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

November, 2025

Abstract: While the geographical agglomeration of inventors has been theorized and shown to increase raw patent productivity, we model and illustrate how it also reduces exploration of new technologies. Clusters enhance individual incentives to exploit because of knowledge sharing, specialized invention infrastructure, and career security in a local labor market that will value the inventor's 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. Highly saturated fixed effects and IV regressions with U.S. patent data from 1976-2016 support prior results in the literature as well as the model's predictions of increased exploitation and decreased exploration. Consistent with narrower search patterns, we also find that, on average, patents originating from more concentrated clusters exhibit lower creativity, reduced financial value, and fewer future citations.

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

---

Currently there is no funding support for this research. The authors thank future readers for their feedback. Errors and omissions remain ours.

1 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 et al. 2010; Ellison et al. 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 unequal over at least the last 40 years (Suss et al. 2024; Forman et al. 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 causal evidence confirms that inventors who reside in, or move into, large clusters produce more patents than comparable peers elsewhere (Moretti 2021). Replicating that result with the universe of U.S. patents, we decompose the 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 all of the cluster-driven productivity gain comes from the first category. In contrast, proximity to same-field peers significantly reduces the probability that an inventor enters a new technology class. The novelty, citation impact, and average value of inventions also decreases within clusters.

We begin by formalizing the intuition in an explore–exploit model (March, 1991; Manso, 2011). At the start of each period an inventor has previously invented in what should now be familiar technology classes. 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. Cluster size enters the model in two ways. First, routine work becomes cheaper in thicker clusters of similar inventors—shared suppliers, deeper labor pools, and thicker knowledge sharing within the cluster's field drive the marginal cost of creating an incremental patent down, however that cost reduction applies mainly to same-field exploitation. Second, the spillover from a breakthrough in a brand-new class either rises with cluster size due to local complementarities or decreases with cluster size due to congestion and 'cognitive lock-in'. As such, the social return to

exploration may rise or decrease with cluster size. The private inventor therefore leans ever more heavily toward exploitation as the cluster grows, even as society's benefit from exploration may rise.

The model delivers three testable predictions that guide the empirics. First, individual patent counts increase with cluster size because cheaper routine costs make exploitation more attractive. Second, the increase in cluster creativity 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—proxied by textual novelty, grant-date market value, or forward citations—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. The fraction and absolute count of new-to-inventor patents fall with cluster concentration. Creativity scores, financial valuations, and forward-citation rates per patent all decline, producing a concave relationship between cluster concentration and inventive productivity: more trees are planted, but each tree is narrower. The pattern is driven almost entirely by (i) the top decile of inventors and (ii) the selective inflow of newcomers into existing hubs, not by incumbent non-movers suddenly becoming more prolific.

Our findings suggest a novel distortion in the process of recombinant search (Weitzmann 1998; Fleming, 2001; 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 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 (Balsmeier et al., 2023)—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 locally popular skills. 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 knowledge access (Giroud et al., 2025). 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, generating 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 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 in the United States and argues that "ideas are getting harder to find" (Bloom et al., 2020). Furthermore, it may also speak to the arguments that the rate of transformative, breakthrough innovation has slowed relative to incremental improvements (Jones, 2009; Arora, Belenzon, & Patacconi, 2019). 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 2 presents the formal model. Section 3 describes data construction and measurement. Section 4 documents the methodology; Section 5 presents empirical results, and Section 6 concludes.

2 Theory

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 novelty, whereas exploration leads to decreased productivity and greater novelty (Manso, 2011; Arts and Fleming 2018; Wang et al. 2025). Clusters (i.e., when geographically proximate inventors work in similar technologies) lower an individual inventor's search costs for exploitation because they facilitate knowledge sharing and provide specialized invention infrastructure. They also increase the incentives to exploit because they provide career security due to a local labor market that will value the inventor's 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.

Agents, actions, private pay-offs, and knowledge spillovers

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 already 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 subclasses are cited more frequently, command higher renewal fees and generate larger stock-market responses than incremental patents, indicating higher social and private value per successful filing (Kolyani et al. 2025). At the same time, incremental projects enjoy shorter development cycles, higher grant probabilities and lower abandonment rates 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 ($\beta > \alpha$) 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.

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.

Proposition 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* $k'(s) < 0$, $P'(s) > 0$.

Proposition 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* $n_{X'}(s) = -n'_{\mathrm{E}}(s) > 0$.

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

Social planner

Exploration generates knowledge spillovers $\lambda(s)$ that are not captured by the individual inventor:

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

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 "hyperfocused" 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

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

Proposition 4 (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 on the gap between socially optimal and privately chosen exploration,* $\Delta(s) = \lambda(s)$*, we can break the analysis into two cases:*

- ***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.

Policy implication

The spillover term $\lambda(s)$ grows or shrinks with cluster size, so any Pigouvian remedy must be explicitly scale-dependent. 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 is not always 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 activity. 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 knowledge spillovers without the full exploitation bias.

Cluster size is obviously not assigned by a social planner; firms and inventors endogenously sort into locations. We model 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 replicate precisely that thought experiment. Inventor fixed effects absorb all time-invariant personal traits, city-year

fixed effects capture local shocks that could jointly influence cluster size and patenting behavior, and an IV based on expansions and contractions of multi-location firms provides 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 correspond to the partial-equilibrium responses predicted by the model despite endogenous sorting in the broader spatial equilibrium.

3 Data

We build our analysis sample from a combination of several USPTO patent datasets as provided by Patentsview. We start with all patents granted between 1976 and 2024. To avoid truncation from grant lags or influence from the pandemic we keep all patents with application years between 1977 and 2019. 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 closely follow Moretti (2021) and Giroud et al. (2025) 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 Cooperative Patent Classification (CPC) system that organizes patents into 9 broad sections, each covering a major area of technology (results are robust to finer CPC subsection levels). The largest section is "Electricity" which includes the booming fields of our time, including semiconductors, microelectronics, digital communication and information technology. We use commuting zones as our city concept (Autor et al., 2013) as prior research on localization of knowledge spillovers suggests that knowledge flows are facilitated by travel time (Agrawal et al. 2017). Further, commuting zones cover the entire US and resemble local labor markets, which creates incentives for inventors to work in fields that are most popular in that area. In total we observe 711 different commuting zones, which we will refer to as cities (results are robust to finer and broader geographic concepts). Following Moretti (2021) we define the size of a cluster as the number of inventors in each city x field of the most popular tech class, excluding the patents of the focal inventor. Like the previous literature we 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 in a given year we assign the most frequent tech class as the main class for the analysis sample. In case of a tie we pick randomly.

We are interested in five main dependent variables. For easier comparison with the literature we start with the same definition as in Moretti (2021) and 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. A patent with one co-inventor would thus count as 0.5, a patent with two more co-inventors as 0.3 etc.

In order to differentiate between exploration and exploitation, we split up the number of patents into those with known versus new to the inventor technologies. Known would be any technology where the inventor had already patented in up to 10 years prior to the focal patent. New would be any patent with a technology the inventor had not patented in during the last 10 years. Again, we weigh both measures with the inverse of the number of co-inventors on a given patent to derive our measures of innovative search, which could also be looked at as measures of individual exploitation (patents in known classes) and exploration (patents in new classes) productivity. As it is ex ante not clear whether inventors' exploration and exploitation efforts are independent or influence each other, we complement these measures with a fractional measure of new to inventor patents scaled by total number of patents per year. It 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 et al. (2025) we also only keep years where an inventor actually applied for a patent since it is unclear what to assume otherwise and what kind of bias would be created by filling up data in years without patenting activity.

To avoid reliance only on technological classifications and their categorizations by the USPTO we add a measure of patent creativity from Kalyani (2025) and Kalyani et al. (2025), which is based only 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. We use 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 as our measure of creativity (results are robust to alternative versions as provided by Kalyani). Table 1 shows descriptive statistics of our main variables based on our inventor-year based analysis sample.

The sample comprises 469,232 inventors with an average individual productivity of 0.77 patents (because a patent typically has more than one inventor and we count fractional). 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 0.69) and many fewer into new to the inventor tech classes (average of 0.08). New to inventor tech class patents are mechanically rarer at the field level than at finer levels of tech classifications like CPC subsections. As we measure clusters at the field level we also use fields as the relevant categories for "new." Finer classifications at the CPC subsection level (123 unique classes) lead to more (fewer) new (known) to the inventor patents and reveal qualitatively similar results as shown below with our field-based classifications.

Table 1 – Descriptive statistics

	N	Mean	Standard deviation	50th percentile	Minimum value	Maximum value
Cluster Size	2334495	2279.77	3738.43	843.00	0.00	20622.00
Individual productivity	2334495	0.77	1.11	0.50	0.01	110.23
Patents in Known Technologies	2334495	0.69	1.10	0.40	0.00	108.23
Patents in New Technologies	2334495	0.08	0.20	0.00	0.00	5.00
Fraction New Technologies	2334495	0.14	0.32	0.00	0.00	1.00
Average creativity	2334495	0.29	0.40	0.16	0.00	11.48

Notes: Analysis sample, unit of observation is an inventor-year. The sample comprises 469,232 inventors in total, observed between 1976 and 2016. Cluster size is the number of inventors within a field x city combination. We use the Cooperative Patent Classification (CPC) system that organizes patents into 9 broad sections, each covering a major area of technology (results are robust to 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 10 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 (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.

4 Methods

Armed with these measures we use OLS to estimate the following equation which closely resembles Moretti's (2021) main model:

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

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, but all results are robust to estimations with Poisson or OLS with raw counts or inverse hyperbolic sine transformations. μ_i are inventor fixed effects that control for unobserved time invariant inventor characteristics like ability. $\gamma_c \times \theta_t$ 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 at al. 2025), changing strategies, or firm specific productivity shocks. δ_f and λ_s are full sets of field (CPC section) and technology class (CPC subsection) fixed effects to control for general time invariant differences across field and technologies, respectively. ε_{icfkt} is the error term. 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 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 moves 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. It is also worthwhile to note 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, financial value and creativity, as it is hard to imagine what incentive inventors or firms would have to select themselves into more negative outcomes.

5 Results

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

Table 2 – Search outcomes

	(1)	(2)	(3)	(4)
	Individual productivity	Patents in Known Tech	Patents in New Tech	Fraction New Tech
Cluster Size	0.010***	0.027***	-0.018***	-0.033***
	(0.001)	(0.001)	(0.000)	(0.001)
R^2	0.569	0.564	0.403	0.454
N	2334491	2334491	2334491	2334491
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. The sample comprises 469,232 inventors in total, observed between 1976 and 2016. $Cluster\ size$ 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. 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 10 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 x year level.

Consistent with Moretti (2021) we find that individual inventor productivity increases with cluster size. Splitting this measure up however into exploitative and explorative productivity reveals opposite effects. As shown in Table 2, column b, the marginal cluster effect in known to the inventor classes is 2.6 times as large the total productivity effect, while explorative 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 effects when it comes to exploitation, but it comes at the expense of less exploration, as predicted by our model. Table 2, column d, shows a decline in the fraction of new to the inventor patents, indicating that the search focus shifts towards exploitation. This narrower search also leads to less creative patents on average as measured by the textual measure of Kalayani (2025). Of note, these baseline results on the full sample are not restricted to moving inventors, star inventors or the largest clusters in the country (see Appendix).

Similar to Moretti (2021), we address remaining endogeneity concerns with an instrumental variable approach that leverages the geographical distribution of multi-location firms' research activities. The logic is that when a firm operates R&D labs in multiple cities, a productivity shock to the firm in one city (e.g., Silicon Valley) generates growth in that firm's cluster there but, because the firm reallocates resources across its labs, it also mechanically increases the size of the same technological field in the firm's other cities (e.g. Boston), regardless of local conditions in those places, conditional on city x year fixed effects.

We thus define our instrumental variable as inventor *i*'s peer firms' patenting in field *k* outside inventor *i*'s commuting zone, excluding inventor *i*'s firm's activities. Peer firms are all firms that are active, i.e. applied for at least one patent, in inventor *i*'s commuting zone and field in year *t.* The key identifying assumption is that any peer firms' increase in outside commuting zone activities in field *f* are unrelated to the focal inventor's innovative search towards or away from field *f* through channels other than a change in local cluster size. Noteworthy, even if that assumption would be violated, it is unlikely that any peer firms' increase in outside activity simultaneously has an influence on exploitation *and* exploration search of the focal inventor, besides its influence through an increase in local cluster size.

Using 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 as presented in our baseline Table 2. We do not fear a weak instrument bias given a F-value of 150.1 in the first stage. The marginal effects come out slightly stronger, which can be explained by the fact the instrument

creates variation in estimated cluster size from larger, cross country active firms, which typically will have a stronger influence than smaller local firms.

Table 3: Search outcomes with second stage outside CZ instrument

	(1)	(2)	(3)	(4)
	Individual productivity	Patents in Known Tech	Patents in New Tech	Fraction New Tech
Cluster Size (est.)	0.034***	0.057***	-0.023***	-0.033***
	(0.004)	(0.004)	(0.002)	(0.001)
N	2334491	2334491	2334491	2334491
First stage F	214.16	214.16	214.16	214.16
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. The sample comprises 469,232 inventors in total, observed between 1976 and 2016. Instrumental variable in the first is an inventor *i*'s peer firms' patenting in field *k* outside inventor *i*'s commuting zone, excluding inventor *i*'s firm's activities. Peer firms are all firms that are active, i.e. applied for at least one patent, in inventor *i*'s commuting zone and field in year *t.* $Cluster\ size$ 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 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 10 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.

Discussion

If narrower search leads to less creative, impactful, and financially valuable patents on average, we should observe these implications using a creativity measure (Kalyani et al. 2025), future prior art citations, and the financial value of patents as derived from stock market reactions to patent's grant event and future cites from other patents to the focal one as provided by (Kogan et al. 2017, updated version 2020). Table 4 below illustrates these measures for correlations and IV regressions. The IV regressions become insignificant for financial value, otherwise, all results remain consistent with the model predictions.

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

Panel A full sample			
	(1)	(2)	(3)
	Average creativity	Average future cites	Average financial value
Cluster Size	-0.006***	-0.012***	-0.008***
	(0.001)	(0.002)	(0.002)
R^2	0.570	0.765	0.812
N	2334491	1288482	1288482
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 IV regression			
	(1)	(2)	(3)
	Average creativity	Average future cites	Average financial value
Cluster Size	-0.029***	-0.094***	-0.003
	(0.005)	(0.017)	(0.012)
First stage F	214.16	150.97	150.97
N	2334491	1288482	1288482
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 estimation of equation (1), split samples, unit of observation is an inventor-year. The full sample comprises 469,232 inventors in total, observed between 1976 and 2016. $Cluster\ size$ 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 (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 and future cites from other patents to the focal one as provided by (Kogan et al. 2017, updated version 2020). Standard errors clustered at the city x year level.

6 Conclusion

Technology clusters are increasingly the drivers of modern innovation, yet our evidence reveals a hidden cost to their celebrated contributions. Using the universe of U.S. patents, we replicate the result that inventors in large technology hubs generate more patents, then show that the boost comes from incremental work in classes they already know. Patenting in brand-new classes, patent creativity scores, market valuations, and forward citations all decline as clusters thicken. A simple explore-exploit model explains the pattern: local density makes routine, within-field invention

cheaper, tilting private incentives toward exploitation even though the social value of exploration rises with cluster size.

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 the slowdown in transformative breakthroughs (Bloom et al., 2020), the growing spatial inequality of inventive activity, and the concentration of talent in a handful of superstar cities. Large, well-capitalized firms partially mitigate the bias by operating multiple R&D labs, but that strategy remains costly and incomplete.

Policy makers therefore face a trade-off. Encouraging clustering can raise near-term patent counts and 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 spillovers without the full exploitation bias of mega-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 at the spatial level. Second, these mechanisms may also reinforce declining business dynamism and the reallocation of talent (Decker et al., 2016; Akcigit & Ates, 2021), since geographically concentrated labor markets encourage inventors to double down on 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 et al., 2012; Acemoglu, 2002, 2023), highlighting another wedge between private incentives and socially optimal innovation.

The results also speak to regional inequality: if clusters both raise productivity and bias inventive search toward exploitation, then the geography of invention not only shapes innovation outcomes

but also regional growth trajectories (Moretti, 2012). It connects to labor economics, as the career security provided by localized skill demand affects inventors' incentives. It also speaks to industrial organization and technology policy, raising questions about whether cluster-promoting policies inadvertently reduce exploration, and whether interventions are needed to counterbalance private underinvestment in radical 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; Jovanovic & Rousseau, 2005; Perez, 2002; Berkes et al., 2025). These waves typically 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.

Several limitations point to fruitful directions 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 from the financing environment that supports them. Venture-capital investors, corporate R&D budgets, and government grant programs may differentially reward incremental versus exploratory projects, 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" is unambiguously good. 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 tech cluster.

Bibliography

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.

Ahuja, G., & Lampert, C. M. (2001). Entrepreneurship in the large corporation: A longitudinal study of how established firms create breakthrough inventions. *Strategic Management Journal*, 22(6–7), 521–543.

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.

Balsmeier, B., Fleming, L., & Lück, S. (2023). Isolating personal knowledge spillovers: Coinventor deaths and spatial citation differentials. *American Economic Review: Insights*, 5(1), 21–34.

Benner, M. J., & Tushman, M. L. (2002). Process management and technological innovation: A longitudinal study of the photography and paint industries. *Administrative Science Quarterly*, 47(4), 676–706.

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.

Fleming, L. (2001). Recombinant uncertainty in technological search. *Management Science*, 47(1), 117–132.

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.

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.

Kerr, W. R., & Kominers, S. D. (2015). Agglomerative forces and cluster shapes. *Review of Economics and Statistics*, 97(4), 877–899.

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

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.

Porter, M. E. (1998). Clusters and the new economics of competition. *Harvard Business Review*, 76(6), 77–90.

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.

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

Appendix

Table A1 – Mover versus non-mover inventors

Panel A – Inventors that moved					
	(1)	(2)	(3)	(4)	(5)
	Indv. prod.	Patents in Known Tech	Patents in New Tech	Fraction New Tech	Av. creativity
Cluster Size	0.011***	0.023***	-0.014***	-0.023***	-0.006***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
R^2	0.549	0.551	0.402	0.455	0.546
N	789514	789514	789514	789514	789514
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)
	Indv. prod.	Patents in Known Tech	Patents in New Tech	Fraction New Tech	Av. creativity
Cluster Size	0.010***	0.033***	-0.024***	-0.046***	-0.007***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
R^2	0.601	0.593	0.423	0.471	0.596
N	1410376	1410376	1410376	1410376	1410376
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), full analysis sample, unit of observation is an inventor-year. The sample comprises 469,232 inventors in total, observed between 1976 and 2016. Panel A Movers are inventors that applied for patents in at least two different commuting zones. Panel B Non-movers are inventors that have patent 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). 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 10 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 (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. Standard errors clustered at the city x year level.

Table A2 – Top 10 cities excluded

Panel A – W/O 10 largest cities					
	(1)	(2)	(3)	(4)	(5)
	Indv. prod.	Patents in Known Tech	Patents in New Tech	Fraction New Tech	Av. creativity
Cluster Size	0.009***	0.029***	-0.021***	-0.036***	-0.007***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
R^2	0.585	0.580	0.426	0.477	0.584
N	1269547	1269547	1269547	1269547	1269547
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 10 largest cities only					
	(1)	(2)	(3)	(4)	(5)
	Indv. prod.	Patents in Known Tech	Patents in New Tech	Fraction New Tech	Av. creativity
Cluster Size	0.009***	0.024***	-0.016***	-0.033***	-0.006***
	(0.001)	(0.002)	(0.001)	(0.002)	(0.001)
R^2	0.593	0.586	0.409	0.458	0.590
N	966087	966087	966087	966087	966087
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), full analysis sample, unit of observation is an inventor-year. The sample comprises 469,232 inventors in total, observed between 1976 and 2016. 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). 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 10 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 (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. Standard errors clustered at the city x year level.