

Organizational Bridges and the Direction of Inventive Search*

Sophia Mo

September 2026

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

*Harvard University, yiningmo@fas.harvard.edu. All errors are my 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 tacit, cumulative, and increasingly burdensome 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 core questions regarding the transmission of firm knowledge on innovation to one’s collaborators. First, does having a strong collaborator inside an outside specialist firm 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 land.

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 (Figure 2). In regressions saturated with $\text{inventor} \times \text{field}$, $\text{field} \times \text{year}$, and $\text{inventor} \times \text{year}$ fixed effects, the cross-firm bridge coefficient is 0.029 ($t = 30$), and, critically, it *survives* 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. The association is also distinct from the peer’s personal knowledge: in a horse race between the peer’s employer’s specialization and the peer’s own patent stock in the field, both enter with essentially undiminished coefficients, and the employer-side signal is roughly sixty percent larger. Absorbing community-level technology shocks with $\text{peer-community} \times \text{field} \times \text{year}$ fixed effects leaves a coefficient of 0.019 ($t = 19$).

These estimates control for a great deal, but they are not causal, and the reason is one we can measure. A standing bridge is often a bridge to a firm the inventor herself once worked at or will later join, so we split the association by the timing of the inventor’s own employment at the bridging firm relative to her entry. The cross-firm coefficient is carried by fields at firms the inventor had already worked at, or moved to in the same year as entry; restricted to entries that occur *before* she has any employment contact with the bridging firm, which is the arm’s-length margin the mechanism is about, it is roughly a fifth of the pooled figure.² Standing bridges are contaminated by the inventor’s own organizational trajectory, a confound no feasible fixed-effects structure removes, and this is what motivates exogenous variation in the network itself.

²Section 4.2.1 reports the decomposition. Roughly a fifth of the bridge weight lies in cells where the inventor had already worked at the bridging firm or moves there in the same year, and those cells carry coefficients an order of magnitude above the pooled estimate.

The second and central part of the paper identifies the causal, arm’s-length effect from displacement-induced rewiring of the network. We link WARN advance-notice mass layoffs to inventor employer spells and study the 217,000 events in which a displaced worker is a strong co-inventor of a focal inventor and moves to a new employer, forming a bridge to the new employer’s fields and severing one to the old employer’s fields while the focal inventor’s own employment is unchanged. A stacked within-event design compares entry into the landing firm’s specialist fields, the losing firm’s specialist fields, and the specialist fields of counterfactual landing firms, meaning the employers absorbing other inventors displaced in the same year. Entry into every side of the comparison decays as the collaboration recedes in time; what the displacement changes is the *slope*. Newly gained bridges arrest the decay of entry into the landing firm’s fields, producing a trend break of +0.08 percentage points per year ($t = 3.3$), while lost bridges continue to decay. To address selection into landing firms, we instrument the landing-field composition with the absorption composition of the displacement market, meaning the fields of the employers hiring *other* same-market, same-broad-technology displaced inventors, excluding the peer’s own firm and the focal inventor’s other displaced peers, and restricting to focal inventors in distant states. The reduced form is 0.030 ($t = 3.7$, clustered by market): a one-standard-deviation increase in the instrument raises the five-year probability of a first patent in the field by 0.21 percentage points against a mean of 0.60 percent, and Anderson–Rubin confidence sets exclude zero in every specification. We emphasize the reduced form and these interval estimates rather than the two-stage point magnitude. When we randomly reassign inventor–collaborator links within field-by-year cells, so that the specific tie is broken but the field composition of the instrument is held fixed, the relationship disappears entirely; whatever the instrument captures is specific to the tie and not a technology boom that would draw in any distant inventor.

The same decomposition that undercuts the association supports the causal flow. Applying the identical timing split to the displacement estimates, the causal effect is concentrated where the association was weakest. Most of it appears *before* the inventor has any employment contact with the peer’s new firm; it *strengthens* when inventors with prior contact are removed, the reduced form rising from 0.030 to 0.034; and the component that coincides with a move by the inventor herself is small and cannot be distinguished from zero, which is not what one would expect if the effect were her own relocation in disguise. The same method thus separates two objects that the raw correlation conflates: a confounded *stock*, in which standing bridges track the inventor’s own organizational history, and a clean *flow*, in which an exogenous new bridge transmits access to a collaborator’s employer’s field-specific knowledge at arm’s length, before and independent of any move by the inventor.

What travels over the tie appears to be organizational rather than personal. Bridges

through collaborators who have never themselves patented in the destination field are roughly twice as strong as bridges through collaborators who are incumbents there, the association rises monotonically with how specialized the employer is, and entry patents cite the bridged firm’s prior art in the field far more often than they cite the prior art of other specialist firms to which the inventor has no connection. That last difference survives restricting the comparison group to firms with portfolios at least as large and to firms that never employed the inventor. These tests do not observe what collaborators say to one another, but they are difficult to reconcile with an account in which the collaborator’s own accumulated expertise is doing the work.

The induced entries are consequential, not cosmetic. Re-estimating the same reduced form with unconditional destination-field outcomes, a newly formed bridge raises not only first entry but subsequent patenting in the connected field and the five-year forward citations it generates, and roughly seventy percent of the induced entries persist rather than lapsing after a single patent. Displacement-induced rewiring therefore causes durable, cited inventive output, not a one-time relabeling. We stop short of a net claim and report destination-field outcomes in gross terms.³

The third part of the paper asks where bridge-induced entry goes. We classify field-years by inventor density and lagged citation rates, and we pay particular attention to sparse fields whose patents are highly cited, which we call gaps for brevity and without any implication that a marginal entrant is more socially valuable there. Such fields account for 2.5 percent of at-risk opportunities but only 0.24 percent of realized specialist landings, and in gap field-years the instrument’s first stage collapses toward zero, because displaced specialists in sparse fields barely exist to be reabsorbed. Our estimates are too imprecise to say whether a bridge, once formed, works differently in such fields; what they establish is that the binding constraint is on formation rather than on transmission. Bridge-mediated diffusion spreads knowledge along the existing geography of specialist organizations. Policies that raise mobility or connectivity will therefore increase the volume of diffusion without necessarily changing its technological composition, and directionality, if it is wanted, has to come from the supply of specialist capacity, which our design does not study.

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](#); [Palom-](#)

³An event-level design that would net this destination-field output against the inventor’s activity in her other fields has a strong first stage and no pre-trends, but at roughly seven thousand events it is too imprecise to separate net creation from reallocation.

eras 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: 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 rewires the fields that are accessible through the network. This perspective shifts 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 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; Mo et al. 2026) 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 bring these literatures together by distinguishing knowledge possessed by a contact from knowledge associated with the orga-

nization 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 Opperi 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: 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 add a distinct mechanism operating later in an inventor’s career: access to specialized organizations through existing collaborators changes which technological fields an inventor enters. This field-entry margin is an individual-level analogue of the extensive margin in growth models that distinguish improving existing product lines from acquiring new ones (Akcigit and Kerr 2018), and models of idea diffusion make opportunities to search, interact with, imitate, or combine the knowledge of others central to productivity growth (Lucas and Moll 2014; Perla and Tonetti 2014; Buera and Oberfield 2020; Akcigit and Ates 2021). Our evidence provides a microeconomic counterpart to this diffusion technology: who gains access to which knowledge depends systematically on the organizational affiliations represented in an inventor’s collaboration network. It also behaves differently from a price or a subsidy in one respect that matters for policy. Because bridges form only where specialist employers exist to absorb displaced workers, raising connectivity increases the volume of diffusion along the existing distribution of specialist capacity without redirecting it toward fields where that capacity is thin (Bloom et al. 2019).

The remainder of the paper is organized as follows. Section 2 sets out a short account-

ing 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: it first documents the relationship between organizational bridges and field entry and why standing-network estimates cannot be read as the effect, then develops 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 sets out the framework the empirical work uses. It has three jobs. It defines organizational access, so that a displacement moving a collaborator from one employer to another becomes a change in the field-specific access available to an inventor who does not move. It states an entry decision under uncertainty, which is what makes the mechanism and heterogeneity results of Sections 4.3 and 4.4 tests of a structure rather than a list of correlations. And it separates the effect of a formed bridge from the process that determines where bridges form, which is the distinction the allocation results in Section 5 turn on. We state at the outset that the framework was written after the estimates and is meant to organize them, not to have predicted them, and we note below the findings it does not account for.


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

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.

Section 4.3.1.

2.2 Entry as a decision 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 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.

Three implications follow in the empirically relevant case in which entry is rare, so that the prior mean sits below the threshold. Entry occurs in 0.4 percent of at-risk cells.

First, entry rises with the employer’s degree of specialization. Because the prior is pessimistic relative to the threshold, entry happens only when signals are favorable, and raising precision widens the distribution of the posterior mean and puts more mass above the cost. The response should therefore be graded in s_{jft} rather than switching on at any particular level, which is what Section 4.3.2 examines, finding the association already positive below the five percent threshold used to build the measure.⁴

Second, returns to additional independent bridges are weakly diminishing rather than increasing. The dispersion of the posterior mean is concave in total precision, so a second bridge into a field moves entry by somewhat less than the first. This is the rule’s sharpest prediction, and it is the one place where the panel and the displacement design disagree. The panel shows fields reached by two or more cross-firm bridges carrying a coefficient roughly eight times that of fields reached by exactly one, which no rule of this kind can produce. The

⁴This is the one implication resting on panel evidence alone. The displacement design cannot separate degrees of specialization, because a market that absorbs displaced workers into strongly specialized employers in a field also absorbs them into marginally specialized ones: splitting the instrument by the landing employer’s field share leaves cross-effects in the first stage as large as the own-effects, so the bands are not separately identified.

displacement design finds the opposite ordering, and mildly: the effect of a formed bridge is -0.041 smaller (standard error 0.038) for an inventor who already reached the field through another specialist employer, against a main effect of 0.219 . Taking the causal estimate at face value, the second bridge is worth about four-fifths of the first, with a confidence interval running from roughly a half to just above one. The prediction is qualitative and the interval is wide, but it excludes amplification of the kind the panel suggests, and Section 4.4 reads the panel gap accordingly.

Third, transmission is approximately homogeneous across fields over the relevant range. Nothing in equations (2)–(4) makes the response to a formed bridge depend on which field it reaches, and the displacement design supports that: interacting the effect of a formed bridge with an indicator for sparse, highly cited fields gives -0.021 with a standard error of 0.205 . This is the restriction that makes the aggregation in Section 2.4 usable.

We state plainly that the framework was written alongside the estimates rather than before them. Its second implication is the one that has been put at risk and survived, since it contradicted the panel evidence available when the rule was written and the causal evidence later sided with the rule. Its first implication has not been tested against a design, and we do not treat it as established.

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

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

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 first term is a volume effect and the second a composition effect. A change that raises rewiring proportionally across fields moves $\mu(p)$ and leaves $q_f(p)$ fixed, whereas only a change in the distribution of absorption capacity $\{a_f\}$ moves $q_f(p)$. Homogeneous transmission sharpens this considerably. If $\tau_f = \tau$, then $\mathbb{E}[\Delta Y(p)] = M \mu(p) \tau$, which does not involve the routing shares at all, and the composition term vanishes because the shares sum to one. Where bridges form then determines *where* induced entry occurs and not *how much* of it there is: a policy that redistributes specialist capacity across fields moves entry between them, and only a policy that changes the number of bridges changes the total. If p also changes transmission conditional on formation, the derivative carries a further term, $M \mu(p) \sum_f q_f(p) d\tau_f/dp$, which the paper does not estimate.

This accounting also marks the limits of the allocation analysis. The paper estimates bridge formation, destination-field entry, persistence, and citation outcomes. It does not identify the net social value of entry, the resource cost of changing specialist capacity, or the extent to which destination-field output displaces invention in other fields. We therefore use equations (6)–(8) to describe the volume and composition of inventive entry, and not to report a welfare effect.

3 Data and Measurement

3.1 Patent data and technological fields

Patent data come 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.⁵ 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 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 a tie *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 $\times$

⁵See [Hall et al. \(2001\)](#) for the canonical description of patent data in economic research. We use utility patents only throughout.

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.⁶ 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 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:

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

⁶The 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.

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.⁷

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

⁷The 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.

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 citation rates 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 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 The motivating fact

We begin with the raw relationship in the inventor–field–year panel. We restrict attention to cells in which at least one prior co-inventor has already patented in the field, so that the comparison is within an inventor’s active collaborative neighborhood rather than across the whole space of technologies. Figure 2 classifies each such cell by its strongest connection to a specialist employer, and panel (a) shows a steep gradient. Entry occurs in 5.5 per 1,000 cells with no connection to a specialist firm, in 12.9 per 1,000 cells connected only through weak ties, in 16.0 per 1,000 cells connected by a strong tie to an outside specialist, and in 25.7 per 1,000 cells where the specialist employer is also the focal inventor’s own. The last category combines a collaborative connection with the inventor’s own direct organizational access, and its size is a first indication that the two need to be separated.


Panel (b) begins that separation by stratifying cells on how much the collaborators have themselves patented in the field, the most obvious alternative explanation for the gradient in panel (a). Within every quartile of collaborator patenting, entry remains roughly 2.5 to 3 times as common in cells with a strong cross-firm bridge as in cells with none. The bridge measure therefore carries information beyond what the collaborators have visibly done in the field. It does not, however, distinguish transmission through the tie from selection into which inventor–field cells acquire bridges in the first place, and the regressions that follow make that distinction as sharp as the panel allows.

4.1.2 Baseline specification

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

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

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.

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.⁸ 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

⁸Standard 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.

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×field×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.

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 are not 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 Causal evidence: displacement-induced rewiring

4.2.1 Organizational bridges and the inventor’s own employment history

A standing bridge to a specialist firm is frequently a bridge to a firm with which the focal inventor is herself connected over the course of her career, either because she worked there before or because she joins later. This simple fact governs how much of the panel association can plausibly be attributed to the tie. To measure it we split the cross-firm coefficient by the timing of the inventor’s *own* first employment at the bridging firm relative to her entry into the field, using her first-ever year at that firm. Every bridge then falls into one of four mutually exclusive cases: she never works there, she joins only after having entered the field, she joins in the same year that she enters, or she had already worked there before entering. The four components sum to the pooled estimate by construction, and their weights tell us how much of the identifying variation each case supplies.

Table 2 reports the decomposition. Fields at firms the inventor had already worked at carry a coefficient of 0.095 ($t = 34.2$), and fields she joins in the same year she enters carry 0.351 ($t = 18.7$), roughly three and twelve times the pooled estimate of 0.029. Between them these two cases account for only a fifth of the bridge weight. The remaining four-fifths, consisting of firms she never works at and firms she joins only after entering the field, carry 0.008 and -0.082 respectively, so that entries occurring before any employment contact with the bridging firm carry a coefficient of about 0.006, roughly a fifth of the pooled figure. The panel association is therefore concentrated in precisely the cells where the inventor’s own organizational history and her collaborator’s employment coincide, which are also the cells in which the two are least separable. We read this as a measurement result about what standing bridges encode rather than as a bound on transmission, since the confound and any true effect are both concentrated in the same cells. It is what leads us to identify the effect from displacements that change the collaborator’s employer while leaving the inventor’s own career untouched.

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.

4.2.2 Displacement as a source of variation in organizational access

The decomposition above implies that identifying the effect of an organizational bridge requires variation in the collaborator’s employer that is unrelated to the focal inventor’s own trajectory. We obtain it from mass layoffs. The timing of a layoff is determined by conditions at the collaborator’s employer, and it displaces her into a new firm whose specialist fields differ from the old one’s, while the focal inventor’s employment, location, and net-

work are unchanged.⁹ The remaining identification challenge is that displaced collaborators choose where to go next. We address it first with a within-event comparison and then with a leave-out measure of the displacement market’s field-specific absorption pattern.

4.2.3 WARN displacement and bridge formation

The variation we exploit is easiest to see in a single 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 $T + 1$. From i ’s perspective nothing about her own situation has changed, yet the move ends her contemporaneous bridge to the fields in which firm A specializes and creates one to the fields in which firm B specializes, so that the technological content of a relationship she did not alter has been rewritten by her collaborator’s labor-market outcome. Our first design compares these two sets of fields within the same displacement event, which holds fixed everything common to the event, and our second addresses the selection of firm B itself by using variation in the fields that happen to absorb other inventors displaced in the same market.

4.2.4 Stacked event study


For each displacement event, we define three sets of fields. The NEW set contains the landing firm’s specialist fields, calculated leave-one-out with respect to j . The OLD set contains the WARN firm’s specialist fields. The NEVER set contains fields associated with counterfactual landing firms, defined as employers that hire other inventors displaced in the same notice year. We drop fields that belong to both an actual and a counterfactual set. For each event time $\tau \in [-5, 5]$, we compute the first-filing hazard among fields in each set. We then form the within-event difference $d(e, \tau) = h_{\text{SIDE}}(e, \tau) - h_{\text{NEVER}}(e, \tau)$ and estimate event-time profiles by weighted least squares using inverse-variance weights. Standard errors are clustered by the focal inventor and, in a separate specification, by the displaced collaborator.¹⁰

Figure 3 reports the results. The effect appears as a change in trend rather than as a jump, and the shape is worth stating carefully. Entry hazards decline with event time for all three sets of fields, as the collaboration used to define the tie recedes further into the past and the relationship becomes less active. After the move, however, that decline is markedly

⁹We use displacement followed by reemployment rather than the more familiar severance shocks because of power and exclusion. Co-inventor deaths generate too few events involving strong ties to specialist employers to support a well-powered design, and firm closures and acquisitions are more common but plausibly reveal information about the expected value of the affected firm’s technologies.

¹⁰The within-event difference absorbs shocks shared across the field sets of an event, in the spirit of stacked designs (Cengiz et al. 2019). We use notice years 2000–2014 so that every cohort is observed over the full event window.

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.

slower for the landing firm’s fields than for the former employer’s. Panel (c) estimates the change in the new-minus-old trend at 0.083 percentage points per year ($t = 3.32$, clustered by focal inventor), and averaged over event years two through five the new-minus-old difference stands 1.11 entries per 1,000 at-risk cells above its pre-period mean ($t = 2.84$). What a displacement does, on this evidence, is not to trigger a burst of entry in the year it occurs but to arrest the natural decay of a tie’s influence in the fields the collaborator has moved toward.

4.2.5 Landing-composition instrument

The within-event comparison uses j ’s realized destination, which may be selected in ways correlated with i ’s technological trajectory. We therefore construct an instrument from the field composition of employers that absorb other displaced inventors in j ’s labor market. We define a displacement market m as a state $\times$ notice-year $\times$ dominant CPC class of the old firm. For each field f , the instrument is

$$Z_f(i, j, T) = \frac{\#\{\text{displaced movers in } m \text{ landing at } f\text{-specialists}\}^{(-)}}{\#\{\text{displaced movers in } m\}^{(-)}}, \quad (12)$$

where the superscript $(-)$ indicates two exclusions. We omit all movers from j ’s WARN firm and all movers who are themselves strong ties of focal inventor i . The second exclusion prevents another mover’s destination from creating a bridge for i directly. The endogenous

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.

treatment D_f equals one if j ’s destination employer specializes in f . The sample contains displacement-event $\times$ field cells in the leave-out pool, excluding the old firm’s specialist fields. The outcome is first filing in f within five years after the landing. We also require the focal inventor’s modal state to differ from the displacement state. Event fixed effects absorb differences across displacement events, while field $\times$ notice-year effects absorb national changes in field-specific entry. The remaining variation compares the field composition of absorption within a displacement market and event. We cluster standard errors by displacement market and, separately, by focal inventor. Because the market is the level at which the instrument varies, it is also the level with fewest clusters: the sample contains 84 markets, and they are unbalanced enough that the inverse Herfindahl index of their sizes is 18.4, with the largest market holding 11 percent of the cells. Market clustering is accordingly the most conservative of the schemes we can apply, and it is the one we report in the text; Appendix Table A14 collects the alternatives.

Table 3 reports the estimates. A larger leave-out absorption share in field f raises the probability that j joins an f -specialist, with a first-stage coefficient of 0.148 and an effective first-stage F statistic of 11.2, and the corresponding reduced form for the focal inventor’s own entry is 0.030 ($t = 3.72$, clustered by displacement market). In interpretable units, a one-

standard-deviation increase in the instrument raises her five-year probability of first filing in the field by 0.21 percentage points against a mean of 0.60 percent, and an interquartile increase raises it by 0.10 percentage points, so the market-level composition of reemployment moves entry by roughly a third of its baseline rate.

Under the instrumental-variables assumptions, two-stage least squares recovers the local transmission effect in equation (5), and the estimate is 0.204 for landing compliers with an Anderson–Rubin 95 percent confidence set of $[0.119, 0.384]$. This is far larger than the average entry rate, which is what one should expect given that the treatment is rare and that compliers are selected precisely by field-specific absorption, but the first stage is modest and the complier population is not directly observed. We therefore treat the reduced form as the quantity that describes the magnitude of a market-level shift, report the 2SLS estimate only with weak-instrument-robust inference, and do not extrapolate its point value to other populations or to policies that would form bridges on a different margin.

4.2.6 Understanding the instrumental-variables estimate

This subsection lays out the anatomy of the two-stage estimate. First entry into a new field is a rare event: the baseline entry probability in the instrumental-variables sample is 0.6 percent, and a displaced peer lands at a field- f specialist for only 1.4 percent of at-risk cells. Against this base, the reduced form is economically large: entry occurs at 2.8 percent where a peer lands at an f -specialist and 0.6 percent where one does not, a raw factor of nearly five.

Figure 4 sorts at-risk cells into deciles of the landing-composition instrument and plots, in each decile, both the actual specialist-landing rate (the treatment) and the first-entry rate (the outcome). The specialist-landing rate rises steadily across the informative deciles, from near zero to above five percent, and the entry rate tracks it. The estimate is just-identified, so limited-information maximum likelihood coincides exactly with two-stage least squares, and inference rests on the Anderson–Rubin confidence sets reported in Section 4.2.8, which exclude zero throughout.

The two-stage estimate is a local average treatment effect and should be read as one. It applies to compliers, focal inventors whose peer lands at an f -specialist *because* the displacement market was absorbing specialists into f , and a large per-treated effect on a rare outcome is naturally far larger than the panel average. Section 4.2.5 explains why we emphasize the reduced form and the interval estimate over this point magnitude; Table 4 collects the mechanics behind it.

Finally, the estimate is not the artifact of a few large markets. The identifying variation is concentrated in thick displacement markets, where the leave-out landing composition




Figure 4: The First Stage of the Landing-Composition Instrument

Notes: At-risk cells are sorted into deciles of the landing-composition instrument. For each decile the figure plots the share of cells in which the displaced peer lands at a field- f specialist (treatment D , navy) and the share in which the focal inventor first enters f (outcome y , gray). Whiskers are 95 percent binomial confidence intervals. $n = 526,462$.

Table 4: Anatomy of the Instrumental-Variables Estimate

	Value
Baseline first-entry probability (mean y)	0.006
Specialist-landing rate (mean D)	0.014
Raw entry rate: landing vs. no landing	2.8% vs. 0.6% (4.9 $\times$)
First stage $D \sim Z$	0.148 ($t = 3.4$)
Reduced form $y \sim Z$	0.030 ($t = 3.7$)
2SLS $y \sim D$ (= LIML, just-identified)	0.204
Leave-one-market-out: reduced-form range	[0.026, 0.032]

Notes: Instrumental-variables sample, $n = 526,462$; notice years 2000–2014. Reduced-form and first-stage t statistics are clustered by displacement market (state $\times$ old-firm technology class $\times$ notice year). The leave-one-market-out row drops each of the twenty largest markets individually and re-estimates. LIML coincides with 2SLS because the model is just-identified; inference uses the Anderson–Rubin sets of Section 4.2.8.

varies meaningfully; but the effect does not hinge on any single market. Dropping each of the twenty largest markets *one at a time* leaves the reduced form in the narrow range [0.026, 0.032] (Table 4).

4.2.7 Does the instrument capture technology demand?

The leading threat to the exclusion restriction is that a booming technology both absorbs displaced workers into a market and, independently, draws distant inventors into the same field. If that were the mechanism, the landing-composition instrument would proxy for field-level technology demand rather than for the specific collaborative tie, and it would predict entry for *any* distant inventor exposed to the same field-year, not only for the collaborator of the displaced peer. The field-by-year fixed effects already absorb national field shocks, and the distant-state restriction rules out shared local labor markets, but neither directly tests whether the effect runs through the tie itself.

A randomized-network placebo tests exactly this. We reassign each focal inventor’s displacement-market exposure to a random other focal inventor within the same field-by-year block, breaking the specific inventor–peer link while holding fixed the field-year composition of the instrument, and we re-estimate the reduced form. If the effect reflected field-level demand, the reassigned instrument would predict entry just as well; if it runs through the tie, the reassigned instrument should carry no signal. Across three hundred such reassignments the placebo reduced form is centered on zero (mean 0.00005, standard deviation 0.003), while the actual reduced form of 0.030 lies more than ten standard deviations into the right tail: not one of the three hundred reassignments produces an effect as large ($p < 0.01$). The instrument’s predictive power lives entirely in the specific inventor–peer pairing, not in a field-year technology boom that would draw any distant inventor. Together with the field-by-year fixed effects and the distant-state restriction, this makes the technology-demand interpretation difficult to sustain.

4.2.8 Robustness and falsification

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.¹¹

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 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

¹¹The 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.

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 The causal effect is arm’s-length

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, 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: what the bridge transmits

The displacement design establishes that a collaborator’s destination specialization changes where the inventors she leaves behind patent, but it does not observe what passes over the tie. The tests in this section bear on that question. We begin by setting the collaborator’s employer against the collaborator herself, and then ask how the association varies with the degree of the employer’s specialization, what the displacement design itself can say about whose knowledge is at issue, whether the tie has to be geographically local, and whether the entry patents bear any trace of the bridged firm’s prior art. Only Section 4.3.3 returns to the displacement design. The remaining tests are panel estimates that inherit the selection concerns of Section 4.1.3, so each can establish that the evidence is consistent with

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.

organization-mediated access without ruling out other channels correlated with employer specialization.¹²

4.3.1 Peer expertise versus employer specialization

Table 6 enters the employer-based bridge measure jointly with $PeerOwnSpec_{ift}$, the tie-weighted cumulative share of the collaborators’ own patents in field f . Entered separately, each measure predicts entry (columns 1 and 2). Entered jointly, neither absorbs the other: the employer-based coefficient falls only from 0.0288 to 0.0269, and personal patent specialization from 0.0174 to 0.0168 (column 3). What a collaborator has patented herself therefore

¹²None of them can rule out unobserved 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.

does not account for the association between an inventor’s entry and where that collaborator works. The employer-based coefficient is about 60 percent larger, though the two measures are distributed differently and the comparison should not be pressed beyond the ordering.

4.3.2 Employer specialization beyond observed personal patenting

Two further panel results bear on the same comparison. The first is that the employer-based association is largest where the collaborator has no patent of her own in the destination field, 0.0447 against 0.0217 where she is herself an incumbent there (Table A7). A collaborator without a patent in a field is not a collaborator without knowledge of it, and we do not read the split that way, but the ordering is the reverse of what an account resting on her personal expertise would predict, and it establishes that the measure is not carried by the collaborators with observed inventive experience in the field.

The second is dose, the first implication of Section 2.2. If a bridge conveys a signal whose precision rises with how specialized the employer is, the response should be graded rather than switch on at a threshold. Figure A1 divides the employer’s leave-one-out field share into four mutually exclusive categories and enters the tie-weighted measure separately for each, and the coefficients rise monotonically from 0.0045 for employers whose field share lies between 2 and 5 percent to 0.0112 for employers above 20 percent. The estimate is already positive below the 5 percent threshold imposed everywhere else in the paper, which indicates that the main result is not an artifact of that cutoff.

4.3.3 Whose knowledge? Evidence from the displacement design

The results above are panel estimates, and there is a limit on what any of them can settle. A displacement moves the collaborator to a new employer, which changes both the organization she is embedded in and what she learns by working there. Computing employer specialization while leaving out her own patents removes her prior contribution to the measure, but it does nothing about knowledge she acquires after arriving, so the question remains whether what travels is the employer’s knowledge or the collaborator’s, newly acquired and relayed in the ordinary way. Two features of the displacement design bear on it, and neither conditions on anything realized after the layoff.¹³

The first is whether the collaborator had already patented in the field before the layoff. That is fixed in advance of the displacement and can therefore be interacted with the instrument, and it holds for 2.9 percent of event–field cells. If what travels were her own expertise, the effect of a formed bridge should be concentrated where she has it. It is not. The effect of a formed bridge is 0.209, and its interaction with the collaborator’s prior expertise is -0.016 with a standard error of 0.044, so the confidence interval of $[-0.10, 0.07]$ leaves no substantial role for it in either direction (Table 7, Panel B). A collaborator need not have worked in a field herself for her move into a specialist employer to redirect the inventors she left behind.

The second is timing. The collaborator lands at $T + 1$ and the outcome window runs from $T + 2$ to $T + 6$. If she must first learn the field and then relay it, the response should build with her time at the new employer. Panel A of Table 7 estimates the reduced form separately for first entry in each year. The effect is already present in the first outcome year, reaches its maximum at $T + 4$, and does not continue to rise thereafter. This is weaker evidence than the first test, since research and filing take time under either account, but a profile that flattens is not what a pure learning-then-relaying story predicts.

Neither test is decisive, and we do not present them as such. Together they rule out the version of the personal account in which the collaborator relays expertise she already had, and they fit poorly with the version in which she acquires it after moving. Neither excludes that second version, which is the main qualification on the mechanism evidence.

4.3.4 Geographic distance

Knowledge diffusion is geographically concentrated, although patent citations do not provide an unambiguous measure of localization (Jaffe et al. 1993; Audretsch and Feldman 1996; Thompson and Fox-Kean 2005; Agrawal et al. 2008; Kerr 2008). We therefore divide cross-

¹³Splitting on the collaborator’s tenure at the landing firm, or on whether she patents in the field after arriving, would condition on outcomes that the displacement itself causes. We therefore use only predetermined characteristics and the timing of the inventor’s own entry.

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`.

firm bridges by the collaborator’s location relative to the focal inventor’s modal state. About half of inventors have no identifiable modal U.S. state, almost all of them because their patents place them outside the United States, and we report that group on its own rather than pooling it with collaborators known to be in a different state. The coefficient is 0.0342 ($t = 17.5$) within a state, 0.0292 ($t = 13.2$) across a state line, and 0.0252 ($t = 20.5$) where the collaborator’s state cannot be identified (Table A4). A bridge that crosses a state line is thus about 85 percent as strong as one that does not, so whatever the tie carries does not require the two inventors to work in the same place.¹⁴

4.3.5 Citations to the bridged firm’s prior art

The entry patents themselves carry a trace of where their prior art came from, and this is the most direct evidence available to us. For each first-entry patent supported by a cross-firm bridge we ask whether it cites earlier field- f patents assigned to the collaborator’s employer, setting that employer against a choice set of other f -specialist employers to which the inventor has no connection. Entry-patent fixed effects hold constant the patent’s overall

¹⁴We treat this as a scope test rather than as separate causal evidence. It does not rule out nonlocal common shocks, and modal state is a coarse measure of where an inventor actually works.

propensity to cite specialist firms in the field, so the comparison runs within a single entry event: whether the bridged employer is cited more often than specialists of the same kind that the inventor had no way to reach. What this establishes is the provenance of the prior art, not the content of any conversation.

An entry patent cites the bridged employer’s field- f prior art in 5.5 percent of cases, against 0.1 percent for other specialist employers, and the within-entry difference is 0.054 ($t = 79$; Table A5). Bridged employers hold much larger field portfolios than the firms in that first comparison group, which by itself raises their chance of being cited. Controlling for the log of field-specific portfolio size reduces the coefficient to 0.038 ($t = 57$), and restricting the comparison group to specialists whose portfolios are at least as large as the bridged employer’s gives 0.021 ($t = 20$), on citation rates of 3.6 against 0.6 percent. Excluding firms that had previously employed the focal inventor gives 0.032 ($t = 50$) across 101,660 entry events. The pattern is what we would expect if inventors entering a field draw in part on prior art belonging to the employer their tie reaches.¹⁵

4.3.6 Interpretation

Taken together, the mechanism tests tie field entry to the specialization and the prior art of the collaborator’s employer, conditional on what the collaborator has patented herself and on the specialization of the inventor’s own employer, and the displacement design adds that a formed bridge works whether or not the collaborator had ever worked in the field. The evidence is consistent with organization-mediated access, and it is inconsistent with the version of the personal account in which the collaborator relays expertise she already held. It does not measure the information or assistance that passes over the tie, and it does not exclude a collaborator who learns the field at her new employer and relays it herself. Section 4.4 turns to whom the association is strongest for.

4.4 Heterogeneity

We split the cross-firm bridge coefficient six ways and report every split in full in Appendix C. These are panel estimates and inherit the selection concerns of Section 4.1.3: a subgroup differential may reflect a difference in transmission or a differently sized confound, and we cannot separate the two.¹⁶ The splits are nonetheless informative about one thing. Those

¹⁵A stronger reading would require separating applicant-supplied from examiner-added citations and ruling out continued co-invention with the displaced collaborator.

¹⁶The displacement samples are not large enough to estimate the corresponding interactions with useful precision, so we do not use these patterns to support the causal or allocation conclusions.

that move the association concern the inventor receiving the bridge, while those that describe the technology being transmitted are close to flat.

Four receiving-side splits are worth stating. First, corroboration appears to matter. Fields reached by two or more independent cross-firm bridges carry a coefficient of 0.0636 against 0.0076 for fields reached by exactly one, a difference of a factor of eight that is far larger than anything additive in bridge count would produce (Table A6). This is the one split we can also check against the displacement design, and there the amplification does not appear: the effect of a formed bridge is no larger for an inventor who already reached the field through another specialist employer, with an interaction of -0.041 (standard error 0.038) against a main effect of 0.219. The gap in the panel therefore reflects which inventor-field cells acquire a second bridge rather than what a second bridge does, which is the second of the two readings this section opened by distinguishing. Second, whose knowledge is at issue runs against the intuitive ordering. Bridges through collaborators who have never personally patented in the destination field carry 0.0447, roughly double the 0.0217 carried by bridges through collaborators who are themselves incumbents there (Table A7). A collaborator without a patent in a field is certainly not without knowledge of it, but an account resting on her personal expertise would predict the reverse ranking, and both patterns point in the same direction as the horse race of Section 4.3.1.

Third, technological proximity enters positively in the binary comparison, as the same complementarity implies, but not monotonically across three bins. Relative to fields with which the inventor has no prior overlap (0.0222), the association is much larger for fields adjacent to her existing portfolio (0.0729) and close to zero in the middle range (0.0034). We report three bins rather than two precisely because the binary split averages over that trough and reads misleadingly clean (Table A8). Fourth, local specialist capacity works in the direction opposite to substitution, which we had not expected but which is the complementarity implied by an adjustment cost that falls in the inventor’s own capacity to act. Bridges are associated with entry most strongly where the inventor’s own state already holds a thick share of national capacity in the field (0.0494) rather than where local capacity is thin (0.0282), and the coefficient is smallest, 0.0182, among inventors with no identifiable U.S. location, who constitute 59 percent of sample weight (Table A10). On this evidence organizational bridges complement local specialist capacity rather than compensating for its absence, which is consistent with the formation result of Section 5.3 and cuts against reading connectivity as an equalizing force.

The content-side splits are comparatively flat. Across the six well-populated CPC sections of the destination field, the coefficient runs from 0.0198 to 0.0336, a spread of 1.7 (Figure A3), so the technology an inventor is entering barely moderates the association. The section she

comes *from* matters somewhat more, from 0.0151 to 0.0351, a spread of 2.3, with computing and electrical origins above the pooled estimate (Figure A2), which is again a property of the receiver rather than of the destination. The one content-adjacent split that does move is the industry of the collaborator’s employer rather than the technology of the patent. Bridges through employers in biopharmaceutical (0.0331) or software and computing (0.0386) industries are stronger than the pooled 0.0288 (Table A11), and the contrast with the flat destination-field result suggests that what varies is how legible an employer’s specialization is, not which technology the entry patent falls in.

Tie recency is the one split with no gradient: ties last active within four years carry 0.0293 and ties dormant five years or more carry 0.0277 (Table A9). We report this cautiously.¹⁷

5 Consequences and Allocation of Bridge-Induced Entry

We next examine whether the estimated field entries persist and generate additional destination-field output. We then study where displacement-induced bridges form. The results distinguish the effect of a formed bridge from the frequency with which displaced collaborators join employers specialized in different fields. All outcomes in this section are measured in gross destination-field terms; the design does not identify changes in the inventor’s total output across fields.

5.1 Persistence and destination-field output

A first patent in a new field is a weak outcome if it is never followed by another. We therefore re-estimate the landing-composition reduced form on a set of outcomes measured over the same five-year post-landing window and defined for every event–field cell, assigning zeros where the focal inventor produces no patent in the field, so that none of the estimates condition on realized entry. Table 8 reports the results, and they indicate that the induced entries persist. A durable entry, defined as a first entry followed by at least one additional patent in the same field within the window, carries a reduced form of 0.0211 ($t = 3.7$). That is roughly 70 percent of the 0.0303 estimate for any first entry, so most induced entrants return to the field rather than patenting there once. The instrument also predicts the number of patents the inventor produces in the destination field (0.140, $t = 2.9$), the number she produces after

¹⁷It is the only one of the six patterns that does not survive when the bridge is restricted to firms the inventor never works at, where the recent and stale coefficients separate (0.0058 against 0.0012, the latter not distinguishable from zero). The remaining five orderings hold or widen under that restriction.

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.

the entry patent itself (0.110, $t = 2.5$), and the five-year forward citations those patents accumulate (0.579, $t = 2.2$), with a similar citation estimate when we winsorize at the 99th percentile. Allocating each patent’s citations fractionally across its CPC subclasses reduces the citation estimate to 0.061 ($t = 1.7$), which remains positive but is no longer precisely estimated.

These outcomes establish that the response is a genuine relocation of inventive effort into the destination field rather than a single relabeled patent, but they measure only where activity lands and not whether the inventor’s total output rises, since a positive destination-field effect could in principle be offset by reduced activity in the fields she would otherwise have worked in. We return to that limitation in Section 5.5.

5.2 Where mobility-induced bridges form

Whether this mechanism changes the aggregate direction of invention depends on two distinct margins, and the remainder of this section separates them. The first is transmission: how strongly a formed bridge moves entry, and whether that strength varies across technological environments. The second is formation: how often the labor market creates a bridge into a given kind of field in the first place. We examine both using the density and citation-rate classification of Section 3.6, referring to sparse, high-citation field-years as “gaps” for brevity. The label describes only observed density and citation rates and carries no implication that an additional entrant is socially more valuable there.

Beginning with transmission, Appendix Table A12 interacts the cross-firm bridge with the four field types in the panel specification (11). The field $\times$ year effects absorb level differences across types, so the interactions are identified from differences in the bridge slope rather than in entry rates. The slopes are remarkably similar across environments, ranging only from 0.0186 in sparse, low-citation fields to 0.0278 in dense, high-citation fields, and the interaction between the bridge measure and the gap indicator is -0.0001 ($t = -0.01$). The panel therefore gives no indication that bridges work differently in sparse, highly cited fields than elsewhere, and the pattern is unchanged 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, but they contain no evidence of the heterogeneity with which a formation result would otherwise be confounded.

5.3 Bridge formation across field types

Turning to formation, Table 9 interacts the landing-composition design with the gap indicator, and the contrast is sharp. The reduced-form coefficient is 0.0306 in non-gap fields, while the gap interaction is -0.0259 ($t = -4.4$), leaving a combined estimate of 0.0047 that is close to zero. This difference originates almost entirely in the first stage rather than in any difference in transmission: the instrument’s first-stage coefficient is 0.1499 in non-gap fields against 0.0170 in gap fields, and correspondingly only 0.14 percent of event-field cells involve a collaborator landing at a gap-field specialist, against 1.46 percent elsewhere. Displaced inventors almost never land at employers specialized in sparse, highly cited fields, for the simple reason that few such employers exist to hire them.

Because the first stage nearly vanishes within this subgroup, the corresponding 2SLS interaction of -0.0211 carries a standard error of 0.2053 and is uninformative about whether the effect of a formed bridge differs across field types. We emphasize that a failure to reject equality is not evidence of equality, and we therefore do not claim that transmission is field-invariant. What the design does support is a statement about formation: displacement reallocates access along the existing distribution of specialist employers, and generates almost no variation in specialist landings for the fields where such employers are scarce.

The results therefore do not establish a field-invariant transmission effect. They show that displacement-induced bridge formation is uncommon in sparse, high-citation fields. This pattern is consistent with the limited employment and absorption capacity of specialist firms in those fields.

Equation (6) makes a restriction that can be checked directly, namely that what routes bridges into a field is total specialist hiring $n_f \bar{h}_f$ rather than the number of specialist firms

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.

on its own. We estimate a field-year model of specialist landings on the count of specialist employers and on their absorption of displaced inventors, with year effects and the count of at-risk cells as an offset, over the 3,031 field-years of the landing-composition sample. Absorption carries essentially all of the variation, with an elasticity of 0.93 ($t = 64.1$), while the number of specialist employers enters with a small negative coefficient once absorption is held fixed (-0.065 , $t = -3.6$). A field with many specialist firms that are not hiring displaced workers receives no more bridges than a field with few. This is why we write capacity as a product and why the constraint in sparse fields is one of employment rather than of the presence of specialists.

5.4 Accounting for bridge formation

Equations (6)–(8) separate a change in the number of bridges from a change in their technological composition, and the data speak to both terms. Across the estimation panel, 81 percent of all bridge weight sits in dense, high-citation field-years and under 1 percent in sparse, high-citation ones, while the average field reached by a bridge carries a lagged citation rate roughly 16 percent higher than a randomly drawn at-risk field. Bridges are therefore not concentrated in technologically unpromising fields. By the citation measure the market routes them toward valuable ones; what they do not reach is the sparse end of the valuable range.

The landing-composition sample measures the formation margin directly, and the disparity there is larger still. Gap fields make up 2.52 percent of at-risk event-field cells but receive only 0.24 percent of realized specialist landings, and their estimated share of total specialist absorption capacity is 0.03 percent, with the underlying denominators reported in Appendix Table A13. The gap is not simply that such fields are rare in the sample: their share of realized bridges is an order of magnitude below their share of at-risk opportunities. Rewiring cannot form a bridge into a field in which no employer specializes. It is this constraint, rather than any difference in how well bridges work, that keeps mobility-induced diffusion out of sparse fields.

The accounting supports one counterfactual, and it is worth being precise about which. Because the routing share depends only on absorption capacity, we can ask how the field composition of bridge formation would respond to a change in that capacity without knowing what a bridge produces once formed. The two policy arms of equation (8) behave very differently. A uniform increase in the rate at which relationships are rewired scales every field’s bridge count by the same factor and leaves q_f unchanged by construction, so connectivity has no composition counterfactual to compute; its invariance is the result. A policy that raises absorption capacity in gap field-years by a factor m gives $q_{\text{gap}}(m) = mA_{\text{gap}}/(mA_{\text{gap}} + A_{\text{non}})$, and the scale required is the finding. Tripling specialist absorption in every gap field-year, which is a large intervention, moves their share of bridge formation from 0.029 to 0.088 percent and closes 3.5 percent of the distance to their 2.52 percent share of opportunities. Matching that share would take a factor of roughly 89. The reason is arithmetic rather than behavioral: the base being multiplied is three hundredths of a percent, so proportional expansions of it remain small. We report this as a statement about the scale of the required change and not as a policy evaluation. It is partial equilibrium, it applies the multiplier uniformly across gap field-years, its capacity measure covers the reabsorption of displaced workers rather than all hiring, and the opportunity share is a benchmark for comparison rather than an optimum. What it does not require is any estimate of what a bridge pro-

duces, which is what distinguishes it from the output and welfare calculations we do not attempt.

5.5 Limits of the allocation exercise

The field-level design identifies where inventive activity occurs, not whether total output rises. Event fixed effects absorb outcomes that are common across an inventor’s fields. We constructed a complementary event-level analysis that aggregates bridge exposure and patenting across fields within nonoverlapping five-year windows. The first stage is strong (an effective F statistic of approximately 16), and the instrument does not predict pre-period patents or citations. However, the sample contains only about 7,700 events, and the estimates for total, new-field, and prior-field output are too imprecise to distinguish net creation from reallocation.

We therefore do not convert the estimates into welfare or policy counterfactuals. Such calculations would require at least three additional objects: the effect on total inventive output, the net social value of entry in each field, and the cost and equilibrium effects of changing specialist employment. Forward citations and inventor density do not identify these objects. The estimates instead support a narrower implication: changes in collaborator mobility diffuse access through the fields represented by the employers that hire those collaborators. A policy that changes mobility without changing the distribution of specialist employment need not change the direction of this diffusion. Our design does not estimate whether or how policy can create specialist capacity in underrepresented fields.

The framework of Section 2 adds two qualitative statements about such policies, and it is worth being clear that they are the only policy content the paper can support. Because transmission is approximately homogeneous, the field composition of bridge formation determines where induced entry falls and not how much of it there is. A policy that redistributes specialist capacity across fields moves entry between them; only a policy that changes how many bridges form changes the total. That is a sharper statement than the associational patterns support on their own, and it is the one the causal evidence backs. What we cannot say is anything about magnitudes, about whether entry in one field is worth more than in another, or about what building specialist capacity would cost, for the reasons given above.

6 Conclusion

Where an inventor searches next is shaped by where her strong collaborators work, and this paper shows that a collaborator’s move to a new employer changes that direction for the

inventors she leaves behind. Using mass layoffs that displace a prior co-inventor into a new firm, we find that focal inventors become more likely to file their first patent in the fields in which that firm specializes. Displacement-induced bridge formation arrests the natural decay of entry into the landing firm’s fields, and a landing-composition instrument built from distant markets confirms the effect under weak-instrument-robust inference. Randomly reassigning inventor–collaborator links within field-by-year cells eliminates the relationship, so what the instrument captures is specific to the tie. Most of the response occurs before the focal inventor has any employment relationship with the destination firm, and excluding inventors who had already worked there raises the estimate rather than lowering it, which is what allows us to describe the effect as an externality of someone else’s mobility.

The signal is organizational. It is separable from and larger than the peer’s personal expertise, strongest when the peer has never patented in the destination field herself, and visible in which prior art the entry patent cites, and it operates without any job change by the inventor. The induced entries persist and generate additional patents and forward citations in the new field, though our event-level design is too imprecise to establish whether this is net new invention or a reallocation from the inventor’s other fields. Finally, mobility-induced diffusion follows the existing geography of specialist employment. Sparse but highly cited fields account for a much larger share of at-risk opportunities than of realized specialist landings, and the instrument’s first stage nearly vanishes there because almost no employers specialized in such fields absorb displaced inventors. Our estimates cannot say whether a formed bridge works equally well across field types; what they establish is that the binding constraint lies in formation rather than in transmission.

Three implications follow. First, organizations shape innovation beyond their own boundaries: the boundary of the firm is not the boundary of its knowledge, and the collaboration network defines who holds an outside claim on it. Second, labor-market shocks rewire opportunity sets and not only employment relationships, so evaluations of displacement that ignore spillovers onto non-displaced collaborators miss a margin of the adjustment. Third, connectivity alone does not steer reallocation. Because transmission requires no further mobility once a tie exists, restricting mobility may be less effective at containing knowledge than a mobility-centric view implies, while liberalizing it spreads knowledge along the existing organizational geography rather than toward fields where specialist capacity is thin.

The picture that emerges is of innovation diffusing through a two-layer structure: a social layer of durable collaborative ties, and an organizational layer of specialist firms that those ties reach into. Changes in the social layer, through connectivity, mobility, or reemployment, govern the volume of diffusion, whereas its direction follows the organizational layer. We do not estimate total inventive output, the social value of entry across fields, or the cost

of building specialist capacity, and we leave to future work the dynamics of how specialist capacity itself emerges in nascent fields and whether targeted entry by anchor organizations can seed the bridge geography that the market, by itself, does not build.

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 $\text{inventor} \times \text{field}$, $\text{field} \times \text{year}$, and $\text{inventor} \times \text{year}$ 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 $\times$ same state	0.0284*** (0.0019)	0.0343*** (0.0020)	0.0342*** (0.0020)
Cross-firm bridge $\times$ different state	0.0247*** (0.0022)	0.0294*** (0.0022)	0.0292*** (0.0022)
Cross-firm bridge $\times$ 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 ($d_{if} >$ median positive)	0.0531*** (0.0025)	0.0586*** (0.0025)	0.0729*** (0.0026)
Cross-firm bridge $\times$ middle ($0 < d_{if} \leq$ 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-