the firm, and the institute at which they acquired education at that time. In general, the information is supplied in a panel structure such that for each firm-position combination, there exists a starting date and an ending date.

This is the universe of global professional profiles, and as such it covers all employees who have such online profiles. Therefore, it is limited to individuals who choose to open an account and become users. That being said, we might be worried that the data is drawn from a non-representative sample of the population as a whole. However, as this paper focuses on high-skilled workers, with a very specific occupation – inventors – this issue should be less of a problem.

3.3 Data Construction

The ultimate goal is to create a panel dataset that encompasses yearly observations, providing crucial information on an inventor's location, employer, and patenting behavior. To achieve this, a merge between the USPTO dataet and Revelio Labs dataset is essential. This linking process is accomplished by utilizing the inventor's home address provided on the patent application. Since online professional profiles were introduced in 2003 and became popular around 2008, I limit my dataset to include only inventors who applied for a patent for the first time after 1990. Focusing on inventors who applied for the first time after 1990 allows me to address the potential bias related to their decision to become users and whether to

[21] For the purpose of this paper, and based on the the geographical units I utilize, I also exclude the American territories of the United States.

[22] These are profiles that are made publicly available online. Anyone can access them, whether or not their are users of the websites they were posted on.

include older information in their profile. Since such profiles are usually more beneficial for individuals who are in the labor force or looking to enter, it is less likely that a retiree or an individual close to retirement will have such an account. Building on Kaltenberg, Jaffe and Lachman (2023) who provide empirical evidence suggesting that most inventors apply for their first patent in their late 20s or early 30s, using 1990 as a threshold seems reasonable. By the time online professional profiles became popular, individuals who patented early on would likely not find it beneficial to have an online account.

The linking procedure involves using the name of the inventor, the state they lived in, and the name of the assignee listed on the application.[23] I ensure that the inventor's name matches the name on their online professional profile, that their state of residence matches their state of employment, and that the firm they report working for matches the firm listed on the patent application. The linking procedure is described in full details in Appendix C. This procedure results in approximately 300,000 successful matches, which represent about 30% of all US-based inventors who have exclusively lived within the US and have never obtained a patent before 1990.

Ultimately, the merged dataset effectively tracks inventors over time, providing annual records of their employers, job positions, and the specific city and state of the establishment where they work. Additionally, the dataset includes valuable information on their patenting activity, such as the number of patents and citations received.[24] The inventor's patenting behavior offers insights into both the quantity and quality of their innovative contributions. The number of patents filed in a given year serves as a measure of quantity, while the number of citations received indicates the quality of the patents, or their importance for innovation. However, it is essential to address potential biases and inaccuracies in measuring patent citations, as highlighted by Hall et al. (2001) and Lerner and Seru (2021). These authors pointed out that a simple citation count may be misleading due to changes in innovative activity over time and across CPC classes. To mitigate these issues, I adopt an adjusted measure for patent citations. This involves normalizing each patent's citation count by the average citation count for all other granted patents in the same year and CPC class. For the citation year, I refer to the application year of the cited patent.[25]

The inventor's location and firm are defined based on their employment history, ensuring

[23]Although online profiles were introduced around 2008, it is up to the individual to decided on the employment history they report. An individual can decide to use of their employment history, dating back as far as possible, and they can also decide to focus on the more recent past.

[24]Since the patent granting process can be lengthy and varies across technology classes, the application year might be significantly closer to the actual time of the innovation, and therefore, I choose to use the application year for these counts.

[25]The reference year is the year the patents were applied for. That is, the average number of citation is calculated relative to the application date rather than the grant date.

accurate and reliable information in the dataset. During the calendar year, if an inventor changes employers or relocates, the employer and location information for each entry in the panel dataset are determined based on the maximum number of days spent in a specific year. This means that the employer entry will reflect the firm where the inventor was employed for the greatest number of days throughout the year. Similarly, the location entry will be based on the place where the inventor spent the majority of their time during that year. In both cases, ties are broken randomly. This approach ensures that the panel dataset provides the most accurate and representative information regarding the inventor's employer and location for each year.

Appendix Table C.1 presents summary statistics related to patent activity and demographics for the final sample, comparing it with the full sample of inventors. In both sets of samples, the distribution of innovation activity is skewed towards higher values. This characteristic is evident in both the quantity and quality aspects of innovation. Moreover, a marginal distinction exists between the two samples, with inventors in the linked sample showcasing slightly higher average productivity than their counterparts in the full sample. This discrepancy could potentially be attributed to the fact that the linking tends to favor inventors with more observations. This is a result of the linking process, which primarily considers inventors with higher accuracy rate.[26]

3.4 Movers, Left Behinds and Sample Construction

Following Moretti (2021), a move in this setting will be defined as a relocation between the U.S. Bureau of Economic Analysis' (BEA) "economic areas." These "economic areas" are determined by the patterns of labor commuting, effectively defining local labor markets. They consist one or multiple MSAs, which serve as hubs of economic activity, along with the counties that share economic interdependencies with them. Therefore, the distinction between "economic areas" is more pronounced, making it less likely for spillovers to naturally occur between these locations.

The United States comprises 179 economic areas that span the entire country, which can vary in size depending on their location. Notably, in larger regions like New York, Boston, or San Francisco, the economic areas tend to be larger than their corresponding MSAs as illustrated in Figure 2 (Johnson, 2004).

[26]The scenario where the patent's assignee matches the employing firm of all inventors involved isn't universally applicable. Due to this, inventors associated with patents where the listed assignee differs from their employing firm won't be considered for matching. Conversely, if their patent history includes multiple instances, some of which under their employing firm, the likelihood of successfully linking them with their online profile increases.

Figure 2: Economic Areas Examples: San Francisco and San Diego

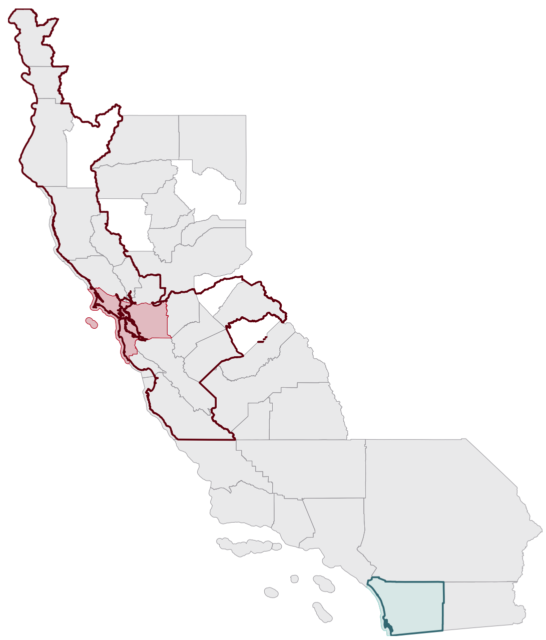

Notes: This map illustrates two economic areas: San Francisco (indicated by a red outline) and San Diego (indicated by a green outline) within the state of California. Notably, in smaller cities like San Diego, the MSA (colored in green) and the economic area (outlined in green) coincide. However, in larger cities such as San Francisco, the corresponding MSA (colored in red) is smaller than the corresponding economic area, delineated in red.

In each inventor-year observation, I determine the economic area by considering the location where the inventor states their work is conducted. Unlike the address provided in patent applications, which often pertains to the corporate headquarters, the city and state reported by the inventor on their professional profile likely accurately represent the physical location of the office where they work. This method ensures that the economic area assigned to each inventor reflects their actual workplace location, providing a more precise and representative measure of the network of inventors they gain access to.

With the information regarding the inventors' locations over time, I identify a group of movers as those who have relocated between economic areas, and I designate the year of their move as the first year in their new location. By examining the employment history provided in professional profiles, I can identify an inventor as a mover even if they have not filed any patents in that specific year or at any point afterward, which is an advantage in my setting.[27] Overall, I identify 49,904 movers who moved at least once. If an inventor moved more than once, I only consider the first move.[28] Moreover, as the network of co-inventors is an important part of the analysis, I only consider moves after the first patent. If the move occurred before the inventor filed their first granted patent, I disregard this move. This is because the network is defined based on patent collaborations, and if no patents were applied for, the set of co-inventors is, by definition, empty.

Table 1 reports the descriptive statistics of the moves. It shows that moves tend to be associated with more productive destinations, as measured by different specification. I follow Moretti (2021) and define the cluster size in a given year as the number of inventors in an economic area X CPC class pair, excluding the mover, as the share of all inventors in that CPC class and year. Specifically, it shows that about 75 percent of the inventors are relocating to areas which are associated with a higher productivity. Moreover, in over 50 percent of cases, the technology class in which the mover patents most frequently is the predominant technology class in their new location.

In accordance with Jaravel et al. (2018), I create a group of "placebo movers." These placebo movers consist of inventors who appear similar to those who actually relocated but, in reality, did not change their location and did not collaborate with any of the movers. To achieve this, I implement an exact matching procedure, matching movers to non-movers based on specific criteria: the cumulative number of patent applications at the time of the move, the year of the first patent, the time of the move, and the CPC class of the last patent before the move. In case of ties, the matched inventor is picked randomly. Accounting for the total number of applications inventors applied for prior to the move allows me to control for some degree of productivity, while considering the year of the first patent serves as a measure of experience in patenting. Utilizing a matching based on the CPC class of the final patent before the relocation enables me to mitigate potential biases that could arise from patenting activities associated with specific technological categories, or industry shocks. For instance, consider the IT revolution that unfolded in the early 2000s. If inventors specializing in IT-related domains were more inclined to relocate across economic areas, those who stayed in the origin location might have appeared more productive due to the influence of this technological shift, rather than their connection to the mover.

Ultimately, I successfully find an exact match for 43,123 movers, accounting for about 87 percent of the identified movers in my dataset. In Table 2, I present the summary statistics

[27] If I had solely relied on patent data, this identification would not have been feasible. The USPTO dataset only includes observations on inventors in years in which they apply for patents, and as most inventors do not file for patents every year, this would have likely resulted in a delayed identification of the inventor's move year and would bias the results.

[28] As a robustness check, I repeat the analysis while including only inventors who relocated only once and I find similar results.

Table 1: Descriptive Statistics on Moves

Move Associated with a More Productive Location by	
Number of Patents in Location (%)	51.78
Weighted Number of Patents by Citations in Location (%)	51.85
Number of Patents in Location X Technology Class (%)	75.15
Weighted Number of Patents by Citations in Location X Technology	75.18
Number of Inventors in Location X Technology Class (%)	74.97
Cluster Size (%)	74.97
Patent in the Same Technology Class (%)	53.88
Patent in the Same Technology Class in the Origin Location (%)	57.09
Patent in the Same Technology Class in the Destination Location (%)	57.17
Moves within the Same Firm (%)	17.55
Continue Patenting after the Move (%)	52.84
Avg. Number of Patents at Time of Move (#)	0.72
Avg. Stock of Citations at Time of Move (#)	1.03
Avergae Distance of the Move (Miles)	980.59

Notes: This table includes information about the moves. It shows the percent of moves that are associated with a more productive destination, based on different measures of productivity, as well as, whether the mover tends to patent in the technology classes that are associated with this location.

of the real and placebo movers' characteristics at the time of their move. The real and placebo movers are perfectly balanced in terms of their first year of patenting, cumulative number of applications, year of the move, and CPC class, as per the matching procedure construction. According to the table, at the time of the move, real and placebo movers have applied for 2.63 patents on average, that were eventually granted. The average number of years between the move and the first patent is 5, while the average year of the move is 2012. Additionally, despite not being directly matched on these characteristics, the real and placebo movers also exhibit balance concerning the number of adjusted citations, averaging 2.75 for the real movers and 2.58 for the placebo mover, and gender distribution, with an average of 86 percent male among the real movers and also the placebo movers. Overall, the table demonstrates that both groups are also evenly balanced in various characteristics that encapsulate their innovation-related activities. This further strengthens the credibility of the matching process, a crucial factor in deriving accurate insights from the findings.

Next, I construct the co-inventor network for both real and placebo movers. This group comprises any inventor who has previously collaborated on a patent with either a real or placebo mover before their respective moves. These inventors are referred to as real and

Table 2: Summary Statistics on Real and Placebo Movers

	Real Movers				Placebo Movers			
	Mean	Median	Std. Dev	# Obs.	Mean	Median	Std. Dev	# Obs.
First Patent Year	2007	2008	9	43,123	2007	2008	9	43,123
Move Year	2012	2014	8	43,123	2012	2014	8	43,123
Patent Stock	2.63	1.00	3.12	43,123	2.63	1.00	3.12	43,123
Average Number of Patents Per Year	0.70	0.50	0.64	43,123	0.70	0.50	0.64	43,123
Adjusted Citations Stock	3.75	0.93	12.43	43,123	3.54	0.89	13.99	43,123
Average Adjusted Citations Per-Year	1.01	0.23	3.68	43,123	1.00	0.22	7.76	43,123
Male	0.86	1.00	0.34	38,991	0.85	1.00	0.36	39,471

Notes: This table presents summary statistics for both the real movers and the matched placebo group that I constructed. The variables are assessed at the time of the real or placebo move and encompass both the total count and the average values over the period before the move. What becomes evident from this table is that, even among variables that were not specifically matched, the real movers and the placebo group appear to be evenly balanced.

placebo "stayers," respectively. To ensure a reliable analysis, I exclude inventors who formerly co-patented with more than one real or placebo mover, leaving me with 23,553 real stayers and 15,401 placebo stayers. The summary statistics for these groups are presented in Table 3, which demonstrates the balance between these groups in terms of their patenting activity and characteristics.

On average, both real stayers and placebo stayers indicate their initial job on their online professional profile as 1997. Their first patent, on average, is applied for about 10 years later, and they experience the move of their collaborate in 2015, on average. Real stayers apply for an average of 6.54 patents before the move, with a cumulative total of 6.97 adjusted citations. In contrast, placebo stayers apply for an average of 7.10 patents prior to the move, accompanied by an average cumulative total of 7.59 adjusted citations. They are also balanced in term of gender, with 84 percent of real stayers and 85 percent of placebo stayers being male.

Remarkably, this balance is achieved despite not conducting the matching procedure at the stayers level. This implies that through the exact matching procedure, which was aimed to establish a foundation estimating more accurately the effects, I created a comparable control group. This will help me to mitigate any biases that could arise from patenting activities and the timing of inventors' relocations. Through CPC class matching, I ensure that the placebo movers are patenting within the same technology category as the real movers. Consequently, any effect driven solely by the patenting category of the movers should be effectively neutralized, as the placebo stayers will be exposed to it in a manner similar to the real stayers.

Table 3: Summary Statistics Real and Placebo Left Behind Pre-Move

	Real Left Behind				Placebo Left Behind			
	Mean	Median	Std. Dev	# Obs.	Mean	Median	Std. Dev	# Obs.
First Year in Sample	1997	1998	10	23,553	1997	1998	10	15,401
First Patent Year	2007	2009	9	23,553	2007	2009	9	15,401
Move Year Mover	2015	2016	6	23,553	2015	2016	6	15,401
Patents Stock	6.54	3.00	11.66	23,553	7.10	3.00	13.41	15,401
Average Number of Patents Per Year	0.40	0.20	0.71	23,553	0.42	0.20	0.73	15,401
Adujsted Citations Stock	5.77	2.27	12.67	23,553	6.32	2.19	14.54	15,401
Average Adujsted Citations	0.36	0.13	0.79	23,553	0.38	0.13	0.88	15,401
Male	0.84	1.00	0.36	21,457	0.85	1.00	0.35	14,036

Notes: The information presented in this table offers summary statistics for both the real stayers and the placebo stayers group. An inventor is considered to be a "stayer" if they have collaborated with an inventor who eventually relocates. As such, the real (placebo) stayers are listed on a patent with real (placebo) movers before their respective moves. The variables are examined at the time of the real (placebo) move and cover both the cumulative count and the average values throughout the period prior to the move. Notably, it's important to observe that although I didn't specifically match characteristics of the stayers, the dataset seems to be well-balanced.

4 The Productivity of Left Behind Inventors

In this section, I outline the methodology I use to estimate the average treatment effect on inventors' productivity following their co-inventor's relocation. This effect is identified through a difference-in-differences research design, where the control group consists of co-inventors who did not experience a relocation of their collaborators but share similar characteristics with the treatment group, as detailed in Section 3. By selecting the co-inventors in this manner, I address the selection bias concern, ensuring that co-inventors who experience a peer's relocation are not substantially different from those who did not experience the relocation of their collaborators. This approach helps mitigate any bias that could arise from productivity disparities between the two groups, enabling a more reliable analysis of the effects of the inventor's move on their co-inventors' productivity.

4.1 Dynamic Effects

Building on the identification strategy in Jaravel et al. (2018), I employ OLS regressions with a full set of leads and lags around the inventor's move. This approach enables me to study how the relocation influences the co-inventor's productivity over time. Furthermore, the methodology helps validate the research design by testing for any observable effects of a collaborator's impending move before the actual relocation event occurs, thereby addressing concerns about a potential common shock. Specifically, using the constructed treatment and control group, I stack the moving events together such that the time index t represents the time relative to the move, which is set to takes place at time $t = 0$, thereby addressing the concerns about estimation of pre-trends raised by Roth (2022). I then estimate the following OLS regression:

$$Y_{it} = \sum_{k=-9}^{9} \beta_k^{\text{Real}} \mathbb{1}_{\left\{L_{it}^{\text{Real}}=k\right\}} + \sum_{k=-9}^{9} \beta^{\text{All}} \mathbb{1}_{\left\{L_{it}^{\text{All}}=k\right\}} + \alpha_i + \alpha_y + \varepsilon_{it} \tag{4}$$

where L_{it}^{Real} are the leads and lags around the time of the relocation for the real stayers. Similarly, L_{it}^{All} are the leads and lags around the time of the relocation for both real and placebo stayers. $\left\{\beta_k^{\text{Real}}\right\}_{k=-9}^{9}$ and $\left\{\beta_k^{\text{All}}\right\}_{k=-9}^{9}$ are the predictive effects associated with the respective leads and lags, where k denotes the time relative to the move. I also include individual (α_i) and year (α_y) fixed effects.[29] To account for a possible serial correlation between inventors who are associated with the same mover, I cluster the standard errors at the mover level as in Jaravel et al. (2018).

If the move is as good as an exogenous variation, the coefficients $\left\{\beta_k^{\text{Real}}\right\}_{k=-9}^{9}$ identify the causal effect of experiencing a relocation of a co-inventor k years relative to the year of the move. If experiencing the relocation of a co-inventor results in an increase in the inventor's productivity k years relative to the move, β_k^{Real} should be positive. Conversely, if it negatively affects the inventor's productivity, the coefficient should be negative, and if there is no effect, it would be zero. However, an identification concerns that needs to be addressed is selection bias. If the inventors who experienced the relocation of their co-inventor display, on average, higher productivity compared to the overall population of inventors, then the estimated effect might not reflect true causality. To mitigate this issue, I use a control group of inventors who share similarities with those experiencing a relocation of a co-inventor but did not experience that themselves. In the analysis below, I demonstrate that there is no evidence of statistically significant pre-trends, effectively addressing this concern. Specifically, the results indicate that the lag terms β_{-9}^{Real} to β_{-1}^{Real} are not statistically significant before the move occurs. This finding suggests that prior to the move, the inventor's productivity is not influenced by the future move, thus supporting the interpretation of the estimates.

In Figure 3, I report the point estimates derived from regression equation (4). No pre-trends are reported. It also reveals that the productivity of stayers, relative to the placebo group, exhibits a positive effect. This improvement becomes statistically significant approximately three years after the initial event when considering the annual number of patents and the weighted metric of annual number of patents, termed adjusted citations. Also, although

[29]These fixed effect also include the time elapsed since the first patent, i.e., experience in the patenting innovation.

the effect weakly increases over time, it appears to stagnate around six years after the move. This suggests that the effect remains for many years.

This outcome aligns logically with the notion that there is a time lag involved. It requires some duration for the relocating inventor to amass new knowledge and subsequently transmit it to their former collaborators in the origin location. Moreover, the delayed adjustment in citations can be attributed to the usual pattern where citations trail behind patents.[30]

4.2 Baseline Regression

As a way of summarizing the results, I utilize a second specification that directly compares the pre-period and post-period to estimate the average treatment effect of the move, instead of examining how the effect evolves over time. In this alternative approach, the dummy variables $PostMove_{it}^{\text{Real}}$ and $PostMove_{it}^{\text{All}}$ turn one after the time of the real or placebo move, respectively. This allows for a more concise assessment of the move's impact without focusing on the dynamic changes over time. This specification is given by:

$$Y_{it} = \beta^{\text{Real}} PostMove_{it}^{\text{Real}} + \beta^{\text{All}} PostMove_{it}^{\text{All}} + \alpha_i + \alpha_y + \varepsilon_{it} \tag{5}$$

I also include individual and year fixed effects, as before. And cluster the standard errors at the mover level.

Table 4 presents the outcomes of the regression analyses. The results indicate a positive and statistically significant impact on both the annual number of patents and the annual number of citations. In column (1), the coefficient β^{Real} equals 0.045 and it is statistically significant at the 5 percent level. This signifies that in comparison to the placebo stayers, those who are truly treated tend to exhibit a higher productivity following their co-inventor's move. When compared to the average annual patents post mean for the control group (0.5), this translates into an approximately 9 percent increase in the number of patents per year compared to the control group.

Similarly, as seen in column (2), the coefficient β^{Real} stands at 0.051 and is highly statistically significant. This suggests that, relative to the placebo group of inventors, real stayers generate a higher annual number of weighted-patents on average, as implied by the number of annual adjusted citations. In percentage terms, this signifies more than 15 percent increase in the annual number of adjusted citations compared to the placebo stayers group.

[30] I use the date of patent application, not the date of patent granting. So while it's generally observed that patent application timing aligns closely with R&D (as indicated by studies such as Griliches (1998) and Griss (1993)), my findings do not conflict with this notion, because the mover should logically not engage in R&D in the new location before the move takes place. Consequently, both sets of results can indeed remain valid, and patent and citation counts can take time to respond.

Figure 3: Dynamic Effects

(a) Annual Number of Patents

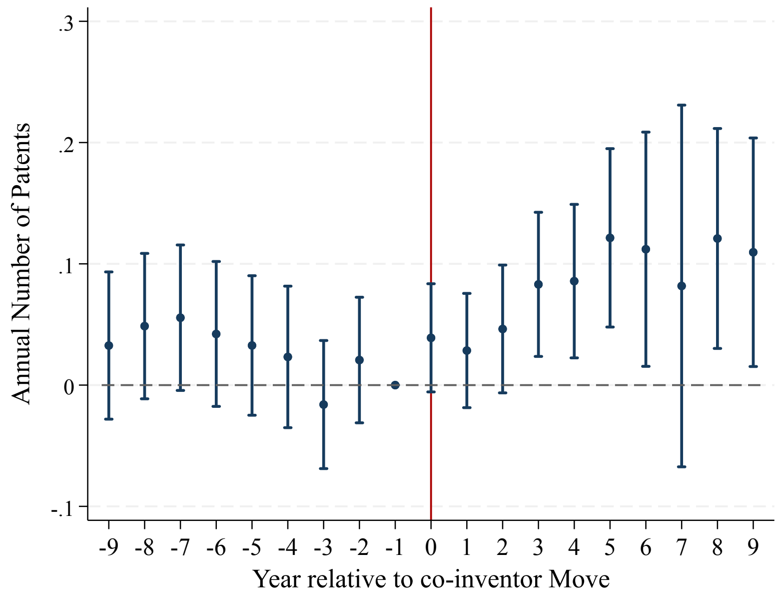

(b) Annual Number of Adjusted Citations

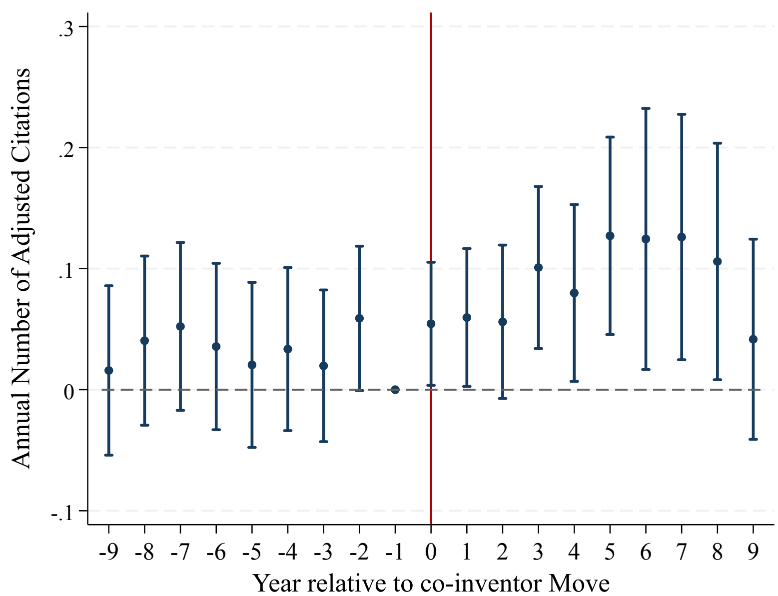

Notes: This plot presents the effect of an inventor's move on their collaborators' productivity. The vertical lines represents a 95% confidence interval, while standard errors are clusters at the mover level. The dependent variables are the annual number of patents and the annual number of adjusted citations, which are defined in Section 3.

A second identification concern relates to a common shock, involving unobserved time-varying productivity shocks at the mover level that could lead to more productive stayers.

Table 4: Baseline Regression Results

	(1) Annual Number of Patents	(2) Annual Number of Adjusted Citations
$PostMove^{\text{Real}}$	0.045^{**}	0.051^{***}
	(0.019)	(0.018)
Control Post Mean	0.5	0.34
Percentage Change	+8.95%	+14.92%
Observations	555815	555815
Individual FE	Yes	Yes
Year FE	Yes	Yes
Experience FE	Yes	Yes

Notes: The information presented in this table presents the results from regression equation (5). The unit of analysis in these regressions is inventor-year. The dependent variable in column (1) is the number of patents per year and in column (2) is the number of adjusted citations per year, as defined in Section 3. Standard errors are clustered at the mover level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

If the mover's opportunity to relocate is influenced by their prior work and success, and the stayer is associated with any of these patents, there is a possibility that the stayer, despite not moving themselves, experiences increased productivity due to the success of these patents and not directly due to the move. This increased productivity may be linked to factors like greater ease in obtaining research grants, which allows them to access better equipment and additional resources. Consequently, the observed effect might be driven by the joint previous success rather than the mere connection to someone who moved. To address this concern, I conduct a series of heterogeneity tests to account for potential confounding factors and ensure the validity of the estimates in the next section.

4.3 Additional Results and Robustness Checks

Within vs. Across Firm Moves. Although inventors within the same company are typically encouraged to collaborate on patents, it is conceivable that when a relocation involves switching employers, these inventors may no longer have the opportunity to co-patent with each other. If this scenario holds true, it could create obstacles to collaboration and, potentially, hinder the exchange of information. To address this possibility, I have confirmed that my results are robust to within and across firm moves, and the results are presented in Appendix Table D.1.

Bad vs. Good Moves. Inventors relocate to various locations, and the characteristics of the location they move to might generate consequences for their co-inventors who remain in the origin location. Research indicates that relocating to a bigger innovation cluster leads to increase in the productivity of the inventor making the move (Moretti, 2021). Building upon this finding, I show that when the mover relocates to a bigger innovation cluster, defined by a higher concentration of inventors patenting in the same modal CPC class, the effect of the productivity of those who stay is larger in a statistically significant way. This implies that the effect is contingent on the nature of the relocation and whether it provides opportunities for heightened patenting activity. The results are presented in Appendix Table D.3.

Additional Robustness Checks. In Appendix D, I report additional robustness checks showing that the results do not depend on the type of firm the inventors move to, and are not driven by movers who relocated more than once.

5 Mechanism: New Information as the Driving Force

In this section, I show that the long-lasting positive effect on the productivity of the mover's co-inventors is a result of the opportunities created by the mover for the stayers to access new information. First, I rule out that the effect is driven by a common shock, such as access to more resources. Second, I show that the effect is not driven by firm or network effect. That is, cases where the whole firm or network is on a higher trajectory before the move. Third, I demonstrate that the effect is primarily attributed to the acquisition of new information, rather than information in general. I show that the stayers cite patents produced in the destination location more extensively, and evidence on network expansion into the destination location, after the move.

I proceed by highlighting the asymmetric nature of this effect, particularly regarding inventors who shared a history of extensive collaboration with the mover prior to the relocation. As anticipated, this subgroup experiences a more pronounced effect, owing, possibly due to the higher likelihood of maintaining some form of relationship in these cases. I conclude by showing two facts: first, I show that the effect remains positive even when I exclude patents produced in collaboration with the mover, indicating that stayers acquire skills that contribute to their independent work. Second, I show that the effect is predominantly prevalent in cases where the mover relocates to a location where the stayer has no prior collaborators.

5.1 Ruling Out the Common Shock Reasoning

Heterogeneity by Sex Differences. The research design I employ, while addressing some potential bias, doesn't completely eliminate the common shock concern. This concern arises when the effect is influenced by individual-level changes over time, such as access to more resources. For instance, these could be resources that were previously occupied by the mover or those resulting from a successful patenting activity for which the mover was "rewarded" with a relocation opportunity. To address this concern, I leverage patterns that common shocks can not produce. I do so by showing that the size of the effect depends on similarities between the mover and the stayers that are not thought to be correlated with success, but rather are similar in nature to homophily.[31]

Table 5: Heterogeneity by Sex Differences

	Same Sex		Opposite Sex	
	(1) Annual Number of Patents	(2) Annual Number of Adjusted Citations	(3) Annual Number of Patents	(4) Annual Number of Adjusted Citations
$PostMove^{\text{Real}}$	0.059***	0.058***	-0.011	0.020
	(0.022)	(0.021)	(0.037)	(0.029)
Control Post Mean	0.501	0.349	0.496	0.304
Percentage Change	+11.72%	+16.7%	-2.28%	+6.42%
P-Value H_0: Diff. = 0			0.09	0.26
Observations	364681	364681	109289	109289
Individual FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Experience FE	Yes	Yes	Yes	Yes

Notes: The information presented in this table presents the results from regression equation (5) on two different samples. Columns (1) and (2) correspond to cases where the mover and the left behind are of the same sex, and columns (3) and (4) cover the cases where the mover and the left behind are of opposite sexes. The unit of analysis in these regressions is still inventor-year. The dependent variable in columns (1) and (3) is the number of patents per year and in columns (2) and (4) is the number of adjusted citations per year, as defined in Section 3. Standard errors are clustered at the mover level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

The findings presented in Table 5 report differential effects which are contingent on shared characteristics between the mover and the stayer. It is evident that when the mover and the stayers are of the same sex, the effect is positive and statistically significant. This effect

[31] An extensive body of literature has emphasized the significance of homophily, a concept characterized by individuals forming relationships with others who share similar attributes. Numerous studies underscore the role of homophily in forging robust connections (see McPherson, Smith-Lovin and Cook (2001) for a survey of the literature in sociology). See also Currarini, Jackson and Pin (2009), Currarini, Jackson and Pin (2010) and Bramoullé, Currarini, Jackson, Pin and Rogers (2012) who model the origins of homophily.

implies that compared to the placebo stayers, the real stayers experience an average of about 12 percent increase in the annual number of patents and almost 17 percent increase in the annual number of adjusted citations. On the other hand, when considering the cases where the mover and the stayer have opposite sexes, the effect becomes negative and statistically insignificant. I can reject that the coefficients are the same at the 10 percent level.[32]

These results dispel concerns related to common shocks, as there is no basis to assume that a common shock would disproportionately impact inventors based on both their own sex and that of their respective movers. If the effects were exclusively driven by a common shock, the logical expectation would be for the results to remain consistent across all sex combinations between the mover and the stayer, as Proposition 3 suggests. However, the findings in Table 5 demonstrate that this is not the scenario.

I find qualitatively similar results when conditioning on the race similarities between the mover and the stayer. I report these results in Appendix Table D.4.

Heterogeneity by the Distance of the Move. Another evidence that the effect is not driven by a common shock comes in the form of the asymmetric effect of the distance between the mover and the stayer after the move.[33]

Specifically, I calculate the distance between the actual mover and the real stayer by utilizing the location coordinates of the city where the inventor's workplace is situated.[34] It is defined as the distance between the mover's destination and the location of the stayer. For the placebo mover, I impose the same distance between the origin and the destination of the mover they match to.

The outcomes presented in Table 6 outline the estimation findings obtained by splitting the dataset based on the distance the mover travels in their relocation. The threshold for determining the distance between the mover and the left behind is set at the 50th percentile of the distance distribution, which is equivalent to 714 miles.[35]

The table indicates that in this scenario, the significance of network novelty outweighs physical proximity. The effect is substantial and statistically significant, especially when

[32]These outcomes align with the findings in Cullen and Perez-Truglia (2023) which indicate that positive outcomes are observed when men interact with men, but no similar effect is found for women interacting with men.

[33]It corresponds to the geographical distance between the mover's initial location (and thus the location of the stayer) and the mover's new destination. In simpler terms, it represents the spatial gap between the mover and the stayer subsequent to the relocation.

[34]The coordinates correspond to the nearest city listed by the inventor as their workplace location. These coordinates indicate a random location within that city, and are persistent across cities. These coordinated align with the coordinates provided by Google Maps.

[35]For a better understanding, consider the following comparisons: the distance between Boston and Detroit spans 613 miles, while Boston to Chicago covers 851 miles. Boston to New York is a distance of 191 miles, whereas Boston to San Francisco spans an extensive 2699 miles.

Table 6: The Heterogeneity by Distance from the Mover

	Short Distance (Bottom 50%)		Long Distance (Top 50%)	
	(1) Annual Number of Patents	(2) Annual Number of Adjusted Citations	(3) Annual Number of Patents	(4) Annual Number of Adjusted Citations
$PostMove^{\text{Real}}$	0.029	0.027	0.061***	0.073***
	(0.033)	(0.028)	(0.020)	(0.021)
Control Post Mean	0.519	0.348	0.481	0.333
Percentage Change	+5.5%	+7.89%	+12.68%	+21.98%
P-Value H_0: Diff. = 0	0.398	0.195		
Observations	279200	279200	276615	276615
Individual FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Experience FE	Yes	Yes	Yes	Yes

Notes: The information presented in this table presents the results from regression equation (5) applied to two distinct subsets of data. The first subset corresponds to left behind inventors who are located in a closer proximity to their respective mover, and the second corresponds to left behind inventors who are located at a greater distance form their respective mover. Close proximity is defined as distance that is shorter than 714 mile (1149 km). The distance between the placebo left behind and placebo mover is defined to be equal to the distance between the real inventor and the real left behind they are matched to. Columns (1) and (2) pertain to scenarios where the left behind and the mover are close to each other, and columns (3) and (4) delve into cases where they are located at a greater distance. The outcome variable in columns (1) and (3) is the number of patents per year, while columns (2) and (4) utilize the number of adjusted citations per year, as elaborated in Section 3. Standard errors are clustered at the mover level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

the inventor relocates to a significantly more distant location. It reveals that in comparison to the placebo stayers, real stayers experience an about 13 percent increase in their annual number of patents and a 22 percent increase in their annual number of adjusted citations post move.[36] This result might seem to stand in a contradiction to the extensive research that underscores the significance of face-to-face interaction and proximity for fostering productive and successful collaboration (Atkin, Chen and Popov, 2022; Battiston, Blanes i Vidal and Kirchmaier, 2021; Emanuel, Harrington and Pallais, 2023). However, following Saxenian (1994) who suggests that mere proximity is insufficient for information exchange among inventors and that collaboration is necessary, it is possible that the prerequisite of an existing prior connection, which is used in order to form the foundation for potential collaboration, can compensate for the absence of later physical proximity. Not to mention that a distance location makes it more likely for new information to flow.

Given that distance should not be correlated with common shocks, these results provide

[36]I find similar results for alternative thresholds such as the top and bottom 10th percentiles or the top and bottom 25th percentiles are chosen.

additional evidence against the common shock explanation. If the entire effect were attributed to a common shock, the magnitude and presence of the effect should not vary based on the distance of the city to which the mover relocates.

5.2 Ruling Out Firm and Network Effects

To investigate whether the diffuse of network effects is an important channel, I consider the group of real and placebo second degree connections. The results reported in Table 7 are obtained form specification (5) and show that the relocation has no significant effect on the productivity of their second degree connections. Hence, suggesting that network diffusion does not take place.

Table 7: Second Degree Inventors

	(1) Annual Number of Patents	(2) Annual Number of Adjusted Citations
$PostMove^{\text{Real}}$	0.018	0.019
	(0.026)	(0.023)
Control Post Mean	0.658	0.447
Percentage Change	-0.78%	4.3%
Observations	349693	349693
Individual FE	Yes	Yes
Year FE	Yes	Yes
Experience FE	Yes	Yes

Notes: The information presented in this table presents the results from regression equation (5) on the sample of second degree co-inventors. The dependent variable in column (1) is the number of patents per year and in column (2) is the number of adjusted citations per year, as defined in Section 3. Standard errors are clustered at the mover level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

5.3 The Effect is Driven by the Access to New Information

In this subsection, I provide evidence that the improvement in innovation productivity resulting from the relocation of a co-inventor is primarily influenced by the new information the mover exposes the stayer to. I do that in three steps. In the first step, I provide evidence that information sharing between the mover and the stayer takes place. I show that, after the move, the stayer is more likely to cite patents produced by inventors in the destination location or to collaborate with them. Finally, I show that the increase in the productivity is also positive and statistically significant when I exclude the patents on which the stayer and the mover collaborate on, suggesting that the information acquired by the stayer is used even on patents the mover does not collaborate with them on.

As a second step, I show that the effect is larger when the mover and the stayer have a more intense relationship prior to the move. These inventors are potentially more inclined to engage in information sharing, and this result supports that notion.

In the last step, I show that the effect is larger when the stayer does not have any prior collaborators in the location the mover relocates to. This result suggest that it is the access to new information that drives the results rather than information in general.

Share of Citations Made to Mover's Destination. Citations are considered a valuable measure to account for knowledge spillovers. Citations are references made to prior work that is relevant for the current patent scope and was filed before the citing application's filing date. An larger proportion of citations to patents produced in the mover's destination indicates a significant exposure to information in that location.

To examine how the citation behavior of the stayers changes after the mover's relocation, I calculate the share of citations the stayers attributed to patents located in the mover's destination. A patent is said to be produced in a certain location if at least one inventor who is listed on the patent is located there at the time of the application.

The findings in Table 8 display the outcomes when considering the share of citations attributed to patents located in the mover's destination. For the placebo movers, I enforce the destination of the real mover they are matched to.

The results indicate that, regardless of whether the inventors patent in a given year or not, there is a notable increase of approximately 11 percent to 12 percent in the share of citations directed toward inventors situated in the mover's destination, relative to the placebo group.

Share of Collaborators in Mover's Destination. Another way to gauge the interaction between the stayer and the mover, after the relocation, is by examining the proportion or count of collaborators the stayer has in the mover's destination. If the collaboration patterns between the stayer and other inventors in the mover's destination undergo changes after the move, it implies the existence of collaboration spillovers. This suggests that the relocation of a former collaborator to that location opens up opportunities for network expansion, which can be highly valuable, both at the individual level and for the firm involved, as these spillovers can lead to new and beneficial collaborative opportunities, fostering innovation and knowledge exchange beyond the direct interactions between the mover and the stayer.

Table 8: The Effect of Relocation on Citations to the Destination

	(1) Annual Percentage of Citations in Destination	(2) Annual Percentage of Citations in Destination
	All	Only When Patenting
$PostMove^{\text{Real}}$	0.047**	0.247***
	(0.023)	(0.095)
Control Post Mean	0.436	2.065
Percentage Change	+10.71%	+11.95%
Observations	555815	142296
Individual FE	Yes	Yes
Year FE	Yes	Yes
Experience FE	Yes	Yes

Notes: The information presented in this table presents the results from regression equation (5). In column (1), I consider all the observations in the panel, whether or not the inventor patented in that year, while column (2) covers only the years when the inventor patents. The unit of analysis in these regressions is still inventor-year. The dependent is the share of citations given to patents with at lease one inventor in the mover's destination. Standard errors are clustered at the mover level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

I calculate the annual share of collaborators located in the destination per year, and use it as the dependent variable in regression equation (5).[37] Here, I, again, assign the placebo mover the destination of the real mover they are matched to.

The results presented in Table 9 indicate that, in comparison to the placebo stayers group, the real stayers experience an increase in the annual share of collaborators they have in the destination location ranging from 74 percent to 77 percent.

Excluding Patents with the Mover. One might be worried that the estimated effect is solely driven by the increase in the productivity of the mover. If the mover and the stayer maintain their collaborative relationship, and the mover becomes more productive, this can lead to an increase in the observed productivity of the stayer through team production, as suggested by the model. If this is the case, the estimated increase in the stayer's productivity should actually be attributed to the mover. To address this concern, I examine the productivity effects on the real stayers when I exclude the patents the stayer produces in collaboration with the mover. I estimate the effect on the annual number of patents and annual number of

[37]Similar results are observed if, alternatively, I calculate the share of collaborators in the destination per patent and then average the values across all patents in the same year.

Table 9: The Effect of Relocation on Share of Collaborators to the Destination

	(1) Annual Percentage of Collaborators in Destination	(2) Annual Percentage of Collaborators in Destination
	All	Only When Patenting
$PostMove^{\text{Real}}$	1.176***	5.318***
	(0.096)	(0.219)
Control Post Mean	1.525	7.229
Percentage Change	+77.10%	+73.56%
Observations	555815	142296
Individual FE	Yes	Yes
Year FE	Yes	Yes
Experience FE	Yes	Yes

Notes: The information presented in this table presents the results from regression equation (5). In column (1), I consider all the observations in the panel, whether or not the inventor patented in that year, while column (2) covers only the years when the inventor patents. The unit of analysis in these regressions is still inventor-year. The dependent variable is the share of collaborators in the destination location. Standard errors are clustered at the mover level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

citations the stayer produces without the mover, which serves as a measure of productivity that is uncorrelated to the mover's patenting productivity.

The results in Table 10 show that there is a positive and statistically significant effect on both the annual number of patents and the annual number of adjusted citations. This implies that even when considering only the patents which are not produced in collaboration with the mover, the real stayers become more productive on average. A possible explanation is that these inventors acquire information from the mover that they can later apply in their own work, even without the mover's direct involvement in the project.

Heterogeneity by the Intensity of the Collaboration. To account for how the intensity of the collaboration between the stayer and the mover affects the results, I run separate regressions conditioning on this characteristic. Specifically, I calculate the number of patents that the mover and the staying inventor are jointly listed on. The frequency of collaboration between inventors can indicate their reliance on each other, and therefore, frequent collaborators are more likely to maintain even when not co-located. If that is indeed the case, the model predicts that effect should be stronger.

To capture the intensity of the connection, I create a dummy variable ($StrongLink$) which takes the value one for the top 50th percentile of the distribution of joint patents prior to the

Table 10: Baseline Specification Excluding Patents Produced with the Mover

	(1) Annual Number of Patents	(2) Annual Number of Adjusted Citations
$PostMove^{\text{Real}}$	0.035**	0.039***
	(0.016)	(0.015)
Control Post Mean	0.369	0.259
Percentage Change	+9.47%	+15.17%
Observations	555815	555815
Individual FE	Yes	Yes
Year FE	Yes	Yes
Experience FE	Yes	Yes

Notes: The information presented in this table presents the results from regression equation (5), when I exclude all the patents the stayer co-patented with the mover. Column (1) reports the effect on the annual number of patents, while column (2) shows the effect on the annual number of adjuated citations. The unit of analysis in these regressions is still inventor-year. The dependent variable is the share of collaborators in the destination location. Standard errors are clustered at the mover level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

move, indicating a strong connection, and zero for the bottom 50the percentile, representing weak connections.[38] Table 11 presents the results for two separate regression analyses: one for staying inventors with strong links and another for those with weak links.

The results reported in Table 11 indicate that the effect is large and statistically significant when the connections between the mover and the stayer are strong. Conversely, for weak links, the effect diminishes and is no longer statistical significant.

Access to a New Network. I presented evidence supporting the notion that the observed increase in productivity is, in part, influenced by knowledge spillovers and network expansion. Now, I delve further into the analysis and demonstrate that these spillovers have a significant impact primarily when the staying inventor is exposed to a new location. In other words, I show that the effect on the productivity of a staying inventor who already has collaborators in the mover's destination prior to the move is considerably smaller. This finding suggests that it is the access to the information in a new location that plays a crucial role in influencing productivity, rather than the relocation itself and the exposure to inventors in general.

[38]The reason I pick the top and bottom 50th percentile is that the 50th percentile is exactly one joint patent between the mover and the left behind. Since one joint patent is a requirement by the definition of being left behind, any other split will only limit the number of observations for the group with strong links.

Table 11: Effect Size and The Strength of the Link

	Strong Link		Weak Link	
	(1) Annual Number of Patents	(2) Annual Number of Adjusted Citations	(3) Annual Number of Patents	(4) Annual Number of Adjusted Citations
$PostMove^{\text{Real}}$	0.314**	0.294**	0.027	0.030
	(0.154)	(0.121)	(0.023)	(0.021)
Control Post Mean	0.971	0.483	0.438	0.305
Percentage Change	+32.36%	+60.86%	+6.07%	+9.85%
P-Value H_0: Diff. = 0			0.06	0.03
Observations	18380	18380	393254	393254
Individual FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Experience FE	Yes	Yes	Yes	Yes

Notes: The information presented in this table presents the results from regression equation (5) on two different samples. Columns (1) and (2) correspond to cases where the mover and the left behind are connected through a strong link, and columns (3) and (4) cover the cases where the mover and the left behind are connected through a weak link. A strong links is defined as collaborating on more than one patent prior to the move. The unit of analysis in these regressions is still inventor-year. The dependent variable in columns (1) and (3) is the number of patents per year and in columns (2) and (4) is the number of adjusted citations per year, as defined in Section 3. Standard errors are clustered at the mover level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

In order to account for the collaboration patterns in the mover's location, I create a dummy variable ($NewNetwork$). This variable takes the value one for the top 10th percentile of the distribution of the share of collaborators in the mover's destination prior to the move, indicating that the inventor gains access to a new network through the mover. Conversely, it takes the value zero for the bottom 10th percentile, indicating that the inventor did not access a new network through the mover.

Table 12 presents the results for two separate regression analyses: one for inventors who gain access to a new network through the mover and another for those who did not, as defined by the $NewNetwork$ dummy variable. The estimation indicates that the effect is primarily noticeable among inventors who lack collaborators in the destination location before the move. In particular, when compared to the placebo left behinds, the real left behind experience a 10 percent increase in the annual number of patents and a 21 percent increase in the annual number of adjusted citations following the move. Conversely, when an inventor has previously gained access to the network in the destination location through collaborations with inventors already situated there before the move, the effect is diminished and lacks statistical significance.

These results also contradict the common shock concern, as the common shock would

Table 12: The Effect of Relocation and the Access to a New Network

	New Network		Old Network	
	(1) Annual Number of Patents	(2) Annual Number of Adjusted Citations	(3) Annual Number of Patents	(4) Annual Number of Adjusted Citations
$PostMove^{\text{Real}}$	0.048*	0.069***	0.056	0.054
	(0.029)	(0.027)	(0.059)	(0.065)
Control Post Mean	0.472	0.325	0.604	0.397
Percentage Change	+10.25%	+21.32%	+9.25%	+13.54%
P-Value H_0: Diff. = 0			0.09	0.08
Observations	333731	333731	37373	37373
Individual FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Experience FE	Yes	Yes	Yes	Yes

Notes: The information presented in this table presents the results from regression equation (5). In column (1), I consider all the observations in the panel, whether or not the inventor patented in that year, while column (2) covers only the years when the inventor patents. The unit of analysis in these regressions is still inventor-year. The dependent variable is the share of collaborators in the destination location. Standard errors are clustered at the mover level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

have affected the staying inventor's productivity regardless of the collaboration patterns in the mover's location, as Proposition 3.

6 Conclusion

In this paper, I study the impact of an inventor's relocation on the productivity of their former collaborators. I build a novel dataset that merges patent and inventor information from the USPTO with inventors' online professional profiles. This allows me to identify movers, track the exact timing of their moves, and accurately determine their origin and destination locations, which would not have been possible otherwise.

I develop a model that integrates team-based collaboration and information sharing within a network. This model allows me to explore each channel in isolation and offers insights into the underlying dynamics at play. Consequently, I am able to delve into the mechanisms potentially underpinning the outcomes and formulate hypotheses about these channels, which I subsequently test empirically.

I find a positive effect of the relocation of co-inventors and an increase in the productivity of the inventors they "leave behind." The productivity measures include the annual counts of patents and of adjusted citations, which collectively reflect both the quantity and quality of the patents produced.

Heterogeneous treatment effects underscore the dominant role of information sharing across different locations, resulting in increased annual patent counts and citations. These findings also indicate that information sharing manifests in several ways, including more frequent citations to patents originating from the destination location and expansion of networks into that destination.

The tradeoff between increasing productivity through enhancing agglomeration on the one hand, and addressing spatial inequality on the other is at the core of debates around place-based policies. In this paper, I highlight that, under certain conditions, an inventor's relocation can lead to sizeable spillover effects in the origin locations on their former collaborators. For these beneficial effects to materialize, the degree and content of information exchange are important. This hinges on factors such as the inventor's prior collaboration history with the mover, their informal connections, and the extent to which the mover's relocation provides access to new information networks to the stayer . Essentially, relocations can foster brain gain by facilitating information sharing across geographically distant locations. This aspect should be taken into account when weighing policy decisions regarding innovation and spatial inequality.

References

Agrawal, Ajay, Devesh Kapur, and John McHale (2008) "How do Spatial and Social Proximity Influence Knowledge Flows? Evidence from Patent Data," *Journal of Urban Economics*, 64 (2), 258–269.

Atkin, David, M Keith Chen, and Anton Popov (2022) "The Returns to Face-to-Face Interactions: Knowledge Spillovers in Silicon Valley,"Technical report, National Bureau of Economic Research.

Azoulay, Pierre, Joshua S Graff Zivin, and Jialan Wang (2010) "Superstar Extinction," *The Quarterly Journal of Economics*, 125 (2), 549–589.

Battiston, Diego, Jordi Blanes i Vidal, and Tom Kirchmaier (2021) "Face-to-Face Communication in Organizations," *The Review of Economic Studies*, 88 (2), 574–609.

Bernstein, Shai, Rebecca Diamond, Abhisit Jiranaphawiboon, Timothy McQuade, and Beatriz Pousada (2022) "The Contribution of High-Skilled Immigrants to Innovation in the United States," (30797).

Bloom, Nicholas, Tarek Alexander Hassan, Aakash Kalyani, Josh Lerner, and Ahmed Tahoun (2021) "The Diffusion of Disruptive Technologies,"Technical report, National Bureau of

Economic Research.

Bramoullé, Yann, Sergio Currarini, Matthew O Jackson, Paolo Pin, and Brian W Rogers (2012) "Homophily and Long-Run Integration in Social Networks," *Journal of Economic Theory*, 147 (5), 1754–1786.

Chillemi, Ottorino and Benedetto Gui (1997) "Team Human Capital and Worker Mobility," *Journal of Labor Economics*, 15 (4), 567–585.

Crescenzi, Riccardo, Max Nathan, and Andrés Rodríguez-Pose (2016) "Do Inventors Talk to Strangers? On Proximity and Collaborative Knowledge Creation," *Research Policy*, 45 (1), 177–194.

Cullen, Zoe and Ricardo Perez-Truglia (2023) "The Old Boys' Club: Schmoozing and the Gender Gap," *American Economic Review*, 113 (7), 1703–1740.

Currarini, Sergio, Matthew O Jackson, and Paolo Pin (2009) "An Economic Model of Friendship: Homophily, Minorities, and Segregation," *Econometrica*, 77 (4), 1003–1045.

——— (2010) "Identifying the Roles of Race-Based Choice and Chance in High School Friendship Network Formation," *Proceedings of the National Academy of Sciences*, 107 (11), 4857–4861.

Ellison, Glenn and Edward L Glaeser (1997) "Geographic Concentration in US Manufacturing Industries: a Dartboard Approach," *Journal of Political Economy*, 105 (5), 889–927.

——— (1999) "The Geographic Concentration of Industry: Does Natural Advantage Explain Agglomeration?," *American Economic Review*, 89 (2), 311–316.

Ellison, Glenn, Edward L Glaeser, and William R Kerr (2010) "What Causes Industry Agglomeration? Evidence from Coagglomeration Patterns," *American Economic Review*, 100 (3), 1195–1213.

Emanuel, Natalia, Emma Harrington, and Amanda Pallais (2023) "The Power of Proximity to Coworkers: Training for Tomorrow or Productivity Today?," https://scholar.harvard.edu/sites/scholar.harvard.edu/files/pallais/files/power_of_proximity_july2023.pdf.

Greenstone, Michael, Richard Hornbeck, and Enrico Moretti (2010) "Identifying Agglomeration Spillovers: Evidence from Winners and Losers of Large Plant Openings," *Journal of Political Economy*, 118 (3), 536–598.

Griliches, Zvi (1998) *R&D and Productivity: The Econometric Evidence*: University of Chicago Press.

Griss, Martin (1993) "Hewlett-Packard Software Reuse: from Library to Factory."

Hall, Bronwyn H, Adam B Jaffe, and Manuel Trajtenberg (2001) "The NBER Patent Citation Data File: Lessons, Insights and Methodological Tools," (8498).

Hall, Bronwyn H, Adam Jaffe, and Manuel Trajtenberg (2005) "Market Value and Patent Citations," *RAND Journal of Economics*, 16–38.

Jaffe, Adam B, Benjamin F Jones et al. (2015) *The Changing Frontier: Rethinking Science and Innovation Policy*: University of Chicago Press.

Jaravel, Xavier, Neviana Petkova, and Alex Bell (2018) "Team-Specific Capital and Innovation," *American Economic Review*, 108 (4-5), 1034–1073.

Jarosch, Gregor, Ezra Oberfield, and Esteban Rossi-Hansberg (2021) "Learning from Coworkers," *Econometrica*, 89 (2), 647–676.

Johnson, Kenneth P (2004) "Redefinition of the BEA Economic Areas," https://apps.bea.gov/scb/pdf/2004/11November/1104Econ-areas.pdf, Acessed on January 20, 2023.

Jones, Benjamin F (2009) "The Burden of Knowledge and the "Death of the Renaissance Man": Is Innovation Getting Harder?," *The Review of Economic Studies*, 76 (1), 283–317.

Kaltenberg, Mary, Adam B Jaffe, and Margie E Lachman (2023) "Invention and the Life Course: Age Differences in Patenting," *Research Policy*, 52 (1), 104629.

Kerr, Sari Pekkala, William Kerr, Çağlar Özden, and Christopher Parsons (2016) "Global Talent Flows," *Journal of Economic Perspectives*, 30 (4), 83–106.

——— (2017) "High-skilled Migration and Agglomeration," *Annual Review of Economics*, 9, 201–234.

Kerr, Sari Pekkala and William R Kerr (2018) "Global Collaborative Patents," *The Economic Journal*, 128 (612), F235–F272.

Kerr, William R (2008) "Ethnic Scientific Communities and International Technology Diffusion," *The Review of Economics and Statistics*, 90 (3), 518–537.

Lerner, Josh and Amit Seru (2021) "The Use and Misuse of Patent Data: Issues for Finance and Beyond," *The Review of Financial Studies*, 35 (6), 2667–2704.

Lerner, Josh, Morten Sorensen, and Per Strömberg (2011) "Private Equity and Long-Run Investment: The Case of Innovation," *The Journal of Finance*, 66 (2), 445–477.

Lucas Jr, Robert E (2009) "Ideas and Growth," *Economica*, 76 (301), 1–19.

Lucas Jr, Robert E and Benjamin Moll (2014) "Knowledge Growth and the Allocation of Time," *Journal of Political Economy*, 122 (1), 1–51.

Mailath, George J and Andrew Postlewaite (1990) "Workers versus Firms: Bargaining over a Firm's Value," *The Review of Economic Studies*, 57 (3), 369–380.

McPherson, Miller, Lynn Smith-Lovin, and James M Cook (2001) "Birds of a Feather: Homophily in Social Networks," *Annual review of sociology*, 27 (1), 415–444.

Moretti, Enrico (2021) "The Effect of High-Tech Clusters on the Productivity of Top Inventors," *American Economic Review*, 111 (10), 3328–3375.

Moser, Petra, Alessandra Voena, and Fabian Waldinger (2014) "German Jewish Émigrés and US Invention," *American Economic Review*, 104 (10), 3222–3255.

Prato, Marta (2022) "The Global Race for Talent: Brain Drain, Knowledge Transfer, and Growth," https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4287268.

Roth, Jonathan (2022) "Pretest with Caution: Event-Study Estimates after Testing for Parallel Trends," *American Economic Review: Insights*, 4 (3), 305–322.

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

Stein, Jeremy C (2008) "Conversations among Competitors," *American Economic Review*, 98 (5), 2150–2162.

Trouvain, Florian (2022) "Technology Adoption, Innovation, and Inequality in a Global World," https://florian-trouvain.com/jmp_florian-trouvain.pdf.

U.S. Patent and Trademark Office, "Data Download Tables," https://patentsview.org/download/data-download-tables, Accessed 2022-10-22.

Waldinger, Fabian (2012) "Peer Effects in Science: Evidence from the Dismissal of Scientists in Nazi Germany," *The Review of Economic Studies*, 79 (2), 838–861.

Wuchty, Stefan, Benjamin F Jones, and Brian Uzzi (2007) "The Increasing Dominance of Teams in Production of Knowledge," *Science*, 316 (5827), 1036–1039.

Zacchia, Paolo (2018) "Benefiting Colleagues but not the City: Localized Effects from the Relocation of Superstar Inventors," *Research Policy*, 47 (5), 992–1005.

——— (2020) "Knowledge Spillovers through Networks of Scientists," *The Review of Economic Studies*, 87 (4), 1989–2018.

Appendix

A Proofs

Proof of Proposition 1. The difference in the output over the two periods, ΔY_i is given by:

$$\begin{aligned}\Delta Y_i &= Y_i^2 - Y_i^1 \\ &= \left\{ I_i\left(2\right) - I_i\left(1\right) + \lambda \sum_{l=1}^{n} w_{il} \left\{ g_{il}\left(2\right)\left[\alpha_l + I_l\left(2\right)\right] - g_{il}\left(1\right)\left[\alpha_l + I_l\left(1\right)\right] \right\} \right\} \\ &= I_i\left(2\right) - I_i\left(1\right) + \lambda w_{ij}\left[g_{ij}\left(2\right)\left(\alpha_j + I_j\left(2\right)\right) - g_{ij}\left(1\right)\left(\alpha_j + I_j\left(1\right)\right)\right] \qquad (6)\end{aligned}$$

where the third equality follows from the assumption that the relocation of inventor j affects inventor i only in a direct way.

When inventor i and inventor j cut all types of their links – $g_{ij}\left(2\right) = 0$ and $s_{ij}\left(1\right) = 0$ – inventors i does not acquire any information in the second period and therefore $I_i\left(2\right) = I_i\left(1\right)$. Therefore, equation (6) becomes

$$\Delta Y_i = -\lambda w_{ij}\left(\alpha_j + I_i\left(1\right)\right) < 0$$

On the other hand, as long as $s_{ij} = 1$, and if the information acquired in the second period is high enough, such that

$$I_i\left(2\right) - I_i\left(1\right) \geq \lambda w_{ij}\left(\alpha_j + I_i\left(1\right)\right)$$

it follows that $\Delta Y_i \geq 0$. □

Proof of Proposition 2. Under the assumptions made in this propositions, the effect on the total output of inventor i is given by

$$\Delta Y_i\left(N_j\left(2\right)\right) = I_i\left(2\right) - I_i\left(1\right)$$

and therefore, Under the assumptions made in this propositions, the effect on the total output of inventor i is given by

$$\Delta Y_i\left(N_j\left(2\right)\right) - \Delta Y_i\left(\tilde{N}_j\left(2\right)\right) = \left[I_i\left(2\right) - I_i\left(1\right)\right] - \left[\tilde{I}_i\left(2\right) - \tilde{I}_i\left(1\right)\right] = I_i\left(2\right) - \tilde{I}_i\left(2\right)$$

where the last equality follows since the information gathered in the first period is equal across these two scenarios by definition.

Now, if information sharing between inventor j and inventor i does not take place, then $s_{ij}(2) = 0$ in both these cases, and no information is acquired in the second period, regardless of connections inventor j's forms. Hence,

$$\Delta Y_i(N_j(2)) - \Delta Y_i\left(\tilde{N}_j(2)\right) = 0$$

If, on the other hand, the inventors engage in information sharing, then $s_{ij}(2) = 1$, and

$$\begin{aligned} I_i(2) - \tilde{I}_i(2) &= \sum_{m \in N_j(2) \setminus N_i(1)} \left[\bar{w}_{im}(2) - \bar{w}_{im}(1)\right] k_m - \sum_{m \in \tilde{N}_j(2) \setminus N_i(1)} \left[\bar{w}_{im}(2) - \bar{w}_{im}(1)\right] k_m \\ &= \sum_{m \in N_j(2) \setminus \left(\tilde{N}_j(2) \cup N_i(1)\right)} \left[\bar{w}_{im}(2) - \bar{w}_{im}(1)\right] k_m \end{aligned}$$

Where the first equality follows by the assumption that inventor i's direct connections do not change (besides potentially that with inventor j), and Assumption 1 which implies that as long as the weight placed on the path leading from inventor i to their indirect collaborators did not increase, inventor i does not acquire information through them. This implicitly means that even if inventor j creates a direct link to one of inventor i's direct connection, further information is not acquired through that link.

And as long as $N_j(2) \setminus \left(\tilde{N}_j(2) \cup N_i(1)\right) \neq \varnothing$, the effect on inventor i's output is positive:

$$I_i(2) - \tilde{I}_i(2) > 0$$

$\square$

Proof of Proposition 3. In this case, the innate ability of inventors i and j changes, but the information they acquire does not depend on the set of inventors they interact with after the move and

$$I_i(2) = \tilde{I}_j(2)$$

Therefore,

$$\Delta Y_i(N_j(2)) = \{[\alpha_i(2) - \alpha_i(1)] + \lambda w_{ij}[g_{ij}(2)\,\alpha_j(2) - g_{ij}(1)\,\alpha_j(1)]\} = \Delta Y_i\left(\tilde{N}_j(2)\right)$$

$\square$

Proof of Proposition 4. By definition, the information acquired by inventor i in both periods, depends on the strength of the relationship between inventor i and inventor j. On the other hand, the information acquired by inventor j depends on w_{ji} that can be different than w_{ij}

and hence does not change, by assumption. Therefore,

$$\Delta Y_i\left(w_{ij}\right) = I_i\left(2\right) - I_i\left(1\right) + \lambda w_{ij} \underbrace{\left\{g_{ij}\left(2\right)\left[\alpha_j + I_j\left(2\right)\right] - g_{ij}\left(1\right)\left(\alpha_j + I_j\left(1\right)\right]\right\}}_{X}$$

Similarly,

$$\Delta Y_i\left(\tilde{w}_{ij}\right) = \tilde{I}_i\left(2\right) - \tilde{I}_i\left(1\right) + \lambda \tilde{w}_{ij}\left\{g_{ij}\left(2\right)\left[\alpha_j + I_j\left(2\right)\right] - g_{ij}\left(1\right)\left(\alpha_j + I_j\left(1\right)\right]\right\} = \tilde{I}_i\left(2\right) - \tilde{I}_i\left(1\right) + \lambda \tilde{w}_{ij} X$$

Now, following equation (3) and the assumption that inventor j is the only inventors moving in the network, the new information can only be acquired through new connections inventor j forms, and as long as they are (and were) not directly connected to inventor i. Thus,

$$I_i\left(2\right) = I_i\left(1\right) + \sum_{l=1}^{n} s_{ij}\left(2\right) \cdot s_{jl}\left(2\right) \cdot \left(1 - s_{il}\left(1\right)\right) \cdot \left(1 - s_{jl}\left(1\right)\right) w_{ij} \cdot w_{jl}$$
$$\tilde{I}_i\left(2\right) = \tilde{I}_i\left(1\right) + \sum_{l=1}^{n} s_{ij}\left(2\right) \cdot s_{jl}\left(2\right) \cdot \left(1 - s_{il}\left(1\right)\right) \cdot \left(1 - s_{jl}\left(1\right)\right) \tilde{w}_{ij} \cdot w_{jl}$$

And,

$$I_i\left(2\right) - I_i\left(1\right) = w_{ij} \underbrace{\sum_{l=1}^{n} s_{ij}\left(2\right) \cdot s_{jl}\left(2\right) \cdot \left(1 - s_{il}\left(1\right)\right) \cdot \left(1 - s_{jl}\left(1\right)\right) w_{jl}}_{Z}$$

And similarly for $\tilde{I}_i\left(2\right) - \tilde{I}_i\left(1\right)$.

This means we can express the effect on inventor i output as:

$$\Delta Y_i\left(w_{ij}\right) = w_{ij}\left\{\sum_{l=1}^{n} s_{ij}\left(2\right) \cdot s_{jl}\left(2\right) \cdot \left(1 - s_{il}\left(1\right)\right) \cdot \left(1 - s_{jl}\left(1\right)\right) w_{jl} + \lambda X\right\}$$

Therefore,

$$\begin{aligned}\left|\Delta Y_i\left(w_{ij}\right)\right| - \left|\Delta Y_i\left(\tilde{w}_{ij}\right)\right| &= w_{ij}\left|Z + \lambda X\right| - \tilde{w}_{ij}\left|Z + \lambda X\right| \\ &= \left(w_{ij} - \tilde{w}_{ij}\right)\left|Z + \lambda X\right| \geq 0\end{aligned}$$

with a strict inequality as long as the benefit from the information gained through inventor j and the cost of discontinuing co-patenting with them do not exactly cancel out. In the case where inventor i and inventor j cut all of their links (informational and co-patenting) or when they continue collaborating, the difference will also be strictly positive. □

B Model Extensions

In this section, I present the model without the assumption imposed on the distance the information can travel. I spell the whole model again, although some of the parts did not change to maintain some level of continuity.

The main goal of these extension is to emphasize that the predictions do not depend on this assumption, and that everything else follows.

B.1 Basic Framework

B.1.1 Inventors' Network

A society of n inventors is connected via a directed and weighted network, which has an adjacency matrix $\mathbf{W} \in [0,1]^{n \times n}$. A general element $w_{ij} \in [0,1]$ represents the status and the strength of the relationship between inventor i and inventor j, where a higher w_{ij} is associated with a stronger connection.[39] Specifically, an entry $w_{ij} = 0$ implies that inventor i and inventor j do not collaborate, and therefore are not connected.

Inventors are also endowed with an innate ability level $\alpha_i \in \mathbb{R}_+$ and with a knowledge level $k_i \in \mathbb{R}_+$. These concepts play a central role in patent production and information sharing.

B.1.2 Information Acquisition

Inventors acquire information through their network, not just from their immediate connections, but also from inventors located farther away. To formalize this concept, it is helpful to introduce a notion that measures the distance between any to inventors in the network. Let d_{ij} be the length of the shortest path between inventor i and inventor j in the network $\mathbf{W}$. This distance metric signifies the smallest number of connected inventors forming a sequence that establishes an indirect link between inventors i and j. Formally, a path of length m between inventor i and inventor j is an ordered set $M = \{i_1, i_2, ..., i_{m+1}\}$ such that $w_{i_l i_{l+1}} > 0$ for all $l \in \{1, ..., m\}$, $i_1 = i$ and $i_{m+1} = j$. The length of the shortest path between inventor i and j is the minimal m that satisfies these conditions. The weight of this path w_{ij}^M is given by multiplying the weights of the links composing this path, which can be expressed as

$$w_{ij}^M = \Pi_{l=1}^m w_{i_l i_{l+1}}$$

[39]One way to interpret the weights of the form w_{ij} is through the eyes of patents production relationship. In this case, a higher w_{ij} corresponds to a higher number of patents both inventor i and inventor j are jointly listed on.

The total information held by inventor i is a result of a combination between their initial knowledge and what they acquire through interactions with other inventors. Let $D \in \{1, ..., n-1\}$ be a bound on the distance information can travel on the network.[40] With this concept and the notation introduced earlier, the total information held by inventor i is given by:

$$I_i = k_i + \sum_{d=1}^{D} \sum_{\{j:\, d_{ji}=d\}} \bar{w}_{ji} k_j \tag{7}$$

where $\bar{w}_{ji}$ corresponds to the cumulative weight of the shortest paths connecting j and i, which also meet the condition of being of a minimal distance d.[41] These weights represent the strengths of the paths, and they capture the idea that the information inventor i acquires through inventor j is proportional to a measure of dependence between them.

Figure B.1 provides two examples that illustrate the bound D and the cumulative weight of the paths between inventor i and inventor j, $\bar{w}_{ij}$. The weights above the links represent the clockwise connections, while those below correspond to the counter-clockwise links.[42] In panel (a), the bound on the distance between inventors for information acquisition purposes is equal to one. In this scenario, inventor 1 gains information from inventors 2 and 3, and this information is scaled by $w_{21} = 0.25$ and $w_{31} = 0.25$, respectively. In panel (b), inventor 1 gains information from inventors 2, 3, 4 and 5, and the information acquired through inventor 5 is scaled by 0.5.

[40] This meant to express the idea that when two inventors are positioned far apart within the network, they are less likely to share information with each other. The parameter D specifies what "far apart" means in this context.

[41] When inventors i and j are directly connected the cumulative weight $\bar{w}_{ij}$ is equal to the strength of the link between inventors i and j, w_{ij}. When the shortest path between inventors i and j is of length two, $\bar{w}_{ij}$ is equal to the sum of the weights of the form $w_{il} \cdot w_{lj}$ where inventor l is directly connected to both inventor i and inventor j, but inventors i and j are not connected to each other. Formally, $\bar{w}_{ij} = \sum_{l=1}^{n} w_{il} w_{ij} \cdot \mathbb{1}\{w_{ij} = 0\}$. The formula for longer paths follows the same reasoning, but it becomes more cumbersome as the minimal distance d increases, since there are more conditions to verify a path is indeed the shortest.

[42] For example, in panel (a), the weight of the link going from inventor j to inventor i, denoted as w_{ji}, is 0.5, whereas the weight of the link from inventor i to inventor j, represented as w_{ij}, is 0.25.

Figure B.1: Information Sharing on Network Example

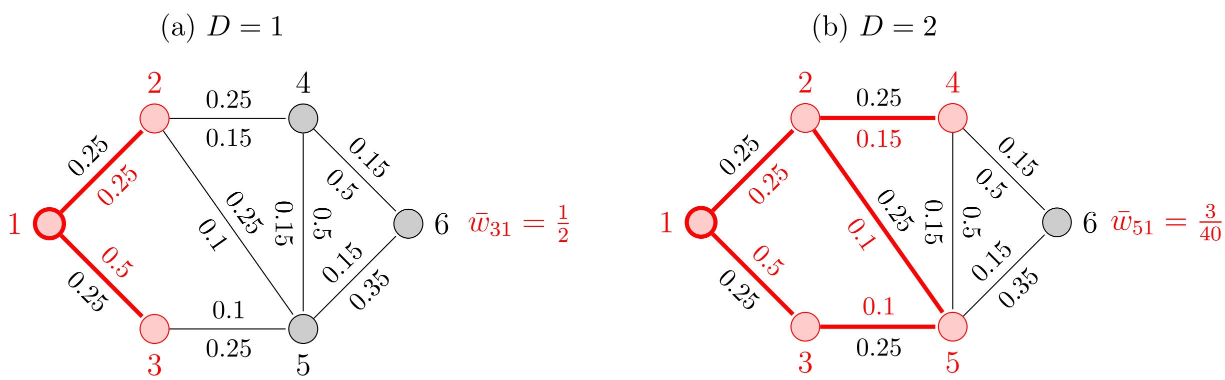

B.1.3 Patent Production Function

Inventors produce patents in teams. Each inventor's total output relies on their individual output, which is determined by their innate ability and total information they hold, as well as the output contributed by their collaborators.

Inventor i's total output is, therefore, given by

$$Y_i = y_i + \sum_{j=1}^{n} w_{ji} y_j \tag{8}$$

$$y_i = \alpha_i + I_i$$

where y_i is inventor i's individual output.

This production function reflects the substitution between the different sources that drive output production. It emphasizes the tradeoff between the information accessed through the network (and one's fixed innate ability) and the direct benefit one gains from co-patenting, which comes in the form of the output contributed by their collaborators. This tradeoff will be the main focus of the next subsection.

B.2 Two Period Model

In general, patents are produce in various geographical locations. As a consequence, inventors may move around. In this section, I study the effect of a relocation on the total output of the mover's former collaborators in the eyes of the model. Specifically, the magnitude and direction of this effect will be contingent upon the specifics of the connections between these inventors, and on how their network changes in response to the relocation.

To start with, consider two time periods and two geographical locations. Let $t = 1$ represent the time before any relocation occurs, and $t = 2$ reflect the time after the relocations have take place, and assume that innate abilities α_i and the weighted adjacency matrix $\mathbf{W}$ are fixed across periods. Since the collaboration network can undergo changes between the two time period, I introduce new notations which capture the state of the network in each one of these periods. Denote by $\mathbf{G}(t) \in \{0,1\}^{n\times n}$ the undirected adjacency matrix at time t. The ij-th entry is the collaboration status between inventor i and inventor j at time t, with the entry equals one if they co-patent, and zero otherwise. This relationship is reciprocal. Additionally, let $\mathbf{S}(t) \in \{0,1\}^{n\times n}$ be the undirected information exchange network at time t. This is a symmetric matrix whose entries equal to one whenever the inventors are engaged in information sharing. In the initial period, information sharing occurs exclusively when inventors co-patent. However, after the relocation, in period $t = 2$, inventors who previously co-patented in $t = 1$ may cease their collaboration in period $t = 2$, and yet still engage in information sharing. Formally,

$$\begin{aligned} g_{ij}(1) \quad &\iff \quad s_{ij}(1) = 1 \qquad\qquad \forall i,j \\ g_{ij}(2) \quad &\implies \quad s_{ij}(2) = 1 \qquad\qquad \forall i,j \end{aligned}$$

Lastly, $I_i(t)$ is the level of the information held by inventor i at time t, where the paths are now measured on the network $\mathbf{S}(t)$.[43]

B.2.1 Information Acquisition

To accommodate some degree of continuity across the two periods, I assume that

Assumption 2. The information inventors acquired in the first period cannot be forgotten, and therefore is not subject to relearning.

This implies that inventor i begins period $t = 2$ with knowledge that is equal to $I_i(1)$, rather than k_i, as it is at the beginning of period $t = 1$.

The idea behind this assumption is that once techniques and ideas are acquired, they can't be unlearned. Once learned, inventors can use them again without relearning.

In particular, equation (1) becomes

$$\begin{aligned} I_i(1) &= k_i + \sum_{d=1}^{D} \sum_{\{j:\, d_{ji}(1)=d\}} \bar{w}_{ji}(1)\, k_j && \forall i \\ I_i(2) &= I_i(1) + \sum_{\{j:\, d_{ji}(2)=d\}} \mathbb{1}\{\bar{w}_{ji}(2) > \bar{w}_{ji}(1)] \cdot [\bar{w}_{ji}(2) - \bar{w}_{ji}(1)]\, k_j && \forall i \qquad (9) \end{aligned}$$

[43]In the first period, the elements of the matrix $\mathbf{G}(1)$ and $\mathbf{S}(1)$ are equivalent to the indicators $\mathbb{1}\{w_{ij} > 0\}$. Therefore, the paths on $\mathbf{W}$, $\mathbf{G}(1)$ and on $\mathbf{S}(1)$ are the same. However, in the second period, since information exchange can take place even when then inventors do not collaborate, it does not necessarily hold.

where $d_{ji}(t)$ denotes the minimal distance on the matrix $\mathbf{S}(t)$ and $\bar{w}_{ji}(t)$ corresponds to the cumulative paths weights on matrix $\mathbf{S}(t)$.[44] The multiplication by the elements $\mathbb{1}\{\bar{w}_{ji}(2) > \bar{w}_{ji}(1)\}$ imposes the restriction that information can only be acquired in the second period, and cannot be forgotten.

B.2.2 Patent Production

Production in both periods follows the same reasoning as in equation (2), with total output depending on both individual and scaled collaborative outputs. That is,

$$\begin{aligned} Y_i(t) &= y_i(t) + \sum_{j=1}^{n} g_{ji}(t)\, w_{ji} y_j && \forall i \\ y_i(t) &= \alpha_i + I_i(t) && \forall i \end{aligned}$$

[44] Although the matrix $\mathbf{W}$ is fixed across the two time periods, the potential differences in the matrices $\mathbf{S}(1)$ and $\mathbf{S}(t)$ can lead to differences in $\bar{w}_{ji}(t)$ across the time periods, if new paths are formed and/or old paths are severed.

C Data

C.1 Description of Patent Data

The information about the patents I use is from the USPTO PatentsView. Besides supplying information about inventors and patents, they also conduct a disambiguation procedure where each inventor gets a unique identifier number, and all of their patents and the information provided at the time of the application are attached to it. This is not trivial as inventors might use different names at the time of the application and are not necessarily patenting under the same assignee or in the same location.

The data I use is patent applications for patents which were eventually granted between 1976 and 2022. The data is provided in a TSV format, where each inventor, location and patent has a unique identifier. Using these identifiers, I can merge all the information given in an application about a specific inventor. In particular, the inventor's name, the names of the other inventors who are listed with them on the patent, the date of the application as well as the date at which the patent was granted, the number of citations the patent received and which patents cited it, the citations granted by that patent, the CPC classes of the patents, the residential address of the inventor and the assignee and the headquarters's address.

C.2 Construction of the Sample

I restrict the sample of the USPTO patent application to include only inventors whose first patent was applied for on 1990 or after. The reason is that online professional profiles were introduced in 2008, and older people, mainly ones who have already retired or are close to retirement, are less likely to have an account due to its purpose being a device that makes it easier to learn about individuals' work history and potentially ease hiring and recruitment. That way I can ensure, with a higher likelihood, that the linking rates are not biased by the probability of opening an account.

I also exclude a very small number of patent numbers who were "withdrawn," which can be found in https://www.uspto.gov/patents/search/withdrawn-patent-numbers. Patent numbers are assigneed before the patents are granted. If between the date at which the patent number was assigned and the date of the issuance of the patent some information that indicates that the application is not ready to be issued is reveled, the patent is not granted at that date and the number will never be used again.[45] It is important to note, however, that the application can still be issued, and in this case, it will be issued under a different patent number. Therefore, dropping the withdrawn patent numbers have two

[45]This may happen if, for example, the fees were not paid in full or if some prior art was found.

purposes. The first, is to avoid including ideas that were not patented after all. And second, to avoid double counting patents, in cases where they were issued under different patent numbers.

Another point is about how to assign CPC classes to patents. When patents are granted and published, more than one CPC class is usually assigned. Following the general method in the literature, the CPC class I assign to the patent is the one listed first in the sequence.

C.3 Linking Algorithm to Construct the Dataset

Both the USPTO data and the data I receive through Revelio Labs have different advantages. While the USPTO dataset provides me with information about the patenting activity of inventors, and their physical location at the time of the application, the information that includes data from online professional profiles, adds to this demographics, work locations and the company workers are employed by at each point in time. Only the combination of both these datasets will allow me to construct the panel that I need in order to study the effect of inventor mobility on their co-inventor productivity.

I leverage as much information as I can from both these dataset to construct the linking in a way that minimizes mistakes. There are two types of possible mistakes. The first is about not linking an inventor that should have otherwise be linked. These mistakes, although costly in term of number of observations, are less likely to bias the data. The type of mistakes that I try most to avoid are the ones around linking an inventor and a user that should not be linked. And my methodology, focuses mainly on the second, although I try to account for both.

The linking is performed on three identifiers. The first is the inventor and the user names. As a first step I only consider the first and last name of the inventor and the user. An inventor is linked to a user if and only if their first and last name match exactly. This implies that there are some inventors I cannot match due to usage of different names on their patent application and their online professional profile. An example is an inventors who uses their full name on their patent and their nickname on their online professional profile.[46] The second element I link on is the state the inventor lived in at the time of the application, and the state the user lists as the location of their workplace at a six-month window around the application data.[47] Linking on the state level can create issues involving the difference between workplace location and residential address. At the application stage inventors are asked to provide their residential address. That is, they report their home

[46]Some of the common names and nicknames are "Robert" and "Bob" or "William" and "Bill."

[47]Some users take time to update their online professional profiles with their new employers. In order to allow for some flexibility, I allow the state to match even in an earlier or a later stage.

address. On online professional profiles, users provide the state at which their work takes place. That is, the location of their office. When inventors live and work in different states, they will not be linked to their user. The last characteristics I link inventors and users on is their employer. I begin with standardizing the USPTO assignees and the online professional profiles companies following the standardizing method in the NBER's Patent Data Project (https://sites.google.com/ site/patentdataproject/). For example, "AMAZON" and "AMAZON USA" are the same company. I also do not use any abbreviations for companies. That is, whenever a company is known to be abbreviated, such as in "IBM," I change all of its instances to "INTERNATIONAL BUSINESS MACHINES" to avoid mismatched for that reason. I, then, use a fuzzy match, linking between the standardize assignee on the inventor's application and the standardize user's employer at the time of the application. More precisely, I use Jaro-Winkler distance measure and set the threshold for match at 0.99. This is a quite high threshold, ensuring that I so not mismatch between different companies and assignees. Finally, I include only inventors who are linked to exactly one user, and I use middle initials to break ties.

There are at least two issues which follow from the last step. The first is that the assignee is not necessarily the employer of all the inventors on the patent. It is usually the case that at least one inventor is employed by the company listed as an assignee, but it is not necessarily the case that all of the inventors are.[48] In these cases, the inventor and the user will not be matched. The second issue might arise in cases where mergers are acquisitions are involved. In these cases users in online professional profiles, might update their user to include the information about the acquiring firm rather than the name of the firm they were employed in prior to the acquisition. If at the time of the application the merger or acquisition has not taken place yet, there will be inconsistency between the information on the online professional profile and the information on their USPTO application.

Characteristics of Linked and Unlined Inventors

Table C.1 presents the summary statistics of the linked and the full samples. It shows that the linked inventors tend to be more productive, and that is due to the fact that linking is easier the more observations there are as it increases the accuracy.

Adding to the information in Table C.1 the BEA's "economic areas" with the lowest linking rates are Great Falls MT, Aberdeen SD, Cape Girardeau-Jackson MO-IL, Lewiston ID-WA and Panama City-Lynn Haven FL. And the BEA's "economic areas" with the highest linking rates are Seattle-Tacoma-Olympia WA, San Jose-San Francisco-Oakland CA, San Diego-Carlsbad-San Marcos CA, Milwaukee-Racine-Waukesha WI, Minneapolis-St. Paul-St.

[48] There are also cases where more than once assignee is listed, and in this case I check all the combinations.

Table C.1: Summary Statistics on Full and Linked Samples

Panel A: Full Sample

	Mean	Median	Std. Dev	# Obs.
Year of First Patent	2004	2004	8	1,150,368
Total Number of Patents	4.39	2.00	7.06	1,150,368
Total Citations Stock	62.80	11.00	159.30	1,150,368
Total Adjusted Citations Stock	3.19	0.96	6.24	1,150,368
Average Team Size	2.75	2.25	1.85	1,150,368
Male	0.85	1.00	0.35	1,048,732

Panel B: Linked Sample

	Mean	Median	Std. Dev	# Obs.
Year of First Patent	2007	2009	9	229,290
Total Number of Patents	6.41	3.00	9.18	229,290
Total Citations Stock	78.82	8.00	200.24	229,290
Total Adjusted Citations Stock	4.34	1.26	7.75	229,290
Average Team Size	3.01	2.60	1.85	229,290
Male	0.85	1.00	0.35	207,203

Notes: The table provides summary statistics for two datasets. Panel A presents the attributes of inventors based in the United States within the patent data, encompassing inventors who initiated patenting activities for the first time after 1990 and never indicated a residential address outside of the United States. Panel B offers summary statistics for the inventors successfully matched to the Revelio Labs data. All variables represent cumulative values over the entire observed period. The Average Team Size is computed across all patents, with a solo patent being considered as a team size of one.

Cloud MN-WI, Raleigh-Durham-Cary NC and Boston-Worcester-Manchester, MA-NH. The latter are more associated with patenting activity.

Moreover, as anticipated, younger inventors have more incentives and are more likely to setup an online account. For that reason, the years of the first patent which is associated with the lowest linking rates are 1990-2004, where the linking rates are monotonically decreasing with the years. And the highest linking rates correspond to inventors patenting for the first time in the years 2013-2020. Where, in this case, the linking rates is monotonically increasing in the years.

The last characteristic I looked at in the CPC classes the linked and and unlinked inventors patent in. The lowest rated of linking are associated with the following CPC classes:

Table C.2: Lowest Linking Rates CPC Classes

CPC Class	Title
B43	WRITING OR DRAWING IMPLEMENTS; BUREAU ACCESSORIES
B44	DECORATIVE ARTS
B63	SHIPS OR OTHER WATERBORNE VESSELS; RELATED EQUIPMENT
B68	SADDLERY; UPHOLSTERY
A45	HAND OR TRAVELLING ARTICLES
A46	BRUSHWARE

Table C.3: Highest Linking Rates CPC Classes

CPC Class	Title
B06	GENERATING OR TRANSMITTING MECHANICAL VIBRATIONS IN GENERAL
B33	ADDITIVE MANUFACTURING TECHNOLOGY
C07	ORGANIC CHEMISTRY
F15	FLUID-PRESSURE ACTUATORS; HYDRAULICS OR PNEUMATICS IN GENERAL
G06	COMPUTING; CALCULATING OR COUNTING
G16	INFORMATION AND COMMUNICATION TECHNOLOGY [ICT] SPECIALLY ADAPTED FOR SPECIFIC APPLICATION FIELDS
H03	ELECTRONIC CIRCUITRY
H04	ELECTRIC COMMUNICATION TECHNIQUE

Note that these CPC classes are aligned with the locations of the linked or unlinked inventors, respectively. To put differently, the main CPC classes in these areas match the CPC classes with the highest or lowest linking rates, respectively.

D Additional Figures and Tables

Figure D.1: Bureau of Economic Analysis' "Economic Areas" Map

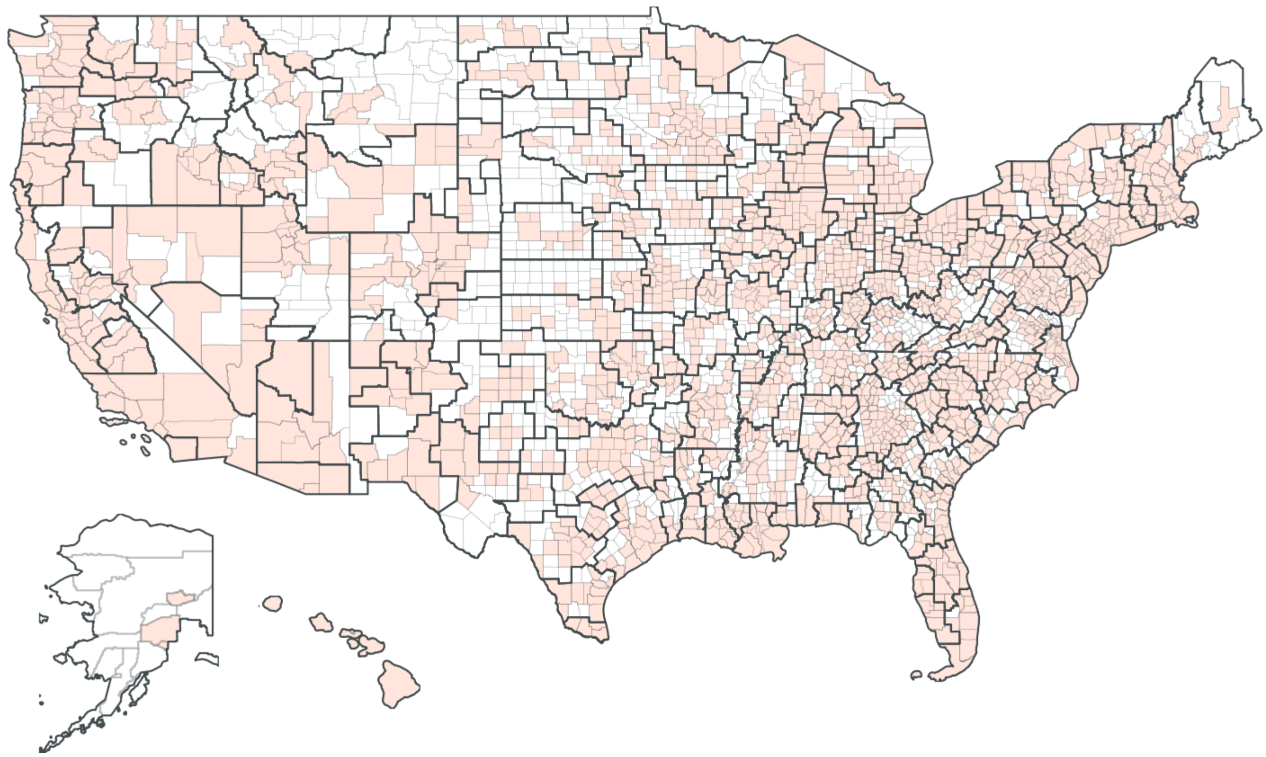

Notes: This figure presents the BEA's economic areas map and compares them the MSA's, painted in pink. One can see that the size of the economic areas changes across locations, and it may coincide with the corresponding MSA, but it can also be larger.

Table D.1: Effect Size and Within Firm vs. Across Firms Move

	Within Firm		Across Firms	
	(1) Annual Number of Patents	(2) Annual Number of Adjusted Citations	(3) Annual Number of Patents	(4) Annual Number of Adjusted Citations
$PostMove^{\mathrm{Real}}$	0.076***	0.085***	0.061**	0.055**
	(0.021)	(0.021)	(0.029)	(0.026)
Control Post Mean	0.501	0.336	0.498	0.348
Percentage Change	+15.08%	+25.23%	+12.18%	+15.78%
P-Value H_0: Diff. = 0			0.67	0.36
Observations	223955	223955	331860	331860
Individual FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Experience FE	Yes	Yes	Yes	Yes

Notes: The information presented in this table presents the results from regression equation (5) on two different samples. Columns (1) and (2) correspond to cases where the mover moves within the same firm, and columns (3) and (4) cover the cases where the mover moves across firms. The unit of analysis in these regressions is still inventor-year. The dependent variable in columns (1) and (3) is the number of patents per year and in columns (2) and (4) is the number of adjusted citations per year, as defined in Section 3. Standard errors are clustered at the mover level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table D.2: Effect Size and the Characteristics of Across Firm Move

	Corp.-Corp.		Non-Corp.-Corp.	
	(1) Annual Number of Patents	(2) Annual Number of Adjusted Citations	(3) Annual Number of Patents	(4) Annual Number of Adjusted Citations
$PostMove^{\mathrm{Real}}$	0.071***	0.047*	-0.015	-0.015
	(0.027)	(0.027)	(0.088)	(0.066)
Control Post Mean	0.441	0.326	0.517	0.345
Percentage Change	+16.15%	+14.29%	-2.95%	-4.42%
Observations	315511	315511	52956	52956
Individual FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Experience FE	Yes	Yes	Yes	Yes

Notes: The information presented in this table presents the results from regression equation (5) on two different samples. Columns (1) and (2) correspond to cases where the mover and the stayer both work in different corporation, and columns (3) and (4) cover the cases where the mover moves into a corporation from a non-corporation and the stayer remains in the non-corporation. The unit of analysis in these regressions is still inventor-year. The dependent variable in columns (1) and (3) is the number of patents per year and in columns (2) and (4) is the number of adjusted citations per year, as defined in Section 3. Standard errors are clustered at the mover level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table D.3: Heterogeneity based on Good vs. Bad Moves

	Bad Move		Good Move	
	(1) Annual Number of Patents	(2) Annual Number of Adjusted Citations	(3) Annual Number of Patents	(4) Annual Number of Adjusted Citations
$PostMove^{\text{Real}}$	-0.008	0.019	0.178***	0.128***
	(0.017)	(0.015)	(0.044)	(0.042)
Control Post Mean	0.442	0.289	0.639	0.464
Percentage Change	+-1.82%	+6.59%	+27.88%	+27.49%
P-Value H_0: Diff. = 0			00	0.01
Observations	149064	149064	406751	406751
Individual FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Experience FE	Yes	Yes	Yes	Yes

Notes: The information presented in this table presents the results from regression equation (5) applied to two distinct subsets of data. Specifically, Columns (1) and (2) pertain to scenarios where the mover relocates to a location where the density of inventors who patent in the same CPC class is higher, and columns (3) and (4) delve into the opposite case. The unit of analysis remains inventor-year in these regression analyses. Columns (1) and (3) encompass the number of patents per year as the dependent variable, while columns (2) and (4) utilize the number of adjusted citations per year, as elaborated in Section 3. Standard errors are clustered at the mover level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table D.4: The Effect of a Relocation and Race Differences

	Same Race		Different Races	
	(1) Annual Number of Patents	(2) Annual Number of Adjusted Citations	(3) Annual Number of Patents	(4) Annual Number of Adjusted Citations
$PostMove^{\text{Real}}$	0.049***	0.048***	0.038	0.057
	(0.018)	(0.017)	(0.043)	(0.039)
Control Post Mean	0.499	0.348	0.503	0.324
Percentage Change	+9.80%	+13.9%	+7.49%	+17.65%
P-Value H_0: Diff. = 0			0.80	0.833
Observations	366726	366726	189089	189089
Individual FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Experience FE	Yes	Yes	Yes	Yes

Notes: The information presented in this table presents the results from regression equation (5) applied to two distinct subsets of data. Columns (1) and (2) address situations where the mover and the left behind inventor have the same race. Columns (3) and (4) report the results when the mover and the left behind inventors have different races. The unit of analysis remains inventor-year in these regression analyses. Columns (1) and (3) encompass the number of patents per year as the dependent variable, while columns (2) and (4) utilize the number of adjusted citations per year, as elaborated in Section 3. Standard errors are clustered at the mover level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table D.5: The Effect of Relocation on Share of Collaborators to the Destination Patent Based Measure

	(1) Annual Percentage of Collaborators in Destination	(2) Annual Percentage of Collaborators in Destination
	All	Only When Patenting
$PostMove^{\text{Real}}$	1.172***	5.291***
	(0.096)	(0.222)
Control Post Mean	1.523	7.216
Percentage Change	+76.94%	+73.32%
Observations	555815	142296
Individual FE	Yes	Yes
Year FE	Yes	Yes
Experience FE	Yes	Yes

Notes: The information presented in this table presents the results from regression equation (5). In column (1), I consider all the observations in the panel, whether or not the inventor patented in that year, while column (2) covers only the years when the inventor patents. The unit of analysis in these regressions is still inventor-year. The dependent variable is the share of collaborators in the destination location. Standard errors are clustered at the mover level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table D.6: The Effect of a Relocation and Gender Differences

Panel A: Same Sex

	Male-Male		Female-Female	
	(1) Annual Number of Patents	(2) Annual Number of Adjusted Citations	(3) Annual Number of Patents	(4) Annual Number of Adjusted Citations
$PostMove^{\text{Real}}$	0.042**	0.043**	-0.021	-0.054
	(0.018)	(0.018)	(0.056)	(0.068)
Control Post Mean	0.517	0.362	0.424	0.289
Percentage Change	+8.07%	+12%	-4.93%	-18.6%
P-Value H_0: Diff. = 0			0.28	0.16
Observations	343092	343092	14969	14969
Individual FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Experience FE	Yes	Yes	Yes	Yes

Panel B: Opposite Sex

	Male-Female		Female-Male	
	(1) Annual Number of Patents	(2) Annual Number of Adjusted Citations	(3) Annual Number of Patents	(4) Annual Number of Adjusted Citations
$PostMove^{\text{Real}}$	0.039	0.064	-0.052	-0.021
	(0.038)	(0.039)	(0.060)	(0.043)
Control Post Mean	0.395	0.252	0.583	0.349
Percentage Change	+9.94%	+25.5%	-8.87%	-6.09%
P-Value H_0: Diff. = 0	0.95	0.62	0.13	0.16
Observations	54593	54593	54696	54696
Individual FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Experience FE	Yes	Yes	Yes	Yes

Notes: The information presented in this table presents the results from regression equation (5) applied to four distinct subsets of data. Each subset corresponds to a different combination of the sexes of both the mover and the left behind inventor. Panel A addresses situations where the mover and the left behind inventor share the same sex. Specifically, Columns (1) and (2) pertain to scenarios where both the mover and the left behind are males, and columns (3) and (4) delve into cases where both the mover and the left behind are females. On the other hand, Panel B delves into scenarios where the sexes of the mover and the left behind inventor differ. Columns (1) and (2) within this panel represent cases where the mover is male and the left behind is female, while columns (3) and (4) encapsulate the reverse situation—where the mover is female and the left behind is male. The unit of analysis remains inventor-year in these regression analyses. Columns (1) and (3) encompass the number of patents per year as the dependent variable, while columns (2) and (4) utilize the number of adjusted citations per year, as elaborated in Section 3. Standard errors are clustered at the mover level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table D.7: The Effect of Relocation and Continued Collaboration

	Continued Collaboration with Mover		No Collaboration with Mover	
	(1) Annual Number of Patents	(2) Annual Number of Adjusted Citations	(3) Annual Number of Patents	(4) Annual Number of Adjusted Citations
$PostMove^{\text{Real}}$	0.055^{**}	0.041^{**}	0.003	0.059
	(0.023)	(0.020)	(0.038)	(0.038)
Control Post Mean	0.34	0.22	0.95	0.66
Percentage Change	+16.01%	+18.15%	+0.3%	+8.87%
P-Value H_0: Diff. = 0	0.243	0.677		
Observations	426153	426153	129662	129662
Individual FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Experience FE	Yes	Yes	Yes	Yes

Notes: The information presented in this table presents the results from regression equation (5) applied to two distinct subsets of data. The first subset corresponds to left behind inventors who continue to collaborate with their respective mover, while the second one corresponds to cases where they stop collaborating with one another after the move. Columns (1) and (2) pertain to scenarios where the left behind and the mover continue to collaborate, and columns (3) and (4) delve into cases where they do not. The outcome variable in columns (1) and (3) is the number of patents per year, while columns (2) and (4) utilize the number of adjusted citations per year, as elaborated in Section 3. Standard errors are clustered at the mover level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.