

Organizational Bridges and the Direction of Inventive Search*

Sophia Mo Jieying Zhang

September 2026

Very preliminary draft

Abstract

How do collaboration networks shape the direction of inventive search? We study organizational bridges: co-inventor ties that connect individual inventors to peers employed outside firms specialized in different fields. Using U.S. patent records matched to employer histories for four million inventors, we show that such bridges influence the propensity with which a given inventor files a first patent in a connected field. Because existing ties are influenced by the inventor’s own employment and innovation history, we exploit inventor displacements during mass layoffs, using an instrument built from which fields absorb other displaced inventors in the same labor market. Organizational bridges raise entry into the fields where the collaborator’s employer specializes, and they do so before the focal inventor has any employment contact with that firm. The employer’s specialization predicts entry conditional on the collaborator’s own patent portfolio. The induced entries persist and generate additional patents and citations in the connected field. Policies that alter worker mobility may therefore affect inventive activity beyond the firms directly involved, but mobility alone does not direct invention toward fields in which few inventors already specialize.

Keywords: Innovation, Co-Inventor Networks, Knowledge Diffusion, Labor Mobility, Technological Search

*Mo: Harvard University, yiningmo@fas.harvard.edu. Zhang: Harvard University, jieyingzhang@g.harvard.edu. All errors are our own.

1 Introduction

The direction of inventive effort shapes the direction of technical change and economic growth. However, redirecting that effort across technological fields is costly.¹ The knowledge required to work in an unfamiliar field is oftentimes tacit and difficult to acquire (Jones 2009; Myers 2020). Therefore, when a new technological opportunity arises, such as machine learning, the inventors who respond are mostly those already working in adjacent fields. The constraint on who can respond is less the visibility of the opportunity than access to the specialized, hard-to-codify knowledge needed to act on it. Understanding what relaxes this access constraint, and for whom the constraint can be relaxed, are therefore central to evaluating policies that aim to redirect innovative activities.

This paper studies a channel through which specialized organizations shape the direction of inventive search beyond their own boundaries, a channel we call organizational bridges. We investigate one specific type of organizational bridge, namely strong collaborative ties that connect an inventor to a peer employed at another firm specialized in a particular field. Since firms are not only producers of innovation but repositories of field-specific knowledge, including problem selection, tooling, and tacit know-how about what works, an inventor’s collaborators who work inside such firms can transmit access to that knowledge to her without a job change by that collaborator.

We address three questions regarding the transmission of firm knowledge on innovation to one’s collaborators. First, does having a strong collaborator inside a firm specializing in certain fields redirect an inventor’s search toward that firm’s fields, over and above the collaborator’s own expertise and the inventor’s own opportunities? Second, is this effect causal and arm’s-length? Specifically, when labor-market shocks redirect where collaborators work, does the resulting change in bridges shape an inventor’s search before she has any contact with the new firm through employment? Third, does bridge-mediated diffusion direct inventive effort towards sparse fields whose patents are highly cited, yet ones in which few inventors already specialize? To answer these questions, we first establish the positive relationship between an inventor’s likelihood to first innovate on a technology and the degree to which her collaborator’s firm specializes in that technology. We then exploit labor-market displacements that exogenously rewire where collaborators work to identify the causal effects of collaborators’ firm specialization on inventor activities. Finally, we investigate where the induced entries go.

¹The difficulty of staffing emerging technologies is a recurring concern in policy and industry. Firms compete for scarce specialists in fields such as artificial intelligence and semiconductors, and shortages of field-specific talent are routinely cited as a binding constraint on where innovative activity can expand (e.g., Bloom et al. 2019).

Our analysis begins by measuring the correlation between the technology an inventor chooses to work on and the specialty of her collaborator’s employer. To do so, we first assemble a panel of first-time field entries covering 8.5 million U.S. utility patents and 4.1 million inventors from PatentsView, matched to inventor employer histories from Revelio Labs. We define fields as CPC subclasses, construct rolling co-inventor networks from collaborations in a window lagged three to seven years, and measure, for each inventor at risk of first entry in a new field during a given year, whether a strong tie connects the inventor to a peer whose current employer, as revealed by its own patent portfolio, specializes in that field.

The descriptive association between an inventor’s decision to invent in a field and her strong ties’ employer specialization in that field is large and robust. An inventor is roughly three times more likely to take a first patent in a field connected by a strong cross-firm bridge than in an unconnected field, holding fixed the amount of her collaborators’ own patenting in the field. In regressions with inventor $\times$ field, field $\times$ year, and inventor $\times$ year fixed effects, the cross-firm bridge coefficient is 0.029. The effect survives after controlling for the inventor’s own employer’s specialization, while the analogous same-firm bridge collapses by nearly half, exactly as it should if same-firm bridges are largely a proxy for own-firm opportunity. When we include both the peer’s employer specialization and the peer’s own patent history, both remain strong predictors of entry. The coefficient on employer specialization is about sixty percent larger. Adding peer-community $\times$ field $\times$ year fixed effects to absorb technology shocks common to similar peer networks leaves a coefficient of 0.019.

Building on the descriptive evidence, we use displacement-induced changes in co-inventors’ employers to identify the causal effect of organizational bridges on inventive search. We link WARN mass-layoff notices to inventor employment histories and identify about 217,000 events in which a displaced worker is a strong prior co-inventor of a focal inventor and moves to a new employer. The move creates new connections to the fields in which the destination firm specializes and removes connections to the fields of the worker’s previous employer, while the focal inventor’s own employment remains unchanged.

We first use a stacked event-study design that compares entry into three sets of fields, including fields in which the co-inventor’s new employer specializes, fields in which the previous employer specializes, and fields associated with employers that hire other displaced inventors in the same year. Entry declines over event time in all three groups as the prior collaboration becomes more distant. However, after the co-inventor moves, this decline becomes significantly slower for fields associated with the new employer. The resulting change in trend is about 0.08 percentage points per year. The destination of the displaced co-inventor may itself be endogenous. To address this concern, we instrument for the fields of the destination

firm using the fields in which other displaced inventors in the same labor market are reemployed. We exclude workers from the same displaced firm and other displaced workers who are directly connected to the focal inventor, and we restrict the sample to focal inventors located in different states. The reduced-form estimate is 0.030, clustered by displacement market). A one-standard-deviation increase in the instrument raises the probability that the focal inventor enters the field within five years by 0.21 percentage points, relative to a mean entry probability of 0.60 percent. As an additional test, we randomly reassign inventor-co-inventor links within field-by-year cells while holding fixed the field composition of the instrument. The relationship disappears under these placebo assignments, suggesting that the result operates through the specific co-inventor relationship rather than through a common technology shock.

We next examine whether these effects reflect the focal inventor’s own employment relationship with the destination firm. Most of the effect occurs before the focal inventor has ever worked at the co-inventor’s new employer. The estimate also becomes slightly larger when we exclude inventors who had previously worked at that firm, increasing from 0.030 to 0.034. In contrast, the component corresponding to entry in the same year that the focal inventor joins the destination firm is small and statistically indistinguishable from zero. These results suggest that the displacement design captures an effect of the co-inventor’s change in employer rather than the focal inventor’s own mobility. This differs from the correlations in the standing network, which are substantially stronger when the focal inventor has her own employment connection to the bridging firm.

We then examine whether the relationship is explained by the co-inventor’s own expertise. Employer specialization continues to predict field entry after controlling for the co-inventor’s prior patenting in the field. The relationship is also stronger for co-inventors who have never themselves patented in the destination field, and it increases with the degree to which the employer specializes in that field. In addition, entry patents are substantially more likely to cite prior patents from the co-inventor’s employer than patents from other specialist firms to which the focal inventor has no connection. This difference remains when we restrict the comparison group to firms with similarly large patent portfolios and to firms that have never employed the focal inventor. These results suggest that the co-inventor’s employer contains information relevant for the focal inventor’s search beyond the co-inventor’s observed prior expertise. However, we do not observe the information exchanged between collaborators and therefore cannot fully distinguish knowledge associated with the organization from knowledge that the co-inventor acquires after joining it.

The resulting field entries also tend to persist. Using the same reduced-form design, we find that newly formed bridges increase not only first entry into the connected field, but also

subsequent patenting and forward citations in that field. Roughly seventy percent of the induced entries are followed by additional patents rather than ending after the first patent. These estimates show that the response extends beyond a one-time field entry. They do not establish whether total inventive output rises, since increased activity in the destination field may be offset by reduced activity in other fields.²

Finally, we study which fields receive new organizational bridges. We classify fields using inventor density and lagged citation rates, focusing in particular on fields with relatively few inventors but highly cited patents. These fields receive relatively few new bridges because few employers specializing in them hire displaced inventors. They account for 2.5 percent of the fields that focal inventors could potentially enter, but only 0.24 percent of realized specialist landings. Consistent with this pattern, the first stage of the displacement instrument is close to zero in these fields. Our estimates are not precise enough to determine whether a bridge, once formed, has a different effect in sparse and highly cited fields. Instead, the evidence shows that such bridges are much less likely to form. More generally, mobility-induced bridges tend to connect inventors to fields in which specialist employers already exist and are hiring. Increasing mobility may therefore increase the amount of knowledge diffusion without necessarily redirecting it toward fields with relatively little existing specialist capacity.

The paper makes three contributions. First, we contribute to research on employee mobility and knowledge transfer by emphasizing mobility as a source of *network rewiring*. Workers carry knowledge across firms and locations when they move (Almeida and Kogut 1999; Song et al. 2003; Rosenkopf and Almeida 2003; Hoisl 2007; Somaya et al. 2008; Palomeras and Melero 2010; Marx et al. 2009; Serafinelli 2019), and a related strategy literature studies when valuable employees move and how mobility affects human capital, competitive advantage, and spin-out formation (Agarwal et al. 2004; Campbell et al. 2012; Ganco et al. 2015; Mawdsley and Somaya 2016), with recent work showing that mobility also affects the productivity of the collaborators a mover leaves behind (Prato 2025). We study a different margin of the mobility process, specifically, what happens to the opportunity set of a worker who does not move when one of her existing collaborators changes organizations. A move changes the organizational knowledge attached to a pre-existing relationship and thereby redirects the fields that are accessible through the network. We therefore shift attention from the human capital carried by the mobile worker to the change in organizational access experienced by that worker’s prior collaborators. Empirically, displacement gives us variation in the organizational content of an existing relationship while holding the prior inventor pair fixed, which

²An event-level design that aggregates patenting across fields has a strong first stage and no evidence of pre-trends, but the sample of roughly seven thousand events is too small to distinguish precisely between net creation and reallocation.

also helps address the standard identification problem in peer-effects settings, where network formation and outcomes are jointly determined (Manski 1993; Waldinger 2012). Our design therefore identifies a network spillover of labor mobility: moving workers affect inventive activity not only at the firms they leave and join, but also among collaborators who remain elsewhere.

Second, the paper contributes to research on how knowledge diffuses through collaborative networks and across organizational boundaries. Knowledge flows are geographically localized (Jaffe et al. 1993), co-invention networks help explain both that localization and the concentration of flows within firms (Singh 2005; Breschi and Lissoni 2009), and a large literature examines how network position and boundary-spanning ties shape what individuals and firms can learn (Powell et al. 1996; Ahuja 2000; Burt 2004; Fleming et al. 2007; Tortoriello et al. 2012; Phelps et al. 2012; Zacchia 2020). Importantly for our setting, relationships formed before a move can continue to transmit knowledge after workers separate, so mobility can affect not only the firms receiving mobile workers but also the collaborators they leave behind (Agrawal et al. 2006; Corredoira and Rosenkopf 2010; Sharoni 2026). These findings are consistent with the view that firms accumulate knowledge and capabilities that are not reducible to the expertise of any one worker (Kogut and Zander 1992), and that the ability to access and use outside knowledge is itself an organizational capability (Cohen and Levinthal 1990). The mechanism also connects to a labor-market literature on coworker networks, in which former coworkers provide information and connections that affect reemployment and reallocation (Glitz 2017; Mo and Lin 2025) and coworkers are important sources of learning and productivity growth (Jarosch et al. 2021; Herkenhoff et al. 2024; Akcigit et al. 2018; Jaravel et al. 2018). We add to these literatures by distinguishing knowledge possessed by a contact from knowledge associated with the organization in which that contact works. An inventor may know little personally about a field while working inside a firm that specializes in it. We therefore measure the collaborator’s own field expertise separately from the specialization of the collaborator’s employer, and use displacement-induced changes in employer affiliation to isolate changes in this organizational access.

Third, the paper contributes to work on the direction of innovation by identifying organizational access as a supply-side determinant of where inventors search. Technological search is path dependent and tends to remain close to existing knowledge, while unfamiliar combinations are more uncertain (Nelson and Winter 1982; March 1991; Stuart and Podolny 1996; Fleming 2001; Katila and Ahuja 2002; Ahuja and Lampert 2001; Carnabuci and Operti 2013), and boundary-spanning relationships, alliances, and inventor mobility can expand this search set (Rosenkopf and Nerkar 2001; Rosenkopf and Almeida 2003; Burt 1992). Existing research shows that economic incentives affect the direction of technological change

(Acemoglu 2002; Popp 2002; Acemoglu et al. 2012; Hanlon 2015), that private incentives can distort the composition of research across projects with different horizons (Budish et al. 2015), and that scientists face substantial costs of redirecting their research toward new topics (Myers 2020). Exposure also matters for technological specialization. In particular, childhood exposure to inventors has technology-specific effects on the fields in which individuals later invent (Bell et al. 2019), the presence or absence of established scientific leaders affects entry by outsiders and the subsequent direction of research (Azoulay et al. 2019), and social experience shapes the products and consumers toward which innovators direct their effort (Einiö et al. 2022). We uncover a distinct mechanism operating later in an inventor’s career, which is access to specialized organizations through existing collaborators changes which technological fields an inventor enters. This field-entry margin is related to the extensive margin in growth models that distinguish between improving existing product lines and entering new ones (Akcigit and Kerr 2018). Models of idea diffusion also make opportunities to search for and learn from others central to productivity growth (Lucas and Moll 2014; Perla and Tonetti 2014; Buera and Oberfield 2020; Akcigit and Ates 2021). We show that these opportunities depend partly on the organizations represented in an inventor’s collaboration network. Because new bridges tend to form where specialist employers already exist and are hiring, greater mobility or connectivity can increase diffusion without necessarily redirecting inventive activity toward fields with little existing specialist capacity.

The remainder of the paper is organized as follows. Section 2 presents a conceptual framework that separates the transmission effect of a formed bridge from the process determining where bridges form. Section 3 describes the patent, employment, collaboration-network, and displacement data and the construction of the bridge measures. Section 4 presents the empirical evidence. We first document the relationship between organizational bridges and field entry, and then describe the displacement event study and landing-composition instrument, and finally examines what the bridge transmits. Section 5 studies the consequences and allocation of bridge-induced entry, including persistence, citation-weighted output, and the fields into which mobility-induced bridges form. Section 6 concludes.

2 Conceptual Framework

This section develops a simple framework to organize the empirical analysis. The framework has three components. First, it defines organizational access as the field-specific knowledge an inventor can reach through her collaborators, so that a collaborator’s move across firms changes the set of fields to which the inventor is connected even when the inventor herself does not move. Second, it describes how this access can affect an inventor’s decision to enter

a new field. Third, it distinguishes the effect of a bridge once it is formed from the process that determines which fields receive bridges in the first place. This distinction is useful for interpreting the allocation results in Section 5.

2.1 Organizational Access

Figure 1 illustrates the object. Inventor i has a strong prior co-invention tie to collaborator j , who is employed at firm e_{jt} . When that firm specializes in field f , the tie connects i to an organization holding field-specific knowledge, and we call the connection an organizational bridge. Let $N_i(t)$ be i 's lagged co-inventor network and $\omega_{ij}(t)$ the normalized weight on repeated collaborations. Access to field f is

$$A_{ift} = \sum_{j \in N_i(t)} \omega_{ij}(t) \mathbf{1}\{e_{jt} \neq e_{it}\} s_{jft}, \quad s_{jft} \equiv \text{Spec}^{(-j)}(e_{jt}, f, t), \quad (1)$$

where s_{jft} is the specialization of j 's employer in f , computed leaving out j 's own patents. Removing j 's contribution is what separates the organization's knowledge from the collaborator's personal expertise, and it is the reason the two can be entered against each other in Section 4.3.1.

2.2 Entry under Uncertainty

An inventor who has never patented in a field does not know what working in it is worth to her, and the framework treats a bridge as supplying information about that. Let V_f denote the payoff to entering, unknown to i , with prior $V_f \sim N(\bar{v}_f, 1/h_0)$, and let her hold a private signal of precision p_0 from her own work and reading. Each strong cross-firm tie delivers a further signal


$$x_j = V_f + \varepsilon_j, \quad \varepsilon_j \sim N(0, \sigma^2/(w_{ij} s_{jft})), \quad (2)$$

independent across collaborators conditional on V_f , since they sit at different organizations. Precision rises in the employer's specialization s_{jft} , because a more specialized organization holds more field-specific knowledge for the collaborator to relay, and in the tie weight w_{ij} , because a closer collaborator relays more of it. The posterior is $V_f \mid \{x\} \sim N(m_{ift}, 1/h_{ift})$ with

$$h_{ift} = h_0 + p_0 + \sigma^{-2} \sum_j w_{ij} \mathbf{1}\{e_{jt} \neq e_{it}\} s_{jft}, \quad (3)$$

so the information a bridge portfolio supplies is the tie-weighted sum of employer specialization; the empirical measure in equation (1) divides that sum by total network weight so

Figure 1: An Organizational Bridge



Note: A stylized illustration, not data. Boxes are employers and circles are inventors. The focal inventor i works at Firm A. A prior co-inventor j works at Firm B, which specializes in field f , and the repeated collaboration between them is the organizational bridge from i to f . Both the descriptive and the causal estimates define a tie as repeated co-invention, at least two shared patents in the window $[t-7, t-3]$. Ties to specialist employers that are also i 's own employer are measured separately as same-firm bridges throughout.

it is comparable across inventors with networks of different sizes, which is a measurement choice rather than something the framework delivers. A same-firm tie is excluded because the collaborator's signal then concerns the inventor's own employer, which she observes directly through $OwnSpec_{ift}$, so it adds nothing beyond what that control already carries; this is why the same-firm coefficient should fall once own employer specialization is held fixed, as Table 1 shows. The inventor enters when the expected payoff clears her adjustment cost,

$$y_{ift} = \mathbf{1} \{ \mathbb{E}[V_f | \{x\}] \geq c(d_{if}) \}, \quad c'(\cdot) < 0, \quad (4)$$

where d_{if} is the proximity of the field to her existing portfolio, so an inventor with more relevant prior knowledge faces a lower cost of acting. A displacement moving collaborator j from employer a to employer b changes her information about field f by $\sigma^{-2}w_{ij}[s(b, f) - s(a, f)]$ and leaves everything else about her unchanged, which is the variation Section 4.2 uses.

Because entry into a new field is rare in the data, we focus on the case in which an

inventor does not enter absent sufficiently favorable information. The framework yields three implications.

First, entry should increase with the degree of employer specialization. A collaborator employed at a more specialized firm provides a more informative signal about the value of entering that field. The effect of organizational access should therefore increase gradually with employer specialization rather than depend on a particular cutoff. Section 4.3.2 examines this prediction and finds a positive relationship even below the five-percent specialization threshold used in the baseline measure.³ Second, the marginal effect of an additional bridge should be weakly diminishing. Additional independent bridges provide more information about a field, but each new signal adds less to total precision than the previous one. The panel estimates do not follow this pattern: fields reached by two or more bridges have a much larger association with entry than fields reached by only one. In the displacement design, however, the effect of a newly formed bridge is slightly smaller when the inventor already has another bridge into the field. The estimated difference is -0.041 (standard error 0.038), relative to a main effect of 0.219 . Although the estimate is imprecise, it provides no evidence of the strong amplification observed in the panel. We return to this contrast in Section 4.4. Third, the framework does not imply that the effect of a formed bridge should differ systematically across fields. In the displacement design, the interaction between a formed bridge and an indicator for sparse, highly cited fields is -0.021 with a standard error of 0.205 . The estimate is too imprecise to establish that transmission is the same across fields, but it provides no evidence of large differences over the range we study. Section 2.4 therefore treats variation in where bridges form separately from variation in what a bridge does once formed.

2.3 The Effect of a Formed Bridge

The displacement design identifies the effect of a change in organizational access. For displacement event r , let $D_{rf} = 1$ when the displaced collaborator joins an employer specialized in field f , and let $y_{rf}(d)$ denote the focal inventor’s subsequent first-entry outcome under $D_{rf} = d$. The instrumental-variables estimand is

$$\tau_f = \mathbb{E}[y_{rf}(1) - y_{rf}(0) \mid \text{landing compliers in field } f]. \quad (5)$$

³This implication is tested using the panel estimates. The displacement design does not separately identify effects by degree of specialization because the corresponding instruments are highly correlated across specialization ranges.

We refer to τ_f as the transmission effect of a formed bridge. Because the instrument shifts the destination of displaced collaborators, this is a local effect for inventors whose collaborator's landing responds to the absorption pattern of the displacement market, and the reduced form is the corresponding intention-to-treat effect of that market-level shift.

2.4 Where Bridges Form

Transmission describes what follows once a bridge exists. A separate process determines which fields receive bridges at all, and it is governed by whether specialist employers in a field are hiring. Let n_f be the number of employers specialized in f and $\bar{h}_f$ their average absorption of displaced inventors, and define field- f *absorption capacity* and the routing share as

$$a_f = n_f \bar{h}_f, \quad q_f(p) = \frac{a_f(p)}{\sum_{\ell} a_{\ell}(p)}, \quad (6)$$

so that q_f is the probability that a rematched collaborator lands at an f -specialist, conditional on landing anywhere. Only the product $n_f \bar{h}_f$ enters, which is a substantive restriction rather than a notational convenience: it says that what routes bridges is total specialist hiring, not the mere existence of specialist firms. Section 5.3 tests that restriction directly.

Writing $\mu(p)$ for the probability that a given relationship is rewired under environment or policy p , and M for the number of relationships that could be rewired, expected induced entry is

$$\mathbb{E}[\Delta Y(p)] = M \mu(p) \sum_f q_f(p) \tau_f. \quad (7)$$

Holding the transmission effects fixed, differentiating separates a change in the number of bridges from a change in their field composition:

$$\frac{d \mathbb{E}[\Delta Y(p)]}{dp} = M \left[\mu'(p) \sum_f q_f(p) \tau_f + \mu(p) \sum_f \tau_f \frac{dq_f(p)}{dp} \right]. \quad (8)$$

The decomposition separates two margins. The first is a volume margin: changes in $\mu(p)$ affect how many relationships are rewired. The second is a composition margin: changes in $q_f(p)$ affect the fields into which new bridges are formed. A change that increases rewiring proportionally across fields raises the number of bridges without changing their field composition. In contrast, a change in the distribution of specialist hiring across fields changes where new bridges are formed.

A useful benchmark is the case in which the effect of a formed bridge is similar across

fields, $\tau_f = \tau$. In that case,

$$\mathbb{E}[\Delta Y(p)] = M\mu(p)\tau,$$

so the distribution of bridges across fields determines where induced entry occurs but not its total amount. Changes in the number of bridges affect total induced entry, while changes in the distribution of specialist capacity shift that entry across fields. If transmission effects also vary with the policy environment, an additional term,

$$M\mu(p) \sum_f q_f(p) \frac{d\tau_f}{dp},$$

enters the decomposition. We do not estimate this margin.

We use this accounting framework to organize the volume and direction of bridge-induced entry, rather than to make welfare claims. In particular, the analysis does not identify whether entry represents net new invention or reallocation from other fields, or the costs and benefits of changing specialist capacity.

3 Data and Measurement

3.1 Patent Data and Technological Fields

We use patent data from the USPTO PatentsView bulk files, which provide disambiguated inventor identifiers, grant and application dates, and Cooperative Patent Classification (CPC) codes for granted U.S. patents.⁴ The data contain 8.45 million utility patents granted between 1976 and 2025. Of 4.26 million disambiguated inventors, 4.06 million have at least one classified technological field. We define a field f as a four-character CPC subclass, such as H04L for the transmission of digital information, of which 643 appear in the regression sample. Patents carry 4.5 CPC codes on average and contribute to every subclass in which they are classified, so an inventor can enter several fields with a single patent. Our outcome throughout is *first entry*, with $y_{ift} = 1$ if year t is the first year in which inventor i patents in subclass f , which we adopt for two reasons. It marks the extensive-margin decision to begin working in an unfamiliar technology, which is the object the paper is about, and it is measurable symmetrically for every inventor and field without requiring a judgment about how much activity constitutes a real move. In the displacement analyses of Section 4.2 we date entry by application year rather than grant year, so that outcomes are timed at the origin of the inventive act rather than at the end of USPTO pendency, which can run to

⁴See [Hall et al. \(2001\)](#) for the canonical description of patent data in economic research. We use utility patents only throughout.

several years and varies systematically across technologies.

3.2 Co-Inventor Networks

For each anchor year t we construct inventor i 's co-inventor network from all patents granted in the window $[t - 7, t - 3]$: j is a peer of i if they share at least one utility patent in the window, and the tie weight w_{ij} counts shared patents. The window ends three years before the anchor year to limit simultaneity between tie formation and the entry decision. We call define a tie as strong if $w_{ij} \geq 2$. About 31 percent of ties are strong, and the median tie involves a single shared patent. The networks comprise 122 million inventor-pair $\times$ year edges over anchor years 2000–2023.

3.3 Employer Histories and Employer Specialization

Employer histories come from Revelio Labs, which assembles professional profiles into employment spells with disambiguated firm identifiers (`rcid`) and provides a crosswalk linking those profiles to PatentsView inventors. The crosswalk contains 7.1 million patent $\times$ inventor-sequence $\times$ candidate-match rows covering 1.16 million distinct disambiguated inventors, with a mean full-name similarity score of 0.982, and it carries several flags indicating whether a matched worker's employer coincides with the patent's assignee. We rely on the strictest of these, which requires not only that the worker's employer match the assignee but that the employment spell actually cover the patent date, and which holds for 55.2 percent of matched rows; where an inventor-year still resolves to more than one employer we retain the spell satisfying that condition.⁵ For each inventor-year we then assign the modal employer, which yields an inventor $\times$ year employer panel that we use both for peers, to establish where j works in $t - 1$, and for the focal inventor, to establish whether a bridge crosses a firm boundary.

We measure a firm's specialization in field f from its own patent portfolio. Let N_{eft} be employer e 's cumulative patent count in subclass f up to year t and N_{et} its cumulative total. Because portfolio counts update only in years when a firm patents in a given subclass, all lookups are *as-of* joins: the specialization of employer e in field f at $t - 1$ uses the latest portfolio observation at or before $t - 1$, with the field numerator and the firm-total denominator tracked separately. To avoid conflating “the employer is a specialist” with “this particular peer patents a lot in f ,” we compute specialization leave-one-out with respect to the peer: j 's own cumulative contribution at employer e is subtracted from both numerator

⁵The weaker flag, requiring only that the worker was employed by the assignee at some point, holds for 81.2 percent of rows and would attribute patents to firms the inventor had not yet joined.

and denominator before applying the threshold. Employer e counts as an f -specialist for peer j at $t - 1$ if the leave-one-out share

$$Spec^{(-j)}(e, f, t - 1) = \frac{N_{eft-1} - N_{eft-1}^{(j)}}{N_{et-1} - N_{et-1}^{(j)}} \quad (9)$$

is at least 5 percent and the leave-one-out total is at least 5 patents. Section 4.2.8 and Appendix Table A17 show robustness to these thresholds.

3.4 Organizational Bridge Measures

The main regressor is the empirical analogue of the access term in equation (1). It aggregates employer specialization across i 's strong ties at other firms:

$$CrossBridge_{ift} = \sum_{j \in P_i(t)} \frac{w_{ij}}{W_i} \mathbf{1}\{w_{ij} \geq 2\} \mathbf{1}\{e_{j,t-1} \neq e_{i,t-1}\} Spec^{(-j)}(e_{j,t-1}, f, t - 1), \quad (10)$$

where $P_i(t)$ is the peer set, $W_i = \sum_j w_{ij}$ is total network weight, and $e_{j,t-1}$ is peer j 's employer in $t - 1$. We construct three related measures. *SameBridge*_{ift} is defined identically except that the peer's specialist employer is also i 's employer. *OwnSpec*_{ift} is the specialization of i 's own employer in f at $t - 1$, calculated as of that year and leave-one-out with respect to i . *PeerOwnSpec*_{ift} is the tie-weighted personal specialization of i 's strong cross-firm peers: each peer's cumulative patent share in f , independent of the peer's employer. Including *PeerOwnSpec*_{ift} and *CrossBridge*_{ift} jointly separates the employer-based measure from the peer's observed patent history. We also control for d_{if} , the cosine similarity between i 's prior patent portfolio and field f , and for *PeerExp*_{ift}, the tie-weighted patenting of i 's peers in f .

The regression panel is the union of two at-risk frames: cells where the field appears in the inventor's collaborative neighborhood (a peer has patented in f), and cells where a strong-tie peer's employer specializes in f even absent any peer patenting. Cells in the second frame with no bridge are sampled at a one-in-five rate and weighted by five in all regressions; because sampling is stratified on a regressor rather than the outcome, the weights recover population-representative magnitudes (Manski and Lerman 1977). The final panel has 136,624,908 inventor $\times$ field $\times$ year cells covering 2000–2023.

3.5 WARN Displacement Data

Displacement events come from Worker Adjustment and Retraining Notification (WARN) Act filings assembled by Revelio Labs, which cover 63,185 notices between 1989 and 2026

and record the filing employer, its state and city, the number of workers affected, the type of action, and both the notice and layoff dates. Because 74.9 percent of notices arrive already linked to a Revelio firm identifier, we require no fuzzy matching of company names and simply drop the remainder, and we key the design on the *notice* date rather than the layoff date. The notice is the legally mandated filing and its timing is fixed at the moment the employer files, whereas the layoff date itself can be revised or postponed afterward, so the notice date is the variable least likely to respond to conditions realized after the decision to lay off.⁶

We define a displaced mover as an inventor employed at a WARN firm in the notice year who is observed at a different employer in the following year, which yields 76,739 displacement-move events across 68,921 distinct inventors. Intersecting these movers with the co-inventor networks of Section 3.2 gives 217,413 events in which the displaced mover is a *strong* tie of some focal inventor, and these are the rewiring events underlying Section 4.2. The event-study and instrumental-variables samples further restrict attention to notice years 2000 through 2014, so that a symmetric five-year outcome window is fully observed for every cohort and no cohort contributes a truncated post-period. Appendix Table A3 reports the sample at each of these steps.

3.6 Field Citations and Inventor Density

For the allocation analysis we classify field-years on two predetermined dimensions: *density*, the number of inventors active in f at t , and *citation rate*, the average five-year forward citations to patents granted in f over $[t - 5, t - 1]$. Splitting both measures at the within-year median across fields yields a two-by-two taxonomy:

Label	Definition
Gap (shorthand)	below-median density, above-median citation rate
Dense, high citation	above-median density, above-median citation rate
Sparse, low citation	below-median density, below-median citation rate
Dense, low citation	above-median density, below-median citation rate

Weighting field-years by the at-risk cells they contain, gap field-years account for 3.1 percent of the estimation panel and 2.5 percent of the at-risk cells in the landing-composition sample. Both dimensions are measured over years strictly preceding the entry decision, so the classification is predetermined, and the results are unchanged when we restrict the sample to

⁶The notice date is populated for 99.3 percent of filings and the number of affected workers for 93.9 percent; the median notice covers 64 workers and the mean 111.

anchor years through 2018, for which the five-year forward-citation windows are complete. We use “gap” throughout as shorthand for the sparse, high-citation cell and nothing more. Neither density nor citation rates identify the marginal social value of an additional entrant, because neither captures duplication, business stealing, knowledge spillovers, or the opportunity cost of the inventor’s effort, and we therefore avoid any language that would treat these cells as underprovided.

3.7 Summary Statistics

Appendix Table A1 summarizes the patent, inventor, employer, and field samples. The median inventor holds two patents, works in two technology subclasses, and has four distinct co-inventors; a smaller group of prolific inventors raises each mean above the corresponding median. Employer patenting is similarly concentrated. Technology fields vary substantially in both inventor density and citation rates, and five-year forward citations are highly dispersed across patents. Appendix A also reports distributions of the outcome and bridge measures and documents the construction of the WARN sample (Tables A2–A3).

4 Empirical Evidence

This section proceeds in three steps. First, we document the association between standing organizational bridges and field entry, which establishes the empirical object and validates the bridge measure. Second, we use displacement-induced changes in collaborators’ employers to estimate the effect of a bridge, since the timing decomposition that opens that step shows the association cannot be read as the effect. Third, we ask what travels over the tie, by examining whether entry is associated with the collaborator’s employer conditional on her own observed expertise in the field.


4.1 Organizational bridges predict field entry

4.1.1 Motivating Fact

We begin by documenting the relationship between organizational bridges and entry into new fields. We restrict the sample to inventor–field–year cells in which at least one prior co-inventor has already patented in the field. This keeps the comparison within fields that are already represented in the inventor’s collaboration network.

Figure 2 groups these cells by the inventor’s strongest connection to a specialist employer. Panel (a) shows a clear gradient. Entry occurs in 5.5 per 1,000 cells with no connection to

Figure 2: Organizational Bridges and Entry into New Fields



Note: Data are inventor $\times$ CPC-subclass $\times$ year cells at risk of first entry with an active collaborative neighborhood ($PeerExp_{ift} > 0$), 2000–2023, weighted by design weights. A cell’s bridge type is the strongest present: *strong, cross-firm* if a co-inventor with ≥ 2 shared patents in $[t - 7, t - 3]$ is employed (at $t - 1$) at an outside firm whose leave-one-out portfolio share in the field is ≥ 5 percent; *strong, same-firm* if that specialist employer is the inventor’s own; *weak-tie* if only single-patent ties connect to specialist employers. Panel (b) splits cells by within-year quartile of tie-weighted peer patenting in the field and plots entry rates for the no-bridge and strong cross-firm categories. Bars show 95 percent confidence intervals; most are smaller than the markers.

a specialist firm, 12.9 per 1,000 cells connected only through weak ties, 16.0 per 1,000 cells connected by a strong tie to an outside specialist firm, and 25.7 per 1,000 cells in which the specialist employer is also the inventor’s own employer. The higher entry rate in the last group suggests that a collaborator’s connection to a specialist firm needs to be distinguished from the inventor’s own direct exposure to that firm.

Panel (b) asks whether the cross-firm relationship simply reflects the collaborator’s own expertise in the field. We divide cells into quartiles based on the collaborators’ prior patenting in the field. Within every quartile, entry is about 2.5 to 3 times as common when the inventor has a strong cross-firm bridge as when she has no bridge. The relationship therefore remains even among inventors whose collaborators have similar levels of prior experience in the field. These comparisons are descriptive, however, and do not distinguish the effect of a bridge from selection into which inventor–field pairs are connected. The regressions that follow address these concerns more directly.

4.1.2 Baseline Specification

Our baseline regression is a linear probability model estimated on the 136.6 million-cell panel,

$$y_{ift} = \beta \text{CrossBridge}_{ift} + \gamma \text{SameBridge}_{ift} + \phi \text{OwnSpec}_{ift} + \delta d_{if} + \alpha_{if} + \alpha_{ft} + \alpha_{it} + \varepsilon_{ift}, \quad (11)$$

estimated by weighted least squares using the design weights described in Section 3.4. Inventor $\times$ field effects α_{if} absorb persistent differences in an inventor’s proximity to each field. Field $\times$ year effects α_{ft} absorb shocks common to a field in a year, and inventor $\times$ year effects α_{it} absorb annual changes in the inventor’s overall propensity to enter new fields. The coefficient is therefore identified from changes in whether a given inventor is connected to a given field, relative to the inventor–field mean and to contemporaneous inventor- and field-level changes.⁷ Column (4) of Table 1 additionally includes peer-community $\times$ field $\times$ year effects. A peer community is defined by the dominant three-character CPC class of the employers of an inventor’s strong ties. These effects absorb field-specific changes common to inventors whose collaborators are located in similar technological communities.

4.1.3 Cross-firm Versus Same-firm bridges

Table 1 reports the estimates. Without controlling for the focal inventor’s employer specialization, the coefficients on cross-firm and same-firm bridges are 0.023 and 0.042 (column 1). Adding the focal inventor’s employer specialization reduces the same-firm coefficient to 0.024 but raises the cross-firm coefficient to 0.029 (column 2), which is what one would expect if same-firm bridges partly proxy for the technological opportunities that the inventor’s own employer provides directly. Adding portfolio similarity has little effect on either estimate (column 3). Peer-community $\times$ field $\times$ year effects, which absorb field-specific shocks common to inventors whose collaborators sit in similar technological communities, reduce the cross-firm coefficient to 0.019 ($t = 19$; column 4).

The conditional association is smaller than the raw difference in Figure 2. The design-weighted standard deviation of CrossBridge_{ift} is 0.0152, and the measure is positive in 5.8 percent of cells (Appendix Table A2). A one-standard-deviation increase is associated with a 0.044 percentage-point increase in entry, approximately 8 percent of the entry rate among bridged cells. Moving from no bridge to the average positive bridge is associated with a 0.095 percentage-point increase, or 17 percent; the corresponding estimate in column (4) is 11 percent. Much of the unconditional difference in Figure 2 therefore reflects which inventor–field cells are connected by bridges.

⁷Standard errors are clustered by inventor throughout this section. Fixed effects are absorbed by iterative weighted demeaning.

Table 1: Organizational Bridges and First Field Entry

	(1)	(2)	(3)	(4)
CrossBridge	0.0233*** (0.0010)	0.0289*** (0.0010)	0.0288*** (0.0010)	0.0189*** (0.0010)
SameBridge	0.0417*** (0.0011)	0.0237*** (0.0013)	0.0230*** (0.0013)	0.0128*** (0.0013)
OwnSpec		0.0352*** (0.0010)	0.0356*** (0.0010)	0.0367*** (0.0010)
d_{if}			-0.0087*** (0.0001)	-0.0090*** (0.0001)
$\alpha_{if}, \alpha_{ft}, \alpha_{it}$	✓	✓	✓	✓
Community×field×year FE				✓
Mean of Dep. Var	0.0039	0.0039	0.0039	0.0039
Observations	136,624,908	136,624,908	136,624,908	136,624,908

Notes: Weighted least squares estimates of specification (11); the dependent variable is an indicator for inventor i 's first patent in subclass f in year t . CrossBridge and SameBridge are defined in equation (10); OwnSpec is the leave-one-out specialization of i 's own employer in f at $t - 1$; d_{if} is the cosine similarity between i 's prior portfolio and f . Column (4) adds community×field×year fixed effects, where a community is the dominant 3-character CPC class of i 's strong peers' employers. Standard errors clustered by inventor in parentheses. *, **, and *** denote significance at the 10, 5, and 1 percent levels.

These conditional associations tell us how strongly organizational bridges are related to entry and which of the two bridge types carries the relationship, but they are not the causal effect of a bridge, for two reasons the fixed effects cannot address. First, no feasible specification absorbs shocks specific to an employer–field pair. A firm may hire into a field it expects to expand at the same time that inventors connected to it independently move toward that field, in which case the bridge measure predicts entry with no transmission over the tie at all. Second, the inventor's own career and her collaborator's employment may not be independent. The second reason we can measure, and it turns out to be large enough to determine the design of the rest of the paper, so we take it up first in the next section.

4.2 Displacement-Induced Rewiring

4.2.1 Organizational bridges and the inventor's own employment history

A cross-firm bridge may be related to the focal inventor's own employment history. In particular, the inventor may have previously worked at the bridging firm or may later join it. To examine this possibility, we split the cross-firm bridge measure according to the timing of the inventor's first employment at the bridging firm relative to her first entry into the field. We distinguish four cases: the inventor never works at the firm, joins only after entering the

field, joins in the same year as entry, or had already worked there before entry.

Table 2 shows that the panel association is concentrated in the cases in which the inventor has her own employment connection to the bridging firm. The coefficient is 0.095 ($t = 34.2$) when the inventor had previously worked at the firm and 0.351 ($t = 18.7$) when she joins in the same year as field entry, compared with a pooled coefficient of 0.029. Together, these cases account for about one-fifth of total bridge weight.

In contrast, the association is much smaller when entry occurs before any employment contact with the bridging firm. The coefficient is 0.008 for firms at which the inventor never works and -0.082 for firms she joins only after entering the field. Combining these cases yields a coefficient of about 0.006, roughly one-fifth of the pooled estimate.

These results suggest that much of the standing-network association reflects the focal inventor’s own relationship with the bridging firm rather than the co-inventor tie alone. Because panel regressions cannot fully separate these channels, the causal analysis that follows uses displacements that change the collaborator’s employer while leaving the focal inventor’s own employment unchanged.

4.2.2 Displacement As a Source of Variation in Organizational Access

The results above motivate using changes in the collaborator’s employer that are not driven by the focal inventor’s own career. We obtain this variation from mass layoffs. A layoff forces the collaborator to leave her current employer and, if she is reemployed elsewhere, changes the set of fields associated with the firm in which she works. From the focal inventor’s perspective, this creates a change in organizational access even though her own employment and prior collaboration network remain unchanged.⁸

The main remaining concern is that the collaborator’s destination after displacement is not random. Displaced inventors may choose new employers based on technological opportunities that are also related to the focal inventor’s subsequent field entry. We address this concern in two steps. We first compare fields associated with the collaborator’s old and new employers within the same displacement event. We then instrument for the fields of the collaborator’s new employer using the field composition of employers that absorb other displaced inventors in the same labor market.

⁸We focus on displacement followed by reemployment because alternative shocks provide less useful variation in this setting. Co-inventor deaths generate too few events involving strong ties to specialist employers, while firm closures and acquisitions may themselves contain information about the expected value of the affected firm’s technologies.

Table 2: Organizational Bridges and the Inventor’s Own Employment History

	(1)	(2)	(3)
Cross-firm bridge (pooled)	0.0288*** (0.0010)		
Before employment contact		0.0058*** (0.0009)	
Never employed at the firm			0.0076*** (0.0009)
Joins only after entering			−0.0823*** (0.0127)
Same year as entry		0.3641*** (0.0191)	0.3512*** (0.0187)
Already worked there		0.0961*** (0.0028)	0.0953*** (0.0028)
SameBridge	0.0230*** (0.0013)	0.0291*** (0.0013)	0.0287*** (0.0013)
OwnSpec	0.0356*** (0.0010)	0.0412*** (0.0010)	0.0409*** (0.0010)
d_{if}	−0.0087*** (0.0001)	−0.0087*** (0.0001)	−0.0087*** (0.0001)
Inventor×field, field×year, inventor×year FE	✓	✓	✓
Mean of Dep. Var	0.0039	0.0039	0.0039
Observations	136,624,908	136,624,908	136,624,908

Notes: Table displays estimates of specification (11) in which the cross-firm bridge is progressively disaggregated by the timing of the focal inventor’s own first-ever employment at the bridging firm relative to her first entry into the field. Column (1) reports the pooled coefficient. Column (2) separates entries occurring before any employment contact with that firm from entries in the same year she joins it and entries after she has already worked there. Column (3) further separates the before-contact component into firms she never works at and firms she joins only after having entered the field. The components sum to the pooled coefficient within each column. Their shares of total bridge weight are 79.9 percent for never employed, 1.2 percent for joins only after entering, 0.7 percent for same year, and 18.3 percent for already worked there. Standard errors are clustered at the inventor level. *, **, and *** denote significance at the 10, 5, and 1 percent levels.

4.2.3 WARN displacement and bridge formation

The identifying variation is easiest to see in a single displacement event. Consider a WARN notice at firm A in year T that displaces inventor j , a strong prior collaborator of focal inventor i , and suppose that j joins firm B in $T + 1$. The move changes the fields to which i is connected through j : i loses a bridge to the fields in which firm A specializes and gains a bridge to the fields in which firm B specializes, while i ’s own employment remains unchanged.

Our first design compares entry into the fields associated with the collaborator’s old and

new employers within the same displacement event. This holds fixed factors common to the event. Our second design addresses the possibility that the collaborator’s destination is selected by using the field composition of employers that hire other displaced inventors in the same labor market.

4.2.4 Stacked event study

We next use a stacked event-study design to examine how field entry changes around a collaborator’s displacement. For each displacement event, we define three sets of fields. The NEW set contains the fields in which the collaborator’s new employer specializes, while the OLD set contains the fields in which the displaced employer specializes. The NEVER set contains fields associated with counterfactual destination firms, defined as employers that hire other inventors displaced in the same notice year. We drop fields that appear in both an actual and a counterfactual set.

We align all displacement events relative to the collaborator’s move and follow field entry from five years before to five years after the event. At each event time $\tau \in [-5, 5]$, we compute the focal inventor’s first-entry rate in each set of fields. We then compare the NEW and OLD fields separately with the NEVER fields within the same event:

$$d(e, \tau) = h_{\text{SIDE}}(e, \tau) - h_{\text{NEVER}}(e, \tau).$$


This within-event comparison removes factors that are common across the field sets of a given displacement event. We then stack these differences across events and estimate the average event-time profiles using weighted least squares, with inverse-variance weights.

The design therefore asks whether entry into the fields associated with the collaborator’s actual old and new employers changes differently around displacement than entry into fields associated with plausible alternative destinations. Because the collaborator’s actual destination may still be selected, we address destination choice separately with the instrumental-variable design below.

Figure 3 reports the event-study results. Entry rates decline over event time for all three sets of fields, consistent with the prior collaboration becoming less predictive of entry as it grows more distant. After the collaborator moves, however, entry into the new employer’s fields declines more slowly than entry into the old employer’s fields.

Panel (c) summarizes this change. The difference between the trends for the new and old fields increases by 0.083 percentage points per year after displacement. Averaged over event years two through five, entry into the new employer’s fields relative to the old employer’s fields is 1.11 entries per 1,000 at-risk cells higher than in the pre-period.

Figure 3: Entry around Displacement-Induced Rewiring



Note: Stacked event study of WARN displacements in which the displaced worker is a strong prior co-inventor of the focal inventor (notice years 2000–2014; 88,732 new-side, 98,281 old-side, and 228,897 counterfactual-side inventor–field cells). Panels (a) and (b) report the within-event difference in first-filing hazards between the indicated fields and counterfactual-landing fields, per 1,000 at-risk cells. Whiskers are 95 percent confidence intervals based on standard errors clustered by focal inventor. Panel (c) reports the within-event difference between the new and old sets, with separate linear trends before and after the move. Entry hazards decline with event time as the collaboration used to define the tie recedes into the past. The estimates therefore compare changes in trends.

The response therefore does not appear as an immediate increase in entry following displacement. Instead, after the collaborator moves, the focal inventor remains more likely to enter fields associated with the collaborator’s new employer than would be expected from the pre-existing decline in the relationship.

4.2.5 Landing-Composition Instrument

The within-event design still uses the collaborator’s realized destination, which may be related to technological opportunities that also affect the focal inventor. We therefore instrument for the fields of the collaborator’s destination firm using the reemployment patterns of other displaced inventors in the same labor market.

We define a displacement market m by state, notice year, and the dominant CPC class of the displaced firm’s patent portfolio. For each field f , the instrument is the share of other displaced inventors in market m who are reemployed at firms specializing in f :

$$Z_f(i, j, T) = \frac{\#\{\text{other displaced movers in } m \text{ who join } f\text{-specialist firms}\}^{(-)}}{\#\{\text{other displaced movers in } m\}^{(-)}}. \quad (12)$$

The leave-out construction excludes all workers displaced from j ’s own WARN firm and all displaced workers who are themselves strong collaborators of focal inventor i . The first exclusion prevents the instrument from being driven by common shocks at the collaborator’s former employer, while the second ensures that the destinations used to construct the

Table 3: The Landing-Composition Instrument

	(1) OLS y on D	(2) First stage D on Z	(3) Reduced form y on Z	(4) 2SLS y on D	(5) RF excl. contact y on Z
Coefficient	0.0132*** (0.0026) [0.0023]	0.1482*** (0.0442) [0.0128]	0.0303*** (0.0081) [0.0059]	0.2044*** (0.0501) [0.0429]	0.0268*** (0.0071) [0.0053]
Event FE, field $\times$ year FE	✓	✓	✓	✓	✓
First-stage F		11.2			11.2
Anderson–Rubin 95% CI				[0.119, 0.384]	[0.096, 0.356]
Observations	526,462	526,462	526,462	526,462	526,462
Displacement events	10,328	10,328	10,328	10,328	10,328

Notes: Sample is displacement-event $\times$ field cells in the leave-out pool (notice years 2000–2014), excluding the old firm’s specialist fields, restricted to focal inventors whose modal state differs from the displacement state. D indicates that the displaced peer’s landing firm is a specialist in the field (as-of, leave-one-out); Z is the doubly-left-out share of same-market, same-class displaced movers absorbed by specialists in the field, equation (12). The outcome is first filing in the field within five years of the landing. Standard errors clustered by market (state $\times$ year $\times$ class) in parentheses; clustered by inventor in brackets. The Anderson–Rubin confidence set is robust to weak instruments. Column (5) repeats column (3) with the outcome redefined to exclude entries whose patent is co-authored with the displaced peer, carries a co-inventor employed at the peer’s landing firm in the filing year, or records the focal inventor herself at that firm; Appendix Table A15 reports the exclusions separately. *, **, *** denote significance at the 10, 5, and 1 percent levels.

instrument do not create additional bridges for i directly.

The endogenous treatment D_f equals one if collaborator j is reemployed at a firm specializing in field f . The outcome is an indicator for whether focal inventor i files her first patent in field f within five years of j ’s move. We estimate the model on displacement-event $\times$ field cells in the leave-out pool, excluding fields in which j ’s former employer specializes. We also restrict the sample to focal inventors whose modal state differs from the displacement state.

Event fixed effects absorb factors common to a given displacement event, while field $\times$ notice-year fixed effects absorb national changes in entry into particular technologies. Identification therefore comes from differences across fields in which other displaced inventors in the same market are reemployed. Standard errors are clustered by displacement market, the level at which the instrument varies. Appendix Table A14 reports alternative clustering schemes and additional inference checks.

Table 3 reports the results. The instrument strongly predicts the collaborator’s destination. A higher leave-out absorption share in field f increases the probability that j is reemployed at a firm specializing in that field. The first-stage coefficient is 0.148, with an effective F statistic of 11.2.

The reduced form shows that the same market-level variation also predicts the focal

inventor’s subsequent field entry. The coefficient is 0.030 ($t = 3.72$, with standard errors clustered by displacement market). A one-standard-deviation increase in the instrument raises the probability that the focal inventor enters field f within five years by 0.21 percentage points, relative to a mean entry rate of 0.60 percent. An interquartile increase raises entry by 0.10 percentage points. Thus, variation in where displaced workers are reemployed is associated with economically meaningful changes in the direction of the focal inventor’s subsequent patenting.

Under the instrumental-variables assumptions, the 2SLS estimate identifies the effect of a formed bridge for inventors whose collaborator’s destination responds to the field-specific absorption pattern of the displacement market. The estimate is 0.204, with an Anderson–Rubin 95 percent confidence set of $[0.119, 0.384]$. The magnitude is large relative to the average entry rate, but the treatment is rare and the estimate applies to a selected complier population. Because the first stage is modest, we place greater emphasis on the reduced-form estimate and weak-instrument-robust confidence intervals than on the 2SLS point estimate itself.

4.2.6 Interpretation

The 2SLS estimate is large relative to the average rate of field entry, but both the treatment and the outcome are rare. In the instrumental-variables sample, the focal inventor enters a given field within five years in 0.6 percent of event–field cells, and the displaced collaborator lands at a firm specializing in that field in 1.4 percent of cells. Entry occurs in 2.8 percent of cells in which the collaborator lands at an f -specialist, compared with 0.6 percent when she does not. These raw differences are descriptive, but they help illustrate the scale of the variation underlying the IV estimate.

Figure ?? shows the first stage graphically. We divide event–field cells into deciles of the landing-composition instrument and plot the probability that the displaced collaborator lands at an f -specialist in each decile. Specialist landings become more common as the instrument increases, consistent with the first-stage estimate in Table 3. The figure also plots the focal inventor’s subsequent entry rate for reference.

The 2SLS estimate should be interpreted as a local effect. It applies to focal inventors whose collaborator’s destination responds to the field-specific reemployment pattern of the displacement market. Because these compliers are selected precisely on the margin shifted by the instrument, the 2SLS estimate need not resemble the average relationship in the full panel. For this reason, and because the first stage is modest, we emphasize the reduced-form estimate and Anderson–Rubin confidence intervals rather than the 2SLS point estimate itself.

Table 4: Interpreting the Instrumental-Variables Estimates

	Value
Mean first-entry probability	0.006
Mean specialist-landing probability	0.014
Entry rate: specialist landing vs. no landing	2.8% vs. 0.6%
First stage: D on Z	0.148 ($t = 3.4$)
Reduced form: y on Z	0.030 ($t = 3.7$)
2SLS: y on D	0.204
Leave-one-market-out reduced-form range	[0.026, 0.032]

Notes: Instrumental-variables sample, $n = 526,462$; notice years 2000–2014. First-stage and reduced-form t statistics are based on standard errors clustered by displacement market. The leave-one-market-out range is obtained by dropping each of the twenty largest markets separately and re-estimating the reduced form. Weak-instrument-robust confidence intervals are reported in Section 4.2.8.

The estimate is also not driven by a single large displacement market. Dropping each of the twenty largest markets one at a time leaves the reduced-form estimate between 0.026 and 0.032. We report these results in Table 4.

4.2.7 Technology Demand

A key concern is that the instrument may capture changes in technology demand rather than the effect of the collaborator’s destination. For example, a growing technology could both attract displaced workers to firms specializing in that field and independently increase the probability that other inventors enter the same field. Field $\times$ year fixed effects absorb national changes in field-specific entry, and we restrict the sample to focal inventors in different states from the displacement market. However, these controls do not directly test whether the reduced-form relationship operates through the specific collaborator tie.

We therefore construct a randomized-network placebo. Within each field $\times$ year cell, we reassign the displacement-market exposure of each focal inventor to another focal inventor. This preserves the field-year distribution of the instrument but breaks the link between the focal inventor and the displaced collaborator. We then re-estimate the reduced form using the reassigned exposure.

If the reduced form were driven by a common technology-demand shock, the reassigned instrument should continue to predict field entry. In contrast, if the relationship depends on the specific collaborator connection, it should disappear once that connection is broken. Across 300 random reassignments, the placebo estimates are centered near zero (mean 0.00005, standard deviation 0.003), compared with an actual reduced-form estimate of 0.030.

None of the placebo estimates is as large as the actual estimate ($p < 0.01$).

These results provide evidence against the interpretation that the reduced form is driven solely by field-level technology demand. Together with field $\times$ year fixed effects and the geographic restriction, the placebo suggests that the relationship depends on the specific inventor–collaborator connection.

4.2.8 Other Robustness Checks

Appendix Table A16 reports alternative definitions of the displacement market and of the estimation sample, and the estimates are stable across them. Requiring the focal inventor to reside outside the displacement market’s census region, a considerably stronger geographic restriction than our baseline, raises the first-stage F statistic from 11.2 to 13.5 while yielding a similar reduced form, and requiring at least 10 or 20 leave-out movers likewise leaves the reduced form intact although the first stage falls to 10.0 and 8.1 respectively. The Anderson–Rubin confidence set excludes zero in every specification, which is the relevant statement given first stages of this magnitude. In the event study, the estimated trend change is similar when we restrict attention to layoffs involving at least 50 workers, and extending the window through $\tau = 8$ for the early cohorts shows the difference persisting rather than reversing.⁹

Two features of the sample bear on how the standard errors should be read, and Appendix Table A14 reports both. The first is the small number of displacement markets. With 84 clusters, and an effective count of 18 once their unevenness is taken into account, asymptotic cluster-robust inference cannot simply be assumed, so we compute a wild cluster bootstrap of the reduced form, imposing the null and resampling markets with Rademacher weights. The bootstrap p value is 0.0005 over 1,999 replications. Two-way clustering by market and focal inventor, and by market and displaced peer, leaves the reduced-form t statistic at 3.68 and 3.72 against 3.72 for market alone, so the repetition of inventors across markets adds essentially nothing to the variance. It is also worth being explicit that the modest first stage is a property of market clustering rather than of the instrument: the same first-stage coefficient carries an effective F of 11.2 clustered by market, 12.9 by WARN firm, 28.4 by displaced peer and 134 by focal inventor. We report the most conservative of these throughout and rest inference on the Anderson–Rubin set.

The second is that the same first entry can enter through several event–field rows, because an inventor may have more than one collaborator displaced and a collaborator may be tied to more than one inventor. The estimate is stable under every deduplication rule we can construct. Keeping only each inventor’s first displacement gives 0.0296 ($t = 3.1$); requiring

⁹The larger contrasts at $\tau \leq -4$ fall inside the $[T-7, T-3]$ window used to define the prior collaboration itself, and should not be read as conventional pre-treatment coefficients.

non-overlapping five-year outcome windows gives 0.0279 ($t = 3.0$); and drawing one displaced collaborator per inventor at random gives a mean of 0.0266 across 200 draws, with every draw positive and a fifth-to-ninety-fifth percentile range of [0.0248, 0.0283]. The one rule that moves the point estimate, keeping a single focal inventor per displaced collaborator, raises it to 0.0818 but leaves only 2,342 events and a first stage of 3.8, which is too weak to interpret.

A distinct concern is that the outcome itself may restate the collaboration rather than measure diffusion through it. The focal inventor’s first patent in the field could be one she writes with the displaced peer, or one staffed by people at the peer’s new firm, or one produced after she has joined that firm herself, and in any of those cases the estimate would describe continued joint production rather than access to an organization she never enters. Because the outcome is built from patents, each of these can be removed. Appendix Table A15 reports the reduced form with entries excluded on each ground in turn and on all three together, holding the sample and the instrument fixed so that only the outcome changes. Co-invention with the peer is rare among entry patents, 2.4 percent of them, and removing those entries leaves the reduced form at 0.0304. Entries with a co-inventor at the landing firm and entries on which the inventor is herself recorded there are more common, about 15 and 14 percent, and removing all three grounds leaves 0.0268 ($t = 3.8$), 88 percent of the baseline, with an Anderson–Rubin set of [0.096, 0.356]. Dropping instead the events in which the inventor is ever employed at the landing firm during the outcome window, which removes 29 percent of events rather than reweighting the outcome, raises the reduced form to 0.0362 ($t = 4.3$). The same exclusions applied to the event study leave the pooled post-period difference at 1.18 per 1,000 against a baseline of 1.11, and the trend break at 0.075 percentage points per year against 0.083. Continued collaboration and destination affiliation therefore account for a modest part of the estimate and none of its sign or significance.

4.2.9 Arm’s-length effect

Applying the same timing split to the displacement design gives the diagnostic of Section 4.2.1 its counterpart. We decompose the landing-composition reduced form by when the focal inventor first works at the peer’s *landing* firm relative to her entry, so that “before” collects entries that occur before she has any employment contact with that firm.

Most of the reduced form is carried by the arm’s-length component (Table 5). Entries before any employment contact with the landing firm carry a reduced form of 0.022 ($t = 3.8$), about 74 percent of the 0.030 total; the same-year component is -0.0001 with a standard error of 0.0002; and the after-contact component is 0.008. Removing events in which the inventor had *already* worked at the landing firm, if anything, *strengthens* the reduced form,

Table 5: Arm’s-Length Timing of the Displacement Effect

	(1) All events	(2) Excl. prior alumni
Reduced form: any first entry	0.0303*** (0.0081)	0.0345*** (0.0083)
Before employment contact	0.0224*** (0.0059)	0.0326*** (0.0079)
Same year as contact	−0.0001 (0.0002)	−0.0002 (0.0003)
After contact	0.0080* (0.0046)	0.0021 (0.0014)
First stage: D on Z	0.1482*** (0.0442)	0.1477*** (0.0454)
2SLS: y on D	0.2044*** (0.0501)	0.2335*** (0.0709)
Event FE, field×notice-year FE	✓	✓
Observations	526,462	390,753

Notes: Table displays the landing-composition reduced form, first stage, and 2SLS estimates of Section 4.2.5, with the outcome split by the timing of the focal inventor’s first employment at the displaced collaborator’s landing firm relative to her first entry into the field. The before-contact, same-year, and after-contact components sum to the any-entry reduced form within each column. Column (2) drops events in which the inventor had already worked at the landing firm before the displacement. Restricting further to inventors who never work at that firm at any point yields a reduced form of 0.0340 (0.0083) over 346,946 event×field cells. The mean of the dependent variable is 0.0060 in column (1). Standard errors are clustered at the displacement-market level. *, **, and *** denote significance at the 10, 5, and 1 percent levels.

from 0.030 to 0.034 ($t = 4.2$). Where the standing association was concentrated in cells that track the inventor’s own organizational trajectory, the causal effect is concentrated in cells where she has none: it operates before, and independently of, any move by the inventor herself. This is the sense in which bridge-induced redirection is arm’s-length.

4.3 Mechanisms

The displacement results show that a collaborator’s destination is related to where the focal inventor subsequently patents, but they do not identify what passes through the collaboration tie. We next examine several mechanisms that may account for this relationship. We first ask whether employer specialization predicts entry beyond the collaborator’s own observed expertise. We then examine how the relationship varies with employer specializa-

tion, what the displacement design implies about the source of the transmitted knowledge, whether geographic proximity matters, and whether entry patents cite prior work from the collaborator’s employer.

Most of these tests use the panel variation from Section 4.1.3 and therefore remain subject to the same selection concerns. We interpret them as evidence consistent with organization-mediated access rather than as separate causal estimates of the mechanism. Section 4.3.3 returns to the displacement design. ¹⁰

4.3.1 Peer Expertise Versus Employer Specialization

We first ask whether the employer-based bridge measure contains information beyond the collaborator’s own observed expertise. Table 6 enters the bridge measure jointly with $PeerOwnSpec_{ift}$, the tie-weighted cumulative share of collaborators’ prior patents in field f .

Entered separately, both measures predict entry (columns 1 and 2). When they are included together, both coefficients change little. Specifically, the coefficient on employer specialization falls from 0.0288 to 0.0269, while the coefficient on collaborators’ own patent specialization falls from 0.0174 to 0.0168 (column 3). Thus, the collaborator’s employer remains predictive of field entry even after conditioning on her observed patent history.

The employer-based coefficient is about 60 percent larger in this specification. Because the two measures have different distributions and scales, however, we do not interpret this difference as evidence that employer specialization is quantitatively more important than the collaborator’s own expertise.

4.3.2 Further Evidence on Employer Specialization

Two additional panel results help distinguish employer specialization from the collaborator’s own observed expertise. First, the employer-based association is actually larger when the collaborator has never patented in the destination field. The coefficient is 0.0447 for collaborators with no prior patent in the field, compared with 0.0217 for collaborators who have previously patented there (Table A7). A collaborator with no patent in a field may still have knowledge of it, so this comparison does not isolate employer knowledge from personal knowledge. It does show, however, that the association is not concentrated among collaborators with observed inventive experience in the destination field.

We next ask whether the relationship becomes stronger as the collaborator’s employer becomes more specialized in the field. Figure A1 divides the employer’s leave-one-out field

¹⁰These tests also cannot rule out unobserved differences or changes in the collaborator’s own knowledge. Detailed estimates appear in Appendix B and Appendix C.

Table 6: Peer Expertise versus Employer Specialization

	(1)	(2)	(3)
CrossBridge (peer’s employer)		0.0288*** (0.0010)	0.0269*** (0.0010)
PeerOwnSpec (peer’s personal)	0.0174*** (0.0004)		0.0168*** (0.0004)
Baseline controls and FE	✓	✓	✓
Mean of Dep. Var	0.0039	0.0039	0.0039
Observations	136,624,908	136,624,908	136,624,908

Notes: All columns include SameBridge, OwnSpec, d_{if} , and the three-way fixed effects of specification (11). PeerOwnSpec aggregates, with the same tie weights and strong-tie restriction as CrossBridge, each cross-firm peer’s own cumulative patent share in field f at $t - 1$, independent of employer. Standard errors clustered by inventor in parentheses. *, **, *** denote significance at the 10, 5, and 1 percent levels.

share into four mutually exclusive ranges and enters the tie-weighted bridge measure separately for each. The coefficients increase monotonically with employer specialization, from 0.0045 for employers with field shares between 2 and 5 percent to 0.0112 for employers with shares above 20 percent. The coefficient is already positive below the 5 percent specialization cutoff used in the main analysis, suggesting that the baseline result is not an artifact of that threshold.

4.3.3 Whose knowledge? Evidence from the Displacement design

The panel results above cannot distinguish knowledge associated with the collaborator’s employer from knowledge held by the collaborator herself. The displacement design does not fully resolve this distinction either. When a collaborator moves to a new employer, she becomes connected to the organization’s knowledge but may also acquire new knowledge herself and transmit it through the existing collaboration tie. Leaving the collaborator’s own patents out of the employer-specialization measure removes her prior contribution to that measure, but does not capture what she learns after joining the new firm.

We use two features of the displacement design to examine this distinction. Both rely on characteristics determined before the collaborator’s move or on the timing of the focal inventor’s subsequent entry.¹¹

First, we ask whether the effect is larger when the collaborator had already patented in the field before displacement. Prior expertise is predetermined and can therefore be

¹¹We do not split the estimates by the collaborator’s subsequent patenting or tenure at the destination firm because these outcomes may themselves respond to the displacement.

interacted with the instrument. It is present in 2.9 percent of event–field cells. The estimated effect of a formed bridge is 0.209, while the interaction with the collaborator’s prior field experience is -0.016 with a standard error of 0.044 (Table 7, Panel B). We therefore find no evidence that the effect is larger when the collaborator herself had previously patented in the destination field. Because prior expertise is uncommon and the interaction estimate is imprecise, however, we cannot rule out meaningful heterogeneity along this dimension.

Second, we examine the timing of the response. The collaborator joins the destination firm in $T + 1$, and we measure the focal inventor’s first entry from $T + 2$ through $T + 6$. Panel A of Table 7 estimates the reduced form separately by outcome year. The effect is already present in the first outcome year, is largest around $T + 4$, and does not continue increasing thereafter. This pattern does not show a long delay before the focal inventor responds to the collaborator’s move. It is therefore consistent with relatively rapid transmission after the move, although it does not distinguish direct access to organizational knowledge from knowledge the collaborator acquires after joining the firm.

Taken together, these results provide little evidence that the displacement effect is driven primarily by expertise the collaborator already possessed before the move. They are also consistent with transmission occurring relatively soon after reemployment. They do not, however, distinguish organizational knowledge from knowledge that the collaborator acquires at the new employer and subsequently relays through the collaboration tie.

4.3.4 Geographic scope

Knowledge diffusion is often geographically concentrated, although patent citations provide an imperfect measure of localization (Jaffe et al. 1993; Audretsch and Feldman 1996; Thompson and Fox-Kean 2005; Agrawal et al. 2008; Kerr 2008). We therefore examine whether the bridge-entry relationship differs with the collaborator’s location relative to the focal inventor’s modal state. About half of inventors have no identifiable modal U.S. state, primarily because their patents place them outside the United States, so we report this group separately rather than combining it with collaborators known to be in another state.

The coefficient is 0.0342 ($t = 17.5$) for collaborators in the same state, 0.0292 ($t = 13.2$) for collaborators in a different state, and 0.0252 ($t = 20.5$) when the collaborator’s state cannot be identified (Table A4). The cross-state coefficient is about 85 percent as large as the within-state coefficient. Thus, the association remains substantial even when the focal inventor and collaborator are located in different states.¹²

¹²We interpret this as a scope test rather than as separate causal evidence. Modal state is a coarse measure of geographic proximity, and the comparison does not rule out nonlocal common shocks.

Table 7: Whose Knowledge Travels: Evidence from the Displacement Design

<i>Panel A. Reduced form by year of first entry</i>					
	$T+2$	$T+3$	$T+4$	$T+5$	$T+6$
Reduced form	0.0036*	0.0063**	0.0079***	0.0069**	0.0056**
	(0.0021)	(0.0026)	(0.0022)	(0.0029)	(0.0024)
Entries	640	652	623	591	627
<i>Panel B. Effect of a formed bridge, by the peer's prior expertise</i>					
	Coefficient		Std. err.	95% CI	
Formed bridge	+0.2092***		(0.0626)	[0.086, 0.332]	
× peer already patented in f	-0.0161		(0.0441)	[-0.103, 0.070]	
Event, field×year FE			✓		
Observations			526,462		

Notes: Both panels use the landing-composition sample of Table 3 (526,462 event–field cells, notice years 2000–2014) with event and field×notice-year fixed effects and standard errors clustered by displacement market. Panel A re-estimates the reduced form with the outcome redefined as first filing in field f in the stated event year; the collaborator lands at $T + 1$. Panel B reports two-stage least squares in which an indicator for the collaborator’s landing at an f -specialist, and its interaction with whether she had herself patented in f before the notice year, are instrumented by the leave-out market composition and its corresponding interaction. Prior expertise is predetermined and holds for 2.9 percent of cells. *, **, and *** denote significance at the 10, 5, and 1 percent levels. Source: `Code/warn_mechanism_timing.py`.

4.3.5 Citations to the bridged employer’s prior work

We next ask whether first-entry patents are more likely to cite prior work from the collaborator’s employer. For each first-entry patent associated with a cross-firm bridge, we compare the bridged employer with other firms that specialize in the same field but are not connected to the focal inventor through the measured bridge network. Entry-patent fixed effects hold fixed the patent’s overall propensity to cite specialist firms, so the comparison is made within the same entry event.

An entry patent cites field- f patents assigned to the bridged employer in 5.5 percent of cases, compared with 0.1 percent for other specialist employers. The within-entry difference is 0.054 ($t = 79$; Table A5). Because bridged employers tend to have larger patent portfolios in the field, part of this difference may simply reflect greater citation opportunities. Controlling for the log of field-specific portfolio size reduces the coefficient to 0.038 ($t = 57$). Restricting the comparison group to specialist firms with portfolios at least as large as that of the bridged employer yields a coefficient of 0.021 ($t = 20$), with citation rates of 3.6 and 0.6 percent, respectively. Excluding firms that previously employed the focal inventor gives a coefficient of 0.032 ($t = 50$) across 101,660 entry events.

These patterns show that patents associated with field entry are disproportionately likely

to cite prior work from the collaborator’s employer. This is consistent with the idea that the bridge provides access to knowledge associated with that organization, although citation patterns alone do not reveal how that knowledge reached the focal inventor.¹³

4.3.6 Summary of Mechanisms

Taken together, the mechanism results show that employer specialization predicts field entry beyond the collaborator’s observed patent history, that the relationship is stronger for more specialized employers, and that entry patents are disproportionately likely to cite prior work from the bridged employer. The displacement results additionally show no evidence that the effect is larger when the collaborator had previously patented in the destination field.

These patterns are consistent with organization-mediated access to knowledge. They are harder to reconcile with an explanation based only on expertise the collaborator had already revealed through her own patenting before the move. At the same time, the evidence does not identify what information or assistance passes through the collaboration tie. In particular, it cannot distinguish direct access to knowledge associated with the employer from knowledge that the collaborator acquires after joining the firm and subsequently transmits to the focal inventor.

Section 4.4 next examines which inventors and fields are most affected by these connections.

4.4 Heterogeneity

We next examine how the bridge–entry relationship varies across inventors, fields, and collaboration ties. These estimates use the panel design of Section 4.1.3 and therefore inherit its selection concerns. Differences across subgroups may reflect heterogeneity in the effect of a bridge, differences in selection into bridges, or both. We therefore treat these results as descriptive and do not use them to identify causal heterogeneity.¹⁴

The largest differences appear across characteristics of the focal inventor and her network. First, the panel association is much larger when an inventor reaches a field through multiple independent cross-firm bridges. The coefficient is 0.0636 for fields reached through two or more bridges, compared with 0.0076 for fields reached through exactly one (Table A6). This pattern could reflect complementarity across independent sources of access, but it could also reflect selection into inventor–field pairs that acquire multiple bridges.

¹³A stronger interpretation would require distinguishing applicant-supplied from examiner-added citations and accounting for continued co-invention with the displaced collaborator.

¹⁴The displacement sample is generally too small to estimate the corresponding interactions precisely. We report the full set of heterogeneity estimates in Appendix C.

The displacement design provides one check on this interpretation. The effect of a formed bridge is not larger when the focal inventor already has another specialist-employer bridge to the field: the interaction is -0.041 (standard error 0.038), compared with a main effect of 0.219 . We therefore find no causal evidence that a second bridge has a larger effect than a first. The much stronger panel association for multiply connected fields may instead reflect selection into where additional bridges form.

Second, the employer-based association is not concentrated among collaborators with prior patenting experience in the destination field. The coefficient is 0.0447 when the collaborator has never patented in the field and 0.0217 when she has (Table A7). As discussed in Section 4.3.1, this does not imply that collaborators without prior patents lack relevant knowledge. It does show that the bridge-entry relationship remains strong when the collaborator has no observed inventive experience in the destination field.

Third, the relationship varies with technological proximity to the focal inventor’s prior work. Relative to fields with no prior overlap with the inventor’s portfolio, for which the coefficient is 0.0222 , the association is larger for technologically adjacent fields (0.0729) but close to zero in the intermediate group (0.0034 ; Table A8). The pattern is therefore non-monotonic, and a binary near-versus-far comparison would obscure the intermediate trough.

Fourth, the association is stronger in fields with greater specialist capacity in the focal inventor’s own state. The coefficient is 0.0494 where local capacity is relatively high, compared with 0.0282 where it is relatively low and 0.0182 among inventors for whom we cannot identify a U.S. state (Table A10). This pattern is consistent with organizational bridges and local specialist capacity acting as complements, although the panel design does not establish that interpretation causally.

Heterogeneity across destination technologies is more limited. Across the six well-populated CPC sections of the destination field, coefficients range from 0.0198 to 0.0336 (Figure A3). There is somewhat more variation by the focal inventor’s originating technology, with coefficients ranging from 0.0151 to 0.0351 (Figure A2). Bridges through employers in biopharmaceutical and software and computing industries also have somewhat larger coefficients than the pooled estimate (Table A11). These industry differences may reflect differences in the organization of inventive activity or in how sharply employer specialization maps into technological fields; we do not distinguish among these explanations.

Finally, we find little difference by the recency of the collaboration tie. Ties last active within four years have a coefficient of 0.0293 , compared with 0.0277 for ties dormant for five years or more (Table A9). This pattern is less robust when we restrict bridges to employers at which the focal inventor never works, so we do not place much weight on it.

5 Persistence and Allocation

We next examine whether displacement-induced field entry is followed by additional patenting in the destination field and where new organizational bridges are formed. These exercises distinguish the consequences of a change in field access from the frequency with which displaced collaborators are reemployed at specialist firms in different fields. All outcomes in this section are measured within the destination field. The design therefore does not identify whether these changes increase the inventor’s total patenting or instead reallocate inventive activity across fields.

5.1 Persistence and Destination-Field Output

A first patent in a new field may represent either a one-time entry or the beginning of continued activity in that field. We therefore re-estimate the landing-composition reduced form using several outcomes measured over the same five-year post-landing window. Each outcome is defined for every event–field cell, with zeros assigned when the focal inventor produces no patent in the field, so the estimates do not condition on realized entry.

Table 8 reports the results. For durable entry, defined as a first entry followed by at least one additional patent in the same field within the five-year window, the reduced-form coefficient is 0.0211 ($t = 3.7$). This is about 70 percent as large as the reduced-form estimate for any first entry, 0.0303. Thus, the increase in field entry associated with the instrument is not concentrated in one-time patenting.

The instrument also predicts subsequent destination-field output. The reduced-form coefficient is 0.140 for the total number of patents produced in the field ($t = 2.9$) and 0.110 for patents produced after the initial entry patent ($t = 2.5$). It also predicts 0.579 additional five-year forward citations to destination-field patents ($t = 2.2$), with a similar estimate after winsorizing citations at the 99th percentile. When citations are allocated fractionally across CPC subclasses, the estimate falls to 0.061 ($t = 1.7$) and is no longer statistically precise.

These outcomes show that the response extends beyond a one-time entry into the destination field. The instrument predicts continued patenting and citations in that field, suggesting a more persistent shift in the direction of inventive activity. These estimates, however, identify only where subsequent activity occurs. They do not show whether the inventor’s total output increases, since greater activity in the destination field could be offset by less activity in fields she would otherwise have pursued. We return to this limitation in Section 5.5.

Table 8: Consequences of Displacement-Induced Bridge Formation

	(1) Any entry	(2) Durable entry	(3) Transient entry	(4) Field patents	(5) After entry	(6) 5-yr cites
Reduced form: Z	0.0303*** (0.0081)	0.0211*** (0.0057)	0.0092*** (0.0035)	0.1402*** (0.0487)	0.1100** (0.0441)	0.5791** (0.2585)
Event FE, field $\times$ notice-year FE	✓	✓	✓	✓	✓	✓
Mean of Dep. Var	0.006	0.004	0.002	0.028	0.022	0.073
Observations	526,462	526,462	526,462	526,462	526,462	526,462

Notes: Table displays the landing-composition reduced form of Table 3 re-estimated on unconditional outcomes measured over the fixed five-year post-landing window, with zeros assigned to cells in which the focal inventor produces no patent in the field. A durable entry is a first entry followed by at least one additional patent in the same field within the window; columns (2) and (3) sum to column (1). Column (5) counts destination-field patents filed after the first entry patent, and column (6) counts five-year forward citations to the inventor’s patents in the field. Winsorizing citations at the 99th percentile gives 0.6099 (0.2542); allocating citations fractionally across a patent’s CPC subclasses gives 0.0614 (0.0367). Standard errors are clustered at the displacement-market level. *, **, and *** denote significance at the 10, 5, and 1 percent levels.

5.2 Where Mobility-Induced Bridges Form

The effect of mobility on the direction of invention depends on two margins. The first is transmission: whether a formed bridge changes entry, and whether that effect differs across fields. The second is formation: how likely mobility is to create a bridge into a particular field in the first place. We examine these margins using the density and citation-rate classification introduced in Section 3.6. For brevity, we refer to sparse, high-citation field-years as “gaps.” This label is descriptive and does not imply that additional invention in these fields has greater social value.

We begin with transmission. Appendix Table A12 interacts the cross-firm bridge measure with the four field types in the panel specification (11). Because the specification includes field $\times$ year fixed effects, the interactions capture differences in the bridge–entry relationship across field types rather than differences in average entry rates.

The estimated coefficients range from 0.0186 in sparse, low-citation fields to 0.0278 in dense, high-citation fields. The interaction between the bridge measure and the gap indicator is -0.0001 ($t = -0.01$). We therefore find no evidence in the panel that the bridge–entry relationship is weaker or stronger in sparse, high-citation fields. The pattern is similar when we restrict the sample to anchor years through 2018, for which the citation windows are complete.

These estimates are associational and do not identify field-specific causal effects. They nevertheless provide no indication that the panel relationship differs systematically for gap

fields.

5.3 Bridge Formation Across Field Types

We next turn to bridge formation. Table 9 interacts the landing-composition design with the gap indicator. The reduced-form coefficient is 0.0306 in non-gap fields, while the interaction for gap fields is -0.0259 ($t = -4.4$), implying a combined reduced-form estimate of 0.0047 in gap fields.

The difference is closely related to the first stage. The first-stage coefficient is 0.1499 in non-gap fields but only 0.0170 in gap fields. Consistent with this, a collaborator lands at a specialist employer in only 0.14 percent of gap-field event-field cells, compared with 1.46 percent in non-gap fields. Thus, the displacement design generates very little variation in new specialist connections to sparse, high-citation fields.

Because the first stage is close to zero within gap fields, the corresponding 2SLS interaction is imprecisely estimated: the interaction is -0.0211 with a standard error of 0.2053. The design therefore does not allow us to determine whether the effect of a formed bridge differs between gap and non-gap fields.

What the results do show more clearly is a difference in formation. New bridges created through displacement are much less likely to point toward sparse, high-citation fields. This is consistent with bridge formation following the existing distribution of specialist employers: fields with few specialist employers provide fewer opportunities for displaced inventors to form new organizational connections.

The results therefore do not establish that transmission effects are the same across fields. They show instead that displacement-induced bridge formation is uncommon in sparse, high-citation fields. This pattern is consistent with limited specialist hiring in those fields.

Equation (6) implies that bridge formation should depend on the amount of specialist hiring in a field, rather than simply on the number of specialist firms. We examine this implication directly using field-year variation. Over the 3,031 field-years in the landing-composition sample, we regress specialist landings on the number of specialist employers and on their absorption of displaced inventors, including year effects and using the number of at-risk cells as an offset.

Specialist absorption is strongly related to bridge formation. Its estimated elasticity is 0.93 ($t = 64.1$). Conditional on absorption, the number of specialist employers enters with a much smaller coefficient of -0.065 ($t = -3.6$). Thus, fields with more specialist firms do not receive more displacement-induced bridges unless those firms also absorb displaced inventors.

Table 9: Bridge Formation and Transmission by Field Type

	(1)	(2)	(3)	(4)	(5)
		Reduced form		First stage	2SLS
	y on Z	y on Z	y on Z	D on Z	y on D
Z	0.0303*** (0.0081)	0.0306*** (0.0081)	0.0282** (0.0128)	0.1499*** (0.0442)	
$Z \times \text{Gap}$		-0.0259*** (0.0058)	-0.0236** (0.0116)	-0.1329** (0.0642)	
$Z \times \text{Dense, high citation}$			0.0033 (0.0123)		
$Z \times \text{Sparse, low citation}$			-0.0902** (0.0412)		
D					0.2043*** (0.0499)
$D \times \text{Gap}$					-0.0211 (0.2053)
Event FE, field $\times$ notice-year FE	✓	✓	✓	✓	✓
Mean of Dep. Var	0.006	0.006	0.006	0.014	0.006
Observations	526,462	526,462	526,462	526,462	526,462

Notes: Table displays the landing-composition design of Table 3 interacted with the field-type classification of Section 3.6, assigned in the landing year. A gap is a sparse, high-citation field-year; the omitted category in column (3) is dense, low-citation field-years. Column (1) reproduces the pooled reduced form, column (2) adds the gap interaction, and column (3) reports the full two-by-two. Column (4) reports the corresponding first stage and column (5) the 2SLS estimates, in which $D \times \text{Gap}$ is instrumented by $Z \times \text{Gap}$. The specialist-landing rate is 0.14 percent in gap cells against 1.46 percent elsewhere, measured over event $\times$ field rows and classified in the landing year, as everywhere in the paper (Appendix Table A13); the large standard error on $D \times \text{Gap}$ reflects the limited first-stage variation in gap fields. Estimates cover 10,328 displacement events. Standard errors are clustered at the displacement-market level. *, **, and *** denote significance at the 10, 5, and 1 percent levels.

This result supports the distinction in Equation (6). In particular, what matters for bridge formation is not simply the presence of specialist employers, but their realized hiring of displaced inventors. Sparse fields therefore appear constrained primarily by the limited absorption of displaced workers into specialist firms.

5.4 Accounting for Bridge Formation

Equations (6)–(8) separate changes in the number of bridges from changes in the fields to which those bridges point. The data show that new bridges are distributed unevenly across fields. In the estimation panel, 81 percent of bridge weight is in dense, high-citation field-years, while less than 1 percent is in sparse, high-citation field-years. The average field reached by a bridge has a lagged citation rate about 16 percent higher than that of a

randomly selected at-risk field. Thus, bridges tend to point toward relatively high-citation fields, but rarely toward high-citation fields in which inventive activity is sparse.

The landing-composition sample provides a more direct measure of the formation margin. Sparse, high-citation fields account for 2.52 percent of at-risk event–field cells but only 0.24 percent of realized specialist landings. Their estimated share of total specialist absorption is still smaller, at 0.03 percent (Appendix Table A13). Thus, their small share of realized bridges does not simply reflect their rarity in the set of potential destinations: they receive substantially fewer specialist landings relative to their share of at-risk opportunities.

These results point to bridge formation as an important constraint on diffusion into sparse, high-citation fields. When few displaced inventors are absorbed by specialist employers in a field, mobility generates few new organizational connections to that field. The estimates do not establish that transmission effects are the same across fields, but they show that there is little scope for the displacement design to generate bridges where specialist absorption is rare.

The accounting framework also allows us to illustrate how large a change in specialist absorption would be required to materially alter the composition of new bridges. Holding the routing relationship fixed, let A_{gap} and A_{non} denote specialist absorption in gap and non-gap fields. If absorption in every gap field-year is multiplied by m , the implied share of new bridges pointing toward gap fields is

$$q_{\text{gap}}(m) = \frac{mA_{\text{gap}}}{mA_{\text{gap}} + A_{\text{non}}}.$$

Because the initial absorption share of gap fields is very small, even large proportional increases produce modest changes in bridge composition. Tripling specialist absorption in gap field-years raises their implied share of bridge formation from 0.029 percent to about 0.088 percent, still far below their 2.52 percent share of at-risk opportunities. Matching the opportunity share would require specialist absorption in gap fields to increase by a factor of roughly 89.

These calculations are mechanical accounting exercises rather than policy counterfactuals. They hold the routing relationship and non-gap absorption fixed, apply the same proportional change across gap field-years, and measure only the reemployment of displaced inventors. The 2.52 percent opportunity share is a benchmark for comparison, not an efficient or socially optimal allocation. The exercise therefore speaks only to the scale of the change in specialist absorption required to substantially alter where mobility-induced bridges form.

5.5 Limitations

The field-level design identifies changes in where inventive activity occurs, not whether the inventor’s total output increases. Because event fixed effects compare fields within the same displacement event, they absorb changes common across the inventor’s fields. We also constructed an event-level analysis that aggregates patenting and bridge exposure across fields. Although the first stage is strong and the instrument does not predict pre-period output, the sample contains only about 7,700 events and the estimates are too imprecise to distinguish net creation from reallocation.

The estimates also do not identify welfare effects. Such an interpretation would require knowing whether bridges increase total invention, how the social value of invention differs across fields, and the costs and equilibrium effects of changing specialist employment. Forward citations and inventor density do not identify these objects.

Our results therefore support a narrower conclusion. Specifically, collaborator mobility creates new connections primarily to fields represented by the employers that absorb displaced inventors. Our design does not establish whether reallocating those connections would increase total innovation or social welfare.

6 Conclusion

We show that where an inventor searches next depends in part on where her close collaborators work. Using mass layoffs that move a prior co-inventor to a new employer, we find that focal inventors become more likely to file their first patent in fields in which the co-inventor’s new employer specializes. Most of this response occurs before the focal inventor has any employment relationship with that firm.

Several results help characterize this relationship. Employer specialization predicts field entry even after conditioning on the collaborator’s own observed patent history, and the relationship remains strong when the collaborator has never herself patented in the destination field. Entry patents are also disproportionately likely to cite prior work from the collaborator’s employer. These patterns are consistent with access to knowledge associated with the collaborator’s organization, although they do not distinguish that channel from knowledge the collaborator acquires after joining the firm and subsequently transmits through the collaboration tie.

The resulting changes in inventive direction are persistent. The displacement instrument predicts not only first entry into the destination field but also subsequent patents and forward citations in that field. These estimates do not establish an increase in total inventive output,

since greater activity in the destination field may reflect reallocation from other fields.

Finally, the fields reached through mobility depend on where new organizational bridges form. Sparse, high-citation fields account for a larger share of potential destination fields than of realized specialist landings. Displaced inventors are much less likely to join specialist employers in these fields, so mobility generates relatively few new connections to them. The evidence does not show that a formed bridge is less effective in such fields; instead, the main observable constraint is that these bridges are rarely formed.

Taken together, the results show that inventive knowledge can diffuse beyond firm boundaries through persistent collaboration ties. Labor-market movements change the organizations represented in those networks and, in turn, the technological fields to which inventors are connected. At the same time, this diffusion follows the existing distribution of specialist employment: mobility can create new channels of access without necessarily redirecting them toward fields in which specialist capacity is thin.

Our analysis does not identify whether these changes raise total invention, whether entry into some fields has greater social value than entry into others, or the costs of expanding specialist employment. An important direction for future work is therefore to understand how specialist capacity develops in emerging fields and how changes in that capacity affect the formation of knowledge connections and the direction of inventive activity.

References

- Daron Acemoglu. Directed technical change. *Review of Economic Studies*, 69(4):781–809, 2002.
- Daron Acemoglu, Philippe Aghion, Leonardo Bursztyn, and David Hémous. The environment and directed technical change. *American Economic Review*, 102(1):131–166, 2012. doi: 10.1257/aer.102.1.131.
- Rajshree Agarwal, Raj Echambadi, April M. Franco, and MB Sarkar. Knowledge transfer through inheritance: Spin-out generation, development, and survival. *Academy of Management Journal*, 47(4):501–522, 2004.
- Ajay Agrawal, Iain Cockburn, and John McHale. Gone but not forgotten: Knowledge flows, labor mobility, and enduring social relationships. *Journal of Economic Geography*, 6(5):571–591, 2006.
- Ajay Agrawal, Devesh Kapur, and John McHale. How do spatial and social proximity influence knowledge flows? evidence from patent data. *Journal of Urban Economics*, 64(2):258–269, 2008. doi: 10.1016/j.jue.2008.01.003.
- Gautam Ahuja. Collaboration networks, structural holes, and innovation: A longitudinal study. *Administrative Science Quarterly*, 45(3):425–455, 2000. doi: 10.2307/2667105.
- Gautam Ahuja and Curba Morris Lampert. Entrepreneurship in the large corporation: A longitudinal study of how established firms create breakthrough inventions. *Strategic Management Journal*, 22(6–7):521–543, 2001. doi: 10.1002/smj.176.
- Ufuk Akcigit and Sina T. Ates. Ten facts on declining business dynamism and lessons from endogenous growth theory. *American Economic Journal: Macroeconomics*, 13(1):257–298, 2021.
- Ufuk Akcigit and William R. Kerr. Growth through heterogeneous innovations. *Journal of Political Economy*, 126(4):1374–1443, 2018.
- Ufuk Akcigit, Santiago Caicedo, Ernest Miguelez, Stefanie Stantcheva, and Valerio Sterzi. Dancing with the stars: Innovation through interactions. Working Paper 24466, National Bureau of Economic Research, 2018.
- Paul Almeida and Bruce Kogut. Localization of knowledge and the mobility of engineers in regional networks. *Management Science*, 45(7):905–917, 1999.

- David B. Audretsch and Maryann P. Feldman. R&D spillovers and the geography of innovation and production. *American Economic Review*, 86(3):630–640, 1996.
- Pierre Azoulay, Christian Fons-Rosen, and Joshua S. Graff Zivin. Does science advance one funeral at a time? *American Economic Review*, 109(8):2889–2920, 2019.
- Alex Bell, Raj Chetty, Xavier Jaravel, Neviana Petkova, and John Van Reenen. Who becomes an inventor in america? The importance of exposure to innovation. *Quarterly Journal of Economics*, 134(2):647–713, 2019.
- Nicholas Bloom, John Van Reenen, and Heidi Williams. A toolkit of policies to promote innovation. *Journal of Economic Perspectives*, 33(3):163–184, 2019.
- Stefano Breschi and Francesco Lissoni. Mobility of skilled workers and co-invention networks: An anatomy of localized knowledge flows. *Journal of Economic Geography*, 9(4):439–468, 2009.
- Eric Budish, Benjamin N. Roin, and Heidi Williams. Do firms underinvest in long-term research? Evidence from cancer clinical trials. *American Economic Review*, 105(7):2044–2085, 2015.
- Francisco J. Buera and Ezra Oberfield. The global diffusion of ideas. *Econometrica*, 88(1):83–114, 2020.
- Ronald S. Burt. *Structural Holes: The Social Structure of Competition*. Harvard University Press, Cambridge, MA, 1992.
- Ronald S. Burt. Structural holes and good ideas. *American Journal of Sociology*, 110(2):349–399, 2004. doi: 10.1086/421787.
- Benjamin A. Campbell, Russell Coff, and David Kryscynski. Rethinking sustained competitive advantage from human capital. *Academy of Management Review*, 37(3):376–395, 2012.
- Gianluca Carnabuci and Elisa Operti. Where do firms’ recombinant capabilities come from? intraorganizational networks, knowledge, and firms’ ability to innovate through technological recombination. *Strategic Management Journal*, 34(13):1591–1613, 2013. doi: 10.1002/smj.2084.
- Doruk Cengiz, Arindrajit Dube, Attila Lindner, and Ben Zipperer. The effect of minimum wages on low-wage jobs. *Quarterly Journal of Economics*, 134(3):1405–1454, 2019.

- Wesley M. Cohen and Daniel A. Levinthal. Absorptive capacity: A new perspective on learning and innovation. *Administrative Science Quarterly*, 35(1):128–152, 1990.
- Rafael A. Corredoira and Lori Rosenkopf. Should auld acquaintance be forgot? The reverse transfer of knowledge through mobility ties. *Strategic Management Journal*, 31(2):159–181, 2010.
- Elias Einiö, Josh Feng, and Xavier Jaravel. Social push and the direction of innovation. Working paper, 2022.
- Lee Fleming. Recombinant uncertainty in technological search. *Management Science*, 47(1):117–132, 2001.
- Lee Fleming, Santiago Mingo, and David Chen. Collaborative brokerage, generative creativity, and creative success. *Administrative Science Quarterly*, 52(3):443–475, 2007. doi: 10.2189/asqu.52.3.443.
- Martin Ganco, Rosemarie H. Ziedonis, and Rajshree Agarwal. More stars stay, but the brightest ones still leave: Job hopping in the shadow of patent enforcement. *Strategic Management Journal*, 36(5):659–685, 2015.
- Albrecht Glitz. Coworker networks in the labour market. *Labour Economics*, 44:218–230, 2017.
- Bronwyn H. Hall, Adam B. Jaffe, and Manuel Trajtenberg. The NBER patent citations data file: Lessons, insights and methodological tools. Working Paper 8498, National Bureau of Economic Research, 2001.
- W. Walker Hanlon. Necessity is the mother of invention: Input supplies and directed technical change. *Econometrica*, 83(1):67–100, 2015. doi: 10.3982/ECTA10811.
- Kyle Herkenhoff, Jeremy Lise, Guido Menzio, and Gordon M. Phillips. Production and learning in teams. *Econometrica*, 92(2):467–504, 2024.
- Karin Hoisl. Tracing mobile inventors—the causality between inventor mobility and inventor productivity. *Research Policy*, 36(5):619–636, 2007. doi: 10.1016/j.respol.2007.01.009.
- Adam B. Jaffe, Manuel Trajtenberg, and Rebecca Henderson. Geographic localization of knowledge spillovers as evidenced by patent citations. *Quarterly Journal of Economics*, 108(3):577–598, 1993.

- Xavier Jaravel, Neviana Petkova, and Alex Bell. Team-specific capital and innovation. *American Economic Review*, 108(4–5):1034–1073, 2018. doi: 10.1257/aer.20151184.
- Gregor Jarosch, Ezra Oberfield, and Esteban Rossi-Hansberg. Learning from coworkers. *Econometrica*, 89(2):647–676, 2021.
- Benjamin F. Jones. The burden of knowledge and the “Death of the Renaissance Man”: Is innovation getting harder? *Review of Economic Studies*, 76(1):283–317, 2009.
- Riitta Katila and Gautam Ahuja. Something old, something new: A longitudinal study of search behavior and new product introduction. *Academy of Management Journal*, 45(6):1183–1194, 2002.
- William R. Kerr. Ethnic scientific communities and international technology diffusion. *Review of Economics and Statistics*, 90(3):518–537, 2008. doi: 10.1162/rest.90.3.518.
- Bruce Kogut and Udo Zander. Knowledge of the firm, combinative capabilities, and the replication of technology. *Organization Science*, 3(3):383–397, 1992.
- Robert E. Lucas and Benjamin Moll. Knowledge growth and the allocation of time. *Journal of Political Economy*, 122(1):1–51, 2014.
- Charles F. Manski. Identification of endogenous social effects: The reflection problem. *Review of Economic Studies*, 60(3):531–542, 1993.
- Charles F. Manski and Steven R. Lerman. The estimation of choice probabilities from choice-based samples. *Econometrica*, 45(8):1977–1988, 1977.
- James G. March. Exploration and exploitation in organizational learning. *Organization Science*, 2(1):71–87, 1991.
- Matt Marx, Deborah Strumsky, and Lee Fleming. Mobility, skills, and the Michigan non-compete experiment. *Management Science*, 55(6):875–889, 2009.
- John K. Mawdsley and Deepak Somaya. Employee mobility and organizational outcomes: An integrative conceptual framework and research agenda. *Journal of Management*, 42(1):85–113, 2016.
- Sophia Mo and Xinyue Lin. Does mobility beget mobility? Coworker networks and the sectoral reallocation of labor. Working paper, Harvard University, 2025.
- Sophia Mo, Xinyue Lin, and Armando Miano. Coworker influence on job choice: Information, connection, and industry switching. Working paper, Harvard University, 2026.

- Kyle Myers. The elasticity of science. *American Economic Journal: Applied Economics*, 12(4):103–134, 2020.
- Richard R. Nelson and Sidney G. Winter. *An Evolutionary Theory of Economic Change*. Belknap Press of Harvard University Press, Cambridge, MA, 1982.
- Neus Palomeras and Eduardo Melero. Markets for inventors: Learning-by-hiring as a driver of mobility. *Management Science*, 56(5):881–895, 2010. doi: 10.1287/mnsc.1090.1131.
- Jesse Perla and Christopher Tonetti. Equilibrium imitation and growth. *Journal of Political Economy*, 122(1):52–76, 2014.
- Corey Phelps, Ralph Heidl, and Anu Wadhwa. Knowledge, networks, and knowledge networks: A review and research agenda. *Journal of Management*, 38(4):1115–1166, 2012. doi: 10.1177/0149206311432640.
- David Popp. Induced innovation and energy prices. *American Economic Review*, 92(1):160–180, 2002. doi: 10.1257/000282802760015658.
- Walter W. Powell, Kenneth W. Koput, and Laurel Smith-Doerr. Interorganizational collaboration and the locus of innovation: Networks of learning in biotechnology. *Administrative Science Quarterly*, 41(1):116–145, 1996. doi: 10.2307/2393988.
- Marta Prato. The global race for talent: Brain drain, knowledge transfer, and growth. *Quarterly Journal of Economics*, 140(1):165–238, 2025. doi: 10.1093/qje/qjae040.
- Lori Rosenkopf and Paul Almeida. Overcoming local search through alliances and mobility. *Management Science*, 49(6):751–766, 2003.
- Lori Rosenkopf and Atul Nerkar. Beyond local search: Boundary-spanning, exploration, and impact in the optical disk industry. *Strategic Management Journal*, 22(4):287–306, 2001.
- Michel Serafinelli. “Good” firms, worker flows, and local productivity. *Journal of Labor Economics*, 37(3):747–792, 2019.
- Brit Sharoni. The effect of inventor mobility on network productivity. Working paper, University of Michigan Ross School of Business, 2026.
- Jasjit Singh. Collaborative networks as determinants of knowledge diffusion patterns. *Management Science*, 51(5):756–770, 2005.

- Deepak Somaya, Ian O. Williamson, and Natalia Lorinkova. Gone but not lost: The different performance impacts of employee mobility between cooperators versus competitors. *Academy of Management Journal*, 51(5):936–953, 2008. doi: 10.5465/AMJ.2008.34789660.
- Jaeyong Song, Paul Almeida, and Geraldine Wu. Learning-by-hiring: When is mobility more likely to facilitate interfirm knowledge transfer? *Management Science*, 49(4):351–365, 2003.
- Toby E. Stuart and Joel M. Podolny. Local search and the evolution of technological capabilities. *Strategic Management Journal*, 17(S1):21–38, 1996.
- Peter Thompson and Melanie Fox-Kean. Patent citations and the geography of knowledge spillovers: A reassessment. *American Economic Review*, 95(1):450–460, 2005. doi: 10.1257/0002828053828509.
- Marco Tortoriello, Ray Reagans, and Bill McEvily. Bridging the knowledge gap: The influence of strong ties, network cohesion, and network range on the transfer of knowledge between organizational units. *Organization Science*, 23(4):1024–1039, 2012. doi: 10.1287/orsc.1110.0688.
- Fabian Waldinger. Peer effects in science: Evidence from the dismissal of scientists in Nazi Germany. *Review of Economic Studies*, 79(2):838–861, 2012.
- Paolo Zacchia. Knowledge spillovers through networks of scientists. *Review of Economic Studies*, 87(4):1989–2018, 2020.

A Sample Composition

This appendix reports distributions of the outcome and bridge measures over the estimation sample, and the construction of the WARN displacement sample used in Section 4.2.

Table A1: Summary Statistics

	Mean	SD	p25	Median	p75
<i>Panel A. Inventors</i> ($N = 4,256,890$)					
Patents	5.58	16.70	1	2	5
Career span (years)	6.87	9.25	1	2	9
Distinct fields entered	3.62	4.51	1	2	4
Co-inventors (network degree)	8.41	14.92	2	4	9
<i>Panel B. Employers</i> ($N = 166,095$ firms)					
Patents (cumulative)	35.7	829.4	1	2	5
Distinct fields	3.44	11.35	1	1	2
Inventors	25.3	395.9	1	3	6
<i>Panel C. Technology fields</i> ($N = 27,477$ field-years, 643 subclasses)					
Active inventors (density)	946.6	3279.4	67	216	685
Five-year citation rate	2.04	2.72	0.10	1.37	3.23
<i>Panel D. Patents</i> ($N = 8,085,002$)					
Forward citations, five-year	3.18	14.48	0	1	3
Forward citations, total	15.79	56.96	2	5	13

Notes: Panel A covers all disambiguated inventors of granted U.S. utility patents; distinct fields and co-inventor degree are available for the 4.06 and 3.72 million inventors, respectively, with at least one classified field or co-inventor. Panel B covers firms (Revelio `rcids` with at least one linked patent); the inventor count is available for the 157,917 firms with at least one matched inventor. Panel C is at the CPC-subclass $\times$ year level over the analysis window; density is the number of active inventors and the citation rate is the mean five-year forward-citation count for patents granted in $[t-5, t-1]$. By field-year, the citation-rate $\times$ density taxonomy of Section 3.6 splits as sparse, high citation (“gap”) 17.3 percent; dense, high citation 32.7 percent; sparse, low citation 32.7 percent; and dense, low citation 17.3 percent. Weighting field-years by at-risk cells reduces the gap share to 3.1 percent. Panel D covers the patents used to construct the citation distributions.

Table A2: Regression Sample and Bridge Measures

	Unweighted			Design-weighted		
	Mean	SD	Share > 0	Mean	SD	Share > 0
CrossBridge	0.0017	0.0142	0.051	0.0019	0.0152	0.058
SameBridge	0.0026	0.0200	0.036	0.0032	0.0226	0.044
OwnSpec	0.0118	0.0440	0.551	0.0121	0.0447	0.532
d_{if}	0.0771	0.1556	0.572	0.0290	0.1024	0.215
PeerExp	0.1118	0.1805	0.578	0.0420	0.1231	0.217

Notes: Distributions over the 136,624,908 at-risk (inventor $\times$ field $\times$ year) cells of `at_risk_union_panel_v3`, spanning 2,289,682 inventors and 643 subclasses. The design weights are those of Section 3.4, so the weighted columns describe the population the regressions estimate and are the moments quoted in Section 4.1.3. Conditional on being positive, CrossBridge averages 0.033 under either weighting, SameBridge 0.071 and 0.073, OwnSpec 0.021 and 0.023, d_{if} 0.135, and PeerExp 0.193. The design-weighted first-entry rate is 0.386 percent. Source: `Code/bridge_moments.py`. Of these cells, 125.4 million are connected only through a peer’s patenting (weighted entry rate 0.35 percent), 7.08 million are connected through both patenting and a strong employer bridge (1.31 percent), and 4.16 million are connected only through an employer bridge (0.22 percent).

Table A3: The WARN Displacement Sample

	Count
WARN notices, 1989–2026	63,185
linked to a firm identifier (%)	74.9
workers per notice (mean / median)	111 / 64
Displaced inventor-movers	76,739
distinct inventors	68,921
Strong-tie rewiring events	217,413
<i>Event-study sample (notice years 2000–2014)</i>	
gained-side (NEW) cells	88,732
lost-side (OLD) cells	98,281
counterfactual-side (NEVER) cells	228,897
Landing-composition IV sample (event $\times$ field)	526,462

Notes: WARN notices are from Revelio Labs (Section 3.5). A displaced mover is an inventor at a WARN firm in the notice year who is at a different employer the following year; a strong-tie rewiring event is such a move in which the mover is a strong co-inventor (≥ 2 shared patents in the $[T - 7, T - 3]$ window) of a focal inventor. The event-study and IV samples restrict to notice years 2000–2014 so the five-year outcome window is fully observed; side-cell counts are inventor $\times$ field cells entering the stacked design of Table/Figure 3.

B Mechanism Tests

This appendix reports the full estimates for the mechanism tests discussed in Section 4.3. Table A4 follows the column structure of Table 1, adding the controls of specification (11) one at a time. All estimates are weighted least squares with inventor×field, field×year, and inventor×year fixed effects on the 136.6-million-cell panel, with standard errors clustered by inventor.

Table A4: Cross-Firm Bridges by Peer Location

	(1)	(2)	(3)
Cross-firm bridge × same state	0.0284*** (0.0019)	0.0343*** (0.0020)	0.0342*** (0.0020)
Cross-firm bridge × different state	0.0247*** (0.0022)	0.0294*** (0.0022)	0.0292*** (0.0022)
Cross-firm bridge × state unknown	0.0196*** (0.0012)	0.0253*** (0.0012)	0.0252*** (0.0012)
SameBridge	0.0418*** (0.0011)	0.0238*** (0.0013)	0.0231*** (0.0013)
OwnSpec		0.0352*** (0.0010)	0.0356*** (0.0010)
d_{if}			-0.0087*** (0.0001)
$\alpha_{if}, \alpha_{ft}, \alpha_{it}$	✓	✓	✓
Mean of Dep. Var	0.0039	0.0039	0.0039
Observations	136,624,908	136,624,908	136,624,908

Notes: The cross-firm bridge is decomposed into three mutually exclusive components by the location of the strong cross-firm peer relative to the focal inventor’s modal state: the same state, a different state, or no identifiable modal U.S. state. The three sum to the undifferentiated cross-firm bridge. Roughly half of inventors have no identifiable modal state, and an earlier version of this table assigned them to the different-state column; they are now reported separately. Columns add the controls of equation (11) one at a time, mirroring Table 1. Standard errors clustered by inventor. *, **, *** denote significance at the 10, 5, and 1 percent levels. Source: `Code/mechanism_splits_geo3.py`.




Figure A1: Dose-Response in Employer Specialization

Notes: Each point is the coefficient on one of four mutually exclusive tiers of the peer-employer's leave-one-out field specialization, entered simultaneously and each measured as the tie-weighted count of strong cross-firm peers whose employer field share falls in the stated range (so the coefficient is the entry response per unit of bridge weight at that tier). The regression includes SameBridge, OwnSpec, d_{if} , and the three-way fixed effects of equation (11); $n = 136,624,908$. Bars are 95 percent confidence intervals from inventor-clustered standard errors.

Table A5: Citations to the Bridged Employer's Prior Art

	(1) Baseline	(2) + log portfolio size	(3) Size-matched controls
Bridging employer (= 1)	0.0541*** (0.0007)	0.0381*** (0.0007)	0.0209*** (0.0010)
Entry-patent (event) FE	✓	✓	✓
Bridge cited rate	0.055	0.055	0.036
Control cited rate	0.001	0.001	0.006
Events	142,498	142,498	50,023
Candidate rows	3,024,358	3,024,358	—

Notes: The unit of observation is an (entry event $\times$ candidate employer) pair. For each first-entry event backed by a cross-firm bridge, candidates are the bridging employer(s) and up to twenty randomly sampled non-bridge f -specialist employers (field share ≥ 0.05 as of $t - 1$). The outcome is one if the inventor's entry patent cites at least one field- f patent affiliated with the candidate employer. All columns include entry-patent (event) fixed effects; column (2) adds the log of the candidate's cumulative field- f portfolio; column (3) restricts controls to specialists whose field portfolio is at least as large as the event's bridging employer. Standard errors clustered by inventor. *, **, *** denote significance at the 10, 5, and 1 percent levels.

C Heterogeneity


This appendix collects the heterogeneity estimates discussed in Section 4.4, one exhibit per split. All are estimated on the same panel and specification; the splits differ only in how the cross-firm bridge is partitioned.

Table A6: Unique versus Redundant Bridges

	(1)	(2)	(3)
Cross-firm bridge $\times$ redundant (≥ 2 peers)	0.0523*** (0.0017)	0.0635*** (0.0018)	0.0636*** (0.0018)
Cross-firm bridge $\times$ unique (1 peer)	0.0053*** (0.0010)	0.0079*** (0.0010)	0.0076*** (0.0010)
SameBridge	0.0443*** (0.0012)	0.0259*** (0.0013)	0.0253*** (0.0013)
OwnSpec		0.0370*** (0.0010)	0.0374*** (0.0010)
d_{if}			-0.0087*** (0.0001)
$\alpha_{if}, \alpha_{ft}, \alpha_{it}$	✓	✓	✓
Mean of Dep. Var	0.0039	0.0039	0.0039
Observations	136,624,908	136,624,908	136,624,908

Notes: A cell is redundant if two or more cross-firm strong-tie peers at distinct specialist employers point to the same field in the same year; the two components sum to the undifferentiated cross-firm bridge. Columns add the controls of equation (11) one at a time, mirroring Table 1. Estimates come from `heterogeneity_nested.py`; the column (3) coefficients reproduce the single-specification estimates reported previously. Standard errors clustered by inventor. *, **, *** denote significance at the 10, 5, and 1 percent levels.

Figure A2: Cross-Firm Bridges by the Inventor's Origin Field




Note: Each point is the coefficient on the cross-firm bridge interacted with an indicator for the inventor's dominant CPC section, measured as the section accounting for the largest share of her cumulative patents as of $t - 1$; the nine indicators are mutually exclusive and estimated jointly with *SameBridge*, *OwnSpec*, d_{if} and the three-way fixed effects ($n = 136,624,908$). Bars are 95 percent confidence intervals from inventor-clustered standard errors. Parentheses give each section's share of sample weight. The dashed line is the pooled estimate of 0.0288. A ninth category, inventors with no prior patents, carries negligible weight (0.002 percent) and is omitted from the figure.

Table A7: Bridges through Incumbent and Non-Incumbent Peers

	(1)	(2)	(3)
Cross-firm bridge $\times$ peer non-expert in f	0.0325*** (0.0018)	0.0463*** (0.0018)	0.0447*** (0.0018)
Cross-firm bridge $\times$ peer expert in f	0.0192*** (0.0011)	0.0212*** (0.0011)	0.0217*** (0.0011)
SameBridge	0.0424*** (0.0012)	0.0246*** (0.0013)	0.0239*** (0.0013)
OwnSpec		0.0360*** (0.0010)	0.0363*** (0.0010)
d_{if}			-0.0087*** (0.0001)
$\alpha_{if}, \alpha_{ft}, \alpha_{it}$	✓	✓	✓
Mean of Dep. Var	0.0039	0.0039	0.0039
Observations	136,624,908	136,624,908	136,624,908

Notes: A peer is an expert in f if she had at least one patent in the field before the anchor year. Columns add the controls of equation (11) one at a time, mirroring Table 1. Estimates come from `heterogeneity_nested.py`; the column (3) coefficients reproduce the single-specification estimates reported previously. Standard errors clustered by inventor. *, **, *** denote significance at the 10, 5, and 1 percent levels.

Figure A3: Cross-Firm Bridges by Destination Field



Note: As in Figure A2, but splitting on the CPC section of the field being entered. Section levels are absorbed by the field $\times$ year fixed effects, so the regression identifies only differences in the bridge slope. Separate binary comparisons for biology-related and artificial-intelligence-related fields yielded no detectable difference.

Table A8: Technological Proximity to the Destination Field

	(1)	(2)	(3)
Cross-firm bridge $\times$ near ($d_{if} > \text{median}$ positive)	0.0531*** (0.0025)	0.0586*** (0.0025)	0.0729*** (0.0026)
Cross-firm bridge $\times$ middle ($0 < d_{if} \leq \text{median}$)	0.0007 (0.0015)	0.0036** (0.0015)	0.0034** (0.0015)
Cross-firm bridge $\times$ far ($d_{if} = 0$)	0.0199*** (0.0011)	0.0261*** (0.0011)	0.0222*** (0.0011)
SameBridge	0.0417*** (0.0011)	0.0237*** (0.0013)	0.0229*** (0.0013)
OwnSpec		0.0354*** (0.0010)	0.0357*** (0.0010)
d_{if}			-0.0090*** (0.0001)
$\alpha_{if}, \alpha_{ft}, \alpha_{it}$	✓	✓	✓
Mean of Dep. Var	0.0039	0.0039	0.0039
Observations	136,624,908	136,624,908	136,624,908

Notes: d_{if} is the cosine similarity between i 's prior portfolio and field f ; the median is taken over cells with strictly positive overlap. The profile is U-shaped in every column, so we report three bins rather than the binary overlap-versus-none split, which averages over the middle bin. Columns add the controls of equation (11) one at a time, mirroring Table 1. Estimates come from `heterogeneity_nested.py`; the column (3) coefficients reproduce the single-specification estimates reported previously. Standard errors clustered by inventor. *, **, *** denote significance at the 10, 5, and 1 percent levels.

Table A9: Recency of the Collaborative Tie

	(1)	(2)	(3)
Cross-firm bridge $\times$ active within 4 yrs	0.0239*** (0.0011)	0.0295*** (0.0011)	0.0293*** (0.0011)
Cross-firm bridge $\times$ dormant ≥ 5 yrs	0.0222*** (0.0012)	0.0278*** (0.0012)	0.0277*** (0.0012)
SameBridge	0.0417*** (0.0011)	0.0237*** (0.0013)	0.0230*** (0.0013)
OwnSpec		0.0352*** (0.0010)	0.0356*** (0.0010)
d_{if}			-0.0087*** (0.0001)
$\alpha_{if}, \alpha_{ft}, \alpha_{it}$	✓	✓	✓
Mean of Dep. Var	0.0039	0.0039	0.0039
Observations	136,624,908	136,624,908	136,624,908

Notes: Recency is the lag between the anchor year and the last year the pair co-authored. All ties are strong ties (≥ 2 shared patents in the $[t-7, t-3]$ window). Columns add the controls of equation (11) one at a time, mirroring Table 1. Estimates come from `heterogeneity_nested.py`; the column (3) coefficients reproduce the single-specification estimates reported previously. Standard errors clustered by inventor. *, **, *** denote significance at the 10, 5, and 1 percent levels.

Table A10: Local Specialist Capacity in the Inventor's State

	(1)	(2)	(3)
Cross-firm bridge $\times$ thick local capacity	0.0432*** (0.0025)	0.0495*** (0.0025)	0.0494*** (0.0025)
Cross-firm bridge $\times$ middle tercile	0.0243*** (0.0024)	0.0306*** (0.0024)	0.0305*** (0.0024)
Cross-firm bridge $\times$ thin local capacity	0.0228*** (0.0022)	0.0283*** (0.0022)	0.0282*** (0.0022)
Cross-firm bridge $\times$ no U.S. state (non-U.S.)	0.0134*** (0.0012)	0.0183*** (0.0012)	0.0182*** (0.0012)
SameBridge	0.0419*** (0.0011)	0.0239*** (0.0013)	0.0232*** (0.0013)
OwnSpec		0.0353*** (0.0010)	0.0357*** (0.0010)
d_{if}			-0.0087*** (0.0001)
$\alpha_{if}, \alpha_{ft}, \alpha_{it}$	✓	✓	✓
Mean of Dep. Var	0.0039	0.0039	0.0039
Observations	136,624,908	136,624,908	136,624,908

Notes: Local capacity is the inventor's modal state's share of national inventor capacity in field f as of $t - 1$, binned at row-weighted terciles carrying 13.7, 13.8 and 13.7 percent of sample weight. The fourth row covers inventors with no identifiable modal U.S. state (58.8 percent of weight), 99.6 percent of whom have non-U.S. modal patent locations; it is reported rather than dropped so the split stays exhaustive. Columns add the controls of equation (11) one at a time, mirroring Table 1. Estimates come from `heterogeneity_nested.py`; the column (3) coefficients reproduce the single-specification estimates reported previously. Standard errors clustered by inventor. *, **, *** denote significance at the 10, 5, and 1 percent levels.

Table A11: Industry of the Peer's Employer

	(1)	(2)	(3)
Cross-firm bridge $\times$ bio employer	0.0271*** (0.0033)	0.0331*** (0.0033)	0.0331*** (0.0033)
other employers (bio)	0.0230*** (0.0010)	0.0285*** (0.0010)	0.0283*** (0.0010)
Cross-firm bridge $\times$ AI/software employer	0.0306*** (0.0053)	0.0387*** (0.0053)	0.0386*** (0.0053)
other employers (AI)	0.0231*** (0.0010)	0.0285*** (0.0010)	0.0284*** (0.0010)
$\alpha_{if}, \alpha_{ft}, \alpha_{it}$	✓	✓	✓
Mean of Dep. Var	0.0039	0.0039	0.0039
Observations	136,624,908	136,624,908	136,624,908

Notes: Industry is the peer-employer's Revelio NAICS code: bio is 3254* and 541714; AI/software is 5415* and 511210. The two cuts are estimated as separate regressions, each splitting the cross-firm bridge against its own out-group, because a peer employer may be neither and the two classifications are not mutually exclusive. Columns add the controls of equation (11) one at a time; *SameBridge*, *OwnSpec* and d_{if} are included as in Table A6 but omitted here for space. Contrast with Figure A3, which finds no differential by the destination field's technology area. Standard errors clustered by inventor. *, **, *** denote significance at the 10, 5, and 1 percent levels.

D Allocation Details

Table A12: Cross-Firm Bridge Slope by Field Type

	(1)	(2)	(3)	(4)
CrossBridge	0.0269*** (0.0010)	0.0269*** (0.0010)		0.0269*** (0.0010)
CrossBridge $\times$ Gap		-0.0001 (0.0054)		-0.0020 (0.0053)
$\times$ Gap (sparse, high citation)			0.0251*** (0.0054)	
$\times$ Dense, high citation			0.0278*** (0.0010)	
$\times$ Sparse, low citation			0.0186*** (0.0043)	
$\times$ Dense, low citation			0.0218*** (0.0019)	
OwnSpec	0.0408*** (0.0009)	0.0408*** (0.0009)	0.0408*** (0.0009)	0.0409*** (0.0009)
OwnSpec $\times$ Gap				-0.0131*** (0.0034)
d_{if}	-0.0087*** (0.0001)	-0.0087*** (0.0001)	-0.0087*** (0.0001)	-0.0087*** (0.0001)
$\alpha_{if}, \alpha_{ft}, \alpha_{it}$	✓	✓	✓	✓
Mean of Dep. Var	0.0039	0.0039	0.0039	0.0039
Observations	136,624,908	136,624,908	136,624,908	136,624,908

Notes: Table displays estimates of specification (11) in which the cross-firm bridge is interacted with the field-type classification of Section 3.6. Column (1) reports the pooled slope, column (2) adds an interaction with the gap indicator, column (3) reports separate slopes for all four field types, and column (4) adds the corresponding interaction for the inventor's own employer specialization as a benchmark. *SameBridge* is included throughout. Field $\times$ year effects absorb the field-type indicators, so the interactions are identified from differences in the bridge slope across types rather than from differences in entry levels. Standard errors are clustered at the inventor level. *, **, and *** denote significance at the 10, 5, and 1 percent levels.

Table A13: Bridge Formation in Sparse, High-Citation Fields

Quantity	Share
Specialist-landing rate in gap cells	0.135%
Specialist-landing rate in non-gap cells	1.457%
Gap share of realized specialist landings	0.240%
Gap share of estimated specialist absorption capacity	0.029%
Gap share of at-risk event-field cells	2.525%

Notes: Statistics are calculated over the 526,462 event-field cells of the landing-composition sample, aggregated by field-year. A gap is a sparse, high-citation field-year as defined in Section 3.6, attached to a cell in the landing year, which is the classification the regressions of Table 9 use. The specialist-landing rate divides cells in which the displaced peer lands at a field specialist by at-risk cells within the indicated field type; there are 18 such landings in gap cells and 7,476 elsewhere. The realized-landing and absorption-capacity shares use all fields as the denominator. Absorption capacity is the product of the number of specialist employers and their average absorption of displaced inventors in the field-year. Source: `Code/build_gap_ledger.py`.

E Additional Robustness

Table A14: Cluster Structure, Inference, and Repeated Events

	Clusters	Median size	Largest share	Effective clusters	RF t	FS F
<i>Panel A. One-way clustering</i>						
Displacement market	84	1,678	11.4%	18.4	3.72	11.2
WARN firm $\times$ year	186	748.5	9.6%	34.0	3.85	12.9
Displaced peer	2,342	96	1.7%	533.9	5.80	28.4
Focal inventor	7,606	49	0.2%	3,421.4	5.10	134.4
Event	10,328	46	0.0%	6,355.4	6.14	171.1
<i>Panel B. Two-way clustering</i>						
Market $\times$ focal inventor					3.68	11.2
Market $\times$ displaced peer					3.72	11.2
<i>Panel C. Repeated events</i>						
	Cells	Events	Reduced form		t	FS F
Baseline	526,462	10,328	0.0303		3.72	11.2
First displacement per inventor	475,398	9,405	0.0296		3.07	11.3
Non-overlapping windows	392,197	7,654	0.0279		2.99	14.2
One inventor per peer	115,676	2,342	0.0818		3.38	3.8
One random peer, 200 draws			0.0266		[0.0248, 0.0283]	

Notes: Panel A describes the dimensions along which event–field observations repeat, over the 526,462 cells of Table 3; the displacement market is a state $\times$ old-firm technology class $\times$ notice year cell and the event is an inventor $\times$ displaced peer $\times$ year triple. “Largest share” is the share of cells in the biggest cluster and “effective clusters” is the inverse Herfindahl index of cluster sizes, which discounts for unevenness. The last two columns give the reduced-form t statistic and the effective first-stage F under that clustering; for a single instrument the effective F is the cluster-robust first-stage F . Panel B applies the Cameron–Gelbach–Miller variance $V_A + V_B - V_{A \cap B}$. A wild cluster bootstrap of the reduced form, drawing Rademacher weights at the market level with the null imposed over 1,999 replications, gives $p = 0.0005$, and the Anderson–Rubin 95 percent set for the 2SLS estimate is $[0.119, 0.387]$. Panel C re-estimates the reduced form under rules that remove repeated events, re-absorbing the fixed effects within each subsample; the final row reports the mean across 200 random draws of one displaced collaborator per focal inventor with the fifth and ninety-fifth percentiles beside it, and every draw was positive. Source: `Code/warn_landing_iv_inference.py`.

Table A15: Continued Collaboration and Destination-Firm Affiliation

	(1)	(2)	(3)	(4)	(5)
<i>Panel A. Landing-composition reduced form</i>					
Reduced form: y on Z	0.0303*** (0.0081) [0.0059]	0.0304*** (0.0081) [0.0059]	0.0282*** (0.0075) [0.0054]	0.0262*** (0.0073) [0.0054]	0.0268*** (0.0071) [0.0053]
Excl. co-invention with peer		✓			✓
Excl. landing-firm co-inventor			✓		✓
Excl. inventor at landing firm				✓	✓
Entries counted	3,133	3,082	2,766	2,801	2,677
AR 95% CI	[0.12, 0.39]	[0.12, 0.39]	[0.10, 0.37]	[0.09, 0.34]	[0.10, 0.36]
Observations	526,462	526,462	526,462	526,462	526,462
<i>Panel B. Stacked event study, new versus old fields</i>					
Post 2–5 vs pre (per 1,000)	1.108 [2.84]	1.183 [2.89]	1.096 [2.81]	1.098 [2.81]	1.182 [2.89]
Trend break (pp per year)	0.083 [3.32]	0.075 [2.99]	0.084 [3.34]	0.083 [3.31]	0.075 [2.99]

Notes: Each column redefines the outcome to exclude entries on the grounds checked in the lower block, holding the sample, the instrument and the first stage (0.1482, $F = 11.2$) fixed, so that only the outcome varies across columns. An entry is excluded when its patent is co-authored with the displaced peer, when another inventor on it is employed at the peer's landing firm in the filing year, or when the focal inventor is herself recorded at that firm; a cell still counts as an entry if any of its first-year patents survives. Standard errors clustered by displacement market in parentheses and by focal inventor in brackets. Panel B applies the same exclusions to the stacked event study of Figure 3, recomputing first-entry years from the surviving patents so that both the hazard and its at-risk denominator adjust; the trend break excludes the move year from both trend segments and its t statistic is clustered by focal inventor. Dropping instead the 29 percent of events in which the inventor is ever employed at the landing firm during $[T+1, T+6]$ leaves 373,148 cells and gives reduced forms of 0.0362 ($t = 4.28$) on the baseline outcome and 0.0352 ($t = 4.37$) under all three exclusions. Source: Code/warn_landing_iv_exclusions.py, Code/warn_event_study_exclusions.py.

Table A16: Robustness of the Landing-Composition Instrument

	(1)	(2)	(3)	(4)
OLS: y on D	0.0132*** [5.15]	0.0135*** [5.07]	0.0139*** [5.31]	0.0125*** [4.46]
First stage: D on Z	0.1482 [3.35]	0.1687 [3.68]	0.1572 [3.17]	0.1585 [2.85]
Reduced form: y on Z	0.0303*** [3.72]	0.0309*** [3.07]	0.0379*** [4.23]	0.0384*** [3.97]
2SLS: y on D	0.2044	0.1834	0.2411	0.2426
Distance restriction	State	Census region	State	State
Min. leave-out movers	≥ 5	≥ 5	≥ 10	≥ 20
Event, field $\times$ year FE	✓	✓	✓	✓
First-stage F	11.2	13.5	10.0	8.1
AR 95% CI	[0.12, 0.38]	[0.09, 0.33]	[0.15, 0.51]	[0.15, 0.60]
Observations	526,462	377,819	500,475	476,548

Notes: Each column re-estimates the four specifications of Table 3 under the sample definition shown in the lower block. The distance restriction requires the focal inventor's modal state, or in column (2) her census region, to differ from the displacement market's; the minimum leave-out movers row raises the number of other displaced movers a market cell must contain. t statistics clustered by displacement market in brackets. The Anderson–Rubin set is robust to weak instruments and excludes zero in every column. *, **, and *** denote significance at the 10, 5, and 1 percent levels.

Table A17: Robustness to Bridge Construction Thresholds

	(1)	(2)	(3)
<i>Panel A. Baseline: specialist share $\geq 5\%$, ≥ 2 shared patents</i>			
CrossBridge	0.0233*** (0.0010)	0.0289*** (0.0010)	0.0288*** (0.0010)
SameBridge	0.0417*** (0.0011)	0.0237*** (0.0013)	0.0230*** (0.0013)
OwnSpec		0.0352*** (0.0010)	0.0356*** (0.0010)
d_{if}			-0.0087*** (0.0001)
<i>Panel B. Looser specialist threshold: $\geq 3\%$</i>			
CrossBridge	0.0253*** (0.0010)	0.0308*** (0.0010)	0.0307*** (0.0010)
SameBridge	0.0454*** (0.0012)	0.0276*** (0.0013)	0.0271*** (0.0013)
OwnSpec		0.0345*** (0.0010)	0.0349*** (0.0010)
d_{if}			-0.0087*** (0.0001)
<i>Panel C. Tighter specialist threshold: $\geq 10\%$</i>			
CrossBridge	0.0196*** (0.0010)	0.0250*** (0.0010)	0.0250*** (0.0010)
SameBridge	0.0347*** (0.0012)	0.0166*** (0.0013)	0.0161*** (0.0013)
OwnSpec		0.0365*** (0.0010)	0.0369*** (0.0010)
d_{if}			-0.0087*** (0.0001)
<i>Panel D. Deeper ties: ≥ 3 shared patents</i>			
CrossBridge	0.0260*** (0.0013)	0.0324*** (0.0013)	0.0325*** (0.0013)
SameBridge	0.0449*** (0.0013)	0.0307*** (0.0014)	0.0305*** (0.0014)
OwnSpec		0.0359*** (0.0009)	0.0362*** (0.0009)
d_{if}			-0.0087*** (0.0001)
<i>Panel E. Deepest ties: ≥ 4 shared patents</i>			
CrossBridge	0.0298*** (0.0017)	0.0367*** (0.0017)	0.0368*** (0.0017)
SameBridge	0.0495*** (0.0017)	0.0351*** (0.0018)	0.0351*** (0.0018)
OwnSpec		0.0366*** (0.0009)	0.0368*** (0.0009)
d_{if}			-0.0087*** (0.0001)

Notes: Each panel rebuilds the bridge measure with the thresh-