

The Geography of Research: The Trade-Off Between Knowledge Production and Access

Sitian Liu*

Yichen Su[†]

August 2026

Abstract

Research activity generates highly localized positive spillovers, yet in the U.S. it has become increasingly spatially misaligned with population and economic activity as people moved away from legacy cities where major research institutions remain anchored. Reallocating researchers toward population centers could broaden local access to knowledge but may reduce knowledge production if legacy institutions or large research clusters raise researcher output. The spatial reallocation of researchers therefore entails a trade-off between gains in knowledge access and losses in knowledge production. To quantify the knowledge production loss from reallocation, we use bibliographic data to separate individual effects from location effects and show that location effects account for substantial differences in research output across locations. Our instrumental-variables estimates provide robust evidence of own-institution agglomeration effects but weaker evidence of external-cluster effects. We then incorporate these estimates into counterfactual reallocations and show that marginally reallocating researchers toward large metropolitan areas with relatively little research activity can plausibly raise aggregate output because the associated knowledge production losses are modest and can be offset by small access gains.

Keywords: Academic Research, Researchers, Knowledge Access, Knowledge Spillovers, Knowledge Production, Agglomeration Economies, Spatial Inequality, Universities, Researcher Reallocation

JEL Codes: O33, J24, R12

*Queen's University. Email: sitian.liu@queensu.ca.

[†]Southern Methodist University. Email: yichensu@outlook.com. We thank seminar participants at Hitotsubashi University and Renmin University of China for their valuable comments.

1 Introduction

Universities and research institutions raise local productivity, innovation, entrepreneurship, and human capital, but many of these benefits decline sharply with distance (Jaffe, 1989; Moretti, 2004; Aghion et al., 2009; Kantor and Whalley, 2014, 2019; Hausman, 2022; Biasi and Ma, 2023; Fieldhouse and Mertens, 2024; Howard et al., 2024). The localized nature of these spillovers provides a rationale for dispersing research capacity beyond established hubs and toward population and economic activity, so that a broader segment of the economy can access frontier knowledge. This rationale has shaped higher-education policy for more than a century, from the U.S. land-grant system to more recent decentralization initiatives in Europe.¹

However, knowledge production may itself benefit from spatial concentration. Researchers can be more productive at some institutions than at others because of location-specific factors, and larger research clusters may further raise productivity through agglomeration effects. Reallocating researchers away from productive hubs may therefore lower movers’ productivity; by shrinking the sizes of major research clusters, it may also reduce the productivity of researchers who remain. These forces create a trade-off: Dispersing researchers toward population and economic activity may improve access to frontier knowledge, but it may also reduce the aggregate amount of knowledge produced.

This paper quantifies the trade-off between knowledge access and knowledge production. Specifically, we ask whether dispersing researchers toward population and economic activity reduces knowledge production enough to offset the resulting gains in knowledge access. Using detailed bibliographic data from 1970 to 2019, we first document the growing geographic misalignment between population and research activity across the U.S. and the resulting widening disparities in local access to knowledge. We then estimate location effects in knowledge production and the causal effects of own-institution and external-cluster sizes on research output using an instrumental-variables strategy. Finally, based on these estimates and a spatial model of researchers and non-researchers, we conduct a series of counterfactual reallocations. To our knowledge, this is the first paper to formally characterize and quantitatively evaluate the trade-off between access benefits of researcher dispersion and knowledge production benefits of researcher concentration.

¹The U.S. land-grant system was established under the Morrill Act of 1862, which supported colleges teaching agriculture and the mechanic arts “to promote the liberal and practical education of the industrial classes” (National Archives, 1862). Land-grant institutions were intended to broaden access beyond traditional elite colleges (Association of Public and Land-Grant Universities, 2012; National Education Association, 2022). For example, Illinois founded its land-grant campus at Urbana–Champaign to expand access for working people statewide. Later examples include the University of California’s expansion into fast-growing regions through campuses such as Irvine and Merced (Pelfrey, 2004; Merritt and Lawrence, 2007; Legislative Analyst’s Office, 2024). More recently, similar efforts have appeared in higher-education policy in Europe. For example, France’s Université 2000 plan aimed to decentralize higher education by creating new university campuses outside traditional centers.

We begin by documenting the geography of researchers and knowledge access. Combining bibliographic data from the Microsoft Academic Graph (MAG) with population and employment, we show that research activity is highly concentrated today. The Northeast Corridor, coastal California, and a number of established research centers in the Midwest account for a disproportionate share of researchers and publications, while many large and fast-growing areas in the South and West have relatively little nearby research activity. This mismatch was less pronounced historically: Research activity was more closely aligned with population in 1970, but the distributions of population and researchers diverged as population shifted toward the South and West while research remained anchored in long-established institutions. As a result, research output per capita became substantially more unequal across metropolitan areas. The same pattern appears when research output per capita is scaled by the college-educated population, employment, business establishments, or high-technology activity, and it holds for both university and non-university research and across nearly all broad fields. Given the highly localized spillovers from research activity, these patterns suggest potential gains from dispersing research resources toward population and economic activity. Such dispersion, however, may reduce knowledge production by moving researchers away from large and productive legacy research clusters.

To quantify the knowledge production loss from dispersing researchers, we estimate an individual-level research production function using a panel constructed from MAG. We model a researcher’s output as the sum of a portable individual effect and an institution-specific location effect. First, we empirically separate the individual effects from the location effects. Separating the two effects is essential for evaluating reallocation: if research output is driven primarily by portable individual effects, moving researchers across locations would have little effect on knowledge production. But if location effects are important, reallocation can substantially change research output, leading to large aggregate effect on knowledge production.

Exploiting the panel structure of the bibliographic data, we use the variation generated by moves across institutions, with researcher fixed effects, to separate location effects from sorting on individual effects. However, limited mobility leaves many institutions weakly connected, making unrestricted location fixed effects imprecise. We therefore parameterize the location effects as a field-specific function of lagged observable characteristics of the institution and its surrounding research environment, and use the fitted values as our location effects. These predicted effects capture the portion of location effects associated with observed characteristics.

A potential concern with using a mover design to distinguish location effects from individual effects is

that researchers with greater expected productivity growth may sort into more productive research clusters. If so, our measured location effects may be biased by capturing movers' pre-existing productivity trajectories. To examine this concern, we conduct an event study around moves across institutions. The estimates show no differential trends before a move and a persistent increase in output after a move to a higher-output institution. We then assess the importance of location effects by decomposing output differences across individuals. Location effects explain a modest but meaningful share of the differences between high- and low-output researchers, but a substantially larger share of the differences between researchers at large and small institutions or in large and small metropolitan research clusters. These sizable location effects imply that reallocating researchers can directly affect the output of those who move.

Second, we estimate agglomeration elasticities for own institutions and external research clusters. Own-institution cluster size is the number of active researchers in the researcher's field at her institution in a given year, while external-cluster size is the number of active researchers in the same field at other institutions within the same metropolitan statistical area (MSA) in a given year. Estimating the agglomeration elasticities is essential for evaluating researcher reallocation because moving researchers changes cluster sizes at both origin and destination institutions and MSAs. These elasticities determine how strongly those size changes affect the output of researchers, including those who do not move.

The main identification challenge is that cluster size may be correlated with unobserved local productivity advantages. We address this concern using Bartik shift-share instruments that combine the 1995 detailed *subfield* composition of each institution and external cluster within a broad field with subsequent *worldwide* growth in those subfields. The instruments predict greater cluster growth for research environments initially specialized in globally expanding subfields. The regressions absorb researcher fixed effects, field-specific life-cycle profiles, persistent differences across institutions within fields, and common field-year shocks. We also control for researchers' and their coauthors' direct exposure to worldwide subfield growth, helping ensure that the instruments do not simply capture expanding opportunities in their own research areas.

The IV estimates provide robust evidence of own-institution agglomeration effects across all three outcomes. In our preferred specification, a 10% increase in own-institution cluster size raises publications by about 2.0%, citations by 8.3%, and patent citations by 1.5%, suggesting that institutional scale may matter especially for research impact. Evidence for external-cluster effects is weaker: A 10% increase in external-cluster size raises publications by about 0.7%, while the estimates for citations and patent citations are small and imprecise. The effects also vary across fields, with particularly strong effects in engineering and the

life sciences. These sizable agglomeration effects imply that, beyond its direct effect on movers, researcher reallocation can also indirectly affect non-movers' output by changing the size of their own-institution and surrounding research clusters.

We then combine the location effects and agglomeration elasticities to compute the counterfactual aggregate research production in counterfactual reallocation scenarios that move researchers from metropolitan areas with more researchers than their population share would imply to areas with fewer: movers experience changes in location effects, while researchers who remain at the origins and destinations are affected by changes in cluster sizes of their institutions and MSAs. Because we do not estimate the elasticity governing the local economic benefit of research access, we do not impose a particular value *a priori*. Instead, for each reallocation, we invert the exercise and solve for the break-even access elasticity at which the access gain exactly offsets the knowledge production loss.

The results are similar across several exercises. Bilateral reallocations of researchers from institutions in Boston to those in Dallas–Fort Worth and from San Francisco to Las Vegas imply break-even access elasticities of 0.018–0.019. Thus, if the local access spillover elasticity is at least this large, the aggregate gains from broader knowledge access would exceed the accompanying loss in knowledge production. Reallocating researchers among the 50 largest MSAs lowers the break-even access elasticity to about 0.01. Extending the exercise to all MSAs produces larger knowledge-production losses because many destinations contain very small institutions with low location effects, but even then, the break-even elasticity remains around 0.04. All of these thresholds fall far below the 0.08 estimate of local university research spillovers in Kantor and Whalley (2014), as well as related estimates in Jaffe (1989), although the margin is narrower when researchers are dispersed across all MSAs containing research institutions. These results suggest if the true access elasticity is 0.08 or above, marginally reallocating researchers to become closer the population and economic activity would raise aggregate output.

The findings are robust across alternative measures of knowledge access and cluster size. Measuring access to research by publications or citations per capita lowers the break-even elasticity because movers from research-intensive origins bring more research output and impact than their headcount alone captures. Measuring cluster size by publication counts rather than researcher counts also reduces the estimated knowledge production loss. Reallocating researchers toward the college-educated population rather than the total population raises the break-even access elasticity only modestly, and restricting the analysis to STEM researchers produces similar results. Taken together, the results suggest that reallocating researchers to align

more closely with population and economic activity—particularly toward large, underserved metropolitan areas—can plausibly raise aggregate output.

The paper makes three contributions. First, it documents a new fact about the geography of research: geographic access to research has become increasingly unequal because research activity remained anchored in legacy cities while population and economic activity moved. This finding is particularly relevant given existing work that shows that universities generate important local spillovers in productivity, innovation, entrepreneurship, and economic activity (Jaffe, 1989; Aghion et al., 2009; Kantor and Whalley, 2014, 2019; Hausman, 2022; Lerner et al., 2024). A related literature shows that proximity to universities also affects educational attainment and access to higher education (Card, 1995; Rouse, 1995; Black et al., 2020; Fu et al., 2022; Mountjoy, 2022; Biasi and Ma, 2023; Acton et al., 2024). Fabre (2023) and Ishimaru (2025) show how the uneven distribution of colleges contributes to persistent spatial inequality in France and the United States. We contribute to this literature by documenting the geography of research and showing how it has diverged from the geography of population and economic activity over time and highlight the potential benefits of reallocating researchers to be more aligned with population and economic activity.

Second, we contribute to the literature on how local environments and spatial concentration affect knowledge production. Within academia, Bosquet and Combes (2017) find that same-field peers raise publication quantity and quality among economists in France, Chandra and Xu (2025) show that institutional factors account for a large share of variation in life-science researchers’ output, and Helmers and Overman (2017) find that proximity to a major UK scientific facility increases related research activity. Evidence on localized peer effects is not uniform: Azoulay et al. (2010) and Waldinger (2012) find limited evidence that geographic proximity to prominent scientists raises peers’ productivity. Beyond academia, a large literature documents agglomeration economies among firms and workers (Glaeser, 1999; Glaeser and Mare, 2001; Rosenthal and Strange, 2003; Charlot and Duranton, 2004; Duranton and Overman, 2005; Arzaghi and Henderson, 2008; Rosenthal and Strange, 2008; Ellison et al., 2010; Greenstone et al., 2010; D’Costa and Overman, 2014; De La Roca and Puga, 2017; Gaubert, 2018; Eckert et al., 2022; Baum-Snow et al., 2024), while research on innovation shows that knowledge spillovers are highly localized and that inventors, startups, and other knowledge workers benefit from proximity to innovative clusters (Jaffe et al., 1993; Belenzon and Schankerman, 2013a,b; Carlino and Kerr, 2015; Moretti, 2021; Atkin et al., 2022; Emanuel et al., 2023; Guzman, 2024).² On the theory side, Davis and Dingel (2019) show how interactions among

²A related literature examines how team structure and communication costs shape research, and how information technologies

knowledge workers can increase more than proportionally with cluster size. We contribute to this literature by separating the location effects from individual effects in knowledge production and estimating agglomeration economies at both the own-institution and external-cluster levels for academic researchers across a broad range of fields.

Third, we connect the literature on the local benefits of research access with the literature on the productivity gains from concentrating knowledge workers. The first provides a rationale for dispersing researchers, while the second implies a potential production benefit from concentration. We formally characterize this trade-off in a spatial framework and take a first step toward quantifying it by embedding our estimates of location and agglomeration effects. In doing so, we provide an initial quantitative assessment of whether alternative spatial allocations of researchers can plausibly generate net gains in aggregate output. This access-concentration trade-off parallels Dingel et al. (2023), who study productivity and geographic access in medical services, and Rossi-Hansberg et al. (2026), who examine spatial allocations in the presence of localized high-skilled spillovers.

The rest of the paper proceeds as follows. Section 2 presents a spatial framework to characterize the trade-off between knowledge access and knowledge production. Section 3 describes the data. Section 4 documents the geography of research and access to knowledge. Section 5 presents our empirical strategies and the corresponding estimates of location effects and agglomeration effects. Section 6 conducts the counterfactual analysis. Section 7 concludes.

2 Trade-off between Knowledge Access and Knowledge Production

To characterize the spatial trade-off between knowledge access and knowledge production, we develop a spatial framework with researchers and non-researchers. The spatial allocation of researchers affects aggregate output through two channels. First, researchers improve nearby non-researchers' access to frontier knowledge. Second, researchers produce the frontier knowledge itself, and their productivity can be higher in larger or more productive research clusters.

Consider J locations, indexed by j . Let N_j denote the quality-adjusted number of researchers in location j , and let M_j denote the number of non-researchers. We define N_j formally below after introducing

alter the process of knowledge production (Hesse et al., 1993; Cohen, 1996; Van Alstyne and Brynjolfsson, 1996; Walsh and Bayma, 1996; Kaminer and Braunstein, 1998; Walsh et al., 2000; Hamermesh and Oster, 2002; Rosenblat and Mobius, 2004; Adams et al., 2005; Wuchty et al., 2007; Agrawal and Goldfarb, 2008; Butler et al., 2008; Jones et al., 2008; Kim et al., 2009; Ding et al., 2010; Winkler et al., 2010; Adams and Clemmons, 2011; Goldstein, 2024).

individual researcher productivity. The baseline model takes the spatial allocation of non-researcher population $\{M_j\}_{j=1}^J$ as given and studies the effect of exogenously reallocating researchers on aggregate output. Appendix A1 endogenizes the location choices of researchers and non-researchers.

Knowledge Access and the Local Economy Economic output in location j is modeled as the product of non-researcher input M_j and labor productivity $S_{0j}^{1-s^N} G_j^{s^N}$:

$$Y_j = S_{0j}^{1-s^N} G_j^{s^N} M_j.$$

Labor productivity combines a location-specific component unrelated to frontier knowledge, S_{0j} , and access to frontier knowledge, G_j , in Cobb-Douglas form, with weight s^N on knowledge access. Access to frontier knowledge in location j is

$$G_j = S(N_j, M_j) Q = \left(S_0 + \left(\frac{N_j}{M_j} \right)^\theta \right) Q, \quad (1)$$

where Q is the location-invariant stock of frontier knowledge relevant to production, endogenized later as the steady-state outcome of knowledge accumulation. The function $S(N_j, M_j)$, abbreviated as S_j , governs how accessible this frontier knowledge is to the local non-researchers.

We adopt the tractable form in Equation (1). The baseline term S_0 captures non-local channels, such as publications and online media, which deliver access regardless of local researchers' presence. The second term captures access mediated by nearby researchers. Holding N_j fixed, local access declines with M_j , reflecting congestion in researchers' attention. We assume $0 < \theta < 1$, so local access is increasing and concave in N_j : an additional researcher improves access, but the marginal gain declines as the local research community expands. For a given stock Q , this concavity creates a force favoring the dispersion of researchers toward economically large locations with relatively few researchers.³

To summarize the relative importance of local and non-local access factors, define

$$\lambda_j \equiv \frac{(N_j/M_j)^\theta}{S_0 + (N_j/M_j)^\theta} \in (0, 1),$$

which is the share of local knowledge access factor attributable to the relative presence of nearby researchers. This term governs how strongly local output responds to the number of researchers: The elasticity of local access with respect to the number of local researchers is $\partial \ln S_j / \partial \ln N_j = \theta \lambda_j$, holding the number of local

³The current specification assumes that local access to knowledge is rival: access mediated by nearby researchers is congestible because researchers have limited time and attention.

non-researchers fixed. Thus, holding the frontier knowledge stock Q and local non-researchers M_j fixed, the elasticity of local output Y_j with respect to N_j is $s^N \theta \lambda_j$.⁴

Knowledge Production The stock of frontier knowledge, Q , is also endogenous to the spatial distribution of researchers across locations. Because knowledge accumulates but can become obsolete over time, we model Q using a perpetual-inventory approach (Romer, 1990; Jones, 1995; Bloom et al., 2020). Researchers add to the stock of knowledge each period, while a constant fraction of existing knowledge depreciates as it moves away from the frontier. We first specify the flow of new research produced in each period, and then describe how these flows accumulate into the stock Q .

The research output of researcher i in location j at time t is

$$F_{ijt} = D_i A_j(N_j) Q_{t-1}^\eta. \quad (2)$$

The term D_i captures the researcher’s portable individual productivity. The term $A_j(N_j)$ captures location productivity, through factors such as institutions, infrastructure, networks, and local research spillovers. Finally, Q_{t-1} is the existing stock of frontier knowledge, capturing the idea that new discoveries may stand on the shoulders of giants.⁵ We assume $0 < \eta < 1$, so the existing knowledge stock raises current research productivity with diminishing returns and a finite steady state exists.⁶

To allow for potential agglomeration economies, we specify location productivity as

$$A_j(N_j) = A_{0j} N_j^\alpha, \quad (3)$$

where A_{0j} captures exogenous local determinants of research productivity that are unrelated to cluster size, while $\alpha > 0$ governs the strength of research agglomeration economies. It can be interpreted as the elas-

⁴When S_0 is small, access comes primarily from local researchers— λ_j approaches 1 and the output elasticity approaches $s^N \theta$. When S_0 is large, non-local channels dominate— λ_j and the output elasticity approach 0. Between these cases, λ_j is small where local researchers are scarce and large where they are abundant, so output becomes more responsive to researcher presence as the existing local research base grows. Without the baseline term S_0 , a constant-elasticity specification would imply very large gains from adding researchers to locations with a very small initial research base. The term S_0 tempers this small-base effect: when researchers are scarce, $(N_j/M_j)^\theta$ is small relative to S_0 , so λ_j is also small. An additional researcher therefore has a limited effect on total access in the thinnest research locations.

⁵Our analysis focuses on steady-state spatial allocations. We therefore suppress time subscripts on D_i and N_j , treating individual productivity and the spatial allocation of researchers as fixed over time. We use time subscripts for F and Q only to describe the law of motion. After deriving the steady-state values, we drop these time subscripts.

⁶We restrict attention to $0 < \eta < 1$, so the knowledge stock converges to a finite steady state, as in Jones (1995). The case $\eta = 1$ corresponds to the linear “standing on shoulders” specification in Romer (1990), while $\eta < 0$ captures a “fishing out” effect in which ideas become harder to find as the stock of knowledge grows. If $\eta > 1$, knowledge production rises more than proportionally with the existing stock, so there is no finite steady state.

ticity of location productivity with respect to the local research cluster size. A larger α implies stronger productivity gains from researcher concentration.

Let $\mathcal{R}_j$ denote the set of researchers working in location j , and define the quality-adjusted number of researchers as $N_j = \sum_{i \in \mathcal{R}_j} D_i$. Total research output in location j is defined as the sum of individual research output across all researchers in that location:

$$\begin{aligned} \sum_{i \in \mathcal{R}_j} F_{ijt} &= \sum_{i \in \mathcal{R}_j} D_i A_j(N_j) Q_{t-1}^\eta \\ &= A_j(N_j) N_j Q_{t-1}^\eta = A_{0j} N_j^{1+\alpha} Q_{t-1}^\eta. \end{aligned}$$

Aggregate research output is the sum of research output across all locations:

$$F_t = \sum_{j=1}^J A_{0j} N_j^{1+\alpha} Q_{t-1}^\eta.$$

The stock of frontier knowledge then evolves according to the standard perpetual-inventory law of motion:

$$Q_t = (1 - \delta) Q_{t-1} + F_t,$$

where $0 < \delta < 1$ is the depreciation rate of the existing knowledge stock. F_t denotes the flow of new knowledge produced by researchers in period t . In steady state, $Q_t = Q_{t-1} = Q$. The steady-state stock of frontier knowledge is therefore

$$Q = \left(\frac{\sum_{j=1}^J A_{0j} N_j^{1+\alpha}}{\delta} \right)^{\frac{1}{1-\eta}}.$$

Because our analysis compares steady states, we drop time subscripts from this point forward. We use Q and F to denote their steady-state values.

Aggregate Output Combining the steady-state knowledge stock with Equation 1 gives aggregate output:

$$Y = \sum_{j=1}^J M_j S_{0j}^{1-s^N} \underbrace{\left(S_0 + \left(\frac{N_j}{M_j} \right)^\theta \right)^{s^N}}_{\text{Access}} \underbrace{\left(\frac{\sum_{j=1}^J A_{0j} N_j^{1+\alpha}}{\delta} \right)^{\frac{s^N}{1-\eta}}}_{\text{Frontier Knowledge}}. \quad (4)$$

Let $y = \ln Y$. Taking logs gives

$$y = \ln B + s^N \ln Q,$$

where

$$B = \sum_{j=1}^J M_j S_{0j}^{1-s^N} \left(S_0 + \left(\frac{N_j}{M_j} \right)^\theta \right)^{s^N}.$$

The term B captures the spatial distributions of researchers and non-researchers and how they determine access factor given the frontier knowledge stock and the spatial distribution of non-research productivity. The term Q is the steady-state stock of frontier knowledge defined above.

2.1 Spatial Reallocation of Researchers

Having derived aggregate output, we now characterize the paper's central trade-off: knowledge access and knowledge production. Consider the *marginal* effect on log aggregate output of moving one quality-adjusted researcher from location j to location j' :

$$\Delta y_{j \rightarrow j'} = \underbrace{s^N \left[\frac{Y_{j'}}{Y} \frac{\theta \lambda_{j'}}{N_{j'}} - \frac{Y_j}{Y} \frac{\theta \lambda_j}{N_j} \right]}_{\text{Effect through knowledge access}} + \underbrace{\frac{s^N (1 + \alpha)}{(1 - \eta) \delta Q^{1-\eta}} [A_{0j'} N_{j'}^\alpha - A_{0j} N_j^\alpha]}_{\text{Effect through knowledge production}}. \quad (5)$$

Equation 5 separates the two channels through which researcher reallocation affects aggregate output. The first term captures the change in access to a given stock of frontier knowledge. The second term captures the change in the steady-state knowledge stock. The net effect is generally ambiguous: access tends to favor dispersion toward economically large, researcher-scarce locations, whereas knowledge production tends to favor locations with high research productivity.

Effect Through Knowledge Access Relocating a researcher from j to j' reduces access in the origin and raises it in the destination, so aggregate output increases only when the access gain at the destination outweighs the loss at the origin. The marginal access effect in each location depends on three factors. The output share Y_j/Y captures the economic importance of the location—improved access matters more in economically larger locations. The access elasticity $\theta \lambda_j$ measures how strongly local access responds to researcher presence. Finally, $1/N_j$ reflects diminishing marginal access as the research base expands.

To make the intuition concrete, suppose that the destination j' is Dallas, which has few researchers relative to its economic size. Let the origin j be Boston, which has a smaller economy but many more researchers. If λ_j and $\lambda_{j'}$ do not differ too sharply, moving a researcher from Boston to Dallas can raise aggregate access because the gain in Dallas exceeds the loss in Boston. This channel creates a force toward a spatial allocation of researchers that is more closely aligned with the distribution of non-research economic

activity.

Effect Through Knowledge Production The second term captures the effect of reallocation on the steady-state stock of frontier knowledge, which depends on the location-specific research productivity, $A_j(N_j) = A_{0j}N_j^\alpha$, of the origin and destination. Moving a researcher from j to j' reduces knowledge production and thus the steady-state knowledge stock when $A_j(N_j) > A_{j'}(N_{j'})$. This loss arises even when $\alpha = 0$, because locations differ in baseline productivity A_{0j} . When $\alpha > 0$, the contraction of the origin cluster and expansion of the destination cluster create an additional agglomeration channel.⁷

Decomposing the Knowledge-Production Effect The effect on knowledge production consists of two channels. First, when a researcher i moves from location j to j' , she experiences a direct change in output because she is exposed to a different research environment. Holding each location's productivity component fixed, researcher i 's contribution to the change in log steady-state knowledge stock is

$$\Delta_i^{direct} \ln Q = \frac{D_i [A_{j'}(N_{j'}) - A_j(N_j)]}{(1 - \eta)\delta Q^{1-\eta}} = \frac{D_i [A_{0j'}N_{j'}^\alpha - A_{0j}N_j^\alpha]}{(1 - \eta)\delta Q^{1-\eta}}.$$

This is the direct effect of reallocation on the mover. It depends on the difference in location productivity between the destination and origin, so identifying it requires separating location productivity, $A_j(N_j) = A_{0j}N_j^\alpha$, from individual researcher productivity, D_i .

Second, the move reduces cluster size at the origin and increases it at the destination, so a move also changes the productivity of researchers who remain in both locations. This indirect contribution to the change in $\ln Q$ is

$$\Delta_i^{indirect} \ln Q = \frac{D_i [N_{j'}A'_{j'}(N_{j'}) - N_jA'_j(N_j)]}{(1 - \eta)\delta Q^{1-\eta}} = \frac{\alpha D_i [A_{0j'}N_{j'}^\alpha - A_{0j}N_j^\alpha]}{(1 - \eta)\delta Q^{1-\eta}}.$$

Intuitively, removing a researcher of quality D_i from origin j lowers its location productivity by $D_iA'_j(N_j)$, reducing the output of incumbent researchers by $D_iN_jA'_j(N_j)$. The corresponding increase in cluster size at the destination raises the output of incumbent researchers at destination j' by $D_iN_{j'}A'_{j'}(N_{j'})$. Under the constant-elasticity specification, $N_jA'_j(N_j) = \alpha A_j(N_j)$, so the indirect effect equals the direct effect scaled

⁷Under the constant-elasticity specification, moving researchers from a location with higher marginal research productivity to one with lower marginal research productivity reduces knowledge production. If agglomeration gains decline sufficiently with cluster size, however, an additional researcher could contribute more in a small cluster than in a large one, and dispersion could instead increase knowledge production. Section 5 examines this possibility by allowing the elasticity to vary with cluster size.

by the agglomeration elasticity α . Section 5 estimates both location productivity at observed cluster size, $A_j(N_j)$, and the agglomeration elasticity α .

2.2 Model Extension: Endogenous Location Choices

The baseline model treats the spatial distributions of researchers and non-researchers as given and evaluates exogenous reallocations of researchers. Appendix A1 endogenizes both groups' location choices. This extension introduces an additional general-equilibrium channel: When researcher reallocation changes local access to frontier knowledge, it changes local productivity and can induce non-researchers to migrate.

Reallocating researchers toward researcher-scarce locations may attract non-researchers toward places with lower baseline productivity, partially offsetting the gains from improved knowledge access. Section 6 shows that this migration channel is quantitatively small relative to the two central forces in the model: the access gains from dispersion and the knowledge production losses from weakening productive research clusters. Thus, we present the extension in the appendix.

3 Data

We combine bibliographic and patent records with supplementary data on higher education and local economic conditions to estimate the research production function and to conduct counterfactual analyses.

3.1 Bibliographic Data: Microsoft Academic Graph (MAG)

Our primary data source is Microsoft Academic Graph (MAG), a large bibliographic database developed by Microsoft Research.⁸ MAG contains records for more than 260 million scientific publications and their citation relationships since 1900. For each publication, the data provide consistent identifiers of authors, institutional affiliations at the time of publication, journals, and fields of study. Institutional affiliations are further matched to geographic coordinates and U.S. state and county identifiers by Pollacci (2022).

MAG assigns publications to detailed subfields. We use its classification of 20 broad fields:

⁸Microsoft Academic was an academic search engine developed by Microsoft Research. It ceased operations at the end of 2021. Its successor, OpenAlex, incorporates and extends MAG. The MAG dataset, compiled by Sinha et al. (2015), is publicly available at <https://zenodo.org/records/2628216>. Appendix A2 discusses MAG's coverage and reliability, showing that, relative to other major bibliographic databases, MAG provides broader coverage of scientific publications and comparable citation counts (Thelwall, 2018; Wang et al., 2020; Martín-Martín et al., 2021; Visser et al., 2021). The section also summarizes MAG's use in recent economics research.

- **Social Sciences:** Economics, Business, Sociology, Political Science, Psychology, Population, and Geography.
- **Humanities:** History, Philosophy, and Art.
- **Natural Sciences:** Mathematics, Physics, Chemistry, Geology, and Environmental Science.
- **Engineering and Computer Science:** Engineering, Computer Science, and Materials Science.
- **Life Sciences:** Biology and Medicine.

We assign each researcher a primary field based on the modal field of her publications. When several fields have the same publication count, we randomly select one of the tied fields, following Moretti (2021) and Lerner et al. (2024).

Using MAG, we construct an annual panel of researchers with at least one U.S. institutional affiliation from 1970 to 2015. We define the beginning of a researcher’s career as the year of her first publication and follow her for 35 years. A researcher’s geographic location is determined by her institutional affiliation. If multiple affiliations are reported in a year, we select the institution that appears most frequently over the researcher’s career.⁹ Researchers do not publish in every year. For years without publications, we record zero publications and carry forward her most recently observed affiliation. We define a move as occurring in the first year in which she publishes with a new institution.¹⁰

Research Outcomes MAG provides two core measures of researcher output. Our baseline quantity measure is the number of publications produced by a researcher in a given year. To measure research impact, we construct a ten-year citation measure. For each paper, we count the citations it receives within ten years of publication. We then sum these citations across all papers published by the researcher in that year.

Research Clusters We measure cluster sizes by field, both within a researcher’s own institution (the internal cluster) and outside it (the external cluster). Internal cluster size is the number of active researchers in the researcher’s field at her own institution. External cluster size is the number of active researchers in the same field at other institutions within the same MSA, excluding the home institution.

We also use alternative measures of cluster size and an alternative geographic definition of the external

⁹A researcher may move between U.S. and non-U.S. institutions during her career; she enters the sample if she has at least one publication associated with a U.S. institution. When several affiliations appear with equal frequency over a researcher’s career, we first select a U.S. institution when applicable and otherwise select the institution with the smallest research cluster to obtain a more conservative estimate.

¹⁰Table A1 reports sample sizes.

cluster. First, we measure internal and external cluster size using the number of publications rather than the number of active researchers. Second, we replace the MSA definition with a distance-based external cluster consisting of same-field researchers within 25 kilometers of the home institution, again excluding the home institution. Table A2 presents the top 10 metropolitan areas by the number of researchers and publications.

3.2 Patent Data

To measure the influence of scientific research on patented innovation, we link MAG to U.S. patent records from PatentsView, an open-access database maintained by the U.S. Patent and Trademark Office. It provides linked records on patents, inventors, assignees, inventor locations, technology classes, and patent citations.

We supplement PatentsView with the patent-to-science linkage developed by Marx and Fuegi (2020), which identifies citations from patents to *scientific* publications and matches the cited papers to MAG records. We aggregate these linked citations to the researcher-year level to construct our third research outcome: the number of patent citations received by papers published by the researcher in that year.

3.3 Additional Data Sources

Integrated Postsecondary Education Data System We use the Integrated Postsecondary Education Data System (IPEDS), administered by the National Center for Education Statistics, to measure tenured and tenure-track faculty across institutions and MSAs. These counts allow us to track changes in the geography of university faculty, a group distinct from the researchers identified in our bibliographic data.

American Community Survey We use the American Community Survey (ACS) accessed through IPUMS NHGIS to measure population and demographic composition at detailed geographic levels (Manson et al., 2022). The data provide population counts by educational attainment, income, and occupation. We use census-tract centroids to calculate geographic access to research activity and aggregate these measures to larger geographic units using appropriate population weights. We also use 2015–2019 ACS wage data to construct the measure of non-research productivity used in the counterfactual analysis.

Employment and Business Establishments We use the Quarterly Census of Employment and Wages and County Business Patterns to measure local employment and business establishments by industry and

geography (U.S. Census Bureau, 2023). These data allow us to examine geographic access to research from the perspective of workers and firms, including those in high-technology industries.

4 The Geography of Researchers and Access to Knowledge

Current Spatial Distribution Research activity is highly uneven across the United States, generating large spatial differences in access to knowledge. Figure 1 maps nearby research papers per 1,000 residents between 2014 and 2019 at the county level, with values ranging from essentially zero to 95 papers per 1,000 residents. Research activity is concentrated along the Northeast Corridor, coastal California, and around several established research centers in the Midwest, reflecting the legacy of nineteenth- and early-twentieth-century research institutions. In contrast, many areas in the South, interior West, and Great Plains have relatively little nearby research activity despite substantial population growth in recent decades. Thus, the current geography of research does not match that of population and economic activity.

Historical Divergence The current mismatch was not always as pronounced. Research activity was more closely aligned with population historically, but the two distributions diverged as population shifted toward the South and West while research remained anchored in long-established institutions in legacy cities. Figure 2 illustrates this pattern for seven large MSAs. In 1970, research papers per resident were broadly similar across these metropolitan areas. Over the next five decades, Boston and San Francisco pulled far ahead, while research output per resident in Dallas/Fort Worth, Miami, and Phoenix remained much closer to its earlier levels despite substantial population growth. The pattern is similar when research output per capita is scaled by the college-educated population, indicating that the divergence is not simply driven by increasing concentration of college-educated workers in legacy cities (Diamond, 2016).

The spatial persistence of research activity helps explain this divergence. Figure A1 shows that, in 1970, the ten largest research clusters accounted for similar shares of the population and researchers. Their shares of researchers and research papers remained broadly stable through 2019, while their population share declined substantially. Thus, population moved away from established research centers, while research activity remained anchored to long-lived institutions.

Figure 3 shows that the selected-MSA examples reflect a broader national pattern. Both the 90/50 ratio and the Gini coefficient of research papers per resident increased substantially between 1970 and 2019. The

same measures also rose when research output was scaled by the number of college graduates, indicating that geographic access to research has become increasingly unequal even relative to the distribution of high-skilled workers. This increase appears in research produced by both universities and non-university institutions (Figure A3) and in every broad field except the humanities (Figure A4).¹¹

The appendix complements these inequality measures with evidence from the full distributions. The distribution of research output per resident and per college graduate became substantially more dispersed between 1970 and 2019 (Figure A5).¹² The distributions of research output relative to local employment and business establishments, including high-technology sectors only, also widened (Figure A7). Similar widening occurred separately for research produced by universities and non-university institutions (Figure A8). Finally, the distribution of faculty per capita remained comparatively stable over time, suggesting that the growing geographic dispersion of research access reflects more than the broader distribution of faculty employment (Figure A9).

5 Estimation of the Research Production Function

The previous section documents substantial geographic disparities in research access, suggesting potential access gains from reallocating researchers toward underserved locations.¹³ The key remaining question is how much such reallocation would change the production of frontier knowledge. In this section, we estimate the research production function used to answer this question.

We proceed in two steps. First, we measure the field-specific location effects across institutions, separately from the portable individual effects that researchers carry with them. These location effects capture baseline differences in research output associated with moving across institutions. Second, we estimate the *causal* effects of own-institution and external-cluster size on researchers' output. These elasticities determine how reallocation affects researchers' output through changes in cluster sizes, including that of researchers who remain in the origin and destination clusters.

¹¹Figure A2 reports the composition of types of research institutions by field. Universities remain the dominant host of researchers, but some fields show notable shares elsewhere: A large share of researchers in the Life Sciences work at medical centers, and a meaningful share of those in Engineering and the Natural Sciences are based at private enterprises.

¹²Figure A6 shows that the rising access inequality is primarily the result of population migration away from the existing research-intensive locations.

¹³Appendix A3 provides a back-of-the-envelope calculation of the potential aggregate output gains from improving geographic access to research, holding the stock of frontier knowledge fixed.

5.1 Individual vs. Location Effects

The research production function in Equation 2 specifies that a researcher's output depends on both an individual-specific component and a location-specific component. Taking logs gives

$$\ln F_{ijt} = \underbrace{\ln D_i}_{\text{Individual-specific}} + \underbrace{\ln A_j(N_j)}_{\text{Location-specific}} + \eta \ln Q_{t-1}. \quad (6)$$

Equation 6 clarifies why location effects are central to the counterfactual analysis. When a researcher is reallocated, she carries her individual productivity with her but her output may change because she enters a different research environment. Thus, to quantify the direct effect of reallocation on the productivity of researchers who move (hereafter “movers”), we need to empirically separate location effects from individual researcher effects.

Because our data contain information on researchers' fields, institutions, and years of publication, we enrich the baseline model to allow for field- and time-specific variation. Let F_{ijft} denote the research output of researcher i , affiliated with institution j , in field f , and year t . Let a denote academic age, measured as years since the researcher's first publication. As an empirical counterpart of Equation 6, we write research output as the additive sum of an individual effect, a field-specific location effect, a field-year effect, and an idiosyncratic component:

$$\ln F_{ijft} = \underbrace{d_i + d_{fa}}_{\text{Individual Effect}} + \underbrace{d_{jft}}_{\text{Location Effect}} + d_{ft} + u_{ijft}.$$

The individual researcher effect d_i captures time-invariant researcher productivity, and the field-by-academic-age effect d_{fa} captures field-specific life-cycle profiles. Together, they correspond to the portable individual productivity component $\ln D_i$ in the model. The term d_{jft} is the empirical counterpart of the location-specific productivity component $\ln A_j(N_j)$, allowing location productivity to vary over time and across fields within an institution. The field-year effect d_{ft} absorbs the frontier-knowledge component $\eta \ln Q_{t-1}$, along with other nationwide field-specific time shocks.

Inverse Hyperbolic Sine (IHS) Transformation of F_{ijft} Because research output can be zero, in the empirical implementation we replace $\ln F_{ijft}$ with the inverse hyperbolic sine (IHS) transformation of F_{ijft} , which is defined and behave linearly at zero and behaves like a logarithm for large positive values. To preserve the connection with the model, we continue to express the specification in logarithmic form. In

the counterfactual analysis, we use the IHS transformation in the research production function to remain consistent with our empirical estimates. In Section 5.2.3, we assess and demonstrate the robustness of the IHS specification.

Identification of Location Effects The annual researcher panel allows us to empirically separate individual effects from location effects. Identification comes from researchers who move across institutions. Because a mover carries her individual effect with her, the change in her output following a move identifies the difference between the location effects of the origin and destination institutions. Movers therefore identify location effects throughout the network of institutions they connect. These location effects, in turn, allow us to recover the individual effects of non-movers at those institutions as each researcher’s output net of the relevant location and field-year effects.

Empirical Challenge: Limited Mobility A natural approach would estimate d_{jft} using institution-field-year fixed effects.¹⁴ However, in our context, only 44,009 of 187,865 researchers—approximately 23%—move at least once. Thus, many institution-field cells are weakly connected or disconnected from the mobility network. A fixed-effects approach would yield imprecise estimates. Restricting the analysis to a sufficiently connected mobility network would exclude many institutions and result in a substantially smaller selected sample.

Predicted Location Effects To address the limited-mobility problem, we model the location effect using observable characteristics of the institution and its surrounding research environment:

$$d_{jft} = X'_{jft}\beta_f + \xi_{jft},$$

where X_{jft} is a vector of observed local characteristics and ξ_{jft} is the portion of location effects not captured by these observed characteristics. Substituting this expression into the empirical production function gives

$$\ln F_{ijft} = \underbrace{d_i + d_{fa}}_{\text{Individual Effect}} + \underbrace{X'_{jft}\beta_f}_{\text{Predicted Location Effect}} + d_{ft} + \epsilon_{ijft}, \quad (7)$$

where $\epsilon_{ijft} = u_{ijft} + \xi_{jft}$.

¹⁴Because a complete set of institution-field-year effects spans the field-year effects, such an approach would first require an additional restriction, such as holding location effects fixed within multi-year periods.

The vector X_{jft} includes lagged log measures of active researchers, publications, ten-year citations, and star researchers in field f , measured both at institution j and in the external MSA. These variables are constructed using the preceding five years, excluding the current year, so that the researcher's current output does not directly enter the local characteristics. We allow the coefficients β_f to vary across fields.

We define the predicted location effect as

$$\hat{d}_{jft}^{pred} = X'_{jft} \hat{\beta}_f,$$

which captures the portion of the full location effect predicted by observed characteristics of the institution and its surrounding research environment. The portion associated with unobserved local characteristics enters the residual through ξ_{jft} . These predicted location effects provide the location-productivity inputs used in the counterfactual analysis. When a researcher is reallocated, her individual effect is held fixed, while the predicted location effect changes from that of her origin to that of her destination.

Identification of β_f and the Predicted Location Effects Importantly, the coefficients β_f should *not* be interpreted as the causal effects of location characteristics, because those characteristics may be correlated with unobserved features of the research environment. The predicted location effects $X'_{j,t} \beta_f$ nevertheless *do* have a causal interpretation. Coefficients in β_f , though predictive and non-causal, captures how the location effects, which are causal, vary with observed location characteristics.

The identification of β_f follows the same logic as the mover design rather than relying on exogenous variation in any particular location characteristic. Changes in a researcher's output following a move identify differences in location effects across institutions, net of her portable individual effect and field-year shocks. The coefficients β_f then describe how these mover-identified location effects covary with the observed characteristics $X_{j,t}$. If researchers experienced no systematic change in output when moving between institutions with different characteristics, the predicted difference $(X_{j,t} - X_{j',t})' \beta_f$ would be close to zero. Conversely, if researchers systematically experienced larger output gains when moving to larger institutions, the coefficient on institution size would be positive, generating larger predicted location effects for larger institutions. This relationship does not imply that institution size causally increases location effects; it only describes how mover-identified location effects covary with institution size.

5.1.1 Event-Study Evidence Around Moves

A key identification concern in the mover design is that researchers anticipating future productivity growth may sort into more productive research clusters. In that case, the predicted location effects could partly reflect researchers' pre-existing productivity trajectories rather than the effects of their new research environments. To assess its likelihood, we examine research output around moves and test whether changes in output occur only after migration or instead reflect trends already underway before the move. We do not use the event-study estimates to construct the location effects; the event study provides complementary evidence on whether changes in research environments are followed by changes in the same researcher's output.

We follow the event-study approach in Finkelstein et al. (2016) and Lerner et al. (2024).¹⁵ We restrict the sample to researchers who change institutions once, so that pre-move outcomes are not affected by earlier moves. To account for the direction and magnitude of each move, we scale the event-study coefficients by the destination-origin difference in average research output. Let $o(i)$ and $d(i)$ denote researcher i 's origin and destination institution, respectively, and let $f(i)$ denote her field. Define

$$\Delta_i = \overline{\ln F}_{d(i)f(i)} - \overline{\ln F}_{o(i)f(i)},$$

where $\overline{\ln F}_{jf}$ is the average log research output in institution j and field f over all researcher-years (including non-movers). We then estimate

$$\ln F_{ijft} = \sum_{s \neq -1} \delta_s \Delta_i \mathbf{1}\{t - t_i^m = s\} + \sum_{s \neq -1} \lambda_s \mathbf{1}\{t - t_i^m = s\} + d_i + d_{fa} + d_{ft} + u_{ijft}, \quad (8)$$

where t_i^m is researcher i 's move year, and s indexes years relative to the move. The omitted period is the year before the move ($s = -1$).

The coefficients of interest are δ_s , which measure how much a mover's research output changes around the move, scaled by the average output gap between the destination and origin. If moves are driven by anticipated output growth, we would expect δ_s to rise before the move. In the absence of sorting on expected productivity growth, the pre-move coefficients should be close to zero.

Event-Study Results Figure 4 shows no statistically significant pre-trend in any outcome during the 15 years before a move. Output rises after researchers move to higher-output institutions and remains persis-

¹⁵A standard event study that plots average output around all moves is not informative in this setting because movers have different origins and destinations. Moves from less to more productive locations and moves in the opposite direction could offset each other, even if location effects are large.

tently higher. The initial increase may partly reflect short-run institutional support, but the persistence of the effect is consistent with location-specific productivity advantages. The estimates for patent citations are less precise, likely because patent citations are rare and relevant to fewer fields. Overall, the results are consistent with adjustment to the destination research environment rather than a trend already underway.¹⁶

5.1.2 Additive Decomposition: Individual vs. Location Effects

We next use the estimates from Equation 7 to quantify the importance of individual and predicted location effects for differences in research output across researchers. We conduct the decomposition over 2010–2015. Let $g \in \{H, L\}$ denote the high and low groups in a given comparison. For comparisons by researcher productivity, researchers are ranked within their fields by their average value of the corresponding research output over the period. For comparisons by researchers' cluster size, institutions and MSAs are ranked within each field by their average number of active researchers.

For each researcher represented in group g , we average the outcome and each component estimated in Equation 7 over the years in which she is observed in that group and then average across researchers. Let $\overline{\ln F}_g$, $\bar{d}_g^{indiv}$, $\bar{d}_g^{location}$, $\bar{d}_g^{field-year}$, and $\bar{\epsilon}_g$ denote the resulting group averages, where $\bar{d}_g^{indiv}$ is the group average of $\hat{d}_i + \hat{d}_{fa}$, $\bar{d}_g^{location}$ is the group average of $X'_{jft}\hat{\beta}_f$, $\bar{d}_g^{field-year}$ is the group average of $\hat{d}_{ft}$, and $\bar{\epsilon}_g$ is the group average of the estimated residuals.

Therefore, the difference in average output between the high and low groups can be written as

$$\begin{aligned} \overline{\ln F}_H - \overline{\ln F}_L &= \left(\bar{d}_H^{indiv} - \bar{d}_L^{indiv} \right) + \left(\bar{d}_H^{location} - \bar{d}_L^{location} \right) \\ &\quad + \left(\bar{d}_H^{field-year} - \bar{d}_L^{field-year} \right) + (\bar{\epsilon}_H - \bar{\epsilon}_L). \end{aligned}$$

The shares of the observed output difference attributable to the difference in individual effects, predicted location effects, and unobserved factors are

$$\begin{aligned} \text{Share}_{H,L}^{indiv} &= \frac{\bar{d}_H^{indiv} - \bar{d}_L^{indiv}}{\overline{\ln F}_H - \overline{\ln F}_L}, \\ \text{Share}_{H,L}^{location} &= \frac{\bar{d}_H^{location} - \bar{d}_L^{location}}{\overline{\ln F}_H - \overline{\ln F}_L}, \\ \text{Share}_{H,L}^{unexplained} &= \frac{\bar{\epsilon}_H - \bar{\epsilon}_L}{\overline{\ln F}_H - \overline{\ln F}_L}. \end{aligned} \tag{9}$$

¹⁶Appendix Figure A10 examines moves to higher- and lower-output institutions separately and shows similar adjustment patterns, suggesting that research output adjusts in the direction predicted by the destination-origin output gap. This symmetry also supports the additive separability assumption in Equation 7.

The unexplained share may partly capture location effects associated with unobserved local characteristics. We expect the field-year component to be small because groups are defined within fields over a short period.

Decomposition Results Table 1 presents the decomposition by individual researcher productivity. Its three panels correspond to IHS publications, citations, and patent citations; within each panel, the columns compare researchers above and below the median, in the top and bottom quartiles, and in the top and bottom deciles. The upper block reports the overall output difference and the corresponding component differences, while the lower block reports each component as a share of the overall difference. Researcher effects account for 63–67% of the publication gap, 71–76% of the citation gap, and 75% of the patent-citation gap. Location effects account for 8%, 7%, and 5–10% of these gaps, respectively. Although these shares are smaller, they are still quantitatively important.¹⁷ Since we are comparing *individuals* based on their productivity, it is unsurprising that the researchers’ individual effects account for the largest explanatory share.

More directly relevant for assessing the effects of reallocating researchers across institutions and MSAs, Table 2 groups researchers by institution or MSA size within fields. Location effects explain a substantial share of the research output differences: 38–80% for publications, 34–64% for citations, and 82–94% for patent citations.¹⁸ The sizable variation in location effects across institutions and MSAs suggests that reallocating researchers across locations may have substantial effects on research output, which we formally quantify in Section 6.

5.2 Estimating Agglomeration Effects

Section 5.1 separates individual from location effects to quantify the direct effect of reallocation on researchers who move. We now estimate how research output responds to changes in own-institution and external-cluster size. These elasticities determine the indirect effect of reallocation on researchers who remain in the origin and destination clusters.

Recall from Equation 3 that location productivity depends on research cluster size: $\ln A_j(N_j) = \ln A_{0j} + \alpha \ln N_j$, where A_{0j} captures location-specific productivity unrelated to cluster size and α is the

¹⁷The remaining 16–29% of the output gap is unexplained and may include location productivity associated with unobserved local characteristics. The field-year component is not reported because its contribution is negligible.

¹⁸These large explanatory shares may partly arise mechanically because the predicted location effects are constructed from observed characteristics of the local research environment, including the grouping variables themselves, namely, institution and MSA size. Grouping researchers by these same characteristics may generate greater separation in predicted location effects by construction. Figures A11–A13 report the corresponding field-specific results. The broad patterns are similar across fields, although the magnitude of the shares varies substantially.

agglomeration elasticity. To capture agglomeration flexibly, we distinguish between internal and external cluster size. Internal cluster size, N_{jft}^I , is the number of active researchers in field f at institution j in year t . External cluster size, N_{jft}^E , is the number of active researchers in the MSA containing j within the same field, excluding the home institution. Specifically, we estimate

$$\ln F_{ijft} = \underbrace{d_i + d_{fa}}_{\text{Individual Effect}} + \underbrace{\alpha^I \ln N_{jft}^I + \alpha^E \ln N_{jft}^E + d_{jf} + d_{ft}}_{\text{Location Effect}} + u_{ijft}, \quad (10)$$

where α^I and α^E are the internal and external agglomeration elasticities, respectively.¹⁹ The institution-field fixed effect d_{jf} absorbs persistent productivity differences across institution-field cells that are unrelated to contemporaneous cluster size, corresponding to the empirical counterpart of $\ln A_{0j}$.

Section 5.1 explains that limited mobility makes it difficult to estimate the levels of location fixed effects precisely. This is less of a concern here because we do not use the estimates of d_{jf} —they enter only as controls. Our parameters of interest are α^I and α^E , which are identified from changes in cluster size within institution-field cells over time, controlling for individual effects and common field-year shocks.

As a further robustness check, we add MSA-by-year fixed effects to absorb time-varying shocks common to all institutions in the same metropolitan area. These fixed effects also reduce the identifying variation in external-cluster size, so we treat this specification as a robustness check rather than the baseline.

Identification Challenge The main identification challenge is that cluster size is endogenous. Cluster size may be correlated with unobserved local productivity advantages, such as the generosity of funding or the quality of research infrastructure. Institution-field fixed effects absorb time-invariant advantages of each institution-field pair, but they do not absorb time-varying institution-field productivity shocks. Thus, OLS estimates of α^I and α^E may be biased.

5.2.1 Bartik Instrument for Cluster Size

To address this identification problem, we construct Bartik shift-share instruments using detailed *subfields* nested within each broad field f . The instruments combine each institution or external cluster’s predeter-

¹⁹Because cluster size can be zero, in the empirical implementation, we enter both cluster size measures as $\ln(N + c)$, with $c = 1$ in the baseline analysis, although we retain the notation $\ln N$ to maintain the connection with the model. Adding one is the minimal adjustment needed to define the regressor at zero and keeps the specification closest to the model’s log form. The implied elasticity with respect to N is $\alpha N / (N + c)$, which approaches α when N is large relative to c . When N is small, the implied elasticity is more sensitive to the additive constant. Section 6 examines this sensitivity in counterfactuals involving very small clusters.

mined composition across detailed subfields with subsequent *worldwide* growth in those subfields. The idea is that institutions and local research clusters initially specialized in globally expanding subfields should experience larger predicted growth in the number of researchers. This predicted growth is likely driven by the global evolution of scientific fields rather than by contemporaneous productivity shocks at any single institution or local cluster.²⁰

Construction of the Bartik Instruments To allow worldwide subfield growth sufficient time to generate variation while smoothing annual noise in research output, we organize the IV construction around three periods: 1995, 2005, and 2015. Each period pools data over the focal year and the two preceding years: For example, the 1995 period uses data from 1993–1995. We use 1995 to measure base-period cluster sizes and subfield shares. We then predict cumulative growth from the base period to each later period, 2005 and 2015, using worldwide growth in detailed subfields.

Let h index a detailed subfield within broad field f . Let $N_{jh,0}^I$ and $N_{jh,0}^E$ denote the 1995 base-period number of active researchers in subfield h at institution j and in its external cluster, respectively. The corresponding baseline shares of subfield h within broad field f at institution j are

$$\omega_{jh(f)0}^I = \frac{N_{jh(f)0}^I}{\sum_{\ell \in f} N_{j\ell 0}^I} \quad \text{and} \quad \omega_{jh(f)0}^E = \frac{N_{jh(f)0}^E}{\sum_{\ell \in f} N_{j\ell 0}^E}.$$

Using worldwide publication records, let g_{ht}^w denote the growth rate in the number of active researchers in subfield h from the base period to period t , where $t \in \{2005, 2015\}$. The predicted internal and external cluster growth rates are share-weighted sums of worldwide subfield growth:

$$\hat{g}_{jft}^I = \sum_{h \in f} \omega_{jh(f)0}^I g_{ht}^w \quad \text{and} \quad \hat{g}_{jft}^E = \sum_{h \in f} \omega_{jh(f)0}^E g_{ht}^w.$$

Because we estimate the production function in levels rather than first differences, we convert these predicted growth rates into predicted cluster-size *levels*. Let N_{jf0}^I and N_{jf0}^E denote base-period internal and

²⁰For example, a statistics department or regional statistics community with a large preexisting share of researchers in machine-learning-related areas would be predicted to grow more as machine learning expands worldwide. For a researcher in that department or cluster, controlling for characteristics of her own subfields, this predicted growth shifts local cluster size through channels that are plausibly unrelated to the direct productivity shocks of any individual researcher.

external cluster sizes. The internal and external Bartik instruments for periods 2005 and 2015 are

$$\begin{aligned}\ln \hat{N}_{jft}^{I,\text{Bartik}} &= \ln [N_{jfo}^I (1 + \hat{g}_{jft}^I)] , \text{ and} \\ \ln \hat{N}_{jft}^{E,\text{Bartik}} &= \ln [N_{jfo}^E (1 + \hat{g}_{jft}^E)] ,\end{aligned}\tag{11}$$

for $t \in \{2005, 2015\}$. The predicted cluster size in 1995 is simply its baseline value and does not provide meaningful shift-share variation. Because all IV specifications include institution-field fixed effects, these base-period cluster sizes are absorbed by these fixed effects. Identification therefore comes from differences across institution-field cells in Bartik-predicted growth over 1995–2005 and 1995–2015.

Identifying Assumptions The Bartik instruments must satisfy two conditions:

- **Inclusion Restriction:** Baseline exposure to rapidly growing detailed subfields must predict subsequent growth in actual cluster size. We examine this restriction directly in the first-stage regressions.
- **Exclusion Restrictions:** Conditional on the fixed effects and controls in Equation 10, the Bartik-predicted changes in cluster size must affect a researcher’s output only through internal and external cluster size. To ensure the validity of this restriction, one of the following two conditions should hold.
 - *Exogeneity of worldwide subfield growth.* Following Borusyak et al. (2022), worldwide subfield growth g_{ht}^w must be uncorrelated with institution-field-specific productivity shocks. This requires worldwide subfield growth to reflect broad scientific trends rather than shocks originating from any particular institution or local research cluster. This condition is plausible because no single institution or cluster accounts for a large share of the worldwide stock of researchers in a subfield.
 - *Exogeneity of baseline shares.* Following Goldsmith-Pinkham et al. (2020), base-period subfield shares must be uncorrelated with subsequent institution-field productivity shocks. The condition is plausible because the shares are measured before the growth they predict, and persistent local advantages associated with the shares are absorbed by institution-field fixed effects. The remaining concern is that base-period shares may proxy for differential institution-field trends, which we examine using the pre-trend tests below.

Potential Threats to Identification We address three potential threats to the exclusion restriction.

Direct exposure to worldwide subfield growth. Worldwide growth in a researcher’s own subfields may directly affect her output by increasing publication opportunities or citation demand, rather than solely

through changes in local cluster size. To address this concern, we control for the researcher's and her coauthors' direct exposure to worldwide subfield growth. Let $\mathcal{H}_i$ denote the set of subfields associated with researcher i , and let $\omega_{ijh,0}^A$ denote the base-period share of subfield h among subfields in $\mathcal{H}_i$ at institution j . Researcher i 's predicted *direct* exposure to worldwide subfield growth is thus

$$\widehat{g}_{ijt}^A = \sum_{h \in \mathcal{H}_i} \omega_{ijh,0}^A g_{ht}^w.$$

We construct the author's exposure control as

$$\ln \widehat{N}_{ijft}^A = \ln [N_{jfo}^I (1 + \widehat{g}_{ijt}^A)]. \quad (12)$$

We construct an analogous coauthor exposure control, $\ln \widehat{N}_{ijft}^C$, using the detailed subfields associated with researcher i 's coauthors. Including both controls helps absorb the direct effects of worldwide growth in the researcher's own research areas and collaboration network, and thus isolates the variation in cluster size generated by the Bartik instruments.

Concentrated weights. A second concern is that the instruments may be driven by a small number of rapidly growing subfields whose baseline shares are correlated with unobserved institution-field-level productivity shocks. To assess this concern, we reconstruct the Bartik instruments, omitting one detailed subfield at a time, and re-estimate the preferred specification. If the estimates remain stable across these exercises, the results are unlikely to be driven by any single subfield.

Pre-existing productivity trends. A third concern is that institutions initially specialized in rapidly growing subfields may already have been on different productivity trajectories before 1995. If future Bartik-predicted growth is correlated with output growth before the base period, the instruments could capture these pre-existing trends rather than exogenous changes in cluster size. To assess this concern, we test whether Bartik-predicted cluster growth after 1995 is associated with institution-field output growth before the base period. Let $\overline{\ln F}_{jfr}$ denote average IHS research output for institution j in field f over the three-year period ending in year r , where $r \in \{1985, 1990\}$. We estimate

$$\overline{\ln F}_{jf,1995} - \overline{\ln F}_{jfr} = \pi_0 + \pi^I \widehat{g}_{jff,2005}^I + \pi^E \widehat{g}_{jff,2005}^E + \pi_f + \varepsilon_{jf}, \quad r \in \{1985, 1990\}, \quad (13)$$

where π_f denotes field fixed effects. We repeat the same regression using predicted cluster growth over the longer horizon, $\widehat{g}_{jff,2015}^I$ and $\widehat{g}_{jff,2015}^E$. If predicted growth in an institution-field cell is not preceded by an

upward trend in research output, the coefficients π^I and π^E should not be positive.

5.2.2 Main Results

OLS Estimates Appendix Table A3 presents the OLS estimates of Equation 10 using the full sample from 1970–2015. Own-institution size is positively associated with all three research outcomes across specifications. The estimates for external-cluster size are smaller and vary across specifications. Because these estimates are descriptive and do not enter the counterfactual analysis, we focus on the IV estimates below.

First-Stage Estimates Before discussing our main IV estimates, we assess the relevance of the instruments in Table 3. The IV sample is organized around three periods: 1995, 2005, and 2015; each period pools data over the focal year and the two preceding years. Columns 1–3 use log own-institution size as the dependent variable, and Columns 4–6 use log external cluster size. Columns 1 and 4 include the baseline fixed effects from Equation 10. Columns 2 and 5 add the author and coauthor controls for direct exposure to worldwide subfield growth. Columns 3 and 6 further add MSA-by-year fixed effects.

The results show that own-institution size is predicted almost entirely by the own-institution instrument, whereas external cluster size responds to both instruments. This cross-loading is plausible because an institution’s initial subfield composition may reflect broader local scientific specialization and therefore predict growth both within the institution and in the surrounding research environment. As a result, the external-cluster elasticity is identified from part of predicted external-cluster growth that is distinct from predicted own-institution growth. This limited independent variation may contribute to the lower precision of the external-cluster estimates.

Adding the author and coauthor controls has little effect on the first stage. MSA-by-year fixed effects absorb additional geographic variation, particularly variation generated by the external-cluster instrument, but the instruments remain jointly relevant. The Kleibergen–Paap rk Wald F-statistics reported in Table 4 range from 12.97 to 15.88 across specifications.

Main OLS and IV Estimates Table 4 reports OLS and 2SLS estimates using the IV sample and follows the same sequence of fixed effects and controls as in Table 3. We use Column 5 as our preferred specification because it includes the author and coauthor controls for direct exposure to worldwide subfield growth, while MSA-by-year fixed effects are included as a more demanding robustness specification in Column 6.

The OLS estimates provide a within-sample benchmark. Own-institution size is positively associated with all three research outcomes, whereas the coefficients on external-cluster size are close to zero.

The IV estimates show a robust positive effect of own-institution size. The coefficient is positive and statistically significant for all three outcomes and across all specifications. In our preferred specification, the estimates imply that a 10% increase in own-institution size raises output by approximately 1.99% for publications, 8.28% for citations, and 1.48% for patent citations.²¹ The effect is largest for citations, suggesting that agglomeration may matter more for research impact than for publication quantity.

The evidence for external-cluster effects is weaker. The external-cluster coefficient is positive and statistically significant for publications, but is small and insignificant for citations and patent citations. Adding MSA-by-year fixed effects further reduces the precision of the estimates. The weaker evidence for external-cluster effects is consistent with the limited independent variation discussed in the first-stage results.²²

Overall, the results indicate a robust *own-institution* agglomeration effect. For the baseline counterfactual analyses, we use the publication agglomeration elasticities $\hat{\alpha}^I = 0.199$ and $\hat{\alpha}^E = 0.074$.

Heterogeneity Across Fields The strength of agglomeration forces may differ across fields because the processes of knowledge production and dissemination vary substantially across fields. In addition, fields differ in the extent to which research relies on team production, specialized infrastructure, and commercialization, all of which may affect the gains from geographic clustering.

Table 5 re-estimates our preferred IV specification separately for the five broad field groups. Because splitting the sample weakens the first stage, we interpret these estimates as suggestive. The estimates broadly mirror the pooled results. The own-institution coefficient is positive for publications in every field group, and the estimates are of the same order of magnitude, although only statistically significant for engineering and computer science. For citations, the coefficient is positive and statistically significant in all groups except the humanities, with the largest estimate in the life sciences. This is consistent with the possibility that institution scale may be particularly important for research impact in fields that rely heavily on teams,

²¹One concern with the patent-citation results is that patent citations are relevant to fewer fields and are zero for many observations. As a robustness check, we restrict the sample to fields in which more than 30% of papers receive at least one patent citation. The estimates remain similar. Due to space constraints, these results are not reported but are available upon request.

²²The IV estimates are larger than the corresponding OLS estimates, especially for citation-based outcomes. This pattern is consistent with attenuation bias in OLS due to measurement error in cluster size. It may also reflect heterogeneous treatment effects: The Bartik instrument identifies effects for institution-field cells whose size grows because they were initially exposed to globally expanding subfields. These cells may have especially high returns to scale, since larger institutions in fast-growing research areas may be particularly effective at increasing visibility, improving research quality, and strengthening connections to other researchers and inventors. This may generate especially large effects on citations and patent citations.

infrastructure, and professional networks. For patent citations, the larger positive estimates are concentrated in the natural and life sciences, where chemistry, physics, biology, and medicine are closely linked to patents.

The external-cluster coefficients are generally smaller and less precisely estimated. They are positive and statistically significant for publications in the social sciences and life sciences, but show no consistent pattern for citations or patent citations. Overall, the field-group-specific estimates are consistent with the pooled results: Agglomeration effects appear to operate primarily through the size of researchers' own institutions, while the evidence for broader external-cluster effects is weaker. Because the subgroup estimates are less precisely identified, the counterfactual analysis uses the pooled estimates from Table 4.

Leave-One-Out Test To assess whether a small number of subfields drive the IV estimates, we reconstruct both Bartik instruments 364 times after omitting one detailed subfield at a time and re-estimate the preferred specification for each outcome. Appendix Figure A14 displays, for each outcome and both internal and external elasticities, the ten leave-one-out estimates with the largest absolute deviations from the corresponding baseline estimate. Even among these most influential omissions, the estimates remain nearly unchanged for both elasticities and all three outcomes. Moreover, every leave-one-out estimate of the own-institution elasticity remains statistically significant across all three outcomes, and every estimate of the external-cluster elasticity remains statistically significant for publications.

Pre-Trend Tests Table A4 reports the estimates of Equation 13, relating pre-1995 changes in institution-field research output to Bartik-predicted growth in own-institution and external-cluster size. We examine output growth over 1990–1995 and 1985–1995, using predicted cluster growth over 1995–2005 and 1995–2015. Predicted own-institution growth is generally uncorrelated or negatively correlated with pre-1995 output growth. These patterns would tend to *attenuate* rather than generate our positive own-institution estimates. Predicted external-cluster growth is positively associated with some pre-period outcomes, so we interpret the external-cluster estimates more cautiously, consistent with the generally weaker evidence for this channel.²³

Anderson-Rubin Confidence Sets Although the Kleibergen–Paap rk Wald F-statistics do not indicate a serious weak-instrument problem, we report Anderson–Rubin confidence sets because they provide infer-

²³The positive pre-trends associated with predicted external-cluster growth are less relevant to the own-institution elasticity, because in the first stage, the estimated effect of the external Bartik instrument on own-institution size is close to zero.

ence that remains valid under weak identification. Figure A15 shows that the 95% confidence sets exclude all parameter pairs with non-positive values of α^I for all three outcomes.

5.2.3 Additional Robustness Checks

Heterogeneity by Cluster Size Our model and baseline empirical specification assume *constant* internal and external agglomeration elasticities. However, if agglomeration gains decline sufficiently with cluster size, an additional researcher could contribute more in a small cluster than in a large one, so dispersion might raise rather than reduce knowledge production. To assess this possibility, we assign each institution-field cell to a tertile based on its 1995 cluster size and interact current log cluster size with the corresponding tertile indicators:

$$\ln F_{ijft} = d_i + d_{fa} + \sum_{q=1}^3 \alpha_q^I (\ln N_{jft}^I \times Q_{jf,q}^I) + \sum_{q=1}^3 \alpha_q^E (\ln N_{jft}^E \times Q_{jf,q}^E) + d_{jf} + d_{ft} + u_{ijft}, \quad (14)$$

where $Q_{jf,q}^I$ indicates that institution-field cell (j, f) belongs to tertile q of the institution size distribution, and $Q_{jf,q}^E$ is defined analogously using external-cluster size; α_q^I and α_q^E are the internal and external agglomeration elasticities that vary across cluster-size groups.

Table A5 reports IV estimates of α_q^I and α_q^E . For publications, the own-institution elasticity is similar in the smallest and largest tertiles. For citations and patent citations, the own-institution elasticity declines with cluster size but remains positive and sizable even in the largest tertile. The external-cluster estimates show no systematic decline with size. Overall, the results do not indicate a sharp decline in agglomeration gains, so we retain the pooled constant-elasticity estimates in the counterfactual analysis.

Alternative Measures of Cluster Size Table A6 considers alternative measures of cluster size—the number of active researchers and publications—and alternative definitions of external cluster—the MSA or a 25-km radius. All cluster measures are field-specific, and the home institution is excluded from the external cluster in all specifications. Own-institution effects remain positive and are especially stable for citations. However, for publications, the estimate becomes statistically insignificant when cluster size is measured by paper counts rather than author counts. The external-cluster effect on publications is positive and statistically significant under the MSA definition, but we find no statistically significant external-cluster effect for any outcome when the cluster is defined using a 25-km radius.

Alternative Approaches for Zero-Valued Outcomes Our baseline analysis uses the IHS transformation of a research outcome as the dependent variable because research outcomes contain zeros. Chen and Roth (2024) show that estimates using log-like transformations with zero-valued outcomes are not unit-invariant and the unit of the outcome implicitly determines the weight placed on the extensive margin. Thus, we use their approach by explicitly calibrating the value placed on the extensive margin.²⁴ We transform a research outcome using $m(F) = \log(F)$ for $F > 0$ and $m(0) = -x$, considering $x = 0.1$ and $x = 1$, so that moving from zero to one is valued as equivalent to 10- and 100-log-point changes along the intensive margin, respectively. Table A7 shows that the estimated effects using $x = 0.1$ are slightly smaller than the baseline estimates, while the estimates using $x = 1$ are slightly larger. Overall, the estimates are close to the baseline results and are not very sensitive to the choice of x .

6 Costs and Benefits of Spatial Reallocation of Researchers

Section 5 provides the inputs from the research-production side needed to evaluate researcher reallocation: location effects and agglomeration effects. We now move to the counterfactual analysis based on Equation 4. For any counterfactual reallocation, the change in log aggregate output is

$$\Delta \ln Y = \Delta \ln B + s^N \Delta \ln Q,$$

where $\Delta \ln B$ is the change in the knowledge access component of aggregate output and $\Delta \ln Q$ is the change in the steady-state knowledge stock.

In principle, we could evaluate each reallocation by directly computing both terms. However, in practice, we estimate the knowledge-production-side inputs in the previous section but not the access elasticity that governs $\Delta \ln B$. Rather than imposing a particular estimate from the prior literature, we invert the exercise. For each reallocation, we solve for the smallest access elasticity that sets $\Delta \ln Y = 0$, which we refer to as the *break-even access elasticity*. If estimates from the prior literature exceed this threshold, the gain from improved access would be large enough to offset the loss in knowledge production.

²⁴Chen and Roth (2024) also recommend estimating the ATE in levels using Poisson regression and expressing it as a percentage effect. However, we do not pursue this approach because they note that Poisson IV regression with a continuous IV may not have a local average treatment effect interpretation.

6.1 Parameter Calibrations

To calibrate the knowledge production parameters, we use the predicted location effects, $\hat{d}_{jft}^{pred}$, and internal and external agglomeration elasticities, $\hat{\alpha}^I$ and $\hat{\alpha}^E$, estimated in Section 5. On the knowledge-access side, we take s^N , the importance of frontier knowledge in economic output; δ , the depreciation rate of knowledge; and η , the dependence of new research on the existing knowledge stock, from the literature. We calibrate λ_j , the share of local knowledge access attributable to nearby researchers to several plausible constant values. We do not impose a value for θ , the elasticity of knowledge access with respect to researchers per capita, but instead leave it free and solve for its break-even value in each counterfactual.

Production-Side Inputs The internal and external agglomeration elasticities are

$$\hat{\alpha}^I = 0.199 \text{ and } \hat{\alpha}^E = 0.074$$

in the baseline counterfactual analyses. They are from our preferred IV specification in Column 5 of Table 4 using publications as the research output. As a robustness check, we also use elasticities in Table A6, where cluster size is measured by paper counts, rather than author counts.

Knowledge-Stock Parameters We set $s^N = 0.2$, near the upper end of estimates of the elasticity of aggregate productivity with respect to research activity in the macro and cross-country literature, roughly ranging from 0.08–0.23 (Coe and Helpman, 1995; Valero and Van Reenen, 2019; Fieldhouse and Mertens, 2024). Because s^N scales the contributions of both knowledge access and knowledge production, this choice does not mechanically favor either side of the trade-off.

We set the depreciation rate of the frontier knowledge stock $\delta = 0.15$, the conventional value in the R&D-stock literature (Hall et al., 2010).²⁵ We set the dependence of new research on the existing knowledge stock $\eta = 0.5$, a value in the range implied by Jones (1995).²⁶

Access-Side Calibration The share of local knowledge access attributable to nearby researchers is

$$\lambda_j = \frac{(N_j/M_j)^\theta}{S_0 + (N_j/M_j)^\theta}.$$

²⁵Because δ enters the steady-state knowledge stock as an additive constant in $\ln Q$, it cancels from the log change in output that our reallocation exercise computes, so our results are insensitive to its value.

²⁶There is a debate about the magnitude of η in the growth literature. More details are provided in footnote 6.

Because the non-local access component S_0 is unobserved, we do not attempt to recover the location-specific values of access share λ_j . Instead, we approximate λ_j using common values $\bar{\lambda} = 0.2, 0.5$, or 0.7 , assuming that the local and non-local channels each account for a fixed share of knowledge access everywhere. For each value, we solve for the break-even value θ^* and calculate the corresponding access elasticity $s^N \theta^* \bar{\lambda}$. This is the object that we compare with estimates of local research spillovers in the literature.

6.2 Counterfactual Reallocation Procedure

We conduct each counterfactual separately within each field f , reallocating researchers from MSAs with more researchers than their population share implies to MSAs with fewer. We then aggregate the resulting outcomes across fields. To simplify notation, we suppress the field subscript below.

1. **Measure researcher surpluses and deficits.** Let $\bar{r} = \sum_j N_j / \sum_j M_j$ denote the aggregate researcher-to-population ratio. The researcher allocation proportional to population is $N_j^{\text{eq}} = \bar{r} M_j$. Define MSA j 's researcher *gap* as

$$g_j = N_j^{\text{eq}} - N_j.$$

An MSA has a researcher surplus if $g_j < 0$ and a deficit if $g_j > 0$.

2. **Choose the reallocation intensity.** Let $\kappa \in [0, 1]$ denote the fraction of the gap to be closed, i.e., the reallocation intensity. The counterfactual number of researchers in MSA j is

$$N_j(\kappa) = N_j + \kappa g_j = (1 - \kappa) N_j + \kappa \bar{r} M_j.$$

At $\kappa = 0$, the observed allocation is unchanged. At $\kappa = 1$, researchers are allocated in proportion to population.

3. **Select and assign movers.** From each MSA with a researcher surplus, we randomly select researchers without replacement until the required outflow, $N_j - N_j(\kappa)$, is reached. We assign each mover to a destination MSA with a researcher deficit with probability proportional to its remaining researcher deficit, $g_j \kappa$. Within a destination MSA, we assign movers to institutions in proportion to their sizes.
4. **Compute counterfactual outcomes.** We calculate the counterfactual access $\ln B$ and the counterfactual steady-state knowledge stock $\ln Q$, and the implied changes $\Delta \ln B$ and $\Delta \ln Q$. Because individual assignments are random, we repeat each exercise 50 times and report the average outcome.

6.3 Counterfactual Knowledge Access

A counterfactual reallocation affects population's access to knowledge through each MSA's counterfactual researcher count. In each counterfactual, we hold fixed the non-research population M_j , baseline non-research productivity S_{0j} , and the calibrated non-local access component S_0 . Only the researcher count changes from its observed value N_j to its counterfactual value $N_j(\kappa)$. Under reallocation intensity κ , the counterfactual access term is

$$\widehat{B}(\kappa) = \sum_{j=1}^J M_j S_{0j}^{1-s^N} \left[S_0 + \left(\frac{N_j(\kappa)}{M_j} \right)^\theta \right]^{s^N}.$$

Non-Researcher Migration The baseline analysis holds the non-research population fixed across MSAs. As a robustness exercise, we allow non-researchers to respond to counterfactual wage changes using the location-choice model in Appendix A1.3.3. Researcher reallocation raises knowledge access and wages in destination MSAs, which can attract non-researchers toward locations with lower baseline productivity. This response only modestly offsets the direct access gain and does not materially alter our conclusion. The spatial misallocation due to migration is similar to the efficiency cost of place-based policies in general.

6.4 Counterfactual Knowledge Production

Throughout this subsection, we retain the log notation for log research outcome $\ln F$ used in the model, but the empirical calculations use the IHS-transformed F instead, consistent with our estimation.

Non-Movers Based on Equation 4, researchers who are not relocated are indirectly affected because a reallocation changes the size of their cluster. Thus, for a researcher who remains at institution j , her output changes through an agglomeration adjustment:

$$\Delta \ln \widehat{F}_{ijt}^{non-mover} = \hat{\alpha}^I \left[\ln(N_{jft}^{I,cf}) - \ln(N_{jft}^{I,0}) \right] + \hat{\alpha}^E \left[\ln(N_{jft}^{E,cf}) - \ln(N_{jft}^{E,0}) \right]$$

where $N_{jft}^{I,cf}$ and $N_{jft}^{E,cf}$ are the counterfactual internal and external cluster sizes, and $N_{jft}^{I,0}$ and $N_{jft}^{E,0}$ are the corresponding initial values.²⁷

²⁷As noted in footnote 19, we use $\ln(N+1)$ to estimate agglomeration elasticities, so in counterfactual analyses, we apply the same transformation to both observed and counterfactual cluster sizes. Because this transformation is steep near zero, predictions can be sensitive when a counterfactual changes very small clusters by only a few researchers. Thus, for the exercise in Section 6.5.3 that involves very small clusters, we apply a larger value of c to make the transformation less curved near zero, reducing the influence of changes involving only a few researchers. Specifically, we re-estimate the elasticities using $c = 5$ and apply the same transformation in the counterfactual. Results are similar if we use $c = 10$.

Movers For a researcher who is relocated from institution j to j' , the change in her output reflects the change in the full location effect between the destination and origin. This is empirically approximated using the difference in the predicted location effect, augmented by the change in agglomeration effects in the destination induced by the reallocation:

$$\Delta \ln \hat{F}_{ijft}^{mover} = \hat{d}_{j'ft}^{pred} - \hat{d}_{jft}^{pred} + \hat{\alpha}^I [\ln(N_{j'ft}^{I,cf}) - \ln(N_{jft}^{I,0})] + \hat{\alpha}^E [\ln(N_{j'ft}^{E,cf}) - \ln(N_{jft}^{E,0})]$$

In sum, non-movers are affected only through changes in cluster size, while movers experience a change in the full location effect. Individual fixed effects, academic-age effects, and field-year effects are held fixed.

We add the corresponding change to each researcher's initial IHS research output and apply the standard hyperbolic sine function to convert the *counterfactual* IHS research output back to levels, $\hat{F}_{ijft}^{cf}$. After aggregating counterfactual *flow* research output $\hat{F}_{ijft}^{cf}$ across researchers, the change in the steady-state frontier knowledge *stock* is

$$\Delta \ln \hat{Q} = \ln \left(\frac{\sum_i \hat{F}_{ijft}}{\delta} \right)^{\frac{1}{1-\eta}} - \ln \left(\frac{\sum_i F_{ijft}}{\delta} \right)^{\frac{1}{1-\eta}}.$$

After completing the reallocation separately within each field, we aggregate counterfactual research output across all researchers and fields to obtain the national change in the knowledge stock. For each value of θ , we compute $\Delta \ln \hat{Y} = \Delta \ln \hat{B}(\theta) + s^N \Delta \ln \hat{Q}$. We identify θ^* as the smallest value of θ for which $\Delta \ln \hat{Y} \geq 0$ in each simulation, and report the implied access elasticity $s^N \theta^* \bar{\lambda}$.

6.5 Counterfactual Results

Tables 6–8 present the counterfactual results. The left block reports changes in *flow* research output for three groups: researchers who remain in the origin (“Origin Stayers”), researchers already located in the destination (“Dest. Incumb.”), and researchers who move (“Movers”). The group-specific percentages use each group's own baseline output as the denominator and therefore do not add directly. The “Net Change” column reports the corresponding aggregate change. The final two columns convert this production loss into the break-even θ^* and the implied break-even knowledge-access elasticity $s^N \theta^* \bar{\lambda}$.

6.5.1 Bilateral Case Studies

We begin with two transparent bilateral counterfactual reallocations: Boston to Dallas–Fort Worth and San Francisco to Las Vegas. These exercises move researchers from established research hubs to large or rapidly

growing MSAs with relatively few researchers per capita. In each case, we set $\kappa = 0.1$.²⁸

Table 6 shows that shrinking the origin lowers output among stayers at institutions in the origin MSAs, expanding the destination raises output among incumbents in the destination MSAs, and movers lose output because destination institutions have lower location effects. The declines for movers are moderated because they carry their individual effects. Moreover, the convergence in cluster sizes narrows the gap in location effects between the destination and the origin. On net, aggregate knowledge production flow (i.e., publications) would fall by around 0.6% in each case.

The implied break-even access elasticity ranges from 0.0176 to 0.0193, well below existing estimates of localized university spillovers. Kantor and Whalley (2014), for example, estimate an elasticity of local noneducation income with respect to university spending of 0.08, while Jaffe (1989) report similarly large or larger elasticities of local patenting with respect to university R&D. These estimates do not correspond exactly to our parameter $s^N \theta \bar{\lambda}$, which measures the elasticity of local non-research productivity with respect to access to researchers. Nevertheless, they provide useful benchmarks for its plausible magnitude. Unless $s^N \theta \bar{\lambda}$ is substantially smaller than these estimates, the access gains from the bilateral reallocations would outweigh the associated loss in knowledge production.

6.5.2 Reallocation Among Large MSAs

We next extend the bilateral exercises to the 50 most populous MSAs. Within this group, we move researchers from the ten MSAs with the most researchers per capita to the ten with the fewest, again setting $\kappa = 0.1$. Restricting the exercise to large MSAs avoids reallocations into very small research clusters, where counterfactual predictions are more sensitive to functional-form assumptions.²⁹

Panel A of Table 7 shows that reallocation among the 50 largest MSAs lowers aggregate flow research output by only 0.32%. The break-even access elasticity is about 0.010 and rises modestly to 0.012 when non-researcher migration is allowed.³⁰ Thus, the conclusion from the bilateral exercises extends to a broader set of large MSAs: only a modest access gain is required to offset the loss in knowledge production.

²⁸Under $\kappa = 0.1$, Boston and San Francisco would lose around 5% and 3% of their researchers, while Dallas–Fort Worth and Las Vegas would gain approximately 31% and 70%, respectively.

²⁹See footnotes 19 and 27 for a detailed discussion of this sensitivity.

³⁰Appendix A4 provides more details on how we simulate non-researchers' migration.

6.5.3 Reallocation Across All MSAs

We finally extend the exercise to all MSAs that host research institutions. This broader reallocation includes many very small research clusters, for which the log transformation is particularly sensitive to changes involving only a few researchers. To limit the influence of this near-zero curvature, we use transformations that are less sensitive at small cluster sizes.³¹

Panel B of Table 7 shows that extending the exercise to all MSAs raises the production loss to 1.14%, largely because many movers are assigned to very small institutions with lower location effects. The break-even access elasticity rises to 0.039–0.043, still below the 0.08 benchmark but by a narrower margin. This contrast suggests that reallocation toward large underserved MSAs could be less costly than broad dispersion into small, isolated research clusters.

6.5.4 Alternative Specifications

Table 8 considers four alternatives to the baseline large-MSA reallocation, all under $\bar{\lambda} = 0.5$.

Alternative Measures of Knowledge Access Panel A considers the same reallocation as Table 7 Panel A but allows local access to depend on publications or citations per capita rather than researcher counts. The break-even elasticity falls from 0.0096 to 0.0068 and 0.0034, respectively, because movers from research-intensive origins bring more research output and impact to destinations than their headcount alone captures.

Agglomeration Through Research Activity Panel B measures cluster size by publication output (a quality-weighted scale) rather than researcher counts and applies the corresponding agglomeration estimates from Table A6. Because movers are relatively productive, their relocation changes publication-based cluster size more than headcount-based size, narrowing the origin–destination gap in location productivity. The net decline in flow research output is only 0.10%, reducing the break-even elasticities to 0.0008–0.0031.

College-Educated Population If college-educated workers benefit more from frontier knowledge, the baseline exercise may overstate access gains because these workers already co-locate more closely with researchers. Panel C therefore reallocates researchers according to the spatial distribution of college graduates

³¹ As discussed in footnote 27, the baseline specification uses $\ln(N + 1)$. For this exercise, we re-estimate the agglomeration coefficients using $\ln(N + 5)$ and apply the same transformation to observed and counterfactual cluster sizes. Although the coefficient estimates vary with the additive constant (as shown in Table A8), the counterfactual results are stable when the same transformation is used in estimation and simulation. Counterfactual results using $\ln(N + 10)$ are omitted for brevity and available upon request.

and measures access relative to that population. The break-even elasticities rise only modestly and remain between 0.0042 and 0.0117 across the alternative access measures.

STEM Researchers If economically relevant local spillovers are concentrated in technical fields, an exercise covering all fields may overstate the access gains. Panel D restricts the reallocation to STEM researchers—those in the life sciences, physical sciences, and engineering—and again uses the college-educated population as the reallocation target and access base. The results are nearly identical to Panel C, consistent with STEM fields accounting for most research volume in our data.

In conclusion, across all alternatives, the break-even access elasticities remain below the 0.08 benchmark, reinforcing the conclusion that a moderate reallocation toward large underserved MSAs can plausibly raise aggregate output.

7 Conclusion

This paper studies the geography of research in the U.S. and its implications for the broader economy. Research activity generates highly localized spillovers, providing a rationale for dispersing researchers toward population and economic activity so that a broader segment of the economy can access frontier knowledge. Yet dispersion may move researchers away from productive institutions and shrink the size of productive research clusters, reducing the aggregate stock of knowledge. We formally characterize and quantify this trade-off between the access benefit of dispersion and the production benefit of concentration.


Using detailed bibliographic data, we first show that research activity is highly concentrated across the U.S. and has become increasingly misaligned with population. As population shifted toward the South and West, research activity remained anchored in long-established institutions, making geographic access to research increasingly unequal across metropolitan areas.

We then estimate an individual-level research production function to quantify the potential production cost of dispersion. Location effects account for meaningful differences in research output across researchers, suggesting that researchers' productivity can change when they move between research environments. Our IV estimates also provide robust evidence that own-institution cluster size raises research output, while evidence of external-cluster effects is weaker. These estimates imply that researcher reallocation can affect knowledge production both directly through movers' location productivity and indirectly through changes

in the cluster sizes faced by researchers who remain.


Finally, we combine these estimates on the research-production side in a spatial framework to evaluate counterfactual reallocations and calculate the knowledge-access elasticity required for access gains to offset production losses. This provides a first quantitative assessment of whether alternative spatial allocations of researchers can generate net gains in aggregate output. The results suggest that marginally reallocating researchers toward large metropolitan areas with relatively limited research capacity would lead to modest production losses and require relatively small local access gains to raise aggregate output.

Figure 1: County-Level Access to Research Papers, 2014–2019




Notes: The map shows county-level access to nearby research. For each census tract, we calculate the number of research papers produced between 2014 and 2019 by institutions within 50 miles per 1,000 residents living within the same radius. County-level values are population-weighted averages of the tract-level measures.

Figure 2: Trends in Access to Research Papers: Selected MSA Case Studies




(a) Research Papers per 1,000 Residents




(b) Research Papers per 1,000 College Graduates

Notes: Panel (a) reports research papers per 1,000 residents, and Panel (b) reports research papers per 1,000 college graduates. For each MSA, publication counts include papers produced by authors affiliated with institutions in that MSA and are divided by the corresponding population measure. Publication counts are aggregated over 1965–1970, 1975–1980, 1985–1990, 1995–2000, 2005–2010, and 2014–2019. Population data come from the 1970, 1980, 1990, and 2000 Censuses and the 2005–2009 and 2015–2019 ACS.

Figure 3: Rising Inequality in Access to Research Papers: 1970–2019




(a) 90/50 and 95/50 Ratios




(b) Gini Coefficients

Notes: Panel (a) plots the 90/50 and 95/50 ratios of access to research papers across MSAs from 1970 to 2019, and Panel (b) plots the corresponding Gini coefficients. The black series use research papers per resident, while the blue series use research papers per college graduate.


Figure 4: Event Study: Changes in Research Output Around a Move



(a) Publications



(b) Citations



(c) Patent Citations

Notes: Each panel plots the estimated coefficients δ_s for $-15 \leq s \leq 15$ ($s \neq -1$) from Equation 8 using the sample of researchers who changed institutions once. The dependent variable is IHS publications (Panel (a)), IHS citations (Panel (b)), and IHS patent citations (Panel (c)). The x -axis denotes years relative to the move. The relative year -1 is omitted and normalized to zero, and the vertical line marks the year before the move. The shaded bands show 95% confidence intervals based on standard errors clustered at the institution level.

Table 1: Additive Decomposition of Research Outcomes Across Researchers:
By Researcher Productivity Group

Top vs. Bottom	Researcher Productivity Groups		
	50%	25%	10%
<i>Panel A: Publications</i>			
Differences in IHS Publications			
Overall	1.15	1.74	2.38
Individuals	0.77	1.15	1.50
Locations	0.09	0.13	0.18
Unexplained	0.29	0.46	0.69
Share of Difference Attributable to			
Individuals	0.67	0.66	0.63
Locations	0.08	0.08	0.08
Unexplained	0.26	0.27	0.29
<i>Panel B: Citations</i>			
Differences in IHS Citations			
Overall	2.68	3.97	5.17
Individuals	2.03	2.92	3.68
Locations	0.19	0.28	0.36
Unexplained	0.48	0.78	1.14
Share of Difference Attributable to			
Individuals	0.76	0.74	0.71
Locations	0.07	0.07	0.07
Unexplained	0.18	0.20	0.22
<i>Panel C: Patent Citations</i>			
Differences in IHS Patent Citations			
Overall	0.67	0.67	1.01
Individuals	0.51	0.51	0.75
Locations	0.07	0.07	0.05
Unexplained	0.11	0.11	0.21
Share of Difference Attributable to			
Individuals	0.75	0.75	0.75
Locations	0.10	0.10	0.05
Unexplained	0.16	0.16	0.21

Note: This table reports the additive decomposition of differences in IHS publications, citations, and patent citations across researcher productivity groups during 2010–2015. Researchers are ranked within their field by their average research output for the corresponding outcome. The 50%, 25%, and 10% columns compare researchers above and below the median, in the top and bottom quartiles, and in the top and bottom deciles, respectively. The researcher component includes author fixed effects and field-specific academic-age effects. The location component is predicted using lagged observable characteristics of the institution–field and external MSA–field research environments. The unexplained component includes differences in residuals. The field-year component is not reported because its contribution to the overall difference is negligible. Component shares are calculated as the difference in each reported component divided by the overall difference in the corresponding outcome; because the field-year component is omitted from the table, the reported shares may not sum exactly to one.

Table 2: Additive Decomposition of Research Outcomes Across Researchers:
By Institution and MSA Size Group

Top vs. Bottom	Institution Size Group			MSA Size Group		
	50%	25%	10%	50%	25%	10%
<i>Panel A: Publications</i>						
Differences in IHS Publications						
Overall	0.40	0.62	0.77	0.11	0.32	0.56
Individuals	0.10	0.23	0.33	0.01	0.10	0.27
Locations	0.26	0.32	0.34	0.09	0.18	0.21
Unexplained	0.03	0.07	0.09	0.01	0.03	0.07
Share of Difference Attributable to						
Individuals	0.25	0.36	0.42	0.09	0.32	0.48
Locations	0.66	0.51	0.44	0.80	0.57	0.38
Unexplained	0.07	0.10	0.12	0.06	0.08	0.13
<i>Panel B: Citations</i>						
Differences in IHS Citations						
Overall	0.95	1.45	1.82	0.31	0.80	1.30
Individuals	0.42	0.77	1.07	0.15	0.43	0.84
Locations	0.53	0.65	0.71	0.20	0.39	0.44
Unexplained	-0.01	0.05	0.07	-0.03	-0.02	0.06
Share of Difference Attributable to						
Individuals	0.45	0.53	0.58	0.47	0.55	0.64
Locations	0.56	0.45	0.39	0.64	0.49	0.34
Unexplained	-0.01	0.03	0.04	-0.10	-0.02	0.04
<i>Panel C: Patent Citations</i>						
Differences in IHS Patent Citations						
Overall	0.10	0.14	0.17	0.07	0.11	0.13
Individuals	0.02	0.04	0.06	0.02	0.03	0.05
Locations	0.09	0.12	0.14	0.06	0.10	0.11
Unexplained	-0.01	-0.02	-0.02	-0.01	-0.02	-0.02
Share of Difference Attributable to						
Individuals	0.23	0.31	0.35	0.28	0.32	0.40
Locations	0.94	0.86	0.82	0.92	0.90	0.83
Unexplained	-0.13	-0.11	-0.10	-0.13	-0.14	-0.13

Note: This table reports the additive decomposition of differences in IHS publications, citations, and patent citations across researchers by institution- or MSA-size group during 2010–2015. Institutions and MSAs are ranked separately within each field by their number of active researchers. The 50%, 25%, and 10% columns compare researchers at institutions or MSAs above and below the median, in the top and bottom quartiles, and in the top and bottom deciles, respectively. The researcher component includes author fixed effects and field-specific academic-age effects. The location component is predicted using lagged observable characteristics of the institution-field and external MSA-field research environments. The field-year component is not reported because its contribution to the overall difference is negligible. Component shares are calculated as the difference in each reported component divided by the overall difference in the corresponding outcome; because the field-year component is omitted from the table, the reported shares may not sum exactly to one.

Table 3: First-Stage Estimates for the Bartik Instruments

	Ln(Institution Size)			Ln(External Size)		
	(1)	(2)	(3)	(4)	(5)	(6)
$\ln \hat{N}^{I,Bartik}$	0.851*** (0.162)	0.816*** (0.161)	0.675*** (0.145)	0.604*** (0.208)	0.590*** (0.207)	0.623*** (0.174)
$\ln \hat{N}^{E,Bartik}$	-0.026 (0.024)	-0.034 (0.023)	-0.000 (0.019)	0.387*** (0.028)	0.385*** (0.028)	0.190*** (0.036)
Author Control		0.175*** (0.019)	0.140*** (0.018)		0.003 (0.038)	-0.023 (0.038)
Coauthor Control		0.004 (0.016)	0.005 (0.012)		0.036** (0.017)	0.030* (0.016)
Observations	798,548	798,548	797,434	798,548	798,548	797,434
R-squared	0.964	0.964	0.967	0.937	0.937	0.940
Baseline FEs	Yes	Yes	Yes	Yes	Yes	Yes
Author/Coauthor Controls	No	Yes	Yes	No	Yes	Yes
MSA $\times$ Year FE	No	No	Yes	No	No	Yes

Note: This table reports first-stage estimates for the two endogenous variables in Equation 10. The dependent variable is log own-institution size in Columns 1–3 and log external-cluster size in Columns 4–6. The excluded instruments are the Bartik-predicted own-institution size and Bartik-predicted external-cluster size, constructed from base-period subfield shares and worldwide subfield growth as described in Equation 11. The sample consists of author-year observations in three periods: 1995, 2005, and 2015, where each period pools the focal year and the two preceding years. Columns 1 and 4 include the baseline fixed effects from Equation 10, including author fixed effects, academic-age-by-field fixed effects, field-by-year fixed effects, and affiliation-by-field fixed effects. Columns 2 and 5 add the author and coauthor direct-exposure controls. Columns 3 and 6 further add MSA-by-year fixed effects. Standard errors are in parentheses and clustered at the institution level: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 4: OLS and IV Estimates of Agglomeration Elasticities

	OLS			IV		
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: Publications</i>						
Ln(Institution Size)	0.174*** (0.007)	0.173*** (0.007)	0.175*** (0.007)	0.234*** (0.072)	0.199*** (0.070)	0.300*** (0.103)
Ln(External Size)	0.003 (0.002)	0.003 (0.002)	0.002 (0.002)	0.082*** (0.020)	0.074*** (0.020)	0.037 (0.052)
Author Control		0.001 (0.007)	-0.004 (0.006)		-0.005 (0.015)	-0.021 (0.017)
Coauthor Control		0.096*** (0.010)	0.101*** (0.010)		0.091*** (0.011)	0.097*** (0.011)
<i>Panel B: Citations</i>						
Ln(Institution Size)	0.375*** (0.016)	0.379*** (0.017)	0.384*** (0.016)	0.898*** (0.192)	0.828*** (0.190)	1.028*** (0.271)
Ln(External Size)	0.002 (0.004)	0.000 (0.004)	-0.001 (0.004)	0.054 (0.062)	0.056 (0.060)	-0.022 (0.134)
Author Control		-0.157*** (0.015)	-0.158*** (0.016)		-0.238*** (0.038)	-0.248*** (0.044)
Coauthor Control		0.281*** (0.027)	0.296*** (0.024)		0.269*** (0.031)	0.288*** (0.027)
<i>Panel C: Patent Citations</i>						
Ln(Institution Size)	0.022*** (0.004)	0.022*** (0.004)	0.024*** (0.004)	0.153*** (0.042)	0.148*** (0.043)	0.192*** (0.062)
Ln(External Size)	-0.002** (0.001)	-0.002** (0.001)	-0.002** (0.001)	-0.013 (0.012)	-0.011 (0.012)	-0.030 (0.029)
Author Control		-0.011*** (0.003)	-0.007* (0.004)		-0.033*** (0.008)	-0.031*** (0.010)
Coauthor Control		0.025*** (0.005)	0.026*** (0.005)		0.023*** (0.006)	0.025*** (0.006)
Observations	798,548	798,548	797,434	798,548	798,548	797,434
Baseline FEs	Yes	Yes	Yes	Yes	Yes	Yes
Author/Coauthor Controls	No	Yes	Yes	No	Yes	Yes
MSA $\times$ Year FE	No	No	Yes	No	No	Yes
Kleibergen–Paap Wald F	N/A	N/A	N/A	15.88	14.90	12.97

Note: This table reports OLS and IV estimates of Equation 10 using the IV sample. The dependent variable is IHS publications in Panel A, IHS citations in Panel B, and IHS patent citations in Panel C. The sample consists of author-year observations in three periods: 1995, 2005, and 2015, where each period pools the focal year and the two preceding years. Columns 1–3 report OLS estimates. Columns 4–6 report IV estimates that instrument log own-institution size and log external-cluster size using the Bartik instruments in Equation 11. Columns 1 and 4 include the baseline fixed effects from Equation 10, including author fixed effects, academic-age-by-field fixed effects, field-by-year fixed effects, and affiliation-by-field fixed effects. Columns 2 and 5 add the author and coauthor direct-exposure controls. Columns 3 and 6 further add MSA-by-year fixed effects. The Kleibergen–Paap rk Wald F statistic is reported for the IV specifications. Standard errors are in parentheses and clustered at the institution level: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 5: IV Estimates of Agglomeration Elasticities by Field Group

	Social Sciences (1)	Humanities (2)	Natural Sciences (3)	Engineering (4)	Life Sciences (5)
<i>Panel A: Publications</i>					
Ln(Institution Size)	0.118 (0.121)	0.123 (0.198)	0.216 (0.143)	0.210*** (0.066)	0.292 (0.186)
Ln(External Size)	0.082*** (0.022)	-0.018 (0.033)	0.074 (0.058)	0.019 (0.046)	0.118*** (0.040)
<i>Panel B: Citations</i>					
Ln(Institution Size)	0.920** (0.377)	-0.637 (0.589)	0.685* (0.352)	0.654*** (0.192)	1.516*** (0.515)
Ln(External Size)	0.019 (0.074)	-0.142 (0.125)	0.119 (0.135)	-0.160 (0.126)	0.161 (0.117)
<i>Panel C: Patent Citations</i>					
Ln(Institution Size)	0.012 (0.025)	-0.013 (0.035)	0.138* (0.071)	0.003 (0.028)	0.499*** (0.169)
Ln(External Size)	-0.004 (0.004)	-0.007 (0.008)	-0.008 (0.028)	-0.024 (0.026)	-0.003 (0.038)
Observations	118,826	21,133	173,316	152,814	332,459
Kleibergen-Paap Wald F	8.48	2.37	4.01	20.45	4.66

Note: This table reports IV estimates of Equation 10 using the IV sample separately for each broad field group. The dependent variable is IHS publications in Panel A, IHS citations in Panel B, and IHS patent citations in Panel C. The sample consists of author-year observations in three periods: 1995, 2005, and 2015, where each period pools the focal year and the two preceding years. All columns include the baseline fixed effects from Equation 10 and the author and coauthor direct-exposure controls. Standard errors are in parentheses and clustered at the institution level: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 6: Counterfactual Reallocations of Researchers: Bilateral Cases

Knowledge Production Changes (%)				Access Calibration	Break-Even θ^*	Implied Access Elast. $s^N \theta^* \bar{\lambda}$
Origin Stayers	Dest. Incumb.	Movers	Net Change			
<i>Panel A: Boston → Dallas–Fort Worth</i>						
−1.5%	+8.65%	−4.91%	−0.59%	$\bar{\lambda} = 0.2$	0.482	0.0193
				$\bar{\lambda} = 0.5$	0.186	0.0186
				$\bar{\lambda} = 0.7$	0.132	0.0185
<i>Panel B: San Francisco → Las Vegas</i>						
−0.78%	+15.54%	−5.33%	−0.56%	$\bar{\lambda} = 0.2$	0.462	0.0185
				$\bar{\lambda} = 0.5$	0.178	0.0178
				$\bar{\lambda} = 0.7$	0.126	0.0176

Notes: Each panel reallocates researchers within one origin–destination MSA pair, with $\kappa = 0.1$. Researchers are assigned following the procedure in Section 6.2, and results are averaged across 50 simulations. The first three columns report percentage changes in output among researchers remaining in the origin, incumbent researchers in the destination, and movers, each relative to that group’s baseline output. “Net Change” reports the percentage change in aggregate flow research output. The change in knowledge production depends only on research output, so the first four columns are invariant to the access calibration. For each value of $\bar{\lambda}$, the break-even θ^* solves $\Delta \ln Y = 0$, and the final column reports the corresponding access elasticity $s^N \theta^* \bar{\lambda}$. Parameter calibrations are described in Section 6.1: $s^N = 0.2$, $\eta = 0.5$, $\delta = 0.15$, $\alpha^I = 0.199$, $\alpha^E = 0.074$.

Table 7: Counterfactual Reallocations of Researchers: Across Metropolitan Areas

Knowledge Production Changes (%)				Access Calibration	Break-Even θ^*	Implied Access Elast. $s^N \theta^* \bar{\lambda}$
Origin Stayers	Dest. Incumb.	Movers	Net Change			
<i>Panel A: Top 50 MSAs (10 Highest- to 10 Lowest-Researchers-per-Capita MSAs)</i>						
-0.99%	+8.78%	-1.42%	-0.32%	$\bar{\lambda} = 0.2$	0.245	0.0098
				$\bar{\lambda} = 0.5$	0.096	0.0096
				$\bar{\lambda} = 0.7$	0.069	0.0096
				$\bar{\lambda} = 0.5$, w/ migration	0.120	0.0120
<i>Panel B: All MSAs with Research Institutions</i>						
-1.94%	+6.13%	-14.54%	-1.14%	$\bar{\lambda} = 0.2$	1.075	0.0430
				$\bar{\lambda} = 0.5$	0.393	0.0393
				$\bar{\lambda} = 0.7$	0.277	0.0388

Notes: Panel A reallocates researchers from the ten MSAs with the highest researchers per capita to the ten with the lowest among the 50 most populous MSAs. Panel B extends the reallocation to all MSAs that host research institutions. In both panels, $\kappa = 0.1$; assignments follow Section 6.2 and results are averaged across 50 simulations. The first three columns report percentage changes in output among researchers remaining in the origin, incumbent researchers in the destination, and movers, each relative to that group's baseline output. "Net Change" reports the percentage change in aggregate flow research output. The change in knowledge production depends only on research output, so the first four columns are invariant to the access calibration. For each value of $\bar{\lambda}$, the break-even θ^* solves $\Delta \ln Y = 0$, and the final column reports the corresponding access elasticity $s^N \theta^* \bar{\lambda}$. The migration row in Panel A allows non-researchers to relocate as described in Appendix A4. Parameter calibrations for Panel A are the same as for Table 6. Panel B re-estimates and applies the agglomeration coefficients using $\ln(N + 5)$ to make the transformation less curved near zero, reducing the influence of changes involving only a few researchers: $\alpha^I = 0.537$, $\alpha^E = 0.043$ (Table A8 Column 2).

Table 8: Counterfactual Reallocations of Researchers: Alternative Specifications ($\bar{\lambda} = 0.5$)

Knowledge Production Changes (%)						
Origin Stayers	Dest. Incumb.	Movers	Net Change	Access Measure	Break-even θ^*	Implied Access Elast. $s^N \theta^* \bar{\lambda}$
<i>Panel A: Alternative Measures of Knowledge Access</i>						
-0.99%	+8.78%	-1.42%	-0.32%	Researchers per capita	0.096	0.0096
				Publications per capita	0.068	0.0068
				Citations per capita	0.034	0.0034
<i>Panel B: Publication-Based Agglomeration</i>						
-0.55%	+6.24%	+1.07%	-0.10%	Researchers per capita	0.031	0.0031
				Publications per capita	0.024	0.0024
				Citations per capita	0.008	0.0008
<i>Panel C: Reallocation Toward the College-Educated Population</i>						
-0.87%	+6.33%	-3.97%	-0.29%	Researchers per capita	0.117	0.0117
				Publications per capita	0.086	0.0086
				Citations per capita	0.042	0.0042
<i>Panel D: Reallocation of STEM Researchers Toward the College-Educated Population</i>						
-0.85%	+6.40%	-4.04%	-0.30%	Researchers per capita	0.121	0.0121
				Publications per capita	0.090	0.0090
				Citations per capita	0.043	0.0043

Notes: All panels reallocate researchers from the ten MSAs with the highest researchers per capita to the ten with the lowest among the 50 most populous MSAs, with $\kappa = 0.1$ and $\bar{\lambda} = 0.5$. Panel A considers the same reallocation as Table 7 Panel A, but allows local access to depend on publications or citations per capita, in addition to researchers per capita. Panel B measures cluster size by paper counts (rather than author counts) and applies the corresponding agglomeration estimates reported in Table A6. Panel C defines researcher surpluses, deficits, and knowledge access relative to the college-educated population. Panel D repeats the exercise in Panel C but reallocates STEM researchers only. Within each panel, the rows measure knowledge access using researcher counts, publications, or citations per capita. The first three columns report percentage changes in output among researchers remaining in the origin, incumbent researchers in the destination, and movers, each relative to that group's baseline output. "Net Change" reports the percentage change in aggregate flow research output. The change in knowledge production depends only on research output, so the first four columns are invariant to the access calibration. The break-even θ^* solves $\Delta \ln Y = 0$, and the final column reports the corresponding access elasticity $s^N \theta^* \bar{\lambda}$. Parameter calibrations for Panels A, C, and D are the same as for Table 6. For Panel B: $\alpha^I = 0.101$, $\alpha^E = 0.049$.

References

- Acton, R., Cortes, K., and Morales, C. (2024). Distance to Opportunity: Higher Education Deserts and College Enrollment Choices. In *Financing Institutions of Higher Education, Chapter 2*.
- Adams, J. D., Black, G. C., Clemmons, J. R., and Stephan, P. E. (2005). Scientific teams and institutional collaborations: Evidence from US universities, 1981–1999. *Research Policy*, 34(3):259–285.
- Adams, J. D. and Clemmons, J. R. (2011). The Role of Search in University Productivity: Inside, Outside, and Interdisciplinary Dimensions. *Industrial and Corporate Change*, 20(1):215–251.
- Aghion, P., Boustan, L., Hoxby, C., and Vandenbussche, J. (2009). The Causal Impact of Education on Economic Growth: Evidence from the United States. *Brookings Papers on Economic Activity*, 2009(2):1–73.
- Agrawal, A. and Goldfarb, A. (2008). Restructuring Research: Communication Costs and the Democratization of University Innovation. *American Economic Review*, 98(4):1578–1590.
- Airoldi, A. and Moser, P. (2024). Inequality in Science: Who Becomes a Star? Technical report, National Bureau of Economic Research.
- Arzaghi, M. and Henderson, V. (2008). Networking off Madison Avenue. *Review of Economic Studies*, 75(4):1011–1038.
- Association of Public and Land-Grant Universities (2012). The Land-Grant Tradition. APLU report.
- Atkin, D., Chen, K., and Popov, A. (2022). The Returns to Face-to-Face Interactions: Knowledge Spillovers in Silicon Valley. *Working Paper*.
- Azoulay, P., Graff Zivin, J. S., and Wang, J. (2010). Superstar Extinction. *The Quarterly Journal of Economics*, 125(2):549–589.
- Baum-Snow, N., Gendron-Carrier, N., and Pavan, R. (2024). Local Productivity Spillovers. *American Economic Review*, 114(4):1030–69.
- Belenzon, S. and Schankerman, M. (2013a). Spatial Heterogeneity in the Diffusion of Technologies: Evidence from Patent Citations. *Journal of Industrial Economics*, 61(3):733–758.
- Belenzon, S. and Schankerman, M. (2013b). Spreading the word: Geography, policy, and knowledge spillovers. *The Review of Economics and Statistics*, 95(3):884–903.
- Biasi, B. and Ma, S. (2023). The Education-Innovation Gap. *Working Paper*.
- Black, S., Cortes, K., and Lincove, J. A. (2020). Apply Yourself: Racial and Ethnic Differences in College Application. *Education Finance and Policy*, 15(2):209–240.
- Blair, P. Q. and Smetters, K. (2021). Why Don’t Elite Colleges Expand Supply? Working Paper 29309, National Bureau of Economic Research.
- Bloom, N., Jones, C. I., Van Reenen, J., and Webb, M. (2020). Are Ideas Getting Harder to Find? *American Economic Review*, 110(4):1104–1144.
- Borusyak, K., Hull, P., and Jaravel, X. (2022). Quasi-Experimental Shift-Share Research Designs. *Review of Economic Studies*, 89(1):181–213.

- Bosquet, C. and Combes, P.-P. (2017). Sorting and Agglomeration Economies in French Economics Departments. *Journal of Urban Economics*, 101:27–44.
- Butler, D. M., Butler, R. J., and Rich, J. T. (2008). The Equalizing Effect of the Internet on Access to Research Expertise in Political Science and Economics. *PS: Political Science & Politics*, 41(3):579–584.
- Card, D. (1995). Using Geographic Variation in College Proximity to Estimate the Return to Schooling. In *Aspects of Labour Economics: Essays in Honour of John Vanderkamp*, edited by Louis Christofides, E. Kenneth Grant and Robert Swindinsky. University of Toronto Press.
- Carlino, G. and Kerr, W. R. (2015). Agglomeration and Innovation. *Handbook of Regional and Urban economics*, 5:349–404.
- Chandra, A. and Xu, C. (2025). Person and Place Effects in Scientific Discovery. *Working Paper*.
- Charlot, S. and Duranton, G. (2004). Communication Externalities in Cities. *Journal of Urban Economics*, 56(3):581–613.
- Chen, J. and Roth, J. (2024). Logs with Zeros? Some Problems and Solutions. *The Quarterly Journal of Economics*, 139(2):891–936.
- Coe, D. T. and Helpman, E. (1995). International R&D spillovers. *European Economic Review*, 39(5):859–887.
- Cohen, J. (1996). Computer Mediated Communication and Publication Productivity Among Faculty. *Internet Research*, 6(2/3):41–63.
- Davis, D. and Dingel, J. (2019). A Spatial Knowledge Economy. *American Economic Review*, 109(1):153–170.
- D’Costa, S. and Overman, H. (2014). The Urban Wage Growth Premium: Sorting or Learning? *Regional Science and Urban Economics*, 48:168–179.
- De La Roca, J. and Puga, D. (2017). Learning by Working in Big Cities. *Review of Economic Studies*, 84(1):106–142.
- Diamond, R. (2016). The Determinants and Welfare Implications of US Workers’ Diverging Location Choices by Skill: 1980-2000. *American Economic Review*, 106(3):479–524.
- Ding, W. W., Levin, S. G., Stephan, P. E., and Winkler, A. E. (2010). The Impact of Information Technology on Academic Scientists’ Productivity and Collaboration Patterns. *Management Science*, 56(9):1439–1461.
- Dingel, J., Gottlieb, J., Lozinski, M., and Mourot, P. (2023). Market Size and Trade in Medical Services. *Working Paper*.
- Duranton, G. and Overman, H. G. (2005). Testing for Localization using Micro-Geographic Data. *The Review of Economic Studies*, 72(4):1077–1106.
- Eckert, F., Hejlesen, M., and Walsh, C. (2022). The Return to Big City Experience: Evidence from Refugees in Denmark. *Journal of Urban Economics*, (103454).
- Ehrenberg, R. G., Rizzo, M. J., and Condie, S. S. (2003). Start-Up Costs in American Research Universities. Working Paper 33, Cornell Higher Education Research Institute.

- Ellison, G., Glaeser, E., and Kerr, W. (2010). What Causes Industry Agglomeration? Evidence from Coagglomeration Patterns. *American Economic Review*, 100(3):1195–1213.
- Emanuel, N., Harrington, E., and Pallais, A. (2023). The Power of Proximity. *Working Paper*.
- Fabre, A. (2023). The Geography of Higher Education and Spatial Inequalities. *Working Paper*.
- Fieldhouse, A. and Mertens, K. (2024). The Returns to Government R&D: Evidence from U.S. Appropriations Shocks. *Working Paper*.
- Finkelstein, A., Gentzkow, M., and Williams, H. (2016). Sources of Geographic Variation in Health Care: Evidence from Patient Migration. *Quarterly Journal of Economics*, 131(4):1681–1726.
- Fu, C., Guo, J., Smith, A., and Sorensen, A. (2022). Students’ Heterogeneous Preferences and the Uneven Spatial Distribution of Colleges. *Journal of Monetary Economics*, 129:49–64.
- Gaubert, C. (2018). Firm Sorting and Agglomeration. *American Economic Review*, 108(11):3117–3153.
- Glaeser, E. (1999). Learning in Cities. *Journal of Urban Economics*, 46(2):254–277.
- Glaeser, E. and Mare, D. (2001). Cities and Skills. *Journal of Labor Economics*, 19(2):316–342.
- Goldsmith-Pinkham, P., Sorkin, I., and Swift, H. (2020). Bartik Instruments: What, When, Why, and How. *American Economic Review*, 110(8):2586–2624.
- Goldstein, E. G. (2024). Communication Costs in Science: Evidence from the National Science Foundation Network. *Industrial and Corporate Change*, 33(4):785–807.
- Greenstone, M., Hornbeck, R., and Moretti, E. (2010). Identifying Agglomeration Spillovers: Evidence from Winners and Losers of Large Plant Openings. *Journal of Political Economy*, 118(3):536–598.
- Guzman, J. (2024). Go West Young Firm: The Impact of Startup Migration on the Performance of Migrants. *Management Science*, 70(7):4824–4846.
- Hall, B. H., Mairesse, J., and Mohnen, P. (2010). Measuring the returns to R&D. In Hall, B. H. and Rosenberg, N., editors, *Handbook of the Economics of Innovation*, volume 2, chapter 24, pages 1033–1082. North-Holland, Amsterdam.
- Hamermesh, D. S. and Oster, S. M. (2002). Tools or Toys? The Impact of High Technology on Scholarly Productivity. *Economic Inquiry*, 40(4):539–555.
- Hausman, N. (2022). University Innovation and Local Economic Growth. *Review of Economics and Statistics*, 104(4):718–735.
- Helmers, C. and Overman, H. G. (2017). My Precious! The Location and Diffusion of Scientific Research: Evidence from the Synchrotron Diamond Light Source. *The Economic Journal*, 127(604):2006–2040.
- Hesse, B. W., Sproull, L. S., Kiesler, S. B., and Walsh, J. P. (1993). Returns to Science: Computer Networks in Oceanography. *Communications of the ACM*, 36(8):90–101.
- Howard, G., Weinstein, R., and Yang, Y. (2024). Do Universities Improve Local Economic Resilience? *Review of Economics and Statistics*, 106(4):1129–1145.
- Hoxby, C. M. (2012). Endowment Management Based on a Positive Model of the University. Working Paper 18626, National Bureau of Economic Research.

- Hsieh, C.-T. and Moretti, E. (2019). Housing Constraints and Spatial Misallocation. *American Economic Journal: Macroeconomics*, 11(2):1–39.
- Ishimaru, S. (2025). Geographic Mobility of Youth and Spatial Gaps in Local College and Labor Market Opportunities. *Journal of Labor Economics*, 43(4):1251–1294.
- Jaffe, A. (1989). Real Effects of Academic Research. *American Economic Review*, 79(5):957–970.
- Jaffe, A., Trajtenberg, M., and Henderson, R. (1993). Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations. *Quarterly Journal of Economics*, 108(3):577–598.
- Jones, B. F., Wuchty, S., and Uzzi, B. (2008). Multi-University Research Teams: Shifting Impact, Geography, and Stratification in Science. *Science*, 322(5905):1259–1262.
- Jones, C. I. (1995). R&D-Based Models of Economic Growth. *Journal of Political Economy*, 103(4):759–784.
- Kaminer, N. and Braunstein, Y. M. (1998). Bibliometric Analysis of the Impact of Internet Use on Scholarly Productivity. *Journal of the American Society for Information Science*, 49(8):720–730.
- Kantor, S. and Whalley, A. (2014). Knowledge Spillovers from Research Universities: Evidence from Endowment Value Shocks. *Review of Economics and Statistics*, 96(1):171–188.
- Kantor, S. and Whalley, A. (2019). Research Proximity and Productivity: Long-Term Evidence from Agriculture. *Journal of Political Economy*, 127(2):819–854.
- Kim, E. H., Morse, A., and Zingales, L. (2009). Are Elite Universities Losing Their Competitive Edge? *Journal of Financial Economics*, 93(3):353–381.
- Kim, S. and Moser, P. (2025). Women in Science. Lessons from the Baby Boom. *Econometrica*, 93(5):1521–1560.
- Koffi, M. and Marx, M. (2026). Cassatts in the Attic: Is There a Gender Gap in the Commercialization of Science. *American Economic Journal: Applied Economics*, 18(2):266–298.
- Krugman, P. R. (1980). Scale Economies, Product Differentiation, and the Pattern of Trade. *The American Economic Review*, 70(5):950–959.
- Legislative Analyst’s Office (2024). UC Merced at 20: Campus Developments and Key State-Level Takeaways. Technical report, California Legislative Analyst’s Office, Sacramento.
- Lerner, J., Manley, H., Stein, C., and Williams, H. (2024). The Wandering Scholars: Understanding the Heterogeneity of University Commercialization. Technical report, National Bureau of Economic Research.
- Manson, S., Schroeder, J., Van Riper, D., Kugler, T., and Ruggles, S. (2022). IPUMS National Historical Geographic Information System: Version 17.0. <https://www.nhgis.org>.
- Martín-Martín, A., Thelwall, M., Orduna-Malea, E., and Delgado López-Cózar, E. (2021). Google Scholar, Microsoft Academic, Scopus, Dimensions, Web of Science, and OpenCitations’ COCI: A Multidisciplinary Comparison of Coverage via Citations. *Scientometrics*, 126(1):871–906.
- Marx, M. and Fuegi, A. (2020). Reliance on Science: Worldwide Front-Page Patent Citations to Scientific Articles. *Strategic Management Journal*, 41(9):1572–1594.

- Merritt, K. and Lawrence, J. F., editors (2007). *From Rangeland to Research University: The Birth of University of California, Merced*. Number 139 in New Directions for Higher Education. Jossey-Bass, San Francisco.
- Moretti, E. (2004). Estimating the Social Return to Higher Education: Evidence from Longitudinal and Repeated Cross-Sectional Data. *Journal of Econometrics*, 121(1-2):175–212.
- Moretti, E. (2021). The Effect of High-Tech Clusters on the Productivity of Top Inventors. *American Economic Review*, 111(10):3328–3375.
- Mountjoy, J. (2022). Community Colleges and Upward Mobility. *American Economic Review*, 112(8):2580–2630.
- National Archives (1862). Morrill act (1862). U.S. National Archives, Milestone Documents.
- National Education Association (2022). Land Grant Institutions: An Overview. NEA brief.
- Pelfrey, P. A. (2004). *A Brief History of the University of California*. University of California Press, Berkeley, 2nd edition.
- Pollacci, L. (2022). EMAKG: An Enriched Version of the Microsoft Academic Knowledge Graph. <https://github.com/LauraPollacci/EMAKG>.
- Romer, P. M. (1990). Endogenous Technological Change. *Journal of Political Economy*, 98(5, Part 2):S71–S102.
- Rosenblat, T. S. and Mobius, M. M. (2004). Getting Closer or Drifting Apart? *The Quarterly Journal of Economics*, 119(3):971–1009.
- Rosenthal, S. and Strange, W. (2003). Geography, Industrial Organization, and Agglomeration. *Review of Economics and Statistics*, 85(2):377–393.
- Rosenthal, S. and Strange, W. (2008). The Attenuation of Human Capital Spillovers. *Journal of Urban Economics*, 64:373–389.
- Rossi-Hansberg, E., Sarte, P.-D., and Schwartzman, F. (2026). Cognitive Hubs and Spatial Redistribution. *American Economic Journal: Macroeconomics*, 18(2):72–111.
- Rouse, C. (1995). Democratization or Diversion? The Effect of Community Colleges on Educational Attainment. *Journal of Business and Economic Statistics*, 13(2):217–224.
- Sinha, A., Shen, Z., Song, Y., Ma, H., Eide, D., Hsu, B.-J., and Wang, K. (2015). An Overview of Microsoft Academic Service (MAS) and Applications. *WWW'15 Companion: Proceedings of the 24th International Conference on World Wide Web*, pages 243–246.
- Thelwall, M. (2018). Does Microsoft Academic Find Early Citations? *Scientometrics*, 114(1):325–334.
- Truffa, F. and Wong, A. (2025). Undergraduate Gender Diversity and the Direction of Scientific Research. *American Economic Review*, 115(7):2414–2448.
- U.S. Census Bureau (2023). County Business Patterns, 2021. <https://www.census.gov/programs-surveys/cbp.html>. U.S. Department of Commerce.
- Valero, A. and Van Reenen, J. (2019). The Economic Impact of Universities: Evidence from Across the Globe. *Economics of Education Review*, 68:53–67.

- Van Alstyne, M. and Brynjolfsson, E. (1996). Could the Internet Balkanize Science? *Science*, 274(5292):1479–1480.
- Visser, M., Van Eck, N. J., and Waltman, L. (2021). Large-Scale Comparison of Bibliographic Data Sources: Scopus, Web of Science, Dimensions, Crossref, and Microsoft Academic. *Quantitative Science Studies*, 2(1):20–41.
- Waldinger, F. (2012). Peer Effects in Science: Evidence from the Dismissal of Scientists in Nazi Germany. *Review of Economic Studies*, 79(2):838–861.
- Walsh, J. P. and Bayma, T. (1996). Computer Networks and Scientific Work. *Social Studies of Science*, 26(3):661–703.
- Walsh, J. P., Kucker, S., Maloney, N. G., and Gabbay, S. (2000). Connecting Minds: Computer-Mediated Communication and Scientific Work. *Journal of the American Society for Information Science*, 51(14):1295–1305.
- Wang, K., Shen, Z., Huang, C., Wu, C.-H., Dong, Y., and Kanakia, A. (2020). Microsoft Academic Graph: When Experts are Not Enough. *Quantitative Science Studies*, 1(1):396–413.
- Winkler, A. E., Levin, S. G., and Stephan, P. E. (2010). The Diffusion of IT in Higher Education: Publishing Productivity of Academic Life Scientists. *Economics of Innovation and New Technology*, 19(5):481–503.
- Wuchty, S., Jones, B. F., and Uzzi, B. (2007). The Increasing Dominance of Teams in Production of Knowledge. *Science*, 316(5827):1036–1039.

Appendix

A1 Model Extension: Endogenous Location Choices

A1.1 Researchers' Location Choice

For simplicity, researchers have homogeneous location preferences and value locations according to the following reduced-form indirect utility function:³²

$$V_j^N = a_j(N_j) + b_j(N_j, M_j) - \beta_r^N r_j(M_j),$$

where

$$a_j(N_j) \equiv \ln A_j(N_j) = a_{0j} + \alpha \ln N_j, \quad a_{0j} \equiv \ln A_{0j}.$$

The term $b_j(N_j, M_j)$ captures local amenities for researchers, which include any non-productivity-related factors that affect researchers' valuation of a location; $r_j(M_j)$ is the local log rent, which reflects local cost of living. Research productivity may affect utility through several channels: More productive locations may offer higher salaries, greater funding opportunities, or stronger intrinsic research value and thus attract researchers through higher $a_j(N_j)$. The amenity term $b_j(N_j, M_j)$ captures both attractive and unattractive features of institutions and surrounding areas, such as institutional support and the availability of research positions. For simplicity, we assume that the local log rent $r_j(M_j)$ depends only on the non-research population, reflecting the assumption that researchers are a small share of the total population and therefore do not materially affect local living costs.

To capture spatial variation in the availability of research positions, we specify

$$b_j(N_j, M_j) = \beta_{0j}^N + \beta_b^N \ln \left(\frac{M_j}{N_j} \right),$$

where β_{0j}^N captures exogenous variation in the availability of research positions and the desirability of research environments, driven by factors such as the historical presence of research universities, support for local institutions, and traditions of scholarly activity. The parameter β_b^N governs how strongly the avail-

³²In the stylized model, researchers have homogeneous location preferences, while non-researchers draw idiosyncratic location tastes that act as migration frictions. The asymmetry is deliberate: researcher location is the policy instrument we manipulate, not a response margin, whereas non-researchers' migration is the endogenous response at the center of the analysis. Productivity heterogeneity among researchers is consistent with this simplification: since individual productivity D_i is separable from location productivity, it enters indirect utility as a location-invariant constant, so every researcher evaluates locations identically. We can therefore treat quality-adjusted researchers as solving a common location-choice problem.

ability of research positions responds to local population. It is a reduced-form parameter that captures, for example, the responsiveness of local faculty positions to the population served by local universities. A larger β_b^N implies that institutional demand for researchers is more responsive to the size of the local population. To ensure the existence of an interior equilibrium, we assume $\beta_b^N > \alpha$, so that researchers face net decreasing returns to scale in each location. Because researchers are assumed to move frictionlessly, their spatial distribution is determined by equalized utility across locations.

A1.2 Non-Researchers' Location Choice

Non-researchers choose locations based on local wages, non-researcher amenities, and living costs. Based on the production function in each location, the local wage equals the marginal product of labor, which depends partially on local access to frontier knowledge. This is the key channel through which the spatial distribution of researchers affects non-researchers' location choices.

Non-researchers are imperfectly mobile across locations. Their indirect utility is

$$V_{ij}^M = (1 - s^N) s_{0j} + s^N \ln \left[S_0 + \left(\frac{N_j}{M_j} \right)^\theta \right] + b_j^M(M_j) - \beta_r^M r_j(M_j) + \varepsilon_{ij},$$

where the first two terms capture the location-varying component of log wages. Here, $s_{0j} \equiv \ln S_{0j}$ is the component of local productivity that is unrelated to exposure to frontier knowledge, while $S_0 + (N_j/M_j)^\theta$ measures access to frontier knowledge. The elasticity of local productivity with respect to N_j/M_j is $s^N \theta \lambda_j$. The term $b_j^M(M_j)$ is location j 's amenity value for non-researchers, which may depend on the size of the local population (largely represented by the non-research population). As with researchers, non-researchers dislike higher log rent, with sensitivity governed by β_r^M . The term ε_{ij} is an individual-specific idiosyncratic preference for location j , which equals σ times a random variable drawn from the Type-I Extreme Value distribution. The parameter σ governs the strength of idiosyncratic attachment to locations and therefore the degree of migration frictions.

The population of non-researchers in location j is therefore

$$M_j = M \frac{\exp(\bar{V}_j/\sigma)}{\sum_{j'} \exp(\bar{V}_{j'}/\sigma)},$$

where M is the national non-researcher population, and $\bar{V}_j$ is the deterministic component of indirect utility.

The spatial equilibrium consists of the distributions $\{N_j, M_j\}$ and rents r_j that clear local housing markets. With J locations, the equilibrium and its comparative statics generally do not admit closed-form

solutions. To illustrate how shocks to the spatial distribution of researchers affect knowledge access and knowledge production, Section A1.3 studies a tractable two-location case.

A1.3 Spatial Effects on Access to Frontier Knowledge: A Two-Location Analysis

To illustrate the model's implications, we consider a setting with two locations, 1 and 2. Location 1 resembles a city such as Boston: a long-established center of scholarship with a dense cluster of universities built up over centuries. Location 2 resembles a city such as Dallas: a fast-growing regional economic hub that has experienced substantial productivity and population growth in recent decades but lacks the same historical depth in higher education and research. Accordingly, we assume that location 1 has a higher location-specific amenity value for researchers, $\beta_{01}^N > \beta_{02}^N$, and that this advantage is sufficiently large to imply $N_1 > N_2$.

For tractability, we assume that local log rents take the form $r(M_j) = r_0 + e \ln(M_j)$, where e is the inverse housing supply elasticity and is common across locations. To further simplify the equilibrium characterization, we set researchers' housing cost sensitivity to zero, $\beta_r^N = 0$.³³ We also shut down endogenous amenities for non-researchers by setting $b_j^M(M_j) = 0$.³⁴

Finally, because the local access shares λ_j generally differ across locations, we evaluate these comparative statics at a common value $\bar{\lambda}$ when solving the system. This is exact at a symmetric equilibrium and a first-order approximation otherwise.

A1.3.1 “Home-Market” Effect

A location with higher non-research productivity attracts more non-researchers. Through the endogenous provision of research positions, captured by β_b^N , a larger non-research population in turn attracts more researchers. In addition, because researcher productivity rises with local cluster size, a larger non-research population can generate a more than proportional increase in the number of researchers. As a result, the more productive and more populous location may also enjoy greater researcher access per capita. This result echoes the home-market effect in trade models with agglomeration economies (Krugman, 1980) and is also related to the spatial concentration of medical services studied by Dingel et al. (2023).³⁵

³³In this case, researchers' housing costs can be interpreted as being absorbed into the location-specific amenity value β_{0j}^N .

³⁴We reintroduce these mechanisms in the quantitative evaluation of the trade-off later in the paper.

³⁵Places with larger non-research populations attract disproportionately more researchers, similar to the classic home-market result in which producers concentrate in locations with larger consumer markets. The mechanism here differs, however, because

We summarize the effect of local non-research productivity, s_{01} , on the spatial distribution of researcher access in the following proposition.

Proposition 1 *Under Assumption 1, an increase in non-research productivity in location 1, s_{01} , raises M_1 and increases N_1/M_1 relative to N_2/M_2 .*

The proof of Proposition 1 is provided in Appendix A1.4.1.

An increase in s_{01} attracts non-researchers to location 1 through higher wages. The larger local population expands research positions, while agglomeration economies further attract researchers. Researcher employment therefore rises more than proportionally, increasing researchers per capita in location 1 relative to location 2. The model predicts that large population and industry centers have both more researchers and more researchers per capita.

A1.3.2 Spatial Inertia of the Historical Legacy of Universities and Research Institutions

In the postwar period, especially after the 1970s, U.S. economic activity shifted substantially away from legacy cities in the Northeast and Midwest toward the Sun Belt. Based on Proposition 1, this shift in productivity and population toward booming cities in the South and West should, all else equal, also attract researchers and gradually narrow differences in researcher access per capita across locations.³⁶ Yet universities and research institutions remained remarkably spatially inert and did not follow population and economic activity into these expanding locations.

In the model, this persistence is represented by a higher researcher amenity β_{01}^N relative to β_{02}^N , reflecting the anchoring of research positions in legacy cities beyond what population or productivity alone would predict. Several forces may contribute to this anchoring. Alumni networks and institutional ties deepen over time and create persistent advantages for existing universities: donor relationships and academic networks are difficult to replicate in new locations (Hoxby, 2012; Blair and Smetters, 2021). In addition, the fixed costs of building a modern research university from scratch are substantial compared with the land-grant era, when federal land transfers and infrastructure requirements were smaller in scale (Ehrenberg et al., 2003).

We summarize the effect of a higher value of β_{01}^N on spatial inequality in research access as follows:

researchers generate local productivity spillovers rather than traded goods.

³⁶In the model, a large increase in non-research productivity in location 2, such as Dallas, raises its non-research population, which then attracts more researchers. As researchers move in, agglomeration effects further increase local researcher productivity, reinforcing this process. All else equal, per-capita access to researchers in location 2 should therefore increase through the home-market effect.

Proposition 2 *Under Assumption 1, an increase in the researcher amenity in location 1, β_{01}^N , lowers N_2/M_2 relative to N_1/M_1 .*

The proof of Proposition 2 is provided in Appendix A1.4.2.

A1.3.3 Moving Research Access Toward Equality

We next consider a policy that reallocates research resources toward an allocation proportional to population. In the model, this operates through an increase in β_b^N , which governs how strongly the provision of local research positions responds to local population. The following proposition summarizes the effect of an increase in β_b^N .

Proposition 3 *Under Assumption 1, if researcher access per capita is initially higher in location 1 than in location 2, an increase in the linkage between the provision of local researcher positions and local population, β_b^N , narrows the researcher access gap by raising N_2/M_2 relative to N_1/M_1 .*

The proof of Proposition 3 is provided in Appendix A1.4.3.

The persistent attraction of legacy cities for researchers can lead to unequal access to frontier knowledge, while policies that direct research resources toward population centers can help reduce this inequality. However, such policies may shrink existing research clusters and weaken the research productivity gains generated by their scale, particularly when those clusters are large and highly productive. In addition, as knowledge access improves in location 2, non-researchers may adjust their own location choices endogenously. We now characterize the costs and benefits of increasing β_b^N , taking into account all of these adjustments.

Costs and Benefits of Reallocating Researchers in Proportion to Population Taking the derivative of national output with respect to β_b^N , we show that the effect on national output is proportional to the following expression:³⁷

$$\begin{aligned} \frac{\partial y}{\partial \beta_b^N} \propto & P_1^N (1 - P_1^N) N \left[\underbrace{s^N \left(\frac{Y_2}{Y} \frac{\theta \lambda_2}{N_2} - \frac{Y_1}{Y} \frac{\theta \lambda_1}{N_1} \right)}_{\text{Gain in Knowledge Access}} - \underbrace{\frac{s^N (1 + \alpha)}{(1 - \eta) \delta Q^{1-\eta}} (A_{01} N_1^\alpha - A_{02} N_2^\alpha)}_{\text{Loss in Knowledge Production}} \right] \\ & - \underbrace{\frac{s^N \theta \bar{\lambda} (1 - s^N \theta \bar{\lambda})}{B (\sigma + s^N \theta \bar{\lambda} + \beta_r^M e)} P_1^M (1 - P_1^M) M \left(S_{01}^{1-s^N} \left(S_0 + \left(\frac{N_1}{M_1} \right)^\theta \right)^{s^N} - S_{02}^{1-s^N} \left(S_0 + \left(\frac{N_2}{M_2} \right)^\theta \right)^{s^N} \right)}_{\text{Loss in Productivity due to Non-Researcher Migration}}, \end{aligned}$$

³⁷See Appendix A1.4.4 for the full derivation.

where N and M are national populations of researchers and non-researchers, respectively. $P_1^N = N_1/N$ and $P_1^M = M_1/M$. The first two terms capture the effects of researcher migration. As researchers move from location 1 to location 2, access to frontier knowledge for non-researchers rises in location 2 and falls in location 1. Because knowledge access is concave in the number of researchers, holding frontier knowledge production fixed, equalizing access across locations raises aggregate output. The second term captures the offsetting effect on knowledge production. As researchers move out of location 1, they forgo the high research productivity of that location, and the contraction of the larger cluster further lowers productivity by weakening agglomeration effects. Because knowledge production is convex in the number of researchers, the productivity loss from shrinking location 1 exceeds the productivity gain from expanding location 2, so total knowledge production declines.

The third term captures the endogenous migration response of non-researchers. As researcher reallocation raises productivity in location 2 and lowers it in location 1, some non-researchers move toward location 2. Because location 2 remains the lower-productivity location overall, this migration can partially offset the gains from improved knowledge access.

A1.4 Proofs for the Comparative Statics in the Two-Location Case

This section derives the comparative statics for the two-location case and proves Propositions 1–3.

A1.4.1 Effects of s_{01} on the Spatial Distributions of Researchers and Non-Researchers: The “Home-Market” Effect

We derive the effect of s_{01} on researcher and non-researcher populations. Because researchers are homogeneous and move frictionlessly, their equilibrium utility is equalized across the two locations. Researcher utility depends on research productivity and amenities rather than on knowledge access. The utility-equalization condition is:

$$0 = (a_{01} - a_{02}) + (\beta_{01}^N - \beta_{02}^N) \\ + (\alpha - \beta_b^N) [\ln N_1 - \ln(N - N_1)] + \beta_b^N [\ln M_1 - \ln(M - M_1)].$$

Totally differentiating this equation yields

$$(\alpha - \beta_b^N) \left(\frac{1}{N_1} + \frac{1}{N_2} \right) \frac{\partial N_1}{\partial s_{01}} + \beta_b^N \left(\frac{1}{M_1} + \frac{1}{M_2} \right) \frac{\partial M_1}{\partial s_{01}} = 0.$$

We then totally differentiate non-researcher population at location 1. Because a non-researcher's real income depends on local access S_j , the access elasticities $\theta\lambda_j$ enter this row; evaluating them at $\bar{\lambda}$ gives

$$(\sigma + s^N \theta \bar{\lambda} + \beta_r^M e) \left(\frac{1}{M_1} + \frac{1}{M_2} \right) \frac{\partial M_1}{\partial s_{01}} - s^N \theta \bar{\lambda} \left(\frac{1}{N_1} + \frac{1}{N_2} \right) \frac{\partial N_1}{\partial s_{01}} = 1 - s^N.$$

The right-hand side remains $1 - s^N$ because the shock s_{01} operates through the location-specific non-research factor S_{01} and not through the access function, so it carries no $\bar{\lambda}$.

Solving this two-equation system yields:

$$\left(\frac{1}{N_1} + \frac{1}{N_2} \right) \frac{\partial N_1}{\partial s_{01}} = \frac{\beta_b^N (1 - s^N)}{(\beta_b^N - \alpha)(\sigma + \beta_r^M e) - \alpha s^N \theta \bar{\lambda}}$$

and

$$\left(\frac{1}{M_1} + \frac{1}{M_2} \right) \frac{\partial M_1}{\partial s_{01}} = \frac{(1 - s^N)(\beta_b^N - \alpha)}{(\beta_b^N - \alpha)(\sigma + \beta_r^M e) - \alpha s^N \theta \bar{\lambda}}.$$

Since $\frac{\partial \ln\left(\frac{N_1}{N_2}\right)}{\partial s_{01}} = \left(\frac{1}{N_1} + \frac{1}{N_2} \right) \frac{\partial N_1}{\partial s_{01}}$ and $\frac{\partial \ln\left(\frac{M_1}{M_2}\right)}{\partial s_{01}} = \left(\frac{1}{M_1} + \frac{1}{M_2} \right) \frac{\partial M_1}{\partial s_{01}}$, we can rewrite the comparative statics as:

$$\frac{\partial \ln\left(\frac{N_1}{N_2}\right)}{\partial s_{01}} = \frac{\beta_b^N (1 - s^N)}{(\beta_b^N - \alpha)(\sigma + \beta_r^M e) - \alpha s^N \theta \bar{\lambda}} \quad (15)$$

and

$$\frac{\partial \ln\left(\frac{M_1}{M_2}\right)}{\partial s_{01}} = \frac{(1 - s^N)(\beta_b^N - \alpha)}{(\beta_b^N - \alpha)(\sigma + \beta_r^M e) - \alpha s^N \theta \bar{\lambda}}. \quad (16)$$

Given these derivatives, we can characterize how a productivity shock in location 1 affects the distribution of per-capita access to researchers:

$$\frac{\partial \ln\left(\frac{N_1/M_1}{N_2/M_2}\right)}{\partial s_{01}} = \frac{\alpha(1 - s^N)}{(\beta_b^N - \alpha)(\sigma + \beta_r^M e) - \alpha s^N \theta \bar{\lambda}}. \quad (17)$$

A sufficient condition for a stable interior equilibrium is that the common denominator in Equations 15–17 be positive:

Assumption 1 *The parameters satisfy the following condition:*

$$(\beta_b^N - \alpha)(\sigma + \beta_r^M e) - \alpha s^N \theta \bar{\lambda} > 0.$$

This condition imposes net decreasing returns to scale in each location. It holds as long as the agglomeration elasticity α is not too large, implying that $\beta_b^N - \alpha > 0$ must hold.

Proposition 1 *Under Assumption 1, an increase in non-research productivity in location 1, s_{01} , raises M_1 and increases N_1/M_1 relative to N_2/M_2 .*

Proof. Under Assumption 1, the comparative static provided by equation 16 shows that an increase in s_{01} leads to an increase in the population share of non-researchers in location 1 if $\beta_b^N > \alpha$. Since $\beta_b^N > \alpha$ is implied from Assumption 1, we know that an increase in s_{01} must lead to an increase in M_1 .

Under Assumption 1, the comparative static provided by equation 17 shows that location 1's per-capita access to researchers by non-researchers increases relative to that in location 2 in response to an increase in s_{01} . ■

A1.4.2 Effects of β_{01}^N on Research and Non-Research Population Distributions

Next, we examine how a change in the location-specific amenity value for researchers, β_{01}^N , affects the spatial distribution of both researcher and non-researcher populations: $\frac{\partial \ln\left(\frac{N_1}{N_2}\right)}{\partial \beta_{01}^N}$ and $\frac{\partial \ln\left(\frac{M_1}{M_2}\right)}{\partial \beta_{01}^N}$.

Totally differentiating the researchers' utility-equalization condition yields:

$$(\alpha - \beta_b^N) \left(\frac{1}{N_1} + \frac{1}{N_2} \right) \frac{\partial N_1}{\partial \beta_{01}^N} + \beta_b^N \left(\frac{1}{M_1} + \frac{1}{M_2} \right) \frac{\partial M_1}{\partial \beta_{01}^N} + 1 = 0.$$

Totally differentiating the equilibrium condition for the non-researcher population in location 1, again evaluating the access elasticities at $\bar{\lambda}$, yields:

$$(\sigma + s^N \theta \bar{\lambda} + \beta_r^M e) \left(\frac{1}{M_1} + \frac{1}{M_2} \right) \frac{\partial M_1}{\partial \beta_{01}^N} - s^N \theta \bar{\lambda} \left(\frac{1}{N_1} + \frac{1}{N_2} \right) \frac{\partial N_1}{\partial \beta_{01}^N} = 0.$$

Solving the two equations for $\frac{\partial M_1}{\partial \beta_{01}^N}$ and $\frac{\partial N_1}{\partial \beta_{01}^N}$ yields:

$$\left(\frac{1}{N_1} + \frac{1}{N_2} \right) \frac{\partial N_1}{\partial \beta_{01}^N} = \frac{\sigma + s^N \theta \bar{\lambda} + \beta_r^M e}{(\beta_b^N - \alpha)(\sigma + \beta_r^M e) - \alpha s^N \theta \bar{\lambda}}$$

and

$$\left(\frac{1}{M_1} + \frac{1}{M_2} \right) \frac{\partial M_1}{\partial \beta_{01}^N} = \frac{s^N \theta \bar{\lambda}}{(\beta_b^N - \alpha)(\sigma + \beta_r^M e) - \alpha s^N \theta \bar{\lambda}}.$$

Since $\frac{\partial \ln\left(\frac{N_1}{N_2}\right)}{\partial \beta_{01}^N} = \left(\frac{1}{N_1} + \frac{1}{N_2} \right) \frac{\partial N_1}{\partial \beta_{01}^N}$ and $\frac{\partial \ln\left(\frac{M_1}{M_2}\right)}{\partial \beta_{01}^N} = \left(\frac{1}{M_1} + \frac{1}{M_2} \right) \frac{\partial M_1}{\partial \beta_{01}^N}$, we arrive at the following:

$$\frac{\partial \ln\left(\frac{N_1}{N_2}\right)}{\partial \beta_{01}^N} = \frac{\sigma + s^N \theta \bar{\lambda} + \beta_r^M e}{(\beta_b^N - \alpha)(\sigma + \beta_r^M e) - \alpha s^N \theta \bar{\lambda}}. \quad (18)$$

and

$$\frac{\partial \ln \left(\frac{M_1}{M_2} \right)}{\partial \beta_{01}^N} = \frac{s^N \theta \bar{\lambda}}{(\beta_b^N - \alpha)(\sigma + \beta_r^M e) - \alpha s^N \theta \bar{\lambda}}. \quad (19)$$

Combining these derivatives gives the effect of β_{01}^N on researcher access per capita in location 2 relative to location 1:

$$\frac{\partial \ln \left(\frac{N_2/M_2}{N_1/M_1} \right)}{\partial \beta_{01}^N} = \frac{-(\sigma + \beta_r^M e)}{(\beta_b^N - \alpha)(\sigma + \beta_r^M e) - \alpha s^N \theta \bar{\lambda}}. \quad (20)$$

Proposition 2 *Under Assumption 1, an increase in the researcher amenity in location 1, β_{01}^N , lowers N_2/M_2 relative to N_1/M_1 .*

Proof. Under Assumption 1, we can easily see that $\frac{\partial \ln \left(\frac{N_2/M_2}{N_1/M_1} \right)}{\partial \beta_{01}^N} < 0$, which proves the proposition. ■

A1.4.3 Effects of β_b^N on Research and Non-Research Population Distributions

Following the same steps as above, we obtain the effects of β_b^N on N_1 and M_1 :

$$\left(\frac{1}{N_1} + \frac{1}{N_2} \right) \frac{\partial N_1}{\partial \beta_b^N} = - \frac{(\sigma + s^N \theta \bar{\lambda} + \beta_r^M e) \ln \left(\frac{N_1/N_2}{M_1/M_2} \right)}{(\beta_b^N - \alpha)(\sigma + \beta_r^M e) - \alpha s^N \theta \bar{\lambda}}$$

and

$$\left(\frac{1}{M_1} + \frac{1}{M_2} \right) \frac{\partial M_1}{\partial \beta_b^N} = - \frac{s^N \theta \bar{\lambda} \ln \left(\frac{N_1/N_2}{M_1/M_2} \right)}{(\beta_b^N - \alpha)(\sigma + \beta_r^M e) - \alpha s^N \theta \bar{\lambda}}$$

Combining these expressions gives the effect of β_b^N on researcher access per capita in location 1 relative to location 2:

$$\frac{\partial \ln \left(\frac{N_1/M_1}{N_2/M_2} \right)}{\partial \beta_b^N} = \frac{-(\sigma + \beta_r^M e) \ln \left(\frac{N_1/M_1}{N_2/M_2} \right)}{(\beta_b^N - \alpha)(\sigma + \beta_r^M e) - \alpha s^N \theta \bar{\lambda}}. \quad (21)$$

Equation 21 is negative when location 1 initially has higher researcher access per capita. Thus, an increase in β_b^N narrows the access gap between the two locations.

Proposition 3 *Under Assumption 1, if researcher access per capita is initially higher in location 1 than in location 2, an increase in the linkage between the provision of local researcher positions and local population, β_b^N , narrows the researcher access gap by raising N_2/M_2 relative to N_1/M_1 .*

Proof. Under Assumption 1, $\frac{\partial \ln \left(\frac{N_1/M_1}{N_2/M_2} \right)}{\partial \beta_b^N} < 0$ if the per-capita access to researchers is higher in location 1, and vice versa if the per-capita access to researchers is higher in location 2. This means that β_b^N equalizes per-capita access to researchers across the two locations. ■

A1.4.4 The Effect of β_b^N on National Output

Totally differentiating log national output gives $\frac{\partial y}{\partial \beta_b^N}$ as a function of $\frac{\partial N_1}{\partial \beta_b^N}$ and $\frac{\partial M_1}{\partial \beta_b^N}$. The marginal effect of a researcher on log output carries the access elasticity $\theta \lambda_j$, while the marginal effect of a non-researcher carries the marginal-to-average product wedge $1 - s^N \theta \lambda_j$; we evaluate both at $\bar{\lambda}$. The access levels take the additive form, with per-non-researcher access productivity $S_{0j}^{1-s^N} S_j^{s^N}$.

The effect of β_b^N on log national output can be written as follows:

$$\begin{aligned} \frac{\partial y}{\partial \beta_b^N} = & \left[s^N \left(\frac{Y_1}{Y} \frac{\theta \lambda_1}{N_1} - \frac{Y_2}{Y} \frac{\theta \lambda_2}{N_2} \right) + \frac{s^N (1 + \alpha)}{(1 - \eta) \delta Q^{1-\eta}} (A_{01} N_1^\alpha - A_{02} N_2^\alpha) \right] \frac{\partial N_1}{\partial \beta_b^N} \\ & + \frac{1 - s^N \theta \bar{\lambda}}{B} \left(S_{01}^{1-s^N} S_1^{s^N} - S_{02}^{1-s^N} S_2^{s^N} \right) \frac{\partial M_1}{\partial \beta_b^N}. \end{aligned}$$

Substituting the comparative statics for researchers and non-researchers and rearranging yields:

$$\begin{aligned} \frac{\partial y}{\partial \beta_b^N} \propto & P_1^N (1 - P_1^N) N \left[\underbrace{s^N \left(\frac{Y_2}{Y} \frac{\theta \lambda_2}{N_2} - \frac{Y_1}{Y} \frac{\theta \lambda_1}{N_1} \right)}_{\text{Gain in Knowledge Access}} - \underbrace{\frac{s^N (1 + \alpha)}{(1 - \eta) \delta Q^{1-\eta}} (A_{01} N_1^\alpha - A_{02} N_2^\alpha)}_{\text{Loss in Knowledge Production}} \right] \\ & - \underbrace{\frac{s^N \theta \bar{\lambda} (1 - s^N \theta \bar{\lambda})}{B (\sigma + s^N \theta \bar{\lambda} + \beta_r^M e)} P_1^M (1 - P_1^M) M \left(S_{01}^{1-s^N} \left(S_0 + \left(\frac{N_1}{M_1} \right)^\theta \right)^{s^N} - S_{02}^{1-s^N} \left(S_0 + \left(\frac{N_2}{M_2} \right)^\theta \right)^{s^N} \right)}_{\text{Loss in Productivity due to Non-Researcher Migration}}. \end{aligned}$$

A2 Comparison of Microsoft Academic Graph with Other Bibliographic Databases

We use the Microsoft Academic Graph (MAG) because it combines broad cross-disciplinary coverage with linked information on publications, citations, authors, institutional affiliations, and fields of study—all of which are necessary to construct our annual researcher panel. An additional advantage is that MAG permits free bulk access, making the construction of our sample more transparent and reproducible.

Existing comparisons support the comprehensiveness and reliability of MAG. Visser et al. (2021) compare MAG with other major bibliographic data sources, including Scopus, Web of Science, Dimensions, and Crossref, and find that MAG provides more comprehensive coverage of scientific publications than the other databases, with most records unique to MAG corresponding to scientific work. While Scopus includes more journal articles and conference proceedings in some subjects, MAG exhibits the greatest overlap with Scopus relative to other sources.

In terms of citation measures, Thelwall (2018) finds that Microsoft Academic is essentially equivalent

to Scopus in terms of its citation counts, making it a robust source of free citation data. Martín-Martín et al. (2021) show that in most subjects, MAG retrieves more citations than Scopus and WoS, although it has coverage gaps in some fields, particularly physics and parts of the humanities. Visser et al. (2021) show that across all databases, documents with a higher number of citations are disproportionately represented in the overlap between data sources. Thus, differences between data sources may not be consequential in terms of the relative “importance” of publications (Visser et al., 2021). These findings suggest that MAG provides broadly comparable citation measures to those of other bibliographic data sources.

Wang et al. (2020) describe MAG as replicating “the success of Google Scholar, which utilizes the massive document index from a web search engine to achieve comprehensive coverage of contemporary scholarly materials, many of which are not published and distributed through traditional channels and not assigned DOIs.” MAG has consequently been used in recent economics research on scientific careers, inequality, and innovation, such as Airolidi and Moser (2024), Kim and Moser (2025), Truffa and Wong (2025), and Koffi and Marx (2026).

A3 Back-of-the-Envelope Calculation: Gains from Improved Access to Knowledge

The evidence in Section 4 shows that research activity has become increasingly misaligned with population, suggesting potential gains from bringing research closer to where people live and work. This appendix uses the spatial framework in Section 2 to quantify the access channel alone. Specifically, we ask how much aggregate output would increase if researchers were distributed across metropolitan areas in proportion to population, holding the total stock of frontier knowledge fixed.

Using Equation (4), we reallocate researchers so that each MSA’s share of researchers equals its share of the national population. We measure non-research productivity using residualized wages from the 2015–2019 ACS and hold it fixed at its baseline value. We set $\lambda_j = \bar{\lambda} = 0.5$, so that local and non-local channels each account for half of knowledge access. We set $s^N = 0.2$, consistent with estimates of the aggregate elasticity of productivity with respect to frontier knowledge. We calibrate the local access elasticity $s^N \theta \bar{\lambda}$ to 0.08, matching the magnitude of local university spillovers estimated by Kantor and Whalley (2014).

Under these assumptions, reallocating researchers in proportion to population raises aggregate output by 5.3%. When geographic access is instead measured using research papers per capita, the implied gain rises


to 8.8%. These calculations illustrate that the access benefits of a more geographically balanced allocation of research activity could be economically meaningful.

The exercise, however, captures only the gross gain through improved knowledge access. It holds the stock of frontier knowledge fixed and therefore does not account for the possibility that reallocation moves researchers away from productive institutions, reduces research-cluster sizes, and lowers knowledge production. Section 6 incorporates these production-side effects and evaluates the net effect of researcher reallocation.

A4 Counterfactual Analysis: Migration of Non-Researchers


To simulate the migration of the non-research population, we assume a standard multinomial logit framework in which the change in per-capita researchers shifts the local marginal product of labor. We set the unobserved idiosyncratic location taste dispersion based on Hsieh and Moretti (2019). We set the counterfactual wage in each MSA equal to this counterfactual marginal product of labor, and, given this wage shock, recompute each MSA's population; the resulting population serves as the counterfactual non-research population used to compute the counterfactual access term $\ln \hat{B}$. We hold housing rents fixed with respect to the change in local wages. Because a rent response would dampen migration, holding rents fixed makes our counterfactual migration response an upper bound: if even this upper-bound response generates no sizable efficiency cost, the true cost is likely smaller still.

Figure A1: Shares of Population, Researchers, and Publications in the Ten Largest Research Clusters




Notes: The figure plots the shares of U.S. population, researchers, and publications that the ten largest research clusters accounted for from 1970 to 2019. The black lines allow the set of ten largest clusters to change over time, while the blue lines hold the set fixed at the ten largest clusters in 1970.

Figure A2: Shares of Researcher Affiliations by Institution Type and Field




(a) Shares by Institution Type and Field




(b) Shares of University Affiliations over Time

Notes: We use Claude AI to classify research institutions into broad types based on the names of the institutions. Panel (a) shows, for each broad field, the share of researcher affiliations associated with each institution type, pooling publication data from 1970 to 2015. The unit of observation is a researcher–institution affiliation appearing in a publication in a given year. Panel (b) plots the annual shares of researcher affiliations associated with universities within each broad field.

Figure A3: Rising Inequality in Access to Research Papers: Universities vs. Non-Universities




(a) University Research



(b) Non-University Research


Notes: Panel (a) plots population-weighted Gini coefficients of access to research papers produced by university-affiliated authors across MSAs, while Panel (b) reports the corresponding coefficients for non-university-affiliated authors. The black series use papers per resident, and the blue series use papers per college graduate.

Figure A4: Gini Coefficients of Research Papers per Resident by Field, 1970–2019




Notes: The figure plots population-weighted Gini coefficients of research papers per resident across MSAs from 1970 to 2019, separately for five broad fields.

Figure A5: Distributions of Access to Research Papers Across MSAs, 1970, 1990, and 2019




(a) Research Papers per 1,000 Residents




(b) Research Papers per 1,000 College Graduates

Notes: Panel (a) plots the population-weighted distribution of demeaned log research papers per 1,000 residents across MSAs. Panel (b) plots the corresponding distribution of research papers per 1,000 college graduates. Publication counts are aggregated over five-year periods for authors affiliated with institutions in each MSA.

Figure A6: Decomposing Changes in the Distribution of Research Papers per 1,000 Residents




(a) Changing Research Activity (Population Fixed at 1970)




(b) Population Shifts (Publications Fixed at 1970)

Notes: Panel (a) plots the population-weighted distribution of demeaned log research papers per 1,000 residents across MSAs, holding each MSA's population fixed at its 1970 level while allowing publication counts to vary. Panel (b) holds publication counts fixed at their 1970 levels while allowing population to vary. The panels isolate the contributions of changing research activity and population shifts, respectively.


Figure A7: Distributions of Research Papers per Worker and Business Establishment




(a) Research Papers Per 1,000 Workers



(b) Research Papers Per 1,000 High-Tech Workers




(c) Research Papers Per 1,000 Establishments




(d) Research Papers Per 1,000 High-Tech Establishments

Notes: The figure plots distributions of access to research papers across MSAs in 1980, 1990, and 2019. Panels (a) and (b) show demeaned log research papers per 1,000 workers and per 1,000 high-tech workers, respectively, weighted by the corresponding employment counts. Panels (c) and (d) show demeaned log research papers per 1,000 establishments and per 1,000 high-tech establishments, respectively, weighted by the corresponding establishment counts. Publication counts are aggregated over five-year periods. Employment and establishment data come from the QCEW and County Business Patterns, respectively. High-tech sectors are defined using 2017 three-digit NAICS codes, including oil and gas extraction (211), utilities (221), pipeline transportation (486), petroleum and coal products (324), chemicals (325), machinery (333), computer and electronic products (334), electrical equipment and components (335), transportation equipment (336), software (511), telecommunications (517), data processing and hosting (518), other information services (519), wholesale electronic markets and agents and brokers (425), professional, scientific, and technical services (541), and management of companies and enterprises (551).


Figure A8: Distributions of Access to University and Non-University Research




(a) University Research per 1,000 Residents



(b) University Research per 1,000 College Graduates




(c) Non-University Research per 1,000 Residents




(d) Non-University Research per 1,000 College Graduates

Notes: The figure plots distributions of access to research papers across MSAs in 1970, 1990, and 2019. Panels (a) and (b) include papers produced by university-affiliated authors, while Panels (c) and (d) include papers produced by authors affiliated with non-university institutions. Panels (a) and (c) show demeaned log papers per 1,000 residents, weighted by MSA population. Panels (b) and (d) show demeaned log papers per 1,000 college graduates, weighted by the college-educated population. Publication counts are aggregated over five-year periods.


Figure A9: Research Papers, Researchers, and Faculty per 1,000 Residents



(a) Research Papers




(b) Researchers



(c) Tenured and Tenure-Track Faculty


Notes: The figure plots selected percentiles of the MSA-level distributions of demeaned log research papers, active researchers, and tenured and tenure-track faculty per 1,000 residents over time. Research papers and researcher counts come from the Microsoft Academic Graph, while faculty counts come from IPEDS.

Figure A10: Event Study: Moves to Higher vs. Lower-Output Institutions




Notes: Each panel plots the estimated coefficients δ_s for $-15 \leq s \leq 15$ ($s \neq -1$) from Equation 8, using the sample of researchers who changed institutions once (analogous to Figure 4). The dependent variable is the inverse hyperbolic sine of publications (Panels (a) and (b)), citations (Panels (c) and (d)), and patent citations (Panels (e) and (f)). Panels (a), (c), and (e) use movers who move up (i.e., $\Delta_i > 0$). Panels (b), (d), and (f) use movers who move down (i.e., $\Delta_i < 0$).

Figure A11: Additive Decomposition of Research Outcomes By Field:
Top vs. Bottom Productivity Deciles



(a) Publications


(b) Citations



(c) Patent Citations


Notes: This figure reports field-specific additive decompositions of research outcomes across researcher productivity groups during 2010–2015. Panels (a), (b), and (c) report results for IHS publications, citations, and patent citations, respectively. Within each field, researchers are ranked by their average value of the corresponding outcome, and the figure compares researchers in the top and bottom deciles. The researcher component includes author fixed effects and field-specific academic-age effects, while the location component is predicted using lagged observable characteristics of the institution–field and external MSA–field research environments. Component shares are calculated as the difference in each component divided by the overall difference in the corresponding outcome. Subfigure (c) includes only patent-relevant fields, i.e., fields in which more than 30% of papers receive at least one patent citation.

Figure A12: Additive Decomposition of Research Outcomes By Field:
Top vs. Bottom Institution Size Deciles



(a) Publications


(b) Citations



(c) Patent Citations


Notes: This figure reports field-specific additive decompositions of research outcomes across researchers by institution size during 2010–2015. Panels (a), (b), and (c) report results for IHS publications, citations, and patent citations, respectively. Within each field, institutions are ranked by the average number of active researchers, and the figure compares researchers in the top and bottom deciles. The researcher component includes author fixed effects and field-specific academic-age effects, while the location component is predicted using lagged observable characteristics of the institution–field and external MSA–field research environments. Component shares are calculated as the difference in each component divided by the overall difference in the corresponding outcome. Panel (c) includes only patent-relevant fields, i.e., fields in which more than 30% of papers receive at least one patent citation.

Figure A13: Additive Decomposition of Research Outcomes By Field:
Top vs. Bottom MSA Size Deciles



(a) Publications


(b) Citations




(c) Patent Citations

Notes: This figure reports field-specific additive decompositions of research outcomes across researchers by MSA size during 2010–2015. Panels (a), (b), and (c) report results for IHS publications, citations, and patent citations, respectively. Within each field, MSAs are ranked by the average number of active researchers, and the figure compares researchers in the top and bottom deciles. The researcher component includes author fixed effects and field-specific academic-age effects, while the location component is predicted using lagged observable characteristics of the institution–field and external MSA–field research environments. Component shares are calculated as the difference in each component divided by the overall difference in the corresponding outcome. Subfigure (c) includes only patent-relevant fields, i.e., fields in which more than 30% of papers receive at least one patent citation.


Figure A14: Leave-One-Out Test for IV Estimates: Top Contributing Subfields




(a) Publications (α^I)




(b) Publications (α^E)




(c) Citations (α^I)



(d) Citations (α^E)




(e) Patent Citations (α^I)




(f) Patent Citations (α^E)

Notes: This figure reports leave-one-subfield-out estimates from the preferred IV specification. For each outcome, we reconstruct the internal and external Bartik instruments 364 times, each time omitting one detailed subfield, and display the ten estimates with the largest absolute deviations from the corresponding baseline estimate. Panels (a) and (b) report results for IHS publications, Panels (c) and (d) for IHS citations, and Panels (e) and (f) for IHS patent citations. Panels (a), (c), and (e) report the own-institution elasticity, α^I , while Panels (b), (d), and (f) report the external-cluster elasticity, α^E . 95% confidence intervals are displayed based on standard errors clustered at the institution level.

Figure A15: Anderson–Rubin Confidence Sets for IV Estimates




(a) Publications



(b) Citations

Figure A15: Anderson–Rubin Confidence Sets for IV Estimates (continued)



Notes: Each panel reports the 95% Anderson–Rubin confidence set for the own-institution and external-cluster elasticities, (α^I, α^E) , from our preferred IV specification for IHS publications in Panel (a), IHS citations in Panel (b), and IHS patent citations in Panel (c). The shaded region contains parameter pairs that cannot be rejected at the 5% level, and the accompanying surface plots the rejection probability, $1 - p$. Anderson–Rubin inference remains valid under weak identification, with standard errors clustered at the institution level. For all three outcomes, the confidence sets lie entirely in the region with $\alpha^I > 0$.

Table A1: Sample Sizes

Sample	Observations
<i>Panel A: Full Sample</i>	
Researcher-years (with imputation)	3,092,028
Researcher-years (without imputation)	1,191,273
Researchers	187,865
Non-movers	143,856
Movers	44,009
Movers (moved once)	17,482
<i>Panel B: IV Sample</i>	
Researcher-years (with imputation)	798,548
Researcher-years (without imputation)	271,602
Researchers	35,263
Non-movers	29,302
Movers	5,961
Movers (moved once)	2,652

Note: This table reports sample sizes under different sample restrictions. Each observation in the data is a researcher-year. “With imputation” includes years in which no publication is observed; for these years, we impute zero publications and carry forward the most recent affiliation. “Without imputation” restricts the sample to researcher-years with at least one publication. The IV sample includes three periods: 1995, 2005, and 2015; each period pools data over the focal year and the two preceding years.

Table A2: Top 10 Metropolitan Research Clusters

Top 10 Clusters in 2019	Number of Authors	Number of Publications
Boston-Cambridge-Newton	3812	22805
New York-Newark-Jersey City	3536	15137
Washington-Arlington-Alexandria	3102	13520
Los Angeles-Long Beach-Anaheim	1706	8155
San Francisco-Oakland-Hayward	1316	7155
Philadelphia-Camden-Wilmington	1451	6426
Chicago-Naperville-Elgin	1396	6351
Houston-The Woodlands-Sugar Land	850	5127
San Jose-Sunnyvale-Santa Clara	1268	4945
Seattle-Tacoma-Bellevue	862	4527

Note: This table reports the ten largest MSAs in academic research in 2019, ranked by the number of publications authored by researchers affiliated with institutions in each MSA. It also reports the corresponding number of active authors affiliated with institutions in each MSA.

Table A3: OLS Estimates of Agglomeration Elasticities: Full Sample

	(1)	(2)	(3)	(4)
<i>Panel A: Publications</i>				
Ln(Institution Size)	0.134*** (0.006)	0.131*** (0.006)	0.258*** (0.007)	0.196*** (0.005)
Ln(External Size)	-0.039*** (0.004)	0.001 (0.002)	-0.007*** (0.002)	0.016*** (0.001)
<i>Panel B: Citations</i>				
Ln(Institution Size)	0.351*** (0.015)	0.283*** (0.011)	0.547*** (0.015)	0.412*** (0.011)
Ln(External Size)	-0.087*** (0.009)	0.011** (0.005)	-0.017*** (0.004)	0.036*** (0.003)
<i>Panel C: Patent Citations</i>				
Ln(Institution Size)	0.044*** (0.003)	0.027*** (0.001)	0.043*** (0.002)	0.031*** (0.003)
Ln(External Size)	-0.004*** (0.002)	0.002*** (0.001)	-0.007*** (0.001)	-0.002*** (0.001)
Observations	3,092,028	3,088,219	3,086,886	2,755,515
Field $\times$ Year FE	Yes	Yes	Yes	Yes
Author FE	No	Yes	Yes	Yes
Academic Age $\times$ Field FE	No	Yes	Yes	Yes
Affiliation $\times$ Field FE	No	No	Yes	Yes
MSA $\times$ Year FE	No	No	No	Yes

Note: This table reports OLS estimates of Equation 10 using the full annual author-year panel from 1970 to 2015. The dependent variable is the inverse hyperbolic sine (IHS) of publications in Panel A, IHS citations in Panel B, and IHS patent citations in Panel C. Institution size is the number of active researchers in the author's field and institution in a given year. External size is the number of active researchers in the same field and MSA, excluding the home institution. Each column corresponds to a separate regression with the fixed effects listed at the bottom of the table. Academic age is measured as years since the author's first publication. Standard errors are in parentheses and clustered at the institution level: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A4: Pre-Trend Tests for Bartik-Predicted Cluster Growth

	1990–1995 Change in IHS Output					
	Pubs (1)	Citations (2)	Patent Cites (3)	Pubs (4)	Citations (5)	Patent Cites (6)
Pred. Institution Growth	-0.044 (0.044)	0.044 (0.072)	-0.173** (0.075)	-0.017 (0.012)	-0.006 (0.020)	-0.054** (0.023)
Pred. External Growth	0.005 (0.012)	-0.035 (0.025)	0.059*** (0.019)	0.003 (0.004)	-0.014 (0.009)	0.023*** (0.008)
Observations	9,781	9,781	9,781	9,781	9,781	9,781
R-squared	0.052	0.033	0.119	0.052	0.033	0.119
Joint p-value	0.586	0.329	0.002	0.324	0.262	0.003
Predicted Growth Window	95–05	95–05	95–05	95–15	95–15	95–15
	1985–1995 Change in IHS Output					
	Pubs (1)	Citations (2)	Patent Cites (3)	Pubs (4)	Citations (5)	Patent Cites (6)
Pred. Institution Growth	-0.085 (0.064)	-0.018 (0.107)	-0.366*** (0.103)	-0.041** (0.018)	-0.032 (0.029)	-0.159*** (0.034)
Pred. External Growth	0.040** (0.016)	0.024 (0.030)	0.170*** (0.032)	0.015*** (0.006)	0.010 (0.011)	0.066*** (0.012)
Observations	8,657	8,657	8,657	8,657	8,657	8,657
R-squared	0.161	0.097	0.346	0.162	0.097	0.348
Joint p-value	0.024	0.733	0.000	0.004	0.381	0.000
Predicted Growth Window	95–05	95–05	95–05	95–15	95–15	95–15

Note: This table reports estimates of π^I and π^E from Equation 13, which relate pre-base-period growth in IHS research output to Bartik-predicted growth in home-institution and external-cluster size. The dependent variable is the change in IHS research output between 1990 and 1995 in the upper panel and between 1985 and 1995 in the lower panel. In Columns 1–3, the key regressors are Bartik-predicted growth in home-institution and external-cluster size between 1995 and 2005. In Columns 4–6, the key regressors are the corresponding predicted growth measures between 1995 and 2015. All regressions include field fixed effects. Standard errors are in parentheses and clustered at the institution level: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A5: IV Estimates of Agglomeration Elasticities by Cluster-Size Tertile

	Publications (1)	Citations (2)	Patent Cits (3)
Institution Size: Small	0.205** (0.091)	0.960*** (0.226)	0.193*** (0.047)
Institution Size: Medium	0.157* (0.083)	0.801*** (0.213)	0.156*** (0.043)
Institution Size: Large	0.225*** (0.067)	0.793*** (0.169)	0.109*** (0.036)
External Size: Small	0.071*** (0.027)	0.035 (0.058)	-0.011 (0.010)
External Size: Medium	0.056*** (0.020)	0.051 (0.054)	-0.009 (0.010)
External Size: Large	0.090*** (0.025)	0.081 (0.071)	-0.004 (0.013)
Observations	778,733	778,733	778,733
Baseline FEs	Yes	Yes	Yes
Author/Coauthor Controls	Yes	Yes	Yes
Kleibergen–Paap Wald F	5.97	5.97	5.97
Equal Institution p-value	0.057	0.053	0.000
Equal External p-value	0.156	0.703	0.809

Note: This table reports IV estimates of Equation 14 using the IV sample. The sample consists of author-year observations in three periods: 1995, 2005, and 2015, where each period pools the focal year and the two preceding years. The dependent variable is IHS publications in Column 1, IHS citations in Column 2, and IHS patent citations in Column 3. The key regressors are log own-institution size and log external-cluster size interacted with indicators for their respective 1995 size tertiles. These interactions are instrumented with the corresponding Bartik instruments interacted with the same tertile indicators. All columns include the baseline fixed effects from Equation 10 and the author and coauthor direct-exposure controls. The final two rows report p -values for equality of the three own-institution coefficients and the three external-cluster coefficients, respectively. Standard errors are in parentheses and clustered at the institution level: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A6: IV Estimates of Agglomeration Elasticities
Using Alternative Cluster Measures

Cluster Definition Size Measurement	MSA		25-km Radius	
	Author Count (1)	Paper Count (2)	Author Count (3)	Paper Count (4)
<i>Panel A: Publications</i>				
Ln(Institution Size)	0.193*** (0.069)	0.101 (0.097)	0.221*** (0.082)	0.133 (0.110)
Ln(External Size)	0.081*** (0.019)	0.049*** (0.014)	-0.139 (0.119)	-0.062 (0.053)
<i>Panel B: Citations</i>				
Ln(Institution Size)	0.795*** (0.183)	0.898*** (0.281)	0.806*** (0.214)	0.870*** (0.320)
Ln(External Size)	0.084 (0.053)	0.048 (0.053)	-0.199 (0.279)	-0.162 (0.167)
<i>Panel C: Patent Citations</i>				
Ln(Institution Size)	0.140*** (0.041)	0.336*** (0.114)	0.138*** (0.052)	0.335*** (0.127)
Ln(External Size)	-0.008 (0.011)	0.007 (0.020)	0.076 (0.082)	0.000 (0.057)
Observations	794,658	794,658	692,763	692,763
Kleibergen-Paap Wald F	15.33	4.718	3.282	3.926

Note: This table reports IV estimates of Equation 10 using the IV sample. The dependent variable is IHS publications in Panel A, IHS citations in Panel B, and IHS patent citations in Panel C. The sample consists of author-year observations in three periods: 1995, 2005, and 2015, where each period pools the focal year and the two preceding years. In Columns 1–2, the external cluster is defined as institutions in the same field and MSA as the home institution; in Columns 3–4, it is defined as institutions in the same field within 25 km of the home institution. Cluster size is measured by the number of active researchers in Columns 1 and 3 and by the number of papers published in Columns 2 and 4. The size measures are instrumented using the corresponding Bartik instruments. All specifications include author fixed effects, academic-age-by-field fixed effects, field-by-year fixed effects, affiliation-by-field fixed effects, and the author and coauthor direct-exposure controls. The Kleibergen–Paap rk Wald F statistic is reported for each specification. Standard errors are in parentheses and clustered at the institution level: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A7: IV Estimates of Agglomeration Elasticities
Under Explicit Extensive-Margin Calibrations

	(1)	(2)	(3)
<i>Panel A: Publications</i>			
Ln(Institution Size)	0.199*** (0.070)	0.112** (0.054)	0.236*** (0.079)
Ln(External Size)	0.074*** (0.020)	0.060*** (0.015)	0.082*** (0.023)
<i>Panel B: Citations</i>			
Ln(Institution Size)	0.828*** (0.190)	0.737*** (0.171)	0.879*** (0.201)
Ln(External Size)	0.056 (0.060)	0.043 (0.054)	0.062 (0.063)
<i>Panel C: Patent Citations</i>			
Ln(Institution Size)	0.148*** (0.043)	0.117*** (0.034)	0.163*** (0.047)
Ln(External Size)	-0.011 (0.012)	-0.007 (0.009)	-0.013 (0.013)
Extensive Margin Value	Baseline	0.1	1

Note: This table reports IV estimates of Equation 10 using the IV sample. The sample consists of author-year observations in three periods: 1995, 2005, and 2015, where each period pools the focal year and the two preceding years. Column 1 uses the baseline IHS transformation. Columns 2 and 3 use $m(F) = \log(F)$ for $F > 0$ and $m(0) = -x$, with $x = 0.1$ and $x = 1$, respectively. All specifications include author fixed effects, academic-age-by-field fixed effects, field-by-year fixed effects, affiliation-by-field fixed effects, and the author and coauthor direct-exposure controls. Standard errors are in parentheses and clustered at the institution level: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A8: IV Estimates of Agglomeration Elasticities
Using Alternative Log Transformations of Cluster Size

	(1)	(2)	(3)
<i>Panel A: Publications</i>			
Ln(Institution Size)	0.199*** (0.070)	0.537*** (0.188)	1.031** (0.467)
Ln(External Size)	0.074*** (0.020)	0.043 (0.030)	-0.023 (0.066)
<i>Panel B: Citations</i>			
Ln(Institution Size)	0.828*** (0.190)	1.932*** (0.599)	3.527** (1.583)
Ln(External Size)	0.056 (0.060)	-0.071 (0.108)	-0.313 (0.235)
<i>Panel C: Patent Citations</i>			
Ln(Institution Size)	0.148*** (0.043)	0.320*** (0.121)	0.566* (0.290)
Ln(External Size)	-0.011 (0.012)	-0.034* (0.020)	-0.075* (0.042)
Transformations	Ln(N+1)	Ln(N+5)	Ln(N+10)
Kleibergen–Paap Wald F	14.9	5.34	2.34

Note: This table reports IV estimates of Equation 10 using different transformations of own-institution and external-cluster size. Columns 1–3 use $\ln(N+1)$, $\ln(N+5)$, and $\ln(N+10)$, respectively. The dependent variable is IHS publications in Panel A, IHS citations in Panel B, and IHS patent citations in Panel C. The sample consists of author-year observations in three periods: 1995, 2005, and 2015, where each period pools the focal year and the two preceding years. All specifications include author fixed effects, academic-age-by-field fixed effects, field-by-year fixed effects, affiliation-by-field fixed effects, and the author and coauthor direct-exposure controls. The Kleibergen–Paap rk Wald F statistic is reported for each specification. Standard errors are in parentheses and clustered at the institution level: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.