Do Technology Clusters Distort Innovation?

Benjamin Balsmeier [a], Gustavo Manso [b], and Lee Fleming [c]

a) University of Luxembourg: benjamin.balsmeier@uni.lu
b) University of California, Berkeley: manso@berkeley.edu
c) University of California, Berkeley: lfleming@berkeley.edu

July 2026

Abstract: While the geographical agglomeration of inventors appears to increase their raw patent productivity, we show that it also reduces their exploration of new technologies. We interpret this result through a stylized Marshallian model in which larger clusters enhance individual incentives to exploit. If truly novel ideas have greater social value, however, then an increase in cluster size may give rise to a wedge between private and social incentives. Highly saturated fixed effects and Bartik regressions with U.S. patent data from 1980-2010 support the model's predictions of increased exploitation and decreased exploration with increasing cluster size. Consistent with narrower search patterns, we also predict and find that, on average, patents originating from larger clusters exhibit lower creativity, fewer future citations, and reduced financial value.

JEL-Classification: O31, O33
Keywords: Knowledge spillovers, Patents, Agglomeration, Exploitation

---

Currently there is no funding support for this research. The authors thank Po-Hsuan Hsu, Andreas Panagopoulos, and Scott Stern for their feedback, as well as the participants of seminars and conferences at Trento University, TUM, 2025 NBER Place-Based Policies and Entrepreneurship, 2025 SMS Conference, GeoInno 2026, 14th AIEA-NBER Conference, and GEIRC 2026. Errors and omissions remain ours.

Introduction

The benefits of locating in a technology cluster - easier access to specialized production assets, talent, and knowledge spillovers - have long been noted both theoretically (Marshall 1890) and empirically (Jaffe, Trajtenberg, and Henderson 1993; Greenstone, Hornbeck, Moretti 2010; Ellison, Glaeser, Kerr 2010). These mechanisms of local advantage appear increasingly important, as the geographical distributions of wealth, valuable innovation, and technical human capital have all become increasingly concentrated over at least the last 40 years (Suss, Kemeny, Conner 2024; Forman, Goldfarb, Greenstein 2016; Kalyani et al. 2025). Rising spatial inequality is now blamed for political populism and social polarization (Rodríguez-Pose 2018), yet governments continue to promote technology clusters in search of innovation, entrepreneurship, and growth.

Recent work suggests that inventors who reside in - or move into - large clusters produce more patents than comparable peers elsewhere (Moretti 2021). Replicating that result with a larger dataset of U.S. patents (and fixing the coding concerns elaborated in Wiebe 2026), and implementing a shift-share instrument (Bartik 1991), we decompose the outcome measure of raw patent count into (i) patents in classes an inventor has already patented in vs. (ii) patents in classes she has not previously patented in. Results indicate that the cluster-driven productivity gain comes from classes the inventor has already patented in. In contrast, proximity to same-field peers significantly reduces the probability that an inventor enters a new technology class. The average creativity, citation impact, and financial value of inventions also decreases within clusters.

To discipline our empirics, we employ a simple explore–exploit conceptualization of creative search (March 1991; Manso 2011). At the start of each period an inventor has previously invented in what has now become a familiar technology class. She can (i) exploit those familiar classes or (ii) pay a fixed, idiosyncratic cost to explore an unfamiliar class. Exploitation succeeds with higher probability but yields lower per-patent value, whereas exploration succeeds with lower probability and yields higher per-patent value. The model frames the impact of agglomeration through two channels. First, more dominant clusters reduce the marginal cost of routine exploitation (via the classic Marshallian mechanisms of thicker knowledge spillovers, shared suppliers, and deeper labor pools). Second, while the social value of exploration might rise or fall with cluster size (due

to the tension between local complementarities vs. congestion), incentives for the private inventor lean toward exploitation as the cluster grows.

This framework serves to ground our three primary findings. First, individual patent counts increase with cluster size because cheaper search costs make exploitation more attractive. Second, the increase in cluster productivity comes from patents in familiar technology classes; in contrast, both the number and the share of "new-to-inventor" patents decline with cluster size. Third, average patent value—textual creativity, forward citations, and grant-date market value—falls as clusters thicken, reflecting a shift toward incremental combinations.

The data support the predictions; while the absolute number of patents rises as a cluster increases, this effect is driven entirely by exploitation patents. Both the fraction and absolute count of new-to-inventor patents fall with cluster concentration. Textual creativity, forward-citation rates, and financial values also decline, producing a concave relationship between cluster concentration and inventive productivity. Metaphorically more trees are planted, but each tree is narrower. The pattern is mostly driven by (i) the top decile of inventors and (ii) the selective inflow of newcomers into existing hubs—rather than by incumbent non-movers becoming more prolific.

Our simple model and measures suggest a novel distortion in the process of recombinant search (Weitzman 1998; Acemoglu, 2023). In principle, inventive progress arises from recombining diverse technological elements. However, in geographically clustered technology clusters, locally specialized assets, pooled labor markets, and thicker knowledge spillovers reduce the cost of search within the same technological field, thereby discouraging inventors from venturing into more distant domains. In a geographically clustered technology landscape—where for example, knowledge transfer is easier over short distances—this mechanism leads to more systematically narrower search paths than would be socially optimal in a frictionless world of equal access to all technological elements. Technology clusters thus create a wedge between individual and social incentives for exploration, with localized labor markets reinforcing the effect by providing more job opportunities to inventors who develop skills that are locally in high demand. This in turn contributes directly to and reinforces the dominant local trajectory.

Firms might internalize these frictions and partly overcome them by opening additional labs in other technology clusters, thereby broadening their access to more distant knowledge (Giroud, Liu, Mueller 2026). Yet such strategies remain a second-best solution: geographic distance continues to impose frictions, and the ability to maintain multiple R&D hubs is only available to large, well-capitalized firms. Our framework thus highlights a previously underappreciated channel through which agglomeration economies can distort innovative search and generate efficiency losses that grow with the scale and concentration of clusters. It refines the theory of recombinant search by showing how agglomeration economies, while fostering productivity within fields, can simultaneously narrow the breadth of search—producing taller, narrower trees of innovation rather than wider, more exploratory ones. That has implications for the socially optimal geographic distribution of innovative activity. Unlike in classic models of agglomeration, clustering may not unanimously create better outcomes. To achieve the optimal outcome, a social planner may need to balance the productivity gains of clusters with the potential losses from distortions in innovative search paths.

Given the increasing dominance of technology clusters for innovation, the results might also speak to a growing body of work that documents a decline in creativity and innovative dynamism, e.g., "ideas are getting harder to find" (Bloom, Jones, Van Reenen, Webb 2020) and the rate of transformative, breakthrough innovation has slowed relative to incremental improvements (Jones, 2009; Arora, Belenzon, & Patacconi, 2018). Scholars have debated whether this decline is the result of structural shifts in R&D, rising costs of invention, or organizational incentives that favor exploitation over exploration; the possibility that increasingly dense technology clusters may contribute to this phenomenon has received little attention.

The remainder of the paper proceeds as follows. Section I develops a stylized framework that organizes the empirical predictions. Section II describes data construction and measurement. Section III documents the methodology; Section IV presents empirical results and a discussion, and Section V concludes.

I. A Stylized Framework for Search

This section develops a stylized framework for thinking about how local technology clusters shape an inventor's choice between exploration and exploitation. We model individual inventor search behavior as a choice between exploitation (working in technology classes the inventor has previously patented in) vs. exploration (moving into new and unfamiliar technology classes). Exploitation search leads to greater productivity but decreased creativity, whereas exploration leads to decreased productivity and greater creativity (Manso, 2011; Arts and Fleming 2018; Hill et al. 2025). Clusters (i.e., when geographically proximate inventors work in similar technologies) lower an individual inventor's search costs for exploitation, arguably for any of the classic Marshall mechanisms, i.e., richer knowledge spillovers (exploitation knowledge is more available and cheaper to access), specialized invention infrastructure (resources that support invention within the cluster become cheaper and more readily available), or denser labor markets (that provide greater career security because they value the inventor's extant expertise). If truly novel ideas have greater social value, however, then an increase in cluster size may give rise to a wedge between private and social incentives.

I A. *Agents, actions, private pay-offs, and Marshallian mechanisms*

A unit continuum of risk-neutral inventors must choose one of two activities in a given period:

1. **Explore** (E): attempt to patent in a *new* technology class which the inventor has not previously patented in.
2. **Exploit** (X): attempt to patent in an *old* technology class which the inventor has previously patented in.

Exploration yields a patent of value A with probability α, while exploitation yields a patent of value B with probability β, such $B < A$ and $\beta > \alpha$. We make the twin assumptions that (i) exploratory patents are rarer but more valuable (A > B) and (ii) exploitation attempts succeed more often ($\beta > \alpha$). Studies of invention quality show that patents combining novel technological components are more commercially successful (Kalyani et al. 2025). At the same time, incremental projects enjoy shorter development cycles because they reuse existing routines, equipment and know-how accumulated through prior work. These patterns imply that an

exploitation attempt is more likely to yield a granted patent (β > α) but that each of those patents, being closer to the technological frontier already mastered, embodies less novel knowledge and therefore delivers a lower expected value (B < A) than the less frequent successes that arise from exploration.

Exploration entails an idiosyncratic private cost $c_i \sim U[0,1]$. The uniform distribution will deliver closed-form shares while preserving all qualitative insights for any continuous distribution.

Let $s \in [0,1]$ index the exogenous size of the local knowledge cluster. Exploitation entails private cost $k(s)$:

$$k: [0,1] \to \mathbb{R}_+, \qquad \boxed{k'(s) < 0,\ k(0) > 0}.$$

Thicker clusters supply specialized labor, shared infrastructure, and greater knowledge sharing, all of which lower the marginal cost of generating incremental patents (Marshall 1890; Moretti 2021). The sole assumption is monotonicity; no functional form is imposed.

I B. ***Private equilibrium***

Private pay-offs are

$$\pi_E(i) = \alpha A - c_i, \qquad \pi_X = \beta B - k(s).$$

Inventor i chooses E iff $c_i \leq c^*$ with

$$c^* \equiv \alpha A - \beta B + k(s).$$

Because $c_i \sim U[0,1]$, the interior share of explorers equals

$$\boxed{n_E^{\text{priv}}(s) = \alpha A - \beta B + k(s).}$$

Interior-solution condition.

We require $0 < c^* < 1$, i.e., $0 < \alpha A - \beta B + k(s) < 1$, for the relevant range of s. This is guaranteed if $\alpha A - \beta B + k(1) > 0$ and $\alpha A - \beta B + k(0) < 1$.

Comparative static (private).

From the solution $n_E^{\text{priv}}(s)$ above and $k'(s) < 0$:

$$\frac{\partial n_E^{\text{priv}}}{\partial s} = k'(s) < 0.$$

Larger clusters tilt inventors toward exploitation. Empirical predictions of the model follow below.

Prediction 1: Productivity, as measured by number of patents, increases with cluster size.

Proof: *Expected patent count* $P(s) = \alpha\, n_E(s) + \beta\, n_X(s) = \beta + (\alpha - \beta)\, n_E(s)$.

Differentiate: $P'(s) = (\alpha - \beta)\, n'_{\mathrm{E}}(s) = (\alpha - \beta)\, k'(s)$.

Because $\alpha - \beta < 0$ *and* $k'(s) < 0$, $P'(s) > 0$.

Prediction 2: The number of exploitation patents increases with cluster size, while the number of exploration patents decreases with cluster size.

Proof: *Immediate from* $n'_{\mathrm{E}}(s) = k'(s) < 0$ *and* $n_{X'}(s) = -n'_{\mathrm{E}}(s) > 0$.

Prediction 3: The average value of patents decreases with cluster size.

Proof: *Expected total value* $V_{\text{tot}}(s) = \alpha A\, n_E(s) + \beta B\, n_X(s) = \beta B + (\alpha A - \beta B)\, n_E(s)$.

Average value per patent $\bar{V}(s) = \frac{V_{\text{tot}}(s)}{P(s)} = \frac{\beta B + (\alpha A - \beta B)\, n_E}{\beta + (\alpha - \beta)\, n_E}$.

Differentiate using the quotient rule and insert $n'_{\mathrm{E}}(s) = k'(s) < 0$.

Collecting terms yields $\bar{V}'(s) = \frac{k'(s)\, \alpha\beta\, (A - B)}{[\beta + (\alpha - \beta)\, n_E]^2} < 0$, *because* $(A - B) > 0$ *and* $k'(s) < 0$.

II. Data

We build our analysis sample from a combination of several USPTO datasets as provided by Patentsview, Kogan, Papanikolaou, Seru, Stoffman (2017), and Marco et al. (2015). We start with all patents granted since 1976 and having historical information on NBER patent classifications until 2014. To avoid truncation from grant lags and to maintain pre-sample data for IV construction, we keep all patents with application years between 1980 and 2010 in our main analysis sample. We add main technology classes, assignees, inventors and their locations of each patent to create an inventor-patent data set. We then aggregate to the inventor-year level based on the application years of patents.

While using a larger dataset that covers more years and inventors, we follow prior approaches (Moretti 2021; Giroud, Liu, Mueller 2026; Wiebe 2026) in building the analysis sample. Like our predecessors we define technology clusters as the number of inventors within a field x city combination. We use the NBER patent classification (Marco et al. 2015) that organizes patents into six broad technological categories at the highest level, each covering a major area of technology (results are robust to finer NBER subcategory or CPC section and subsection concordances). The largest section is "Computers & Communications" which includes semiconductors, microelectronics, digital communication and information technology. We use 711 commuting zones as our city concept (Autor and Dorn 2013) as prior research on the localization of knowledge spillovers suggests that knowledge flows are facilitated by travel time (Agrawal, Galasso, Oettl 2017; results remain robust to finer geographical definitions). Further, commuting zones cover the entire US and approximate local labor markets. We define the size of a cluster as the number of inventors in each city x field, excluding the focal inventor's patents, and assign our cluster variable to each inventor-year based on the main technology class of the focal inventor's patent. If an inventor applied for more than one patent and/or in multiple cities in a given year we assign the most frequent field and city as the main cluster. In case of a tie we pick randomly. To allow absorption of unobserved heterogeneity at the inventor and firm level, we only keep inventors that patented at least twice during the sampling period and exclude lone inventors (no assignee).

To facilitate comparison with prior literature we define individual inventor productivity as the number of patents applied for by an inventor in a given year, weighted by the inverse of the number

of co-inventors on each patent Moretti (2021). A patent with one co-inventor would thus count as 0.5, a patent with two more co-inventors as 0.3 etc. To differentiate between exploration and exploitation, we split up the number of patents into those in known versus new to the inventor technologies. Known to the inventor patents are those in technology classes the inventor had previously patented in (5 years prior to the focal patent's application year). New to the inventor patents are those in technology classes in which the inventor has not previously patented (results remain robust to alternative lengths and definitions, details in the appendix). As it is ex ante not clear whether inventors' exploration and exploitation efforts are independent of each other, we complement these measures with a fractional measure of new to inventor patents, scaled by the total number of patents per year. This allows us to test if any differential in absolute terms would also be reflected in the relative focus within the exploration-exploitation continuum. As our measures of innovative search are not defined for the first patent of an inventor we drop it from the data set. Following Moretti (2021) and Giroud, Liu, Mueller (2026) we also only keep years where an inventor actually applied for a patent since reasonable assumptions remain unclear, as would the possible bias that would be created by assuming data in years without patenting activity.

To measure creativity, impact, and financial value, we add a lexical measure of patent creativity Kalyani et al. (2025). This measure identifies new technical terminology (mainly bigrams not used in any other patent up to 5 years prior) in each patents' title, abstract, and claims, to derive a measure of technological creativity (these measures are further validated as commercially successful from earnings calls). We use the standardized version so that patent creativity becomes comparable across fields and time and average the observations if an inventor has multiple patents in a given year. We also estimate the average citations per patent, a common measure of patent value (Trajtenberg, 1990), as well as a measure of financial value derived from event studies of firms' stock price reactions following patent allowance (Kogan, Papanikolaou, Seru, Stoffman 2017). Table 1 shows descriptive statistics of our main variables based on our inventor-year based analysis sample.

The sample comprises 402,360 inventors with an average individual productivity of 0.66 patents (because a patent typically has more than one inventor and we count fractional patents). Over the entire sample period inventors produce 12.1 patents on average (Median 7), with the star inventors in the top 10% of the distribution producing 25 or more patents. Most of those patents fall into

known to the inventor tech classes (average of 0.50) and many fewer into new to the inventor tech classes (average of 0.15). The maximum number of known patents belong to John C. Shi and are assigned to The Upjohn Company and the maximum number of new patents belong to Michel K. Bowman-Amua assigned to Accenture LLP (results remain almost unchanged with the removal of the top 1% of outliers). New to inventor tech class patents are mechanically rarer at the field level than at finer levels of tech classifications like NBER subsections. As we measure clusters at the field level we also use fields as the relevant categories for "new." Finer classifications at the NBER subsection level (36 unique classes) lead to more (fewer) new (known) to the inventor patents and reveal qualitatively similar results as shown below with the higher-level field-based classifications.

Table 1 – Descriptive statistics

	N	Mean	Standard deviation	50th percentile	Minimum value	Maximum value
Cluster Size	1629542	916.83	1347.85	449.00	0.00	7305.00
Individual productivity	1629542	0.64	0.99	0.33	0.01	179.50
Patents in Known Technologies	1629542	0.49	0.98	0.25	0.00	178.50
Patents in New Technologies	1629542	0.15	0.36	0.00	0.00	51.00
Fraction New Technologies	1629542	0.31	0.46	0.00	0.00	1.00
Average creativity	1625312	0.92	0.88	0.71	0.00	12.20
Average future cites	1629542	28.40	65.83	10.00	0.00	3297.00
Average financial value (KPSS)	1629542	25.47	61.51	8.50	0.00	3297.00

Notes: Analysis sample, unit of observation is an inventor-year. The sample comprises 402,360 inventors in total, observed between 1980 and 2010. Cluster size is the number of inventors within a field x city combination. We use the NBER classification system that organizes patents into five major areas of technology (results are robust to including the "other" category as a sixth field, NBER subsections, CPC sections and finer CPC subsection levels). Individual inventor productivity is the number of patents applied for by an inventor in a given year, weighted by the inverse of the number of co-inventors on each patent. Patents are split into those with known versus new to the inventor technologies. Known is any technology where the inventor had already patented in up to 5 years prior to the focal patent. New would be any patent with a technology the inventor had not patented in. Both measures weighted with the inverse of the number of co-inventors on a given patent. Fraction new is new to inventor patents scaled by total number of patents per year. As measures of innovative search are not defined for the first patent of an inventor they are dropped from the data set. Patent creativity is from Kalyani et al. (2025), which is based on a textual analysis. This measure identifies new technical terminology (mainly bigrams not used in any other patent up to 5 years prior) in each patents' title, abstract, and claims, to derive a measure of technological creativity. Shown is the standardized version so that patent creativity becomes comparable across fields and time and average per patent applied for by a given inventor in a given year.

III. Methods

Armed with these measures we estimate the following OLS equation:

$$Y_{icfkt} = \alpha_0 + \beta_1 Cluster\ size_{-ikcft} + \mu_i + \gamma_c \times \theta_t + \kappa_k \times \theta_t + \delta_f + \lambda_s + \varepsilon_{icfkt} \qquad (1)$$

where Y_{icfkt} is any of the individual productivity measures introduced above, measured for inventor *i* in city *c*, active in field *f,* working for company *k,* in year *t*. $Cluster\ size_{-ifct}$ is the number of inventors that patented in city *c* and field *f* in year *t*, excluding inventor *i*, measured in log(y+1). We estimate with log(y+1) transformed variables for easier comparison with the literature, but all results are robust to estimations with Poisson or OLS with raw counts or inverse hyperbolic sine transformations (see appendix). μ_i are inventor fixed effects that control for unobserved time invariant inventor characteristics like ability or innate curiosity about new fields. $\gamma_c \times \theta_t$ are city times year fixed effects that control for any unobserved time variant and invariant potential confounders like city size, shape, geography, access to universities etc. We further add company times year fixed effects $\kappa_k \times \theta_t$ to control any unobserved time varying or constant firm characteristics like access to other clusters Giroud, Liu, Mueller (2026) changing strategies, or firm specific productivity shocks. δ_f and λ_s are full sets of field (NBER category) and technology class (NBER subcategory) fixed effects to control for general time invariant differences across field and technologies, respectively (note that field × year and technology × year fixed effects are largely nested in firm × year fixed effects since there are only a few firms who are active in multiple fields and technologies in a given year). ε_{icfkt} is the error term. Since the focal variation in the data is at the cluster level, we treat this is as the key design-based uncertainty (Abadie et al., 2023) and estimate standard errors clustered at the city-field (i.e. cluster) level (results are robust to clustering at city and/or year level). Under the assumption that our set of fixed effects captures all unobserved confounders, β_1 will be an estimate of the marginal impact of cluster size on any of our previously introduced outcome variables.

Two main sources for potential endogeneity bias remain. First, inventors might be able to anticipate future technological developments or strategic shifts at a finer than city or firm level, enabling them to select themselves into the most promising places. Indeed, the results become

more pronounced when including inventors that switched places. Restricting the sample to non-moving inventors reveals consistent findings with similarly negative effect sizes with respect to exploration and patent creativity but generally smaller effects with respect to raw individual patent productivity. While precisely identifying the causal effect of inventor mobility lies beyond the scope of the present study, the results based on non-movers, who likely represent a negative selection, indicate that cluster effects are not solely explainable by inventor selection. Second, firms actively set their research agendas. It is worthwhile to note, however, that the most likely potential sources of endogeneity bias, e.g., anticipation or strategic shifts, is less likely to drive the negative effects on exploration, creativity, citations, and financial value, as it is hard to imagine what incentive inventors or firms would have to select themselves into more negative outcomes.

IV. Results

Table 2 shows the results of our baseline model with the key dependent search variables as described above.

Table 2 – Individual productivity and search outcomes

	(1)	(2)	(3)	(4)
	Individual productivity	Patents in Known Tech	Patents in New Tech	Fraction New Tech
Cluster Size	0.012***	0.032***	-0.021***	-0.041***
	(0.001)	(0.003)	(0.002)	(0.004)
R^2	0.577	0.565	0.446	0.527
N	1629542	1629542	1629542	1629542
Inventor FEs	yes	yes	yes	yes
Field FEs	yes	yes	yes	yes
Technology FEs	yes	yes	yes	yes
City x Year FEs	yes	yes	yes	yes
Firm x Year FEs	yes	yes	yes	yes

Notes: OLS estimation of equation (1), full analysis sample, unit of observation is an inventor-year. $Cluster\ size$ is the number of inventors that patented in city *c* and field *f* in year *t*, excluding inventor *i* and firm *k*, measured in log(y+1). We estimate with log(y+1) transformed variables, but all results are robust to estimations with Poisson or OLS with raw counts or inverse hyperbolic sine transformations. Individual inventor productivity is the number of patents applied for by an inventor in a given year, weighted by the inverse of the number of co-inventors on each patent. Patents are split into those with known versus new to the inventor technologies. Known is any technology where the inventor had already patented in up to 5 years prior to the focal patent. New would be any patent with a technology the inventor had not patented in. Both measures weighted with the inverse of the number of co-inventors on a given patent. Fraction new is new to inventor patents scaled by total number of patents per year. As measures of innovative search are not defined for the first patent of an inventor they are dropped from the data set. Standard errors clustered at the city-field (i.e. cluster) level reported in parentheses.

We find that individual inventor productivity increases with cluster size. The estimated marginal impact of Table 2, model (2) indicates that a 10 percent increase in local cluster size (~92 additional inventors for the average city) corresponds to 0.12 percent increase in individual inventor productivity. This appears to be economically relevant considering the standard deviation of 1348 inventors. When we compare our productivity estimates with the previous literature, we find our estimates to be about half as large (using the same definition as Moretti (2021) we estimate an elasticity of 0.31 versus 0.67). The smaller impact is largely driven by our inclusion of non-star inventors in the analysis sample and firm-year fixed effects to the main specification (for star-inventors and Poisson models we find larger effects as well, see appendix).

Splitting the productivity measure up, however, into exploitative and explorative productivity, reveals two underlying and opposite effects. As shown in Table 2, column b, the marginal cluster effect in known to the inventor classes is 2.7 times as large as the total productivity effect, while exploration efforts decline (column c), indicating that any positive productivity effects come exclusively from exploitation search; exploration appears to be discouraged. Tech clusters thus seem to have even larger than previously estimated productivity effect for exploitation, though it comes at the expense of less exploration, as predicted by our model. Table 2, column d, also shows a decline in the fraction of new to the inventor patents, indicating that the search focus shifts towards exploitation. Finally, these baseline results on the full sample are weaker because they are not restricted to moving inventors, star inventors, or the largest clusters in the country (results in Appendix).

We explore endogeneity concerns with a shift-share instrumental variable (SSIV) approach, following the logic of a Bartik-style design (Bartik 1991). In the original Bartik setting, local employment growth is instrumented by interacting a region's predetermined industry structure with national industry-level employment growth. Regions are predicted to grow more when they were historically specialized in industries that subsequently expanded nationally. We translate this logic to inventor clusters. Cities are predicted to experience larger clusters when they had a larger pre-sample presence in technologies that subsequently expanded globally. We implement a Bartik design with shares captured by historical exposure and measured by a pre-sample of technology-

city shares. Modern shifts are captured by non-U.S. global inventor growth and measured by technology-level shocks that are plausibly external to contemporaneous local U.S. city conditions. Therefore, we define our instrument as:

$$\widetilde{Cluster}_{cft} = \sum_{cs} \frac{I_{cs}^{1960-1979}}{I_{s}^{1960-1979}} \cdot G_{s,t-5,t,leave\ out\ US} \qquad (2)$$

where $I_{cs}^{1960-1979}$ is the total number of inventors in city *c* occurring on patents in technology class *s* between 1960 to 1979 (pre-sample), $I_{s}^{1960-1979}$ is the number of inventors in *s* in the entire U.S. between 1960 to 1979 and $G_{s,t-5,t,leave\ out\ US}$ is defined as the log of global growth in patenting inventors between year *t*-5 and year *t* – excluding U.S. inventors to avoid spurious correlation between global growth in technology *s* and local conditions (results are robust to using longer trends, earlier, shorter, and finer pre-sample shares). The expected cluster size $\widetilde{Cluster}_{ctf}$ in city *c* and field *f* at time *t* is thus constructed as the weighted average of technology-specific inventor growth rates *s* (the "shifts"), where the weights are given by the historical distribution of inventors across technologies within each city (the "shares"). Intuitively, cities with greater initial exposure to technologies that subsequently experienced stronger global expansion are predicted to experience stronger growth in clustering.

This identification strategy follows the recent SSIV literature. The key identifying assumption is not that historical technology shares are exogenous, but rather that the technology-level shocks $G_{s,t-5,t,leave\ out\ US}$ remain orthogonal to unobserved determinants of patenting in a given U.S. city. Identification is therefore obtained from global 5-year variation across more than 1500 fine grained technological areas. In this framework, the shares merely determine a city's exposure to these shocks, while causal identification comes from the exogeneity of the shift in non-US inventive activity. Because the shifts are measured at a fine level and leave out the United States, each of them is unlikely to be driven by contemporaneous local economic conditions in an individual U.S. city. While we cannot rule out potential remaining confounders, they would need to vary within the inventor, city-year, field, subfield and firm-year level. For example, it would have to take the form of a time-varying inventor-specific shock that affects both the inventor's exposure to larger clusters and the rate and direction of patents she produces, while remaining uncorrelated with common shocks at the city and firm level or technology-field and subfield trends. Such confounders appear difficult to reconcile with the rich set of fixed effects, especially when

considering the wide variety of theoretically consistent but oppositely signed outcomes, i.e., the positive effects on overall patenting, simultaneously positive and negative effects on exploitation vs. exploration, as well as the negative effects on creativity, citations, and value. Furthermore, the results are not driven by inventors active in particularly large domains or those exposed to particularly large technological shocks, as the results become even slightly stronger when excluding technologies that fall into the largest 5% of local shares or shocks (see appendix).

Applying this instrument as a predictor of cluster size in the first stage and estimating equation (1) as a second stage yields consistent results across all dependent variables. There should not be weak instrument bias given a F-value of 151.45 in the first stage. The marginal effects come out slightly larger (but still below previous studies, and with wider confidence bands), which could result from the instrument using variation in estimated cluster size from global trends, which will typically have a stronger influence for stronger clusters. The marginally larger positive effects for exploitation and larger negative effects for exploration also implies the possibility of attenuation bias in the OLS estimates.

Table 3: Search outcomes with second stage SSIV instrument

	(1)	(2)	(3)	(4)
	Individual productivity	Patents in Known Tech	Patents in New Tech	Fraction New Tech
Cluster Size (est.)	0.016***	0.035***	-0.020***	-0.040***
	(0.004)	(0.009)	(0.006)	(0.012)
N	1629542	1629542	1629542	1629542
First stage F	151.45	151.45	151.45	151.45
Inventor FEs	yes	yes	yes	yes
Field FEs	yes	yes	yes	yes
Technology FEs	yes	yes	yes	yes
City x Year FEs	yes	yes	yes	yes
Firm x Year FEs	yes	yes	yes	yes

Notes: OLS IV regression, full analysis sample, unit of observation is an inventor-year. Instrumental variable in the first as defined in equation (2). $Cluster\ size$ is the estimated number of inventors that patented in city *c* and field *f* in year *t*, excluding inventor *i*, measured in log(y+1). We estimate with log(y+1) transformed variables, but all results are robust to estimations with OLS with raw counts or inverse hyperbolic sine transformations. Individual inventor productivity is the number of patents applied for by an inventor in a given year, weighted by the inverse of the number of co-inventors on each patent. Patents are split into those with known versus new to the inventor technologies. Known is any technology where the inventor had already patented in up to 5 years prior to the focal patent. New would be any patent with a technology the inventor had not patented in. Both measures weighted with the inverse of the number of co-inventors on a given patent. Fraction new is new to inventor patents scaled by total number of patents per year. First stage F-value is the Kleibergen-Paap Wald F statistic, results shown in appendix Table 9. Standard errors clustered at the city-field (i.e. cluster) level reported in parentheses.

If narrower search leads to less creative, impactful, and financially valuable patents on average, we should observe these implications for new bigrams (Kalyani et al. 2025), future prior art citations (Trajtenberg, 1990), and stock market reactions to a patent's grant event (Kogan, Papanikolaou, Seru, Stoffman 2017). Table 4 below illustrates exactly these outcomes for OLS and IV regressions. The estimated marginal cluster impact comes out larger than the OLS estimates, which could be driven, as discussed above, by avoiding attenuation bias in the OLS estimates or the SSIV being influenced by particularly strong technologies that exert larger spillover effects.

Table 4 – OLS and IV regressions of creativity, future cites and financial value

	Panel A full sample OLS		
	(1)	(2)	(5)
	Average creativity	Average future cites	Average financial value
Cluster Size	-0.014***	-0.018***	-0.011***
	(0.002)	(0.004)	(0.002)
R^2	0.609	0.746	0.887
N	1628655	1633977	926394
Inventor FEs	yes	yes	yes
Field FEs	yes	yes	yes
Technology FEs	yes	yes	yes
City x Year FEs	yes	yes	yes
Firm x Year FEs	yes	yes	yes
	Panel B full sample SSIV		
	(1)	(2)	(3)
	Average creativity	Average future cites	Average financial value
Cluster Size (est.)	-0.017***	-0.042***	-0.031***
	(0.006)	(0.016)	(0.011)
First stage F	151.39	151.45	152.90
N	1624289	1629542	923094
Inventor FEs	yes	yes	yes
Field FEs	yes	yes	yes
Technology FEs	yes	yes	yes
City x Year FEs	yes	yes	yes
Firm x Year FEs	yes	yes	yes

Notes: OLS IV regression, full analysis sample, unit of observation is an inventor-year. Instrumental variable of the first stage as defined in equation (2). $Cluster\ size$ is the number of inventors that patented in city c and field f in year t, excluding inventor i, measured in log(y+1). We estimate with log(y+1) transformed variables, but all results are robust to estimations with Poisson or OLS with raw counts or inverse hyperbolic sine transformations. Patent creativity is from Kalyani et al. (2025), which is based on a textual analysis. This measure identifies new technical terminology (mainly bigrams not used in any other patent up to 5 years prior) in each patents' title, abstract, and claims, to derive a measure of technological creativity. Shown is the standardized version so that patent creativity becomes comparable across fields and time and average per patent applied for by a given inventor in a given year. Financial value of patents are derived from stock market reactions to patent's grant event and future cites from other patents to the focal one as provided by (Kogan, Papanikolaou, Seru, Stoffman 2017, updated version 2020). First stage F-value is the Kleibergen-Paap Wald F statistic, results shown in appendix Table 9. Standard errors clustered at the city-field (i.e. cluster) level reported in parentheses.

IV A. ***Further robustness and sensitivity checks***

In addition to the robustness checks mentioned throughout the text above, we undertake a battery of robustness and sensitivity checks that are shown in the appendix. For easier comparison with the literature, we re-run regressions only with star inventors as defined by Moretti (2021), as those that fall into the top 10% most frequently patenting inventors within the sampling period. We find more pronounced exploration vs. exploitation effects for star inventors, but relative to non-star inventors, the decrease in average patent value is not as pronounced, indicating that star inventors might take greater advantage from working in a cluster. Results remain robust to varying the definition of new versus known to the inventor technologies for 10 vs. all prior years, using just a dummy, unscaled creativity, or using CPC class, subclass, or NBER subcategory classifications. Weighing cluster size by the number of patents an individual inventor applied for in a given year reveals slightly stronger estimates. Although responsible for about 50% of U.S. patenting, results are also not driven by the largest ten cities in the U.S. (e.g., Silicon Valley) as excluding the latter or estimating just the 10 largest ten cities returns similar results. Results also hold with the exclusion of the entire "Computers & Communication" field, which was the largest of the five NBER categories and assumedly responsible for a large part of growth in the sampled period.

IV B. ***Discussion***

Here we develop policy implications theoretically and discuss potential options available to the social planner and policy maker. Building on our conceptual framework, exploration generates knowledge spillovers that are not captured by the individual inventor:

$$\lambda: [0,1] \to \mathbb{R}_+, \qquad \boxed{\lambda'(s) \gtrless 0,\ \lambda(s) > 0}.$$

In the *complementarity* case ($\lambda'(s) > 0$), the marginal social value of exploration rises with s. The denser the local ecosystem, the more potential complementors can absorb and build upon the exploratory idea.

In the *congestion* case ($\lambda'(s) < 0$), the marginal social value of exploration decreases with s. In an overly dense or "hyper-focused" cluster, the high volume of existing (mostly exploitative) work creates "noise," making it difficult for novel ideas to be noticed. Furthermore, the cluster's

specialized infrastructure and labor market may become rigid, lacking the "absorptive capacity" to build upon exploratory ideas that deviate from the local norm.

Aggregate welfare equals private earnings plus non-internalized spillovers:

$$W(n_E, s) = n_E\left(\alpha A - \frac{n_E}{2}\right) + (1 - n_E)(\beta B - k(s)) + \lambda(s)\, n_E.$$

Maximizing welfare over $n_E \in (0,1)$ yields

$$\boxed{n_E^{\text{soc}}(s) = \alpha A - \beta B + k(s) + \lambda(s).}$$

Market failure and deadweight loss

$$\boxed{\Delta(s) \;\equiv\; n_E^{\text{soc}}(s) - n_E^{\text{priv}}(s) \;=\; \lambda(s) \;>\; 0}, \qquad \frac{\partial \Delta}{\partial s} = \lambda'(s) \gtrless 0.$$

Under-investment and comparative statics on cluster size: *For every s satisfying the interior condition* $0 < c^* < 1$*, the economy under-invests in exploration. Moreover, we can break the comparative statics of the gap into two cases, between socially optimal and privately chosen exploration,* $\Delta(s) = \lambda(s)$:

- ***Complementarity case****: the gap* $\Delta(s)$ *is strictly increasing in cluster size;*
- ***Congestion case****: the gap* $\Delta(s)$ *is strictly decreasing in cluster size;*

Comparative static (planner).

$$\frac{\partial n_E^{\text{soc}}}{\partial s} \;=\; k'(s) + \lambda'(s).$$

The sign is ambiguous:

- If $|k'(s)| > \lambda'(s) \Rightarrow k'(s) + \lambda'(s) < 0 \Rightarrow$ the planner, like private agents, prefers *less* exploration in larger clusters.
- If $|k'(s)| < \lambda'(s) \Rightarrow k'(s) + \lambda'(s) > 0 \Rightarrow$ the rising social spillover outweighs cheaper exploitation, so the planner wants *more* exploration as the cluster grows while agents curtail exploration.

The spillover term $\lambda(s)$ grows or shrinks with cluster size, so any Pigouvian remedy must explicitly depend on scale. We therefore propose a menu of "first-in-class" instruments—targeted grants, milestone prizes, or refundable tax credits—whose payout rises or lowers with the local

hub's density. Operationally, let $\tau(s) = \lambda(s)/\alpha$, be the reward per successful exploratory patent; because $\lambda'(s) \gtrless 0$, $\tau(s)$ automatically increases or shrinks in larger clusters. Paying $\tau(s)$ ex-post (when the patent is granted) aligns the private return with the social externality without the over-subsidy that would occur if the full amount were paid ex-ante on every attempt (which succeed only with probability α). In practice, a hybrid scheme works best: a modest up-front grant to cover learning costs, followed by an ex-post prize or tax credit keyed to the issue of a "first-in-class" patent. By tying the payment to demonstrably cross-field breakthroughs and indexing its value to cluster size, such instruments restore the explore–exploit balance exactly where agglomeration economies bias it the most.

If economic development policy can shape the map of invention, pursuing ever-larger "superstar" hubs may not always be the winning strategy. Once a cluster crosses a critical mass, each additional inventor mostly makes the existing trees of innovation taller and slimmer, rather than planting new ones, deepening the private-social wedge that mega-clusters create. Beyond that scale, the planner's goal flips; instead of feeding the dominant hub, policy should lean against further concentration and encourage a more even spread of R&D activities. Nurturing several mid-sized, specialized hubs—or even fostering virtual collaborations that mimic their spillovers—preserves the routine cost advantages of agglomeration while still generating the cross-field experimentation that drives long-run growth. In short, regional policy needs to balance the productivity gains of density against the exploration tax imposed by mega-clusters, tilting support toward places and platforms that can deliver positive agglomeration externalities without exploitation bias.

Cluster size is obviously not assigned by a social planner; firms and inventors endogenously sort into locations. We modeled *s* as exogenous only to derive clean comparative statics—how a marginal, unanticipated change in local same-field density affects an inventor's explore-versus-exploit choice, holding her intrinsic type fixed. In the empirical analysis we replicated that thought experiment. Inventor fixed effects absorbed all time-invariant personal traits, city-year fixed effects capture local shocks that could jointly influence cluster size and patenting behavior, and a Bartik IV provided variation in same-field density that is plausibly orthogonal to any single inventor's latent propensity to explore. Because identification comes from changes in *s* that are external to the focal inventor, the estimated coefficients should correspond to the partial-

equilibrium responses predicted by the model, despite endogenous sorting in the broader spatial equilibrium.

V. Conclusion

Technology clusters are increasingly the drivers of modern innovation, yet our model predicts and our evidence reveals a hidden cost to their celebrated contributions. We replicate the result that inventors are more productive in large technology hubs but then show that this boost comes entirely from work in classes they already know. In contrast, patenting in brand-new classes, patent creativity scores, market valuations, and forward citations all decline as clusters thicken. A simple explore-exploit model provides an explanation to the pattern: local density makes routine, within-field invention cheaper, tilting private incentives toward exploitation even when exploratory search might carry greater social value.

These findings refine classic agglomeration theory (Marshall 1890). Clusters do not merely scale up knowledge production; they systematically reshape its composition, producing "taller, narrower" trees of innovation. The resulting private–social wedge may account for part of the increase in the cost of new ideas (Bloom, Jones, Van Reenen, Webb 2020), the growing spatial inequality of inventive activity, and the concentration of talent in a handful of superstar cities (Forman, Goldfarb, Greenstein 2016). Large, well-capitalized firms can partially mitigate the bias by operating multiple R&D labs, but that strategy remains costly and unavailable to smaller firms and startups.

Policy makers therefore face a trade-off. Encouraging clustering can raise near-term patent counts and possibly also local employment, yet it may also undercut the breadth of technological search that seeds future growth. Instruments that subsidize cross-field experimentation—targeted grants, prizes, or tax credits for "first-in-class" patents—may be needed to restore the balance. Equally important is the design of regional policies: supporting smaller, specialized hubs or virtual collaborations could create diverse spillovers without encouraging the exploitation bias of clusters.

The research is also consistent with a variety of arguments and results in the broader innovation literature. First, it complements theories of the "burden of knowledge" (Jones, 2009), which

emphasize rising individual specialization, and illustrate how local ecosystems could reinforce this specialization. Second, these mechanisms may reinforce declining business dynamism and the reallocation of talent (Decker, Haltiwanger, Jarmin, Miranda 2016; Akcigit & Ates, 2021), since geographically concentrated labor markets encourage inventors to double down on their established expertise rather than venturing into new technological spaces. Third, it provides a cluster and knowledge spillover example to the literature on distorted innovation (Acemoglu & Linn, 2004; Acemoglu & Akcigit, 2012; Acemoglu, Aghion, Bursztyn, Hemous, 2012; Acemoglu, 2002, 2023), highlighting another wedge between private incentives and socially optimal innovation.

Finally, our results can be situated within the broader historical literature on technological waves. Research on long-run growth emphasizes that major technological transformations—such as electricity, semiconductors, or information technology—arrive in successive waves that reshape the economy (Bresnahan & Trajtenberg, 1995; Perez, 2002; Jovanovic & Rousseau, 2005; Berkes Gaetani, Mestieri 2025). These waves are thought to emerge from clusters of inventive activity that generate complementary innovations and enable diffusion across sectors. If agglomeration systematically tilts inventors toward exploiting established domains rather than venturing into new ones, the arrival of future waves may be slowed or weakened, delaying the transition to the next technological paradigm. Our framework thus links individual search behavior in localized settings to the macro-dynamics of technological change, suggesting that cluster-driven incentives may not only reduce creative exploration but also contribute to a broader slowdown in the emergence and diffusion of transformative technologies. These arguments are consistent with the decline of Detroit as a pre-eminent technology cluster in the early to mid-20th century (though it is only one strand of a very complex story). While automakers in Detroit invested heavily to invent a variable speed windshield wiper in the 1950s, they missed the incorporation of recently invented semiconductor transistors into a simple RC (resistor capacitor) circuit (Kearns, 1964). Instead, it was a mechanical engineering faculty at Wayne State University, Robert Kearns - also familiar with electronics - who invented what became the industry standard.

Several limitations point to possibilities for future work. First, patents capture only part of inventive activity; incorporating scientific publications, open-source software, or product launches could test the generality of the mechanism. Second, our analysis focuses on inventors' knowledge choices and abstracts that knowledge from the financing environment that influences those choices. Venture-capital investors, corporate R&D budgets, and government grant programs may reward exploration and exploitation different, potentially amplifying—or dampening—the exploitation bias we uncover. Linking patent data to detailed funding information could illuminate whether financial constraints or investor preferences reinforce the tendency to stay within familiar fields and, in turn, identify levers that make exploration more attractive. Third, remote work and digital collaboration tools may alter the geography of knowledge in ways our historical data cannot observe; tracking post-pandemic cohorts could be informative. Finally, the model and empirics imply that all three Marshallian mechanisms might drive the effect; further work could tease out stronger and weaker mechanisms and estimate the magnitudes of each effect.

Innovation policy often assumes that "more clustering" provides unambiguous benefits. Our results urge a more nuanced view: clustering raises output along existing trajectories but simultaneously narrows the diversity and risky exploration that can initiate future trajectories. Ensuring that the next generation of breakthrough ideas receives sufficient exploration may require pushing some inventors—and some resources—deliberately beyond the comfortable orbit of their local technology cluster.

Bibliography

Abadie, A., Athey, S., Imbens, G. W., & Wooldridge, J. M. (2023). When should you adjust standard errors for clustering? *Quarterly Journal of Economics*, *138*(1), 1-35

Acemoglu, D. (2002). Directed technical change. *Review of Economic Studies*, 69(4), 781–809.

Acemoglu, D. (2023). Distorted innovation: Does the market get the direction of technology right? *AEA Papers and Proceedings*, 113, 1–28.

Acemoglu, D., Aghion, P., Bursztyn, L., & Hemous, D. (2012). The environment and directed technical change. *American Economic Review*, 102(1), 131–166.

Acemoglu, D., & Akcigit, U. (2012). Intellectual property rights policy, competition and innovation. *Journal of the European Economic Association*, 10(1), 1–42.

Acemoglu, D., & Linn, J. (2004). Market size in innovation: Theory and evidence from the pharmaceutical industry. *Quarterly Journal of Economics*, 119(3), 1049–1090.

Agrawal, A., Galasso, A. and Oettl, A., 2017. Roads and innovation. *Review of Economics and Statistics*, *99*(3), pp.417-434.

Akcigit, U., & Ates, S. T. (2021). Ten facts on declining business dynamism and lessons from endogenous growth theory. *American Economic Journal: Macroeconomics*, 13(1), 257–298.

Arora, A., Belenzon, S., & Patacconi, A. (2018). The decline of science in corporate R&D. *Strategic Management Journal*, 39(1), 3–32.

Autor, D. and D. Dorn. "The Growth of Low Skill Service Jobs and the Polarization of the U.S. Labor Market." American Economic Review, 103(5), 1553-1597, 2013.

Bartik, T.J., 1991. Who benefits from state and local economic development policies? *WE Upjohn Institute.*

Berkes, E., Gaetani, R., & Mestieri, M. (2025). Technological Waves, Knowledge Diffusion, and Local Growth. *Journal of Political Economy (Macroeconomics)*, 3(1), 75–121

Bloom, N., Jones, C. I., Van Reenen, J., & Webb, M. (2020). Are ideas getting harder to find? *American Economic Review*, 110(4), 1104–1144.

Bresnahan, T. F., & Trajtenberg, M. (1995). General purpose technologies: “Engines of growth”? *Journal of Econometrics*, 65(1), 83–108.

Cohen, W. M., & Levinthal, D. A. (1990). Absorptive capacity: A new perspective on learning and innovation. *Administrative Science Quarterly*, 35(1), 128–152.

Decker, R., Haltiwanger, J., Jarmin, R., & Miranda, J. (2016). Declining business dynamism: What we know and the way forward. *American Economic Review: Papers & Proceedings*, 106(5), 203–207.

Ellison, G., Glaeser, E.L. and Kerr, W.R., 2010. What causes industry agglomeration? Evidence from coagglomeration patterns. *American Economic Review*, *100*(3), pp. 1195-1213.

Forman, C., Goldfarb, A., & Greenstein, S. (2016). Agglomeration of invention in the Bay Area: Not just ICT. *American Economic Review: Papers & Proceedings*, 106(5), 146–151.

Giroud, X. E. Liu, M. Holger 2026. Innovation spillovers across U.S. tech clusters, *Journal of Financial Economics,* Volume 179, 104264.

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), pp. 536-598.

Jaffe, A. B., Trajtenberg, M., & Henderson, R. M. (1993). Geographic localization of knowledge spillovers as evidenced by patent citations. *Quarterly Journal of Economics*, 108(3), 577–598.

Jones, B. F. (2009). The burden of knowledge and the "death of the renaissance man": Is innovation getting harder? *Review of Economic Studies*, 76(1), 283–317.

Jovanovic, B., & Rousseau, P. L. (2005). General purpose technologies. In P. Aghion & S. Durlauf (Eds.), *Handbook of Economic Growth*, Vol. 1B (pp. 1181–1224). Amsterdam: Elsevier.

Kalyani, A., Bloom, N., Carvalho, M., Hassan, T. A., Lerner, J., & Tahoun, A. (2025). The diffusion of new technologies. *Quarterly Journal of Economics*, 140(2), 1299–1365.

Kearns, R. (1964) Windshield wiper system with intermittent operation. *U.S. Patent* No. 3,351,836.

Kogan, L., Papanikolaou, D., Seru, A., Stoffman, N. (2017). Technological Innovation, Resource Allocation, and Growth. *Quarterly Journal of Economics*, 132(2): 665–712.

Hill, R., Yin, Y., Stein, C., Wang, X., Wang, D. and Jones, B.F., 2025. The pivot penalty in research. *Nature*, 642(8069), pp.999-1006.

Marco, A. and Carley, M., Jackson, S. and Myers, A., 2015. The USPTO Historical Patent Data Files: Two Centuries of Innovation. SSRN working paper, available at **http://ssrn.com/abstract=2616724**.

March, J. G. (1991). Exploration and exploitation in organizational learning. *Organization Science*, 2(1), 71–87.

Marshall, A. (1890). *Principles of Economics*. London: Macmillan.

Moretti, E. (2012). *The New Geography of Jobs*. Boston: Houghton Mifflin Harcourt.

Moretti, E. (2021). The effect of high-tech clusters on the productivity of top inventors. *American Economic Review*, 111(10), 3328–3375.

Perez, C. (2002). *Technological Revolutions and Financial Capital: The Dynamics of Bubbles and Golden Ages*. Cheltenham: Edward Elgar.

Rodríguez-Pose, A. (2018). The revenge of the places that don't matter (and what to do about it). *Cambridge Journal of Regions, Economy and Society*, 11(1), 189–209.

Suss, J., Kemeny, T., & Connor, D. S. (2024). Spatial wealth inequality data for the United States, 1960–2020. *Scientific Data*, 11(1), 253.

Trajtenberg, M., 1990. A penny for your quotes: patent citations and the value of innovations. *The Rand journal of economics*, pp.172-187.

Wiebe, M., 2026. *The Effect of High-Tech Clusters on the Productivity of Top Investors: Comment* (No. 292). I4R Discussion Paper Series.

Weitzman, M.L., 1998. Recombinant growth. *The Quarterly Journal of Economics*, *113*(2), pp. 331-360.

Appendix for "Do Technology Clusters Distort Innovation?"

A1 – Sensitivity of results with respect to fixed effects

A2 – Alternative functional forms: raw counts and inverse hyperbolic sine transformation

A3 – Alternative definitions of new patents and cluster size

A4 – CPC classifications instead of NBER patent classification

A5 – Finer tech classification and sub-field clusters

A6 – Mover versus non-mover inventors

A7 – Star inventors only

A8 – Excluding the top 10 biggest cities in the U.S.

A9 – Excluding IT sector

A10 – IV First stages

A11 – IV sensitivity checks: excluding largest shocks and shares, longer shocks

A1 – Sensitivity of results with respect to fixed effects

Panel A - Individual inventor productivity					
	(1)	(2)	(3)	(4)	(5)
	Individual productivity	Individual productivity	Individual productivity	Individual productivity	Individual productivity
Cluster Size	0.014***	0.013***	0.013***	0.020***	0.011***
	(0.001)	(0.001)	(0.001)	(0.002)	(0.001)
R^2	0.483	0.483	0.483	0.492	0.578
N	1633794	1633794	1633794	1633794	1633794
Inventor + Year FEs	yes	yes	yes	yes	yes
Field FEs		yes	yes	yes	yes
Technology FEs			yes	yes	yes
City x Year FEs				yes	yes
Firm x Year FEs					yes
Panel B - Patents in known tech					
	(1)	(2)	(3)	(4)	(5)
	Patents in known tech	Patents in known tech	Patents in known tech	Patents in known tech	Patents in known tech
Cluster Size	0.024***	0.024***	0.024***	0.045***	0.031***
	(0.001)	(0.001)	(0.001)	(0.003)	(0.003)
R^2	0.470	0.470	0.470	0.481	0.567
N	1633794	1633794	1633794	1633794	1633794
Inventor + Year FEs	yes	yes	yes	yes	yes
Field FEs		yes	yes	yes	yes
Technology FEs			yes	yes	yes
City x Year FEs				yes	yes
Firm x Year FEs					yes
Panel C - Patents in new tech					
	(1)	(2)	(3)	(4)	(5)
	Patents in new tech	Patents in new tech	Patents in new tech	Patents in new tech	Patents in new tech
Cluster Size	-0.010***	-0.011***	-0.011***	-0.026***	-0.022***
	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)
R^2	0.311	0.312	0.312	0.325	0.444
N	1633794	1633794	1633794	1633794	1633794
Inventor + Year FEs	yes	yes	yes	yes	yes
Field FEs		yes	yes	yes	yes
Technology FEs			yes	yes	yes
City x Year FEs				yes	yes
Firm x Year FEs					yes

Notes: OLS estimations of equation (1), full analysis sample, unit of observation is an inventor-year. $Cluster\ size$ is the number of inventors that patented in city *c* and field *f* in year *t*, excluding inventor *i*, measured in log(y+1). Individual inventor productivity is the number of patents applied for by an inventor in a given year, weighted by the inverse of the number of co-inventors on each patent. Patents are split into those with known versus new to the inventor technologies. Known is any technology where the inventor had already patented in the 5 previous years prior to the focal patent's application year. New would be any patent with a technology the inventor had not patented in. Both measures weighted with the inverse of the number of co-inventors on a given patent. Standard errors clustered at the city-field (i.e. cluster) level reported in parentheses.

A2 – Alternative functional forms: raw counts, inverse hyperbolic sine transformation, and Poisson

	OLS raw counts		**OLS IHS transformed**		**Poisson**	
	(1)	(2)	(3)	(4)	(5)	(6)
	Patents in known tech	Patents in new tech	Patents in known tech	Patents in new tech	Patents in known tech	Patents in new tech
Cluster Size	0.057***	-0.028***	0.038***	-0.025***	0.120***	-0.239***
	(0.006)	(0.003)	(0.003)	(0.002)	(0.010)	(0.027)
R^2	0.522	0.402	0.563	0.438		
N	1629542	1629542	1629542	1629542	1496754	1288775
Inventor FEs	yes	yes	yes	yes	yes	yes
Field FEs	yes	yes	yes	yes	yes	yes
Technology FEs	yes	yes	yes	yes	yes	yes
City x Year FEs	yes	yes	yes	yes	yes	yes
Firm x Year FEs	yes	yes	yes	yes	yes	yes

Notes: OLS estimations of equation (1), full analysis sample, unit of observation is an inventor-year. $Cluster\ size$ is the number of inventors that patented in city *c* and field *f* in year *t*, excluding inventor *i*, measured as a count (models 1,2,5,6) or the inverse hyperbolic sine of $Cluster\ size$, (models 3 and 4). Patents are split into those with known versus new to the inventor technologies. Known is any technology where the inventor had already patented in the 5 previous years prior to the focal patent's application year. New would be any patent with a technology the inventor had not patented in. Both measures weighted with the inverse of the number of co-inventors on a given patent. Models 3 and 4 take the inverse hyperbolic sine of known and new tech patents. Standard errors clustered at the city-field (i.e. cluster) level reported in parentheses.

A3 – Alternative definitions of new patents and cluster size

Panel A - Alternative definitions of new patents and creativity					
	(1)	(2)	(3)	(4)	(5)
	Patents in new tech > 0, dummy	Patents in new tech (10yr window)	Patents in new tech (ever)	Patents in new tech (excluding 1st)	Creativity w/o normalization
Cluster Size	-0.061***	-0.020***	-0.021***	-0.024***	-0.003***
	(0.005)	(0.002)	(0.002)	(0.002)	(0.000)
R^2	0.507	0.439	0.453	0.402	0.624
N	1629590	1629590	1629590	1130222	1624311
Inventor FEs	yes	yes	yes	yes	yes
Field FEs	yes	yes	yes	yes	yes
Technology FEs	yes	yes	yes	yes	yes
City x Year FEs	yes	yes	yes	yes	yes
Firm x Year FEs	yes	yes	yes	yes	yes

Notes: OLS estimations of equation (1), full analysis sample, unit of observation is an inventor-year. $Cluster\ size$ is the number of inventors that patented in city *c* and field *f* in year *t*, excluding inventor *i*, measured in log(y+1). Known patents are in any technology where the inventor had already patented in the 5 previous years prior to the focal patent's application year. We vary this by considering only a dummy indicating any positive number of new patents (1), the number new patents but taking the 10 previous years of individual patenting into account (2), the number new patents but taking all available previous years of individual patenting into account (3), excluding the first patent of a given inventor which is otherwise always new (4), and taking the creativity measure of Kalyani (2025) without normalization. Standard errors clustered at the city-field (i.e. cluster) level reported in parentheses.

Panel B - Cluster size weighted by individual inventor productivity					
	(1)	(2)	(3)	(4)	(5)
	Patents in known tech	Patents in new tech	Average creativity	Average future cites	Average financial value
Cluster Size	0.037***	-0.017***	-0.012***	-0.014***	-0.010***
	(0.003)	(0.002)	(0.002)	(0.005)	(0.002)
R^2	0.566	0.446	0.609	0.746	0.889
N	1629590	1629590	1624311	1629590	923199
Inventor FEs	yes	yes	yes	yes	yes
Field FEs	yes	yes	yes	yes	yes
Technology FEs	yes	yes	yes	yes	yes
City x Year FEs	yes	yes	yes	yes	yes
Firm x Year FEs	yes	yes	yes	yes	yes

Notes: OLS estimations of equation (1), full analysis sample, unit of observation is an inventor-year. $Cluster\ size$ is the number of inventors that patented in city *c* and field *f* in year *t*, excluding inventor *i*, either weighted by the number of patents spawn out each year (Panel B), measured in log(y+1). Individual inventor productivity is the number of patents applied for by an inventor in a given year, weighted by the inverse of the number of co-inventors on each patent. Patents are split into those with known versus new to the inventor technologies. Known is any technology where the inventor had already patented in the 5 previous years prior to the focal patent's application year. New would be any patent with a technology the inventor had not patented in. Both measures weighted with the inverse of the number of co-inventors on a given patent. Patent creativity is from Kalyani (2025) and Kalyani et al. (2025), which is based on a textual analysis. This measure identifies new technical terminology (mainly bigrams not used in any other patent up to 5 years prior) in each patents' title, abstract, and claims, to derive a measure of technological creativity. Shown is the standardized version so that patent creativity becomes comparable across fields and time and average per patent applied for by a given inventor in a given year. Financial value of patents are derived from stock market reactions to patent's grant event as provided by (Kogan et al. 2017, updated version 2020). Future cites from other patents to the focal one come from the USPTO Patentsview database. Standard errors clustered at the city-field (i.e. cluster) level reported in parentheses.

A4 – CPC classifications instead of NBER patent classification

	(1)	(2)	(3)	(4)
	Individual productivity	Patents in Known Tech	Patents in New Tech	Fraction New Tech
Cluster Size	0.005***	0.021***	-0.019***	-0.033***
	(0.001)	(0.002)	(0.002)	(0.003)
R^2	0.577	0.565	0.444	0.525
N	1629295	1629295	1629295	1629295
Inventor FEs	yes	yes	yes	yes
Field FEs	yes	yes	yes	yes
Technology FEs	yes	yes	yes	yes
City x Year FEs	yes	yes	yes	yes
Firm x Year FEs	yes	yes	yes	yes

Notes: OLS estimation of equation (1), full analysis sample, unit of observation is an inventor-year. $Cluster\ size$ is the number of inventors that patented in city *c* and field *f* in year *t*, excluding inventor *i* and firm *k*, measured in log(y+1). Field is defined as a CPC section classification instead of NBER category. We estimate with log(y+1) transformed variables, but all results are robust to estimations with Poisson or OLS with raw counts or inverse hyperbolic sine transformations. Individual inventor productivity is the number of patents applied for by an inventor in a given year, weighted by the inverse of the number of co-inventors on each patent. Patents are split into those with known versus new to the inventor technologies based on the first CPC section classification (8 sections in total). Known is any technology where the inventor had already patented in up to 5 years prior to the focal patent. New would be any patent with a technology the inventor had not patented in. Both measures weighted with the inverse of the number of co-inventors on a given patent. Fraction new is new to inventor patents scaled by total number of patents per year. As measures of innovative search are not defined for the first patent of an inventor they are dropped from the data set. Standard errors clustered at the city-field (i.e. cluster) level reported in parentheses.

A5 – Finer NBER tech classification and NBER sub-field clusters

	(1)	(2)	(3)	(4)
	Individual productivity	Patents in Known Tech	Patents in New Tech	Fraction New Tech
Cluster Size	-0.002***	0.002***	-0.005***	-0.006***
	(0.000)	(0.001)	(0.001)	(0.001)
R^2	0.577	0.565	0.445	0.540
N	1629542	1629542	1629542	1629542
Inventor FEs	yes	yes	yes	yes
Field FEs	yes	yes	yes	yes
Technology FEs	yes	yes	yes	yes
City x Year FEs	yes	yes	yes	yes
Firm x Year FEs	yes	yes	yes	yes

Notes: OLS estimation of equation (1), full analysis sample, unit of observation is an inventor-year. $Cluster\ size$ is the number of inventors that patented in city *c* and field *f* in year *t*, excluding inventor *i* and firm *k*, measured in log(y+1). Field is defined as a NBER subcategory. We estimate with log(y+1) transformed variables, but all results are robust to estimations with Poisson or OLS with raw counts or inverse hyperbolic sine transformations. Individual inventor productivity is the number of patents applied for by an inventor in a given year, weighted by the inverse of the number of co-inventors on each patent. Patents are split into those with known versus new to the inventor technologies based on the NBER subcategory. Known is any technology where the inventor had already patented in up to 5 years prior to the focal patent. New would be any patent with a technology the inventor had not patented in. Both measures weighted with the inverse of the number of co-inventors on a given patent. Fraction new is new to inventor patents scaled by total number of patents per year. As measures of innovative search are not defined for the first patent of an inventor they are dropped from the data set. Standard errors clustered at the city-field (i.e. cluster) level reported in parentheses.

Table A6 – Mover versus non-mover inventors

Panel A – Inventors that moved					
	(1)	(2)	(3)	(4)	(5)
	Patents in known tech	Patents in new tech	Average creativity	Average future cites	Average financial value
Cluster Size	0.028***	-0.016***	-0.012***	-0.024***	-0.011***
	(0.003)	(0.002)	(0.002)	(0.006)	(0.003)
R^2	0.559	0.443	0.571	0.712	0.885
N	438928	438928	437580	438928	261554
Inventor FEs	yes	yes	yes	yes	yes
Field FEs	yes	yes	yes	yes	yes
Technology FEs	yes	yes	yes	yes	yes
City x Year FEs	yes	yes	yes	yes	yes
Firm x Year FEs	yes	yes	yes	yes	yes
Panel B - Inventors that never moved					
	(1)	(2)	(3)	(4)	(5)
	Patents in known tech	Patents in new tech	Average creativity	Average future cites	Average financial value
Cluster Size	0.038***	-0.028***	-0.019***	-0.020***	-0.012***
	(0.003)	(0.003)	(0.002)	(0.007)	(0.003)
R^2	0.590	0.471	0.630	0.763	0.893
N	1102129	1102129	1098424	1102129	643290
Inventor FEs	yes	yes	yes	yes	yes
Field FEs	yes	yes	yes	yes	yes
Technology FEs	yes	yes	yes	yes	yes
City x Year FEs	yes	yes	yes	yes	yes
Firm x Year FEs	yes	yes	yes	yes	yes

Notes: OLS estimation of equation (1). Panel A: Movers are inventors that applied for patents in at least two different commuting zones. Panel B: Non-movers are inventors that have patented in only one commuting zone. $Cluster\ size$ is the number of inventors that patented in city *c* and field *f* in year *t*, excluding inventor *i*, measured in log(y+1). Individual inventor productivity is the number of patents applied for by an inventor in a given year, weighted by the inverse of the number of co-inventors on each patent. Patents are split into those with known versus new to the inventor technologies. Known is any technology where the inventor had already patented in up to 5 years prior to the focal patent. New would be any patent with a technology the inventor had not patented in. Both measures weighted with the inverse of the number of co-inventors on a given patent. Fraction new is new to inventor patents scaled by total number of patents per year. Patent creativity is from Kalyani (2025) and Kalyani et al. (2025), which is based on a textual analysis. This measure identifies new technical terminology (mainly bigrams not used in any other patent up to 5 years prior) in each patents' title, abstract, and claims, to derive a measure of technological creativity. Shown is the standardized version so that patent creativity becomes comparable across fields and time and average per patent applied for by a given inventor in a given year. Financial value of patents are derived from stock market reactions to patent's grant event as provided by (Kogan et al. 2017, updated version 2020). Future cites from other patents to the focal one come from the USPTO Patentsview database. Standard errors clustered at the city-field (i.e. cluster) level reported in parentheses.

A7 – Star inventors only

Panel A – Inventors that patented at least 10 times					
	(1)	(2)	(3)	(4)	(5)
	Patents in known tech	Patents in new tech	Average creativity	Average future cites	Average financial value
Cluster Size	0.042***	-0.025***	-0.015***	-0.023***	-0.008***
	(0.004)	(0.002)	(0.002)	(0.006)	(0.003)
R^2	0.498	0.364	0.530	0.694	0.881
N	749414	749414	747487	749414	490286
Inventor FEs	yes	yes	yes	yes	yes
Field FEs	yes	yes	yes	yes	yes
Technology FEs	yes	yes	yes	yes	yes
City x Year FEs	yes	yes	yes	yes	yes
Firm x Year FEs	yes	yes	yes	yes	yes
Panel B - Inventors that patented at least 25 times					
	(1)	(2)	(3)	(4)	(5)
	Patents in known tech	Patents in new tech	Average creativity	Average future cites	Average financial value
Cluster Size	0.054***	-0.024***	-0.018***	-0.023***	-0.008**
	(0.006)	(0.003)	(0.003)	(0.008)	(0.004)
R^2	0.501	0.335	0.528	0.700	0.888
N	273612	273612	273051	273612	194686
Inventor FEs	yes	yes	yes	yes	yes
Field FEs	yes	yes	yes	yes	yes
Technology FEs	yes	yes	yes	yes	yes
City x Year FEs	yes	yes	yes	yes	yes
Firm x Year FEs	yes	yes	yes	yes	yes

Notes: OLS estimation of equation (1). Panel A: Inventors that applied for at least 10 patents during the sampling period 1980 to 2010. Panel B: Inventors that applied for at least 25 patents during the sampling period 1980 to 2010. $Cluster\ size$ is the number of inventors that patented in city *c* and field *f* in year *t*, excluding inventor *i*, measured in log(y+1). Individual inventor productivity is the number of patents applied for by an inventor in a given year, weighted by the inverse of the number of co-inventors on each patent. Patents are split into those with known versus new to the inventor technologies. Known is any technology where the inventor had already patented in up to 5 years prior to the focal patent. New would be any patent with a technology the inventor had not patented in. Both measures weighted with the inverse of the number of co-inventors on a given patent. Fraction new is new to inventor patents scaled by total number of patents per year. Patent creativity is from Kalyani (2025) and Kalyani et al. (2025), which is based on a textual analysis. This measure identifies new technical terminology (mainly bigrams not used in any other patent up to 5 years prior) in each patents' title, abstract, and claims, to derive a measure of technological creativity. Shown is the standardized version so that patent creativity becomes comparable across fields and time and average per patent applied for by a given inventor in a given year. Financial value of patents are derived from stock market reactions to patent's grant event as provided by (Kogan et al. 2017, updated version 2020). Future cites from other patents to the focal one come from the USPTO Patentsview database. Standard errors clustered at the city-field (i.e. cluster) level reported in parentheses.

Table A8 – Excluding the top 10 biggest cities in the U.S.

Panel A – Excluding the top 10 biggest cities in the U.S.					
	(1)	(2)	(3)	(4)	(5)
	Patents in known tech	Patents in new tech	Average creativity	Average future cites	Average financial value
Cluster Size	0.034***	-0.023***	-0.012***	-0.015***	-0.014***
	(0.003)	(0.002)	(0.002)	(0.005)	(0.003)
R^2	0.585	0.468	0.625	0.756	0.891
N	903332	903332	899651	903332	507259
Inventor FEs	yes	yes	yes	yes	yes
Field FEs	yes	yes	yes	yes	yes
Technology FEs	yes	yes	yes	yes	yes
City x Year FEs	yes	yes	yes	yes	yes
Firm x Year FEs	yes	yes	yes	yes	yes
Panel B - Top 10 biggest cities in the U.S. only					
	(1)	(2)	(3)	(4)	(5)
	Patents in known tech	Patents in new tech	Average creativity	Average future cites	Average financial value
Cluster Size	0.030***	-0.023***	-0.021***	-0.034***	-0.003
	(0.007)	(0.006)	(0.004)	(0.011)	(0.005)
R^2	0.581	0.453	0.623	0.758	0.894
N	651695	651695	650320	651695	388686
Inventor FEs	yes	yes	yes	yes	yes
Field FEs	yes	yes	yes	yes	yes
Technology FEs	yes	yes	yes	yes	yes
City x Year FEs	yes	yes	yes	yes	yes
Firm x Year FEs	yes	yes	yes	yes	yes

Notes: OLS estimation of equation (1). Panel A: the largest 10 commuting zones as measure by total patent counts excluded. Panel B: only inventors from the largest 10 commuting zones as measure by total patent counts. $Cluster\ size$ is the number of inventors that patented in city *c* and field *f* in year *t*, excluding inventor *i*, measured in log(y+1). Individual inventor productivity is the number of patents applied for by an inventor in a given year, weighted by the inverse of the number of co-inventors on each patent. Patents are split into those with known versus new to the inventor technologies. Known is any technology where the inventor had already patented in up to 5 years prior to the focal patent. New would be any patent with a technology the inventor had not patented in. Both measures weighted with the inverse of the number of co-inventors on a given patent. Fraction new is new to inventor patents scaled by total number of patents per year. Patent creativity is from Kalyani (2025) and Kalyani et al. (2025), which is based on a textual analysis. This measure identifies new technical terminology (mainly bigrams not used in any other patent up to 5 years prior) in each patents' title, abstract, and claims, to derive a measure of technological creativity. Shown is the standardized version so that patent creativity becomes comparable across fields and time and average per patent applied for by a given inventor in a given year. Financial value of patents are derived from stock market reactions to patent's grant event as provided by (Kogan et al. 2017, updated version 2020). Future cites from other patents to the focal one come from the USPTO Patentsview database. Standard errors clustered at the city-field (i.e. cluster) level reported in parentheses.

A9 – Excluding IT sector

	(1)	(2)	(3)	(4)	(5)
	Patents in known tech	Patents in new tech	Average creativity	Average future cites	Average financial value
Cluster Size	0.031***	-0.021***	-0.010***	-0.006	-0.004
	(0.003)	(0.002)	(0.002)	(0.007)	(0.003)
R^2	0.585	0.474	0.614	0.756	0.898
N	1097565	1097565	1093043	1097565	563972
Inventor FEs	yes	yes	yes	yes	yes
Field FEs	yes	yes	yes	yes	yes
Technology FEs	yes	yes	yes	yes	yes
City x Year FEs	yes	yes	yes	yes	yes
Firm x Year FEs	yes	yes	yes	yes	yes

Notes: OLS estimation of equation (1). Sample as described in section 3 but without NBER category "Computers & Communication". $Cluster\ size$ is the number of inventors that patented in city *c* and field *f* in year *t*, excluding inventor *i*, measured in log(y+1). Individual inventor productivity is the number of patents applied for by an inventor in a given year, weighted by the inverse of the number of co-inventors on each patent. Patents are split into those with known versus new to the inventor technologies. Known is any technology where the inventor had already patented in up to 5 years prior to the focal patent. New would be any patent with a technology the inventor had not patented in. Both measures weighted with the inverse of the number of co-inventors on a given patent. Fraction new is new to inventor patents scaled by total number of patents per year. Patent creativity is from Kalyani (2025) and Kalyani et al. (2025), which is based on a textual analysis. This measure identifies new technical terminology (mainly bigrams not used in any other patent up to 5 years prior) in each patents' title, abstract, and claims, to derive a measure of technological creativity. Shown is the standardized version so that patent creativity becomes comparable across fields and time and average per patent applied for by a given inventor in a given year. Financial value of patents are derived from stock market reactions to patent's grant event as provided by (Kogan et al. 2017, updated version 2020). Future cites from other patents to the focal one come from the USPTO Patentsview database. Standard errors clustered at the city-field (i.e. cluster) level reported in parentheses.

A10 – IV first stages

	(1)	(2)	(3)
	Cluster Size	Cluster Size	Cluster Size
IV	0.261***	0.261***	0.264***
	(0.021)	(0.021)	(0.021)
F-value	151.45	151.39	152.90
N	1629542	1624289	923094
Inventor FEs	yes	yes	yes
Field FEs	yes	yes	yes
Technology FEs	yes	yes	yes
City x Year FEs	yes	yes	yes
Firm x Year FEs	yes	yes	yes

Notes: This table reports the first stages of the IV regressions presented in Tables 3 and 4. Column (1) is the first stage for Table 3, columns (1)-(4) and Table 4, Panel B, column (2). Column (2) is the first stage for Table 4, Panel B, column (1) and Column (3) is the first stage for Table 4, Panel B, column (3). $Cluster\ size$ is the number of inventors that patented in city *c* and field *f* in year *t*, excluding inventor *i* and firm *k*, measured in log(y+1). We estimate with log(y+1) transformed variables, but all results are robust to estimations with Poisson or OLS with raw counts or inverse hyperbolic sine transformations. F-value is the Kleibergen-Paap Wald F statistic. Standard errors clustered at the city-field (i.e. cluster) level reported in parentheses.

A11 – IV sensitivity checks: excluding largest shocks and shares, longer shocks

Panel A – Excluding shocks < 5% and > 95% percentile					
	(1)	(2)	(3)	(4)	(5)
	Patents in known tech	Patents in new tech	Average creativity	Average future cites	Average financial value
Cluster Size (est.)	0.036***	-0.020***	-0.017***	-0.043***	-0.032***
	(0.009)	(0.006)	(0.006)	(0.015)	(0.011)
First stage F-value	134.09	134.09	134.14	134.09	131.47
N	1527459	1527459	1522395	1527459	868689
Inventor FEs	yes	yes	yes	yes	yes
Field FEs	yes	yes	yes	yes	yes
Technology FEs	yes	yes	yes	yes	yes
City x Year FEs	yes	yes	yes	yes	yes
Firm x Year FEs	yes	yes	yes	yes	yes
Panel B – Excluding shares > 95% percentile					
	(1)	(2)	(3)	(4)	(5)
	Patents in known tech	Patents in new tech	Average creativity	Average future cites	Average financial value
Cluster Size (est.)	0.036***	-0.020***	-0.020***	-0.047***	-0.030***
	(0.009)	(0.006)	(0.006)	(0.016)	(0.011)
First stage F-value	177.26	177.26	177.12	177.26	190.39
N	1202180	1202180	1199115	1202180	696665
Inventor FEs	yes	yes	yes	yes	yes
Field FEs	yes	yes	yes	yes	yes
Technology FEs	yes	yes	yes	yes	yes
City x Year FEs	yes	yes	yes	yes	yes
Firm x Year FEs	yes	yes	yes	yes	yes

Notes: IV regression as defined in section 4 and Instrument for cluster size as defined in equation (2). $Cluster\ size$ is the number of inventors that patented in city *c* and field *f* in year *t*, excluding inventor *i*, measured in log(y+1). Panel A excludes all shocks as defined in equation 2 that fall below the 5% percentile or exceed the 95% percentile. Panel B excludes all shares as defined in equation 2 that exceed the 95% percentile. Individual inventor productivity is the number of patents applied for by an inventor in a given year, weighted by the inverse of the number of co-inventors on each patent. Patents are split into those with known versus new to the inventor technologies. Known is any technology where the inventor had already patented in up to 5 years prior to the focal patent. New would be any patent with a technology the inventor had not patented in. Both measures weighted with the inverse of the number of co-inventors on a given patent. Fraction new is new to inventor patents scaled by total number of patents per year. Patent creativity is from Kalyani (2025) and Kalyani et al. (2025), which is based on a textual analysis. This measure identifies new technical terminology (mainly bigrams not used in any other patent up to 5 years prior) in each patents' title, abstract, and claims, to derive a measure of technological creativity. Shown is the standardized version so that patent creativity becomes comparable across fields and time and average per patent applied for by a given inventor in a given year. Financial value of patents are derived from stock market reactions to patent's grant event as provided by (Kogan et al. 2017, updated version 2020). Future cites from other patents to the focal one come from the USPTO Patentsview database. First stage F-value is the Kleibergen-Paap Wald F statistic. Standard errors clustered at the city-field (i.e. cluster) level reported in parentheses.