The Changing Geography of American Innovation

Weiliang Tan*
National Bureau of Economic Research

March 26, 2026

Click here for the latest version

Abstract

The geography of innovation has important implications for place-based policy, yet the forces shaping the changing geography of American innovation remain poorly understood. Using comprehensive data on patents, firms, and inventors from 1976 to 2018, I document that the spatial concentration of innovation remained roughly constant before 1990 but rose sharply thereafter, driven by a small number of regions – with Silicon Valley alone accounting for about half of the post-1990 rise and the top 20 CZs accounting for nearly all of it. These regions share a common characteristic: a high initial share of college-educated workers. To identify the sources of this trend, I develop an empirical framework that progressively decomposes it – first by technology field and subfield, then by firm. About half of the post-1990 increase reflects a compositional shift toward information and communication technologies (ICT) – driven almost entirely by *Digital Communication* and *Computer Technology* – while the remainder reflects broad within-field concentration rises. At the firm level, both channels are driven by between-firm reallocation: the growing patent shares of ICT-oriented firms and firms already anchored in high-skill regions before 1990, rather than firm entry or within-firm changes. I provide descriptive evidence on two mechanisms through which the rapid rise of ICT concentrated innovation in high-skill regions: the colocation of ICT innovation with ICT production, and ICT-enabled firm spatial expansion that paradoxically concentrated R&D further in high-skill regions.

Keywords: geography of innovation; high-tech clusters; information and communication technologies

*Email: wt289@cornell.edu; tanw@nber.org. I am extremely grateful to Samuel Kortum, Peter Schott, Scott Stern, and Adam Jaffe for valuable advice. I thank participants at the NBER Productivity Seminar for helpful comments.

1 Introduction

Governments around the world use place-based innovation policies – R&D tax credits, science parks, and more recently the CHIPS Act and the EDA's Regional Technology and Innovation Hubs program in the United States – to shape where innovative activity occurs. The design of such interventions requires understanding not only why innovation clusters in certain regions, but also why the geography of innovation has been changing over time – which forces are concentrating innovation further, and whether policy can redirect or enhance them. A large literature has studied the first question, documenting the role of localized knowledge spillovers (Jaffe et al., 1993; Audretsch and Feldman, 1996; Feldman and Florida, 1994), agglomeration economies (Moretti, 2021; Bettencourt et al., 2007), and connections to academic science (Hausman, 2022; Bikard and Marx, 2020) in explaining where innovation is located (see Carlino and Kerr, 2015; Kerr and Robert-Nicoud, 2020, for surveys). The second question – why the geography has been *changing* – has received far less attention, despite its direct policy relevance.

Recent work has begun to document that innovation has become increasingly spatially concentrated: Andrews and Whalley (2022) trace long-run trends in concentration from 1866 to 2016, Chattergoon and Kerr (2022) decompose aggregate trends into software versus non-software components across city groups, Forman and Goldfarb (2020) track the parallel rise of firm-level and geographic concentration in IT. Less is known, however, about which regions drove the aggregate trend, whether the rising concentration occurred specifically in high-skill regions or simply in large ones, and what the underlying sources are. How much reflects shifts in industry composition versus rising concentration within fields? Which firms drove the reallocation, and where? The answers matter for policy: if rising concentration reflects the colocation of innovation with production in specific industries, then attracting manufacturing plants without the corresponding R&D base may not generate local innovation; if it reflects the spatial expansion of incumbent firms that keep R&D anchored in high-skill regions, then policies targeting new firm entry may be less effective than those addressing the location decisions of existing firms.

This paper addresses these questions using comprehensive data on patents, inventors, and firms from 1976 to 2018. I first provide a detailed characterization of the post-1990 rise – when it began, which regions drove it, and how it relates to the skill composition rather than the size of regions. I then develop a new empirical framework that progressively decomposes the aggregate trend – first by technology field and subfield, then by firm – to narrow down which explanations are consistent with the data. Finally, I provide descriptive evidence on the underlying mechanisms that likely drove the trend.

I begin with the aggregate trends. The locational Gini coefficient of patents per capita across Commuting Zones (CZs) remained approximately constant between 1976 and 1990 but rose from 0.36 to 0.50 thereafter. I then ask which regions drove the aggregate trend, distinguishing a CZ's patenting footprint – its share of national patents – from its patenting intensity – patents per capita. Silicon Valley alone accounts for about half of the post-1990 Gini rise; dropping the top 20 CZs by patent share increase nearly eliminates the entire increase. Moreover, this rising concentration occurred specifically in high-skill CZs rather than in large or densely populated ones: CZs in the top decile of the 1990 college ratio saw their share of national

patents rise from roughly 20% to over 34%, while the bottom half stagnated. The annual elasticity of patents per capita with respect to a CZ's 1990 college ratio hovered around 0.6 from 1976 to 1990 but rose sharply to approximately 1.6 by 2012, while the corresponding elasticity with respect to population shows no comparable break.

I then decompose this trend by technology field and subfield. ICT patents are consistently the most spatially concentrated and the fastest-growing field, while spatial concentration also rose within most non-ICT fields after 1990. Shapley decompositions by technology field attribute about half of the post-1990 increase in the Gini to a compositional shift toward ICT (the cross-field component) and 40% to rising concentration within fields (the within-field component). A finer subfield decomposition pinpoints Digital Communication and Computer Technology as the core drivers of the compositional shift – two-thirds of the post-1990 increase in the patent elasticity is explained by compositional changes across subfields. The within-field rise is broad-based, spanning most fields and subfields – not just ICT.

I further decompose each channel at the firm level using modified Foster-Haltiwanger-Krizan (FHK) and Shapley decompositions. The compositional shift toward ICT was driven by the rising patent shares of *ICT-oriented firms* – not by firms shifting their patent portfolios toward ICT. These firms disproportionately concentrate all their patents in high-skill CZs: cross-firm differences in patent geography, rather than within-firm differences between ICT and non-ICT patents, account for the bulk of the aggregate ICT–non-ICT elasticity gap. The within-field rise in concentration was driven by the growing patent shares of *high-skill CZ firms* – firms already concentrated in high-skill CZs – for both ICT and non-ICT innovation. Between-firm reallocation accounts for nearly all of the post-1990 rise in the non-ICT patent elasticity and more than two-thirds of the rise for ICT, with pre-1990 incumbents rather than new entrants driving the increase. Decompositions by ICT-oriented versus non-ICT-oriented firm groups confirm that this reallocation operates within both groups, indicating that the within-field rise is not simply a byproduct of the compositional shift toward ICT.

Finally, I provide descriptive evidence that the rapid rise of information and communication technologies (ICT) likely drove the growth of these two groups of firms through two channels. First, the sudden arrival of influential ICT breakthroughs between 1985 and 1993 catalyzed a surge in ICT patenting – the ICT share of all patents increased from roughly 8.5% in 1990 to 33% in 2018, with 78% of ICT patents during this period citing these breakthroughs within two degrees. This surge disproportionately increased innovation in ICT-oriented firms, whose ICT production plants – in Software Publishers, Telecommunications, and Data Processing – had already clustered in high-skill CZs before 1990 due to their deeper pools of high-skill labor. The *colocation* of ICT innovation with ICT production in these regions explains both why ICT-oriented firms concentrate their patents in high-skill CZs and why the compositional shift toward ICT raised aggregate innovation concentration. Second, ICT-induced reductions in communication costs – from the adoption and standardization of foundational digital technologies and the scaling of the NSFNET – likely facilitated firm spatial expansion, which paradoxically concentrated innovation in high-skill CZs even further. Firms with R&D in high-skill CZs expanded into substantially more regions – both high- and low-skill – than firms with R&D in low-skill CZs and, following expansion, increased their R&D employment and patenting activity,

predominantly in high-skill CZs. This phenomenon arises for both ICT and non-ICT oriented firms. I interpret it as an *asymmetric scale effect*: high-skill CZ firms can exploit lower production and service costs in other regions, whereas low-skill CZ firms already operate in relatively low-cost environments and therefore gain less from expansion.

Contributions to the Literature

This paper makes three contributions. First, using comprehensive and updated data on patents, inventors, and firms, I carefully document how the geography of American innovation has changed in recent decades. Relative to existing work on the changing geography of innovation – Andrews and Whalley (2022) on long-run concentration trends, Chattergoon and Kerr (2022) on software versus non-software decompositions across city groups – I make several advances: establishing the post-1990 break using the locational Gini, which captures the full distribution rather than being dominated by the upper tail; distinguishing a CZ's patenting *footprint* (its share of national patents) from its patenting *intensity* (patents per capita), which reveals that Silicon Valley accounts for about half the post-1990 Gini rise and the top 20 CZs nearly all of it; and showing that the rising concentration occurred in high-skill CZs rather than in large ones – a distinction that the urban scaling literature (Bettencourt et al., 2007; Balland et al., 2020) does not make.

Second, I develop a new empirical framework to progressively decompose the aggregate trend – first by technology field, then by subfield, and finally by firm – and use it to narrow down which explanations are consistent with the data. To my knowledge, this is the first paper to trace the aggregate spatial trend through all of these layers in an integrated framework. Forman and Goldfarb (2020) document the parallel rise of firm-level and geographic concentration in IT, and Forman et al. (2016) show that the Bay Area's rise extends beyond ICT to a broader "coagglomeration" of invention, but neither paper formally links firm dynamics to the aggregate spatial trend. On the firm side, Agrawal and Cockburn (2003) and Agrawal et al. (2014) show that large anchor firms amplify local innovation through spin-outs and market thickening, while hUallacháin (2025) demonstrates that individual firms' strategic decisions drive city-level patent dynamics – but this work studies firm-level effects on regional innovation in levels rather than connecting firm reallocation to the aggregate spatial trend. My decomposition framework bridges these literatures, revealing that both the compositional shift toward ICT and the within-field concentration rise are driven by *between-firm reallocation* – the growing patent shares of ICT-oriented firms and high-skill CZ firms – rather than within-firm changes or new entry.

Third, I provide micro-level descriptive evidence on the likely mechanisms through which the ICT shock concentrated innovation in high-skill CZs, informing the policy questions raised in the introduction. The *colocation* of ICT innovation with ICT production in high-skill CZs – a sector-specific extension of the broader colocation of innovation and production documented by Fort et al. (2020) and Liu (2024) – explains why ICT-oriented firms concentrate their patents in high-skill CZs and why the compositional shift toward ICT raised aggregate concentration. ICT-enabled firm spatial expansion, which paradoxically concentrated R&D further in high-skill CZs through an *asymmetric scale effect*, explains why high-skill CZ firms grew

disproportionately, connecting the firm spatial expansion literature (e.g. Hsieh and Rossi-Hansberg, 2023; Jiang, 2024; Kleinman, 2025) to innovation geography.

2 Data

I obtain the universe of patents granted by the USPTO since 1976 from PatentsView, supplemented with bulk files from the US Patent and Trademark Office (USPTO). My sample includes all 3.70 million utility patents where at least one inventor lists a US address. Each patent contains extensive information, including inventor addresses, which typically reflect their home city and state. Using the Google Maps API, I geocode these addresses and map them to various spatial resolutions within the United States. I use these inventor locations to calculate annual patent counts for each region and derive measures of the spatial concentration of innovation across regions over time. Additionally, every patent is assigned a unique primary Cooperative Patent Classification (CPC) technology class. To analyze the role of compositional changes in innovative activity across fields, I map these classes to broader technology fields and subfields by adapting the field classification methodology developed by the World Intellectual Property Organization. Most patents also have one or more assignees, typically US firms, that hold ownership of the intellectual property.

I obtain firm and establishment locations from historical Dun & Bradstreet (D&B) files, which provide the near-universe of US business establishments with information on industry, employment, sales, and geographic location. Regional data come from several sources: annual county-level population and income from the Bureau of Economic Analysis (BEA); county-level employment by industry from the harmonized County Business Patterns (CBP) database produced by Eckert et al. (2020); and the composition of workers by college education from the IPUMS Current Population Survey and IPUMS National Historical Geographic Information System (Flood et al., 2025; Schroeder et al., 2025). See Appendix A for more details.

3 Aggregate Trends on the Geography of American Innovation

This section documents the aggregate trends in the spatial concentration of American innovation from 1976 to 2018, establishing three facts that serve as a foundation for the progressive decomposition in Sections 4 and 5. Throughout the paper, I measure innovation using patent counts. Patents are an imperfect proxy: they capture only patentable inventions, and the propensity to patent may vary across industries, firm types, and time periods in ways that differ across regions.[1] Despite these limitations, patents remain the most comprehensive, geographically granular, and consistently available measure of inventive activity over the four decades I study, and are the standard measure in the geography of innovation literature (Jaffe et al., 1993; Audretsch and Feldman, 1996; Andrews and Whalley, 2022).

[1]For example, if patent propensity rose disproportionately among firms concentrated in high-skill CZs — due to changes in patent law or strategic patenting behavior — this could mechanically inflate measured concentration. I discuss this concern in Section 6.4.

3.1 Fact 1: Innovation became more geographically concentrated after 1990.

I use the locational Gini index (Krugman, 1992) to measure the aggregate spatial concentration of innovation annually from 1976 to 2018. For each commuting zone (CZ), I calculate its share of total patents and share of total population in the US. I then rank CZs by their patent-to-population share ratio and plot the cumulative sum of patent shares against the cumulative sum of population shares to construct the locational Gini curve. The locational Gini index, proportional to the area between this curve and the 45-degree line, measures the concentration of patents across CZs relative to population, with each CZ weighted by its population share. A value of zero indicates perfect equality while a value of one reflects perfect inequality[2].

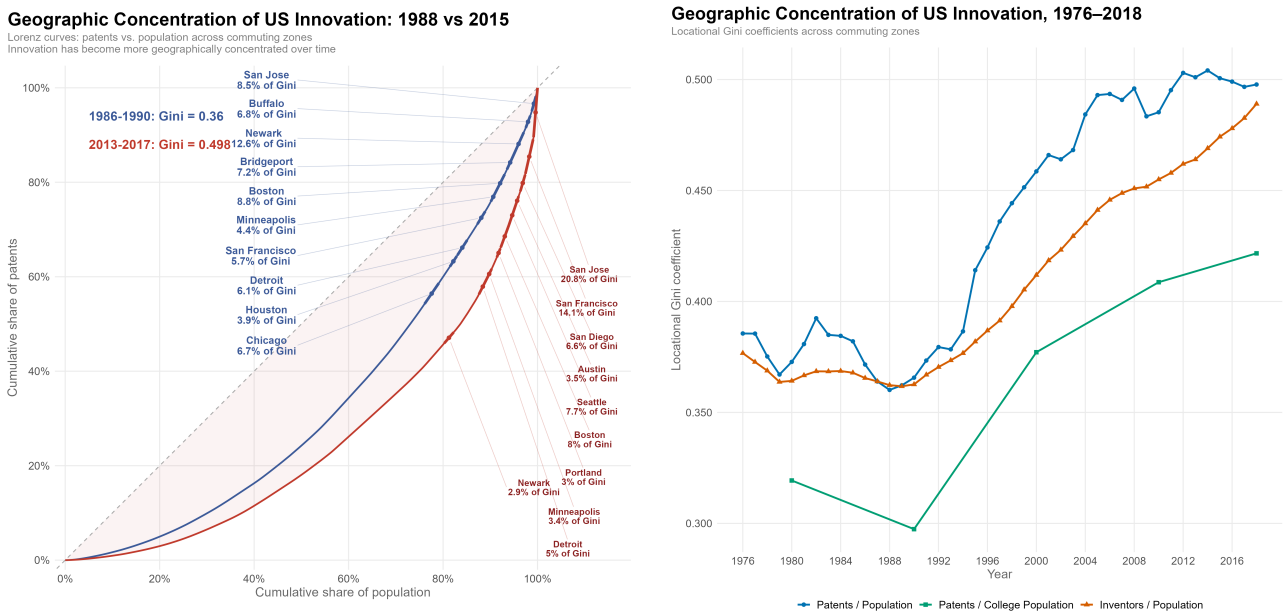


Figure 1: Trends in the spatial concentration of innovation across CZs. The left panel compares locational Gini curves of patents relative to population across CZs in 1988 and 2015, with each CZ's contribution to the overall Gini coefficient labeled for the top contributors. The locational Gini index rose from 0.36 in 1986–1990 to 0.50 in 2013–2017. The right panel plots annual trends in locational Gini coefficients from 1976 to 2018 for three measures: patents relative to population (solid blue), inventors relative to population (dashed orange), and patents relative to college-educated workers (solid green). All three measures remained approximately constant through 1990 and rose substantially thereafter.

Figure 1 documents this trend. The left panel compares locational Gini curves in 1988 and 2015. The Gini index rose from 0.36 to 0.50 over this period, with San Jose's contribution more than doubling from 8.5% to 20.8% of the total Gini, reflecting the explosive growth of Silicon Valley. San Francisco's contribution similarly rose from 5.7% to 14.1%, while previously important innovation centers such as Detroit and Chicago

[2]Formally, the locational Gini index G of patents x with respect to population L across all the N CZs is equivalent to half of the weighted mean absolute deviation of patents across all CZ-pairs o, d:

$$G = \frac{1}{2\mu} \sum_{o=1}^{N} \sum_{d=1}^{N} \frac{L_o}{L} \frac{L_d}{L} |x_o - x_d|, \qquad \mu = \sum_{o=1}^{N} \frac{L_o}{L} x_o.$$

dropped out of the top contributors.

The right panel of Figure 1 plots annual trends in three locational Gini measures. The solid blue line shows that the spatial concentration of patents relative to population remained approximately constant from 1976 to 1990 but increased significantly thereafter. The dashed orange line depicts the trend for inventors, which closely mirrors that for patents; the gap between the two from 1995 to 2018 potentially reflects agglomeration economies in innovation, as documented by Moretti (2021)[3]. The solid green line illustrates trends in the locational Gini index of patents relative to college-educated workers, computed annually by replacing population shares with college-educated worker shares. Relative to college-educated workers, the spatial concentration of patents remained approximately constant in the 1980s and rose significantly after 1990. This finding suggests that the rising spatial concentration of innovation cannot be solely attributed to the geographical sorting of workers by skill. In Appendix B.1, I demonstrate the robustness of these trends to employing alternative measures of spatial concentration.

The most closely related measure in the literature is the dartboard innovation intensity concentration index used by Andrews and Whalley (2022), which builds on the Ellison and Glaeser (1997) index. Because their measure is based on squared deviations of each CZ's patent share from its population share, it is particularly sensitive to the upper tail – a region whose patent share deviates twice as far contributes four times as much. The locational Gini, by contrast, is based on pairwise absolute differences in innovation intensity weighted by population, making it more sensitive to changes throughout the entire distribution rather than being dominated by a few extreme outliers. Moreover, because it is constructed from a Lorenz curve, the Gini provides a visual characterization of the full distribution at each point in time, making it straightforward to identify both the top contributors and how their relative positions shift over time.

3.2 Fact 2: Silicon Valley accounts for about half of the post-1990 rise in concentration, while the top 20 CZs account for nearly all of it

Having established that the spatial concentration of innovation rose sharply after 1990, I now ask: where, specifically, did innovative activity concentrate, and which regions drove the aggregate trend? Figure 2 maps the geography of changes in innovation intensity — measured by the patent-to-population share ratio — between five-year averages around 1990 and 2015. The sharpest gains are concentrated along the West Coast and in a handful of corridors in the Northeast and Mountain West, while much of the Midwest and parts of the Northeast experienced relative declines.

These gains were concentrated in CZs that were already relatively patent-intensive in 1990, consistent with a fanning out of the pre-existing distribution rather than a reshuffling of innovative activity across regions. Figure 3 confirms this pattern: the left panel maps the patent-to-population share ratio across CZs in 1990, while the right panel plots this ratio in 1990 against the same ratio in 2015. The strong positive correlation shows that regions with higher initial intensity generally maintained an even larger lead by 2015, while

[3]Specifically, Moretti (2021) shows that inventors produce more patents when located in regions with more inventors. This finding implies that the spatial concentration of patents exceeds that of inventors.

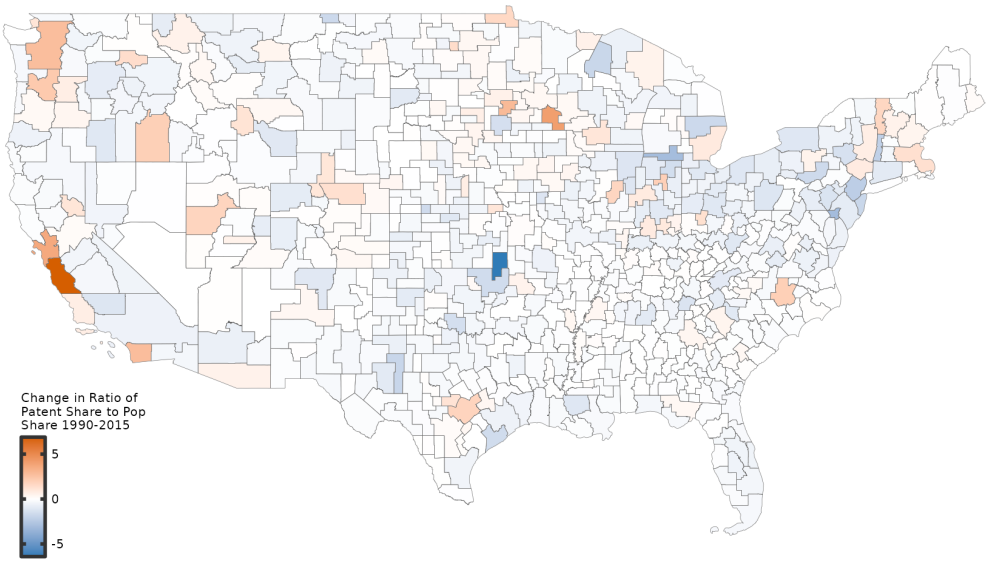


Figure 2: Changes in the geography of innovation intensity from 1990 to 2015, measured using changes in the patent-to-population share ratio on a pseudo-log scale.

regions starting from a lower base fell further behind.

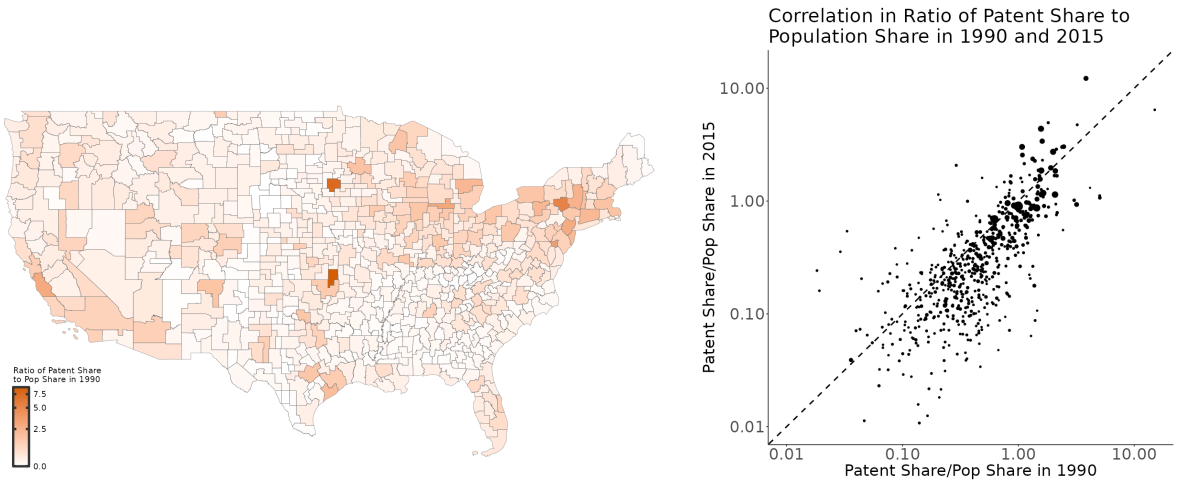


Figure 3: The geography of the patent-to-population share ratio in 1990 (left) and its correlation with the same ratio in 2015 (right).

Which CZs drive the rising concentration? Figure 4 ranks the top 20 CZs by two complementary measures of change between 1986–1990 and 2013–2017. The left panel ranks CZs by the increase in national patent share, capturing changes in a region's absolute *footprint* in national innovation. The right panel ranks CZs by the increase in the innovation location quotient (LQ) — the ratio of a CZ's patent share to its population share — capturing changes in local innovation *intensity* regardless of CZ size. San Jose and San Francisco dominate both rankings by a wide margin, underscoring the outsized role of Silicon Valley in the post-1990 concentration of American innovation. Seattle, San Diego, Austin, and Portland also appear prominently on both measures. The two panels diverge for some CZs because the location quotient divides patent share

by population share: a CZ can experience a large LQ increase with little change in its patent share if its population shrank, while a fast-growing CZ can absorb a large slice of national patents yet see only a modest LQ increase if its population grew proportionally.

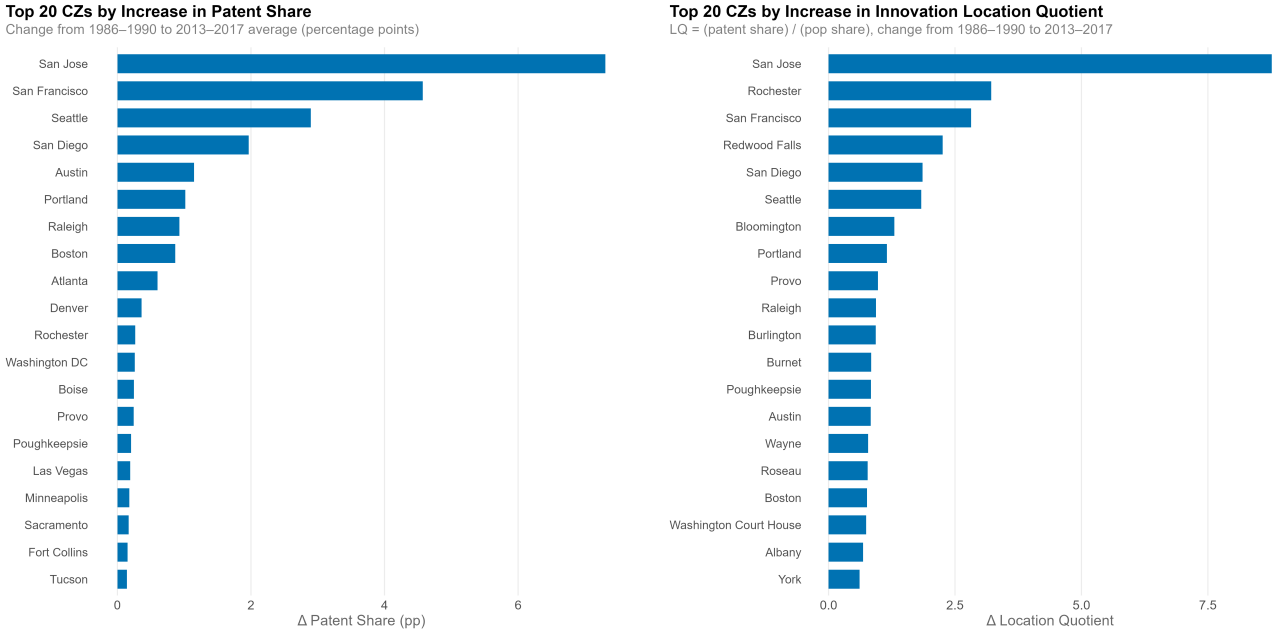


Figure 4: Top 20 CZs by increase in patent share and innovation location quotient. The left panel ranks CZs by the increase in national patent share between five-year averages around 1990 (1986–1990) and 2015 (2013–2017). The right panel ranks CZs by the increase in the location quotient (patent share divided by population share) over the same period.

Although the two measures capture different aspects of the post-1990 shift — innovation intensity versus national footprint — they are strongly positively correlated across CZs: regions that gained the most patent share also tended to experience the largest increases in innovation intensity (Figure 25 in Appendix B.2).

Because patent share changes matter more than intensity changes for the locational Gini coefficient — which is fundamentally a size-weighted measure of the Lorenz curve's deviation from the diagonal — I can quantify how much the aggregate concentration trend depends on the top CZs. Figure 5 recomputes the Gini after sequentially dropping the CZs with the largest increases in patent share and recalculating all shares among the remaining regions. Dropping Silicon Valley alone (San Jose and San Francisco) eliminates roughly half of the post-1990 rise, and dropping the top 20 nearly eliminates it entirely. The post-1990 rise in the spatial concentration of American innovation is thus driven almost entirely by a small number of regions.

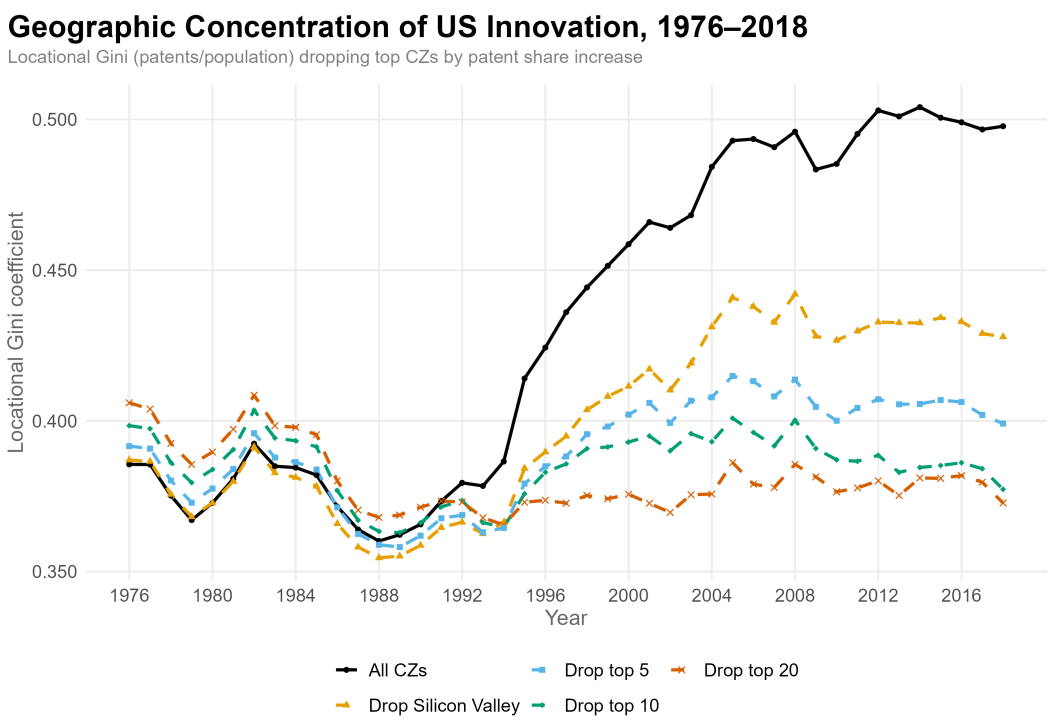


Figure 5: Locational Gini coefficient of patents with respect to population, 1976–2018, recomputed after dropping CZs with the largest increases in patent share between 1986–1990 and 2013–2017. "Drop Silicon Valley" removes San Jose and San Francisco. "Drop top N" removes the N CZs with the largest increase in national patent share; population and patent shares are recomputed among the remaining CZs before calculating the Gini.

3.3 Fact 3: Innovation became more geographically concentrated in *high-skill CZs* after 1990.

The previous subsection identified Silicon Valley and a small number of other regions as the key drivers of rising concentration. What do these regions have in common? Figure 6 reveals a striking pattern by grouping CZs into population-weighted deciles based on their 1990 college ratio — the share of residents aged 25 and older with at least a bachelor's degree — and plotting each decile's share of national patents and its mean innovation location quotient over time.

The left panel of Figure 6 shows that the top college-ratio decile (D10) accounted for roughly 20% of national patents before 1990 — already disproportionate given that each decile contains about 10% of the population. Beginning around 1990, however, D10's patent share rose sharply to over 34% by 2018. The ninth decile also diverged upward, rising from about 11% to 17%. Together, the top two deciles captured over half of all US patents by 2018, up from roughly 30% in 1990. The top CZs identified in Fact 2 share a common characteristic: they are predominantly high-skill regions, and the divergence in patenting follows a clear skill gradient. Meanwhile, the bottom five deciles remained largely flat or declined.

The right panel of Figure 6 tells a similar story using the innovation location quotient introduced in Section

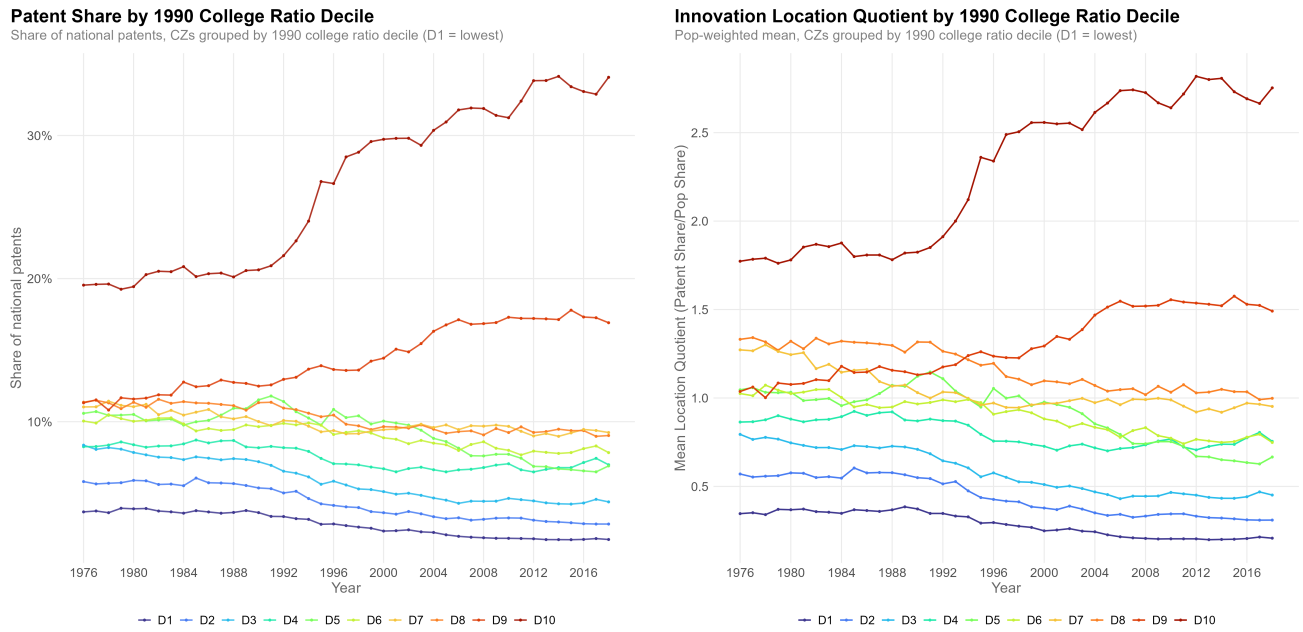


Figure 6: Innovation trends by 1990 college ratio decile. CZs are grouped into population-weighted deciles based on their 1990 college ratio (D1 = lowest, D10 = highest), so each decile contains approximately 10% of the national population. The left panel plots each decile's share of national patents from 1976 to 2018. The right panel plots the population-weighted mean of the innovation location quotient (patent share / population share) for each decile over the same period. Both panels show that the top decile pulled sharply away from the rest after 1990, while lower deciles experienced stagnation or decline.

3.2. A value above one indicates that a CZ is overrepresented in patenting relative to its population, while a value below one indicates a deficit. The top decile's location quotient rose from about 1.8 in 1990 to nearly 2.7 by 2018, and the ninth decile's quotient rose from roughly 1.0 to about 1.5, indicating that a broad set of high-skill CZs became increasingly overrepresented in national patenting activity. Conversely, CZs in the bottom deciles saw their location quotients fall further below one, signaling a growing innovation deficit relative to their population shares.

While Figure 6 documents divergence across discrete decile groups, Figure 7 exploits the continuous variation in college ratios across all CZs to formalize the relationship between initial skill composition and subsequent innovation growth. The left graph plots the absolute change in patenting intensity — measured by patents per capita — between five-year averages around 1990 and 2015 against the 1990 college ratio across CZs. A one-percentage-point higher college ratio in 1990 is associated with a 0.53 greater increase in patents per 1,000 population over this period ($\beta = 0.528$, $p = 0.002$), with San Jose, San Francisco, San Diego, and other technology hubs clearly visible as outliers along the fitted line. To account for regions with zero patents and to exploit the annual frequency of the dataset, I estimate annual elasticities α_t of patent counts with respect to the 1990 college ratio – henceforth "patent elasticity" – using the following Quasi-Poisson specification:

$$\text{Patents}_{r,t} = \exp\left(\alpha_t \cdot \text{Log College Ratio}_{r,1990} \times \text{Year}_t + \beta \cdot \text{Log Population}_{r,t} + \gamma_t + \epsilon_{r,t}\right). \tag{1}$$

where r represents regions and t denotes years. Controlling for log population means the elasticity captures the relationship between patenting *intensity* — patents per capita — and the college ratio, consistent with the intensity focus of the right panel of Figure 6. The Quasi-Poisson estimator accommodates overdispersion in patent counts and naturally handles regions with zero patents. The estimated annual elasticities α_t play an important role in the decomposition analysis that follows: I decompose this elasticity by technology field and subfield in Section 4 to identify which fields drive the sorting of innovation toward high-skill CZs. The right graph of Figure 7 shows that this elasticity hovered around 0.6 from 1976 to 1990 but rose sharply to approximately 1.6 by 2012, where it has since plateaued. This trend reflects the gradual geographical sorting of patents per capita across CZs, driven by differences in their initial college ratios in 1990 – consistent with the divergence across deciles documented in Figure 6. In Appendix B.3, Figure 26 plots trends in the elasticity of CZ patents per capita with respect to the 1990 population and population density, respectively. The absence of a clear break in these trends after 1990 suggests that the skill mix of workers in 1990 is a more important margin for the geographical sorting of patents per capita over time than either population size or density. This distinction matters because the urban scaling literature documents that innovation scales superlinearly with city size (Bettencourt et al., 2007; Balland et al., 2020), yet the post-1990 concentration occurred along the skill dimension rather than the size dimension.

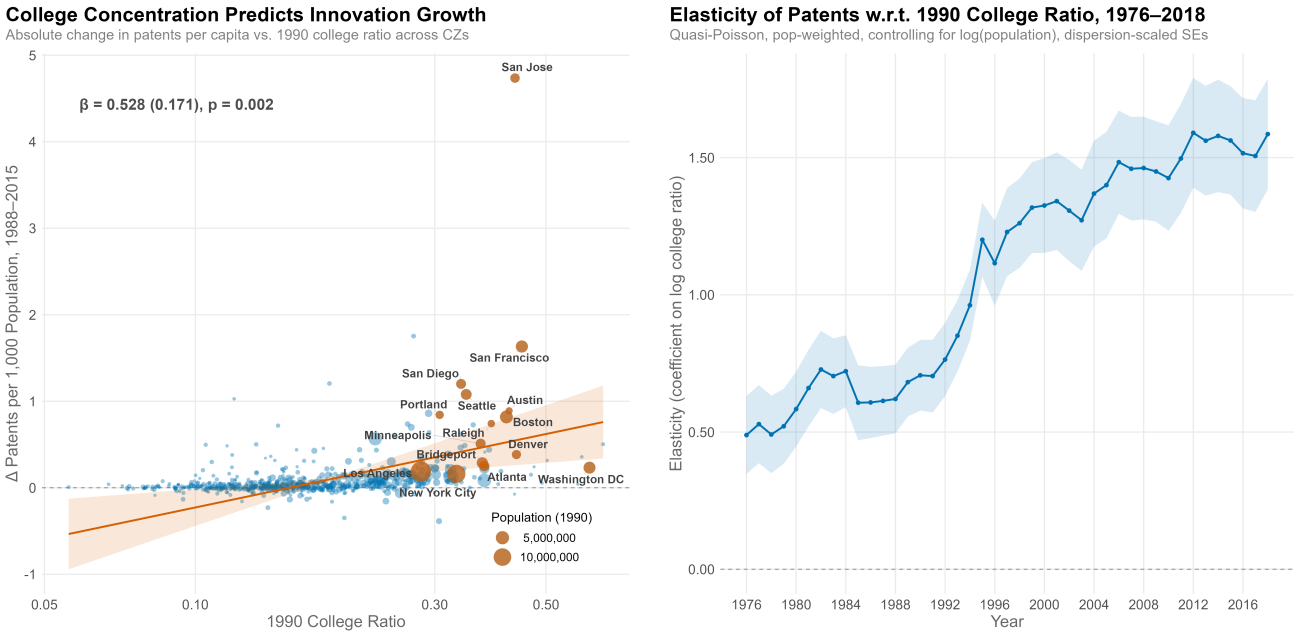


Figure 7: The post-1990 geographical sorting of patents per capita toward high-skill CZs. The left graph plots the absolute change in patents per 1,000 population between five-year averages around 1990 (1986–1990) and 2015 (2013–2017) against the 1990 college ratio across CZs, with dot size proportional to population. The right graph displays trends in the annual patent elasticity α_t of CZ patents per capita with respect to the 1990 college ratio. The confidence intervals in both graphs reflect heteroskedastic-robust standard errors.

4 Decomposition by Technology Field and Subfield

Section 3 established that the spatial concentration of innovation rose sharply after 1990, driven by a small number of high-skill CZs. There are many potential explanations, including the rising importance of knowledge spillovers, agglomeration economies, larger teams given the rising burden of knowledge, connections to university research, worker sorting by education, place-based policies, and changes in patent law. To make progress, I develop an empirical framework that progressively decomposes the aggregate trend – first across technology fields and subfields in this section, then across firms in Section 5 – to narrow down which explanations are consistent with the data. The existing literature has yet to trace the aggregate spatial trend through all of these layers in an integrated framework.

I begin by disaggregating the two measures of spatial concentration from Section 3 – the locational Gini (left panel) and the patent elasticity with respect to the 1990 college ratio (right panel) – by technology field in Figure 8. Two features stand out. First, ICT patents are consistently the most spatially concentrated, with Gini coefficients above 0.7 and elasticities exceeding 3 by 2015 – and as I show next, ICT is also the fastest-growing field, suggesting that a compositional shift toward ICT may be a key source of the aggregate rise. Second, spatial concentration has been rising after 1990 for most fields, indicating that the aggregate trend is not purely a compositional phenomenon. These two observations – a compositional shift toward a spatially concentrated field, and broad within-field concentration increases – suggest two distinct channels through which aggregate concentration rose.

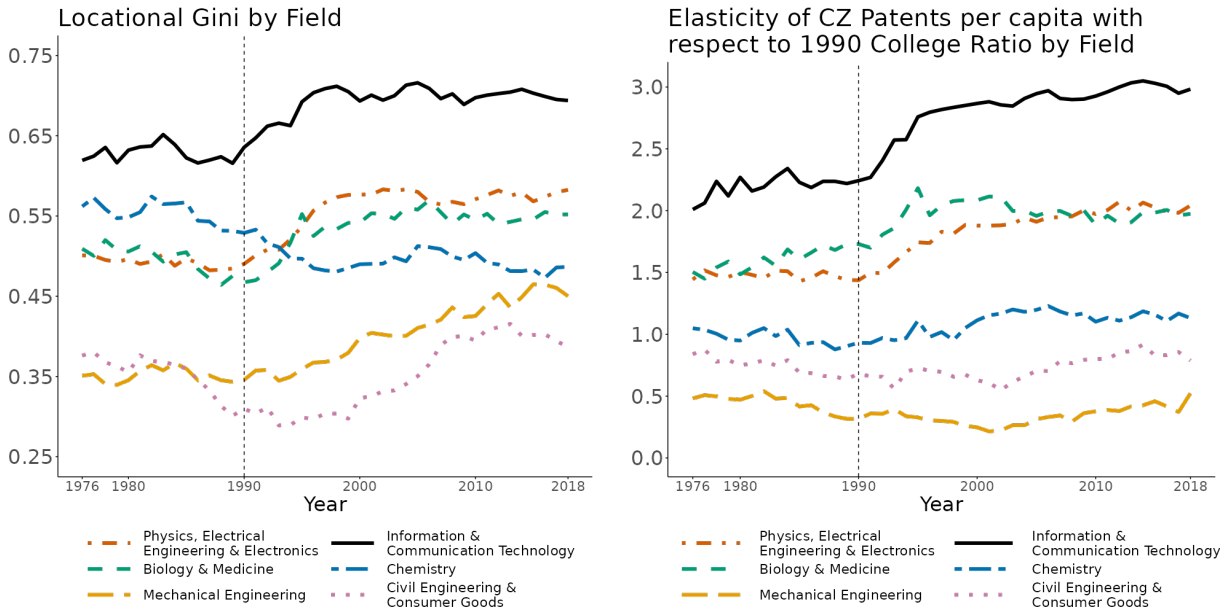


Figure 8: Spatial concentration of patents per capita by technology field: levels. Left: Gini coefficient. Right: elasticity with respect to the 1990 college ratio, estimated using PPML.

To formally quantify these channels, I use Shapley decompositions to separate the post-1990 increase in the spatial concentration of innovation into within-field and cross-field components, using both the Gini

coefficient and the patent elasticity. I decompose changes over time in the overall Gini coefficient G of patents per capita from 1990–2018 into within-field, cross-field, and field colocation components:

$$G_{t^*} - G_{1990} = \sum_{t=1991}^{t=t^*} \Delta G_t = \sum_{t=1991}^{t=t^*} \Bigg[\underbrace{\sum_f \overline{G}_{f,t} \Delta s_{f,t}}_{\substack{\text{changes in field} \\ \text{composition}}} + \underbrace{\sum_f \overline{s}_{f,t} \Delta G_{f,t}}_{\substack{\text{within-field} \\ \text{changes}}} + \underbrace{\Delta \Bigg(G_t - \sum_f s_{f,t} G_{f,t} \Bigg)}_{\substack{\text{changes in the} \\ \text{colocation of fields}}} \Bigg] \tag{2}$$

where $G_{f,t}$ is the Gini coefficient of patents in field f with respect to population across CZs in year t, $s_{f,t}$ is the share of US patents in field f in year t, and $\overline{x}_t = \frac{x_t + x_{t-1}}{2}$, $\Delta x_t = x_t - x_{t-1}$ for any variable x. I conduct an analogous decomposition of the post-1990 increase in the aggregate annual elasticity, α_t, of CZ patents per capita with respect to the 1990 college ratio:

$$\alpha_{t^*} - \alpha_{1990} = \sum_{t=1991}^{t=t^*} \Delta \alpha_t = \sum_{t=1991}^{t=t^*} \Bigg[\underbrace{\sum_k \overline{\alpha}_{k,t} \Delta s_{k,t}}_{\substack{\text{changes in field} \\ \text{composition}}} + \underbrace{\sum_k \overline{s}_{k,t} \Delta \alpha_{k,t}}_{\substack{\text{within-field} \\ \text{changes}}} + \underbrace{\Delta \Bigg(\alpha_t - \sum_k s_{k,t} \alpha_{k,t} \Bigg)}_{\substack{\text{residual: changes in the} \\ \text{colocation of fields}}} \Bigg] \tag{3}$$

where $\alpha_{k,t}$ is the annual elasticity of CZ patents per capita in field k and year t with respect to the 1990 college ratio.

4.1 Fact 4: A compositional shift toward ICT – driven by *Digital Communication* and *Computer Technology* – explains about half of the post-1990 rise in concentration

Figure 9 documents the compositional shift in innovative activity across fields. While patent counts increased across most fields (left panel), ICT grew disproportionately: its share of annual patents rose from about 8% in 1990 to over 32% by 2015, while the shares of most other fields declined or remained flat (right panel). Because ICT innovation is systematically more concentrated in high-skill CZs than non-ICT innovation (Figure 8), this reallocation of innovative activity toward ICT mechanically raised aggregate concentration.

The Gini decomposition (Figure 10) attributes 52% of the increase in the Gini coefficient to changes in field composition (the cross-field component), 40% to rising spatial concentration within fields (the within-field component), and 8% to the increasing colocation of patents across different fields. Most of the rise between 1990 and 1995 is accounted for by changes in field composition, whereas the subsequent increase is more evenly split between the cross-field and within-field components. The elasticity decomposition (Figure 11) yields a similar split: 60% from cross-field changes and 40% from within-field changes. The middle panels of Figures 10 and 11 show that nearly the entire cross-field component is driven by the rising share of ICT patents, with the elasticity decomposition attributing 53% of the overall increase to the reallocation toward ICT alone. An additional 7% reflects the rising share of Biology & Medicine patents. These results are robust to dropping top patenting CZs (Appendix B.4.2).

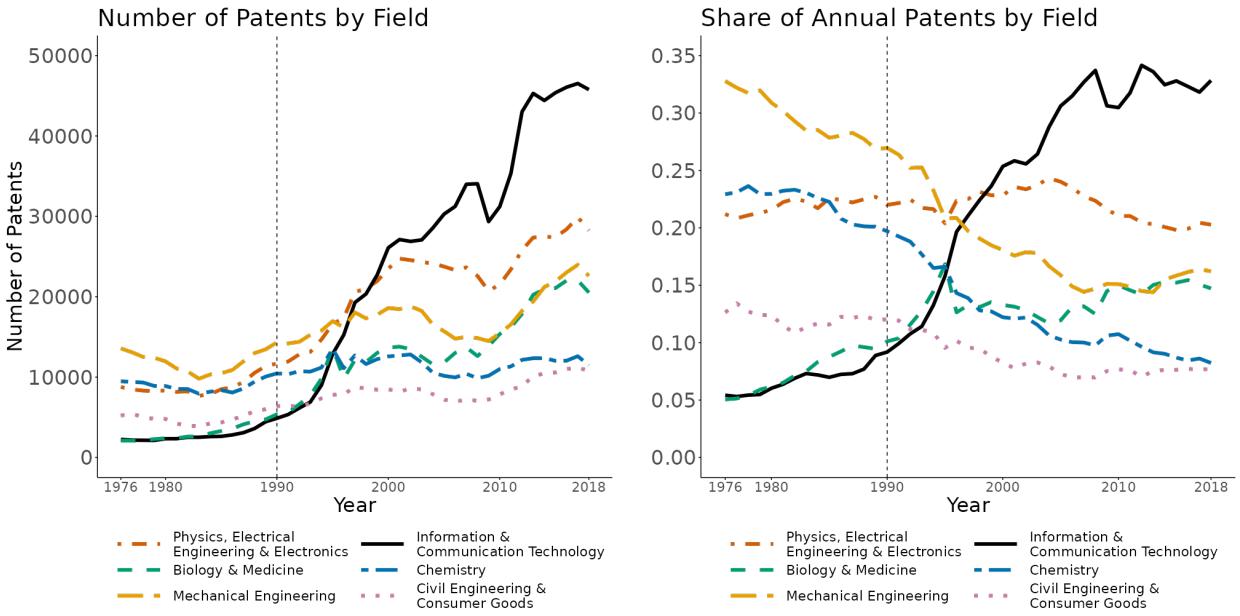


Figure 9: Patent counts and shares by technology field, 1976–2018. The left panel plots the number of patents by field. The right panel plots each field's share of annual US patents. ICT patents grew disproportionately, rising from about 8% to over 32% of national patents between 1990 and 2015.

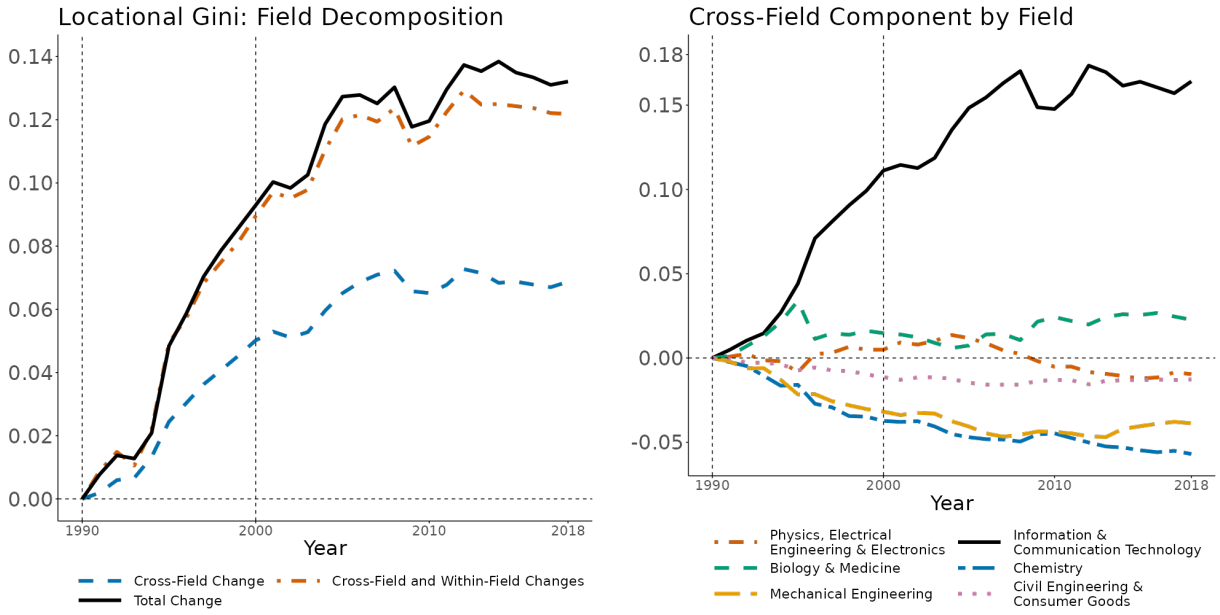


Figure 10: Decomposition of changes in the locational Gini coefficient of patents per capita relative to 1990 into within-field, cross-field, and field colocation components (left), with the cross-field component further decomposed by individual field (right).

A finer subfield decomposition reveals that the compositional shift is even more pronounced at higher granularity, and pinpoints its source. The right panel of Figure 12 shows that two-thirds of the post-1990 increase in the aggregate patent elasticity is explained by compositional changes across subfields – compared to about half at the field level. Table 1 shows that this cross-subfield component is driven predominantly

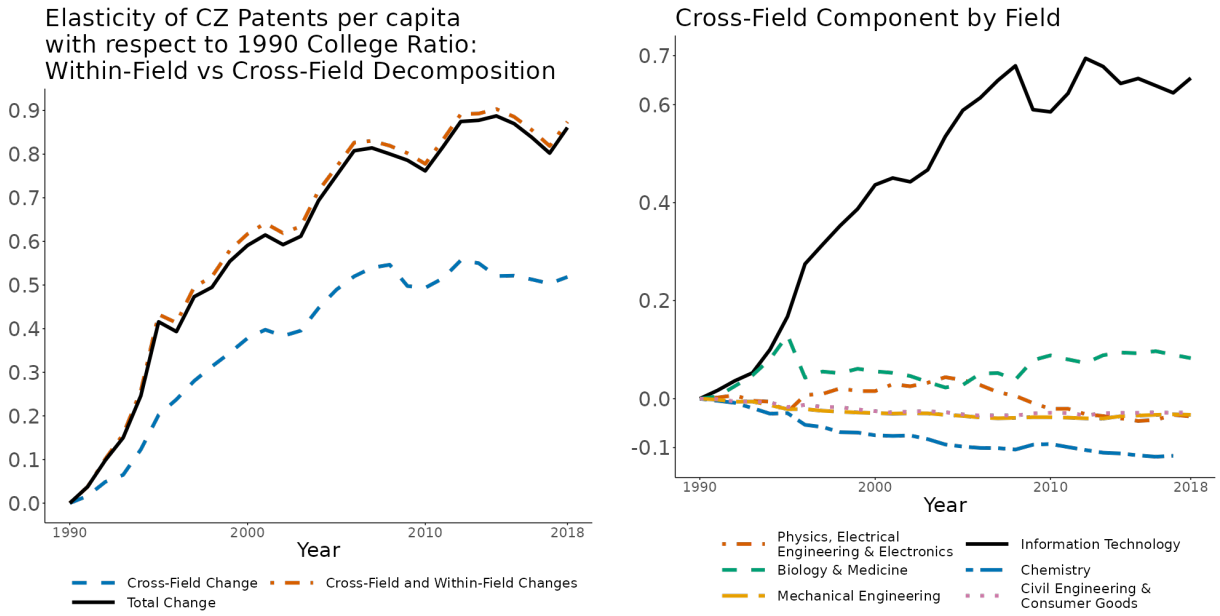


Figure 11: Decomposition of the aggregate elasticity of CZ patents per capita with respect to the 1990 college ratio into within-field and cross-field components (left), with the cross-field component further decomposed by field (right).

by two ICT subfields: *Digital Communication* and *Computer Technology*, which together account for the vast majority of the cross-subfield elasticity increase (Elast. Δ Across column). These two subfields alone grew from a combined 5% of national patents in 1990 to 26% by 2015, while exhibiting the highest patent elasticities of any subfields (above 3.0), making them the primary channel through which the compositional shift raised aggregate concentration in high-skill CZs.

4.2 Fact 5: Within-field concentration rose broadly across fields and subfields, explaining the other half

The compositional shift toward ICT is not the whole story. The remaining 40–48% of the post-1990 rise in concentration reflects increases in spatial concentration *within* fields – innovation becoming more concentrated in high-skill CZs even holding the technology mix fixed.

Figure 13 shows this directly by plotting changes in each field's Gini coefficient and patent elasticity relative to 1990. The Gini coefficient rose for every field except Chemistry, with ICT showing the largest increase. The patent elasticity rose broadly as well – most notably for ICT, but also for Physics, EE & Electronics and Biology & Medicine. Even Mechanical Engineering, whose innovation is anchored in traditional manufacturing hubs, saw a modest increase, and Civil Engineering & Consumer Goods shows a delayed rise beginning around 2000. Chemistry is the one exception to the Gini trend, but its rising elasticity reveals an important distinction: Chemistry innovation was historically concentrated in the New York–New Jersey corridor and has since spread to other high-skill CZs, reducing overall spatial concentration while strengthening the link to skill composition. This broad-based pattern – rising within-field concentration across nearly all fields – indicates that the aggregate trend is not purely a compositional phenomenon.

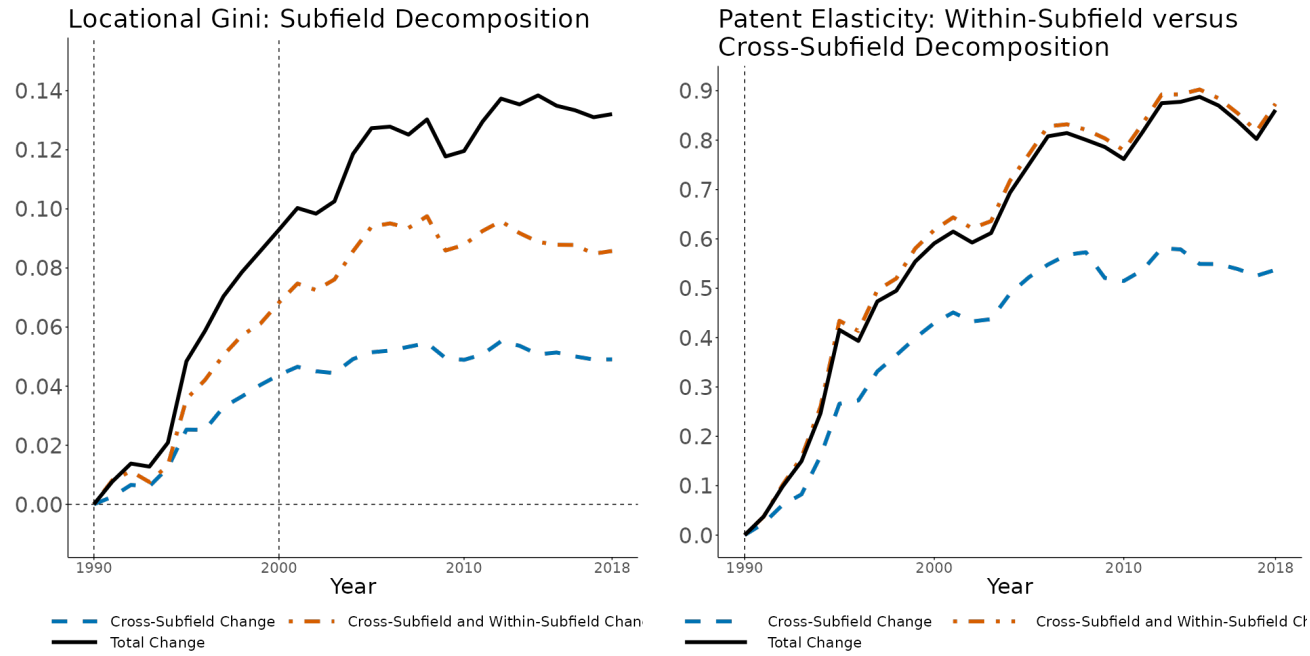


Figure 12: Subfield decomposition: locational Gini decomposition into within-subfield, cross-subfield, and colocation components (left), and patent elasticity decomposition into within-subfield and cross-subfield components (right).

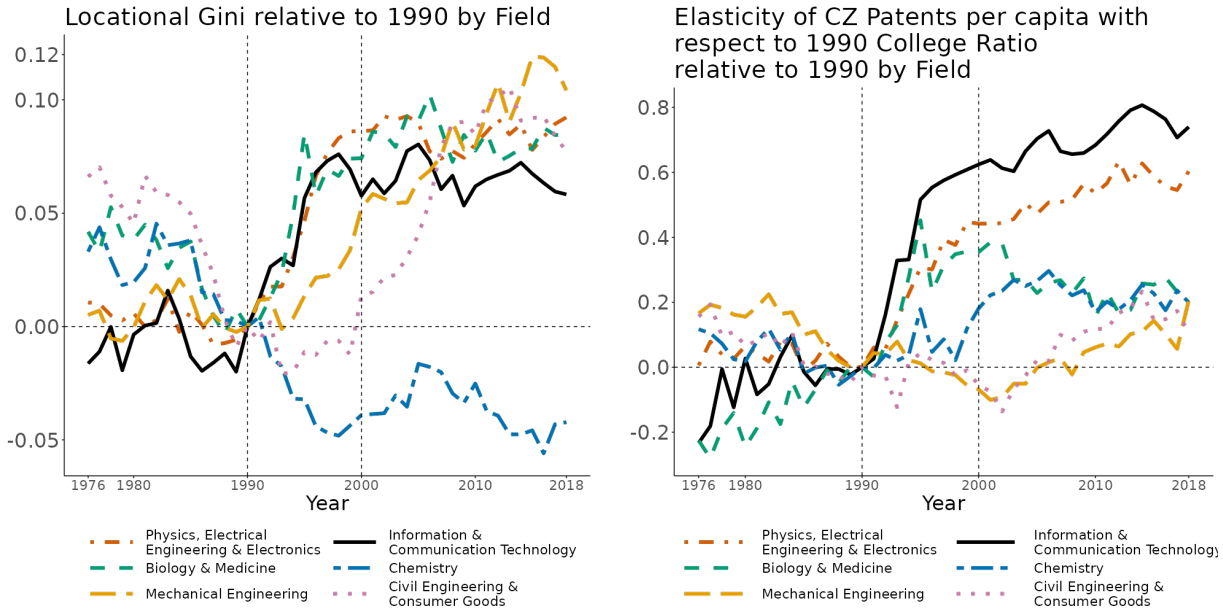


Figure 13: Spatial concentration of patents per capita by technology field: changes from 1990. Left: Gini coefficient. Right: elasticity with respect to the 1990 college ratio, estimated using PPML.

Figure 14 decomposes the within-field component by individual field. The within-field contributions are distributed broadly: non-ICT fields account for 36% of the Gini increase and 27% of the elasticity increase, while within-ICT concentration adds a further 4% and 14%, respectively.

Table 1 confirms this broad-based pattern at an even finer level, reporting the within-subfield contributions

Subfield	Field	Share 2015	Gini 2015	Elast. 2015	Gini Δ Across	Elast. Δ Across	Gini Δ Within	Elast. Δ Within
Computer technology	2	0.153	0.738	3.068	0.091	0.352	-0.002	0.039
Digital communication	2	0.106	0.731	3.184	0.066	0.266	-0.001	0.022
IT methods for management	2	0.030	0.649	2.604	0.017	0.068	0.001	0.007
Medical technology	3	0.093	0.565	1.688	0.014	0.041	0.008	0.020
Pharmaceuticals	3	0.027	0.573	2.230	0.009	0.032	-0.000	0.004
Audio-visual technology	1	0.033	0.704	2.711	-0.001	-0.009	0.003	0.037
Electrical machinery, apparatus, energy	1	0.050	0.546	1.579	-0.003	-0.008	0.003	0.032
Measurement	1	0.041	0.558	2.028	-0.008	-0.027	0.005	0.030
Optics	1	0.019	0.667	2.068	-0.011	-0.029	0.000	0.027
Organic fine chemistry	4	0.019	0.653	2.132	-0.019	-0.053	-0.001	0.025
Telecommunications	2	0.026	0.647	2.585	0.003	0.009	0.000	0.017
Transport	5	0.043	0.585	0.296	0.002	0.003	0.006	0.015
Basic communication processes	2	0.011	0.773	2.788	-0.005	-0.017	0.001	0.011
Control	1	0.025	0.568	1.525	0.007	0.018	0.001	0.010
Other special machines	5	0.028	0.419	0.396	-0.008	-0.005	0.001	0.008
Chemical engineering	4	0.015	0.495	1.206	-0.009	-0.017	0.000	0.007
Biotechnology	3	0.025	0.662	2.702	0.004	0.013	0.000	0.005
Surface technology, coating	4	0.008	0.597	1.386	-0.006	-0.013	0.000	0.005
Handling	5	0.017	0.446	0.758	-0.010	-0.009	0.000	0.005
Furniture, games	6	0.028	0.453	0.984	-0.007	-0.014	0.003	0.004
Civil engineering	6	0.030	0.539	0.653	-0.006	-0.006	0.002	0.004
Machine tools	5	0.012	0.537	0.389	-0.011	-0.011	0.002	0.004
Materials, metallurgy	4	0.007	0.603	0.911	-0.009	-0.010	-0.000	0.003
Textile and paper machines	5	0.006	0.719	0.669	-0.010	-0.009	0.002	0.003
Engines, pumps, turbines	5	0.024	0.702	0.467	-0.004	-0.003	0.002	0.003
Micro-structural and nano-technology	4	0.001	0.835	2.819	0.001	0.001	-0.000	0.001
Analysis of biological materials	3	0.007	0.593	2.150	-0.000	-0.001	-0.001	0.001
Other consumer goods	6	0.018	0.470	1.025	-0.004	-0.009	0.000	0.001
Mechanical elements	5	0.019	0.560	0.183	-0.009	-0.004	-0.000	0.001
Macromolecular chemistry, polymers	4	0.007	0.678	0.513	-0.015	-0.010	-0.001	0.001
Environmental technology	4	0.009	0.515	0.579	-0.004	-0.005	0.000	0.001
Thermal processes and apparatus	5	0.008	0.543	0.680	-0.003	-0.002	0.000	0.000
Basic materials chemistry	4	0.015	0.563	1.055	-0.010	-0.016	-0.002	-0.002
Semiconductors	1	0.032	0.841	2.347	0.007	0.024	0.002	-0.007
Food chemistry	4	0.008	0.775	0.105	0.001	-0.001	0.001	-0.007

Table 1: All subfields: 2015 patent share, Gini, and elasticity levels, and contributions to the within- and across-subfield components of the post-1990 changes in the locational Gini and patent elasticity. Top 5 sorted by elasticity change across subfields; remaining sorted by elasticity change within subfield. Field codes: 1 = Physics, Electrical Engineering & Electronics; 2 = Information & Communication Technology; 3 = Biology & Medicine; 4 = Chemistry; 5 = Mechanical Engineering; 6 = Civil Engineering & Consumer Goods.

to the Gini increase (Gini Δ Within column) across all 35 subfields. These contributions are positive for the large majority of subfields and span all six technology fields. The largest within-subfield contributors include *Medical Technology* (Biology & Medicine), *Transport* (Mechanical Engineering), *Measurement* and *Audio-visual Technology* (Physics, EE & Electronics) – a diverse set with no single field dominating.

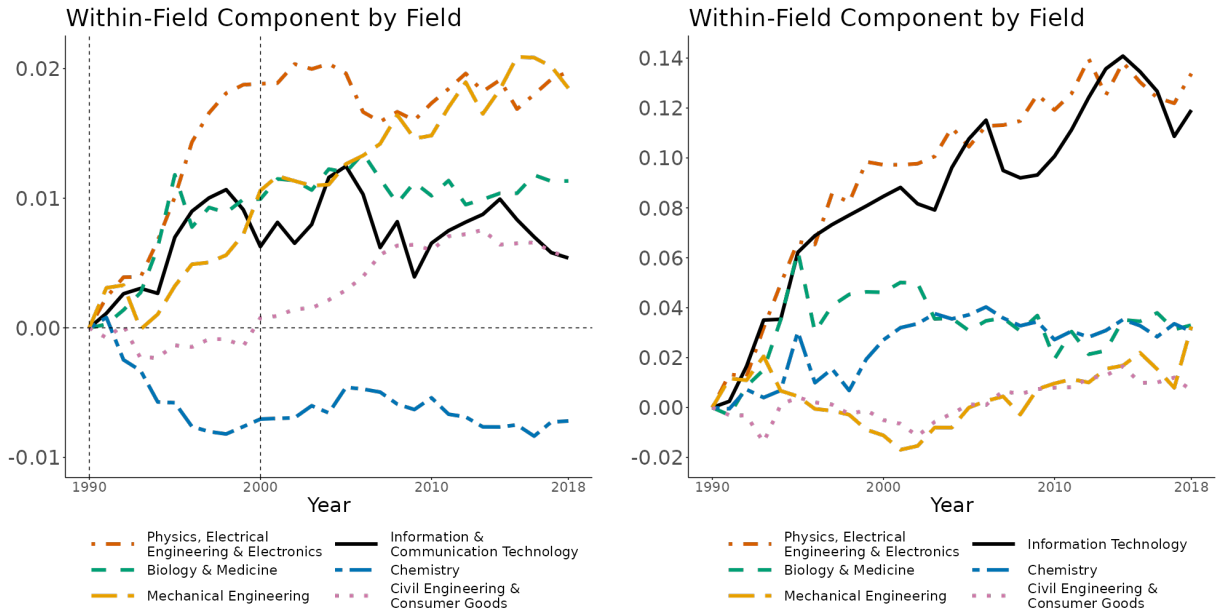


Figure 14: Within-field components of the post-1990 change in spatial concentration, by individual field. Left: locational Gini. Right: elasticity with respect to the 1990 college ratio.

The subfield decomposition also reveals another component of the within-field rise. The left panel of Figure 12 shows that about one-third of the post-1990 increase in the aggregate locational Gini is driven by the increasing colocation of subfields – that is, the tendency for different subfields to concentrate in the same CZs. Because the colocation component is only 8% at the field level (Figure 10), this colocation occurred primarily *between subfields within the same field*: for example, different Chemistry subfields increasingly co-locating in the same CZs, or different Physics subfields doing so. In other words, what appears as within-field concentration at the field level partly reflects cross-subfield colocation within fields. Taken together, the within-field half of the post-1990 concentration rise has two features: it is broad-based across fields and subfields, and a substantial portion reflects the increasing colocation of related subfields within the same regions.

5 Further Decomposition by Firm

Section 4 showed that roughly half of the post-1990 rise in spatial concentration reflects a compositional shift toward ICT (Fact 4), while the other half reflects rising concentration within fields and subfields (Fact 5). This section takes the decomposition one layer further, examining whether each channel was driven by within-firm changes, between-firm reallocation, or the entry of new firms. As I show below, between-firm reallocation is the dominant force in both cases.

5.1 Fact 6: The compositional shift toward ICT (Fact 4) was driven by the growth of *ICT-oriented firms*, which disproportionately concentrate their patents in high-skill CZs

5.1.1 Firm decomposition of the rise in the aggregate ICT patent share

I begin by examining whether the post-1990 rise in the aggregate ICT patent share, s_t, reflects increases in ICT intensity within firms or instead the growing weight of firms that predominantly produce ICT patents. To appropriately account for firm entry and exit, I implement a modified Foster-Haltiwanger-Krizan (FHK) decomposition (Foster et al., 2001):[4]

$$s_{t^*} - s_{1990} = \sum_{t=1991}^{t=t^*} \Delta s_t = \sum_{t=1991}^{t=t^*} \Bigg[\underbrace{\sum_{f \in C} \overline{w}_{f,t} \Delta s_{f,t}}_{\text{within-firm}} + \underbrace{\sum_{f \in C} \left(\overline{s}_{f,t} - \overline{S}_t\right) \Delta w_{f,t}}_{\text{between-firm}} + \underbrace{\sum_{f \in E} w_{f,t} \left(s_{f,t} - \overline{S}_t\right)}_{\text{entry}} + \underbrace{\sum_{f \in X} w_{f,t-1} \left(s_{f,t-1} - \overline{s}_t\right)}_{\text{exit}} + \underbrace{\Delta\Big(s_t - S_t\Big)}_{\text{residual}} \Bigg], \tag{4}$$

where $w_{f,t}$ is firm f's patent weight (its share of total patents in year t), $s_{f,t}$ is the share of firm f's patents in ICT, and $S_t = \sum_f w_{f,t} s_{f,t}$ is the weighted average of firm-level ICT shares in year t. For any variable x, $\overline{x}_t \equiv \frac{x_t + x_{t+1}}{2}$ and $\Delta x_t \equiv x_t - x_{t-1}$. The sets C, E, and X denote continuing, entering, and exiting firms, respectively. The five terms capture: (i) changes in ICT shares within firms, (ii) compositional changes among incumbents, (iii) entry, (iv) exit, and (v) a residual reflecting patents without assignees. The left panel of Figure 15 plots this decomposition and shows that the post-1990 increase in the aggregate ICT patent share is explained almost entirely by the between-firm component (dot-dashed blue line): the ICT share within individual assignees remains relatively stable over time, while ICT-oriented firms expand their overall share of patenting.

To validate the firm-level decomposition, I conduct two complementary group-level Shapley decompositions. First, I classify firms into two groups – ICT-oriented and non-ICT-oriented – based on their long-run ICT intensity. For each firm, I compute the weighted average of its annual ICT patent shares, using as weights the firm's relative national patent share across years. Firms with an average ICT share above the median are designated as ICT-oriented. I then implement the following group-level Shapley decomposition as follows:

$$s_{t^*} - s_{1990} = \sum_{t=1991}^{t=t^*} \Delta s_t = \sum_{t=1991}^{t=t^*} \Bigg[\underbrace{\sum_{g} \overline{w}_{g,t} \Delta s_{g,t}}_{\text{within-group}} + \underbrace{\sum_{g} \overline{s}_{g,t} \Delta w_{g,t}}_{\text{between-group}} + \underbrace{\Delta\Big(s_t - \sum_{g} s_{g,t} w_{g,t}\Big)}_{\text{residual}} \Bigg], \tag{5}$$

[4] I use the FHK decomposition rather than the Shapley decomposition here because the Shapley entry and exit terms comprise the elasticity *level*, making them not directly comparable to the between and within terms, which capture *changes*.

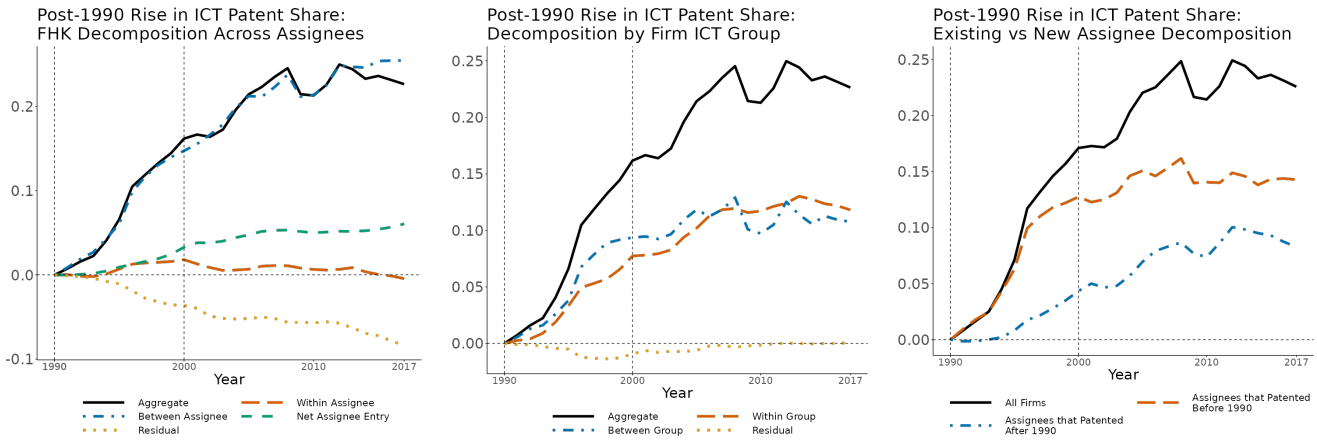


Figure 15: Decomposition of the aggregate rise in the ICT patent share by within versus across firm (left), by firm long-run ICT intensity group (middle) and firm entry cohort group (right).

where $w_{g,t}$ is group g's patent weight (its share of total patents in year t) and $s_{g,t}$ is the share of group g's patents in ICT. The middle panel of Figure 15 shows that, even with this coarse two-group classification, the rising patent share of ICT-oriented firms (dot-dashed blue line) accounts for more than half of the post-1990 increase in the aggregate ICT patent share (solid black line).[5]

Second, I classify firms by entry cohort – those that began patenting before 1990 versus after – and apply the same Shapley framework. I decompose the post-1990 increase in the aggregate ICT share into (i) growth in ICT shares among pre-1990 firms; (ii) growth in ICT shares among post-1990 entrants; and (iii) compositional changes due to the rising patent weight of post-1990 entrants. I attribute (ii) and (iii) jointly to post-1990 entrants. The right panel of Figure 15 shows that about three-quarters of the post-1990 increase in the aggregate ICT patent share is driven by firms that began patenting before 1990.

Taken together, both group-level decompositions confirm the firm-level FHK decomposition results: the post-1990 rise in the aggregate ICT patent share was mostly driven by the growing weight of existing ICT-oriented firms, rather than by within-firm changes or new entrants.

5.1.2 Firm decomposition of the differences in spatial concentration between ICT and non-ICT innovation

The growth of ICT-oriented firms explains the compositional shift, but why does this raise spatial concentration? I next examine whether the greater concentration of ICT innovation in high-skill CZs, relative to non-ICT, reflects within-firm or between-firm differences. Using the firm classification by long-run ICT intensity introduced above, I apply a Shapley decomposition of the aggregate ICT–non-ICT elasticity gap

[5] I also consider an alternative firm classification by ICT intensity. Using annual ICT and non-ICT patent weights for each firm, I define three groups: (i) firms consistently ICT-oriented across all years, (ii) firms consistently non-ICT-oriented, and (iii) mixed firms. Decomposition results are similar under this classification and available upon request.

by year.[6]

$$\alpha_{\text{ICT,t}} - \alpha_{\text{nonICT,t}} = \underbrace{\sum_f \overline{w}_{f,t} \left(\alpha_{f,\text{ICT},t} - \alpha_{f,\text{nonICT},t}\right)}_{\text{within-firm}} + \underbrace{\sum_f \overline{\alpha}_{f,t} \left(w_{f,\text{ICT},t} - w_{f,\text{nonICT},t}\right)}_{\text{between-firm}} + \underbrace{\varepsilon}_{\text{residual}} \tag{6}$$

where $w_{f,k,t}$ denotes firm f's patent weight in sector $k = \{\text{ICT}, \text{non-ICT}\}$ in year t and $\alpha_{f,k,t}$ represents the corresponding elasticity. The first term captures within-firm differences, the second term captures between-firm differences, and the residual reflects patents from firms with undefined elasticities. The left panel of Figure 16 shows that the aggregate ICT–non-ICT elasticity gap is driven more by the between-firm component, reflecting systematic cross-firm differences in patent geography rather than within-firm differences between ICT and non-ICT patents.

Because firm-level elasticities can be noisy, I complement this with group-level decompositions. The middle panel of Figure 16 displays group-level ICT and non-ICT patent elasticities for ICT-oriented versus non-ICT-oriented firms, showing that non-ICT patents produced by ICT-oriented firms are more concentrated in high-skill CZs than ICT patents produced by non-ICT-oriented firms. A group-level Shapley decomposition of the aggregate ICT–non-ICT elasticity gap (right panel) confirms that between-group differences dominate within-group differences. Together, these results indicate that cross-firm differences in patent geography, rather than within-firm differences, account for the bulk of the ICT–non-ICT elasticity gap.

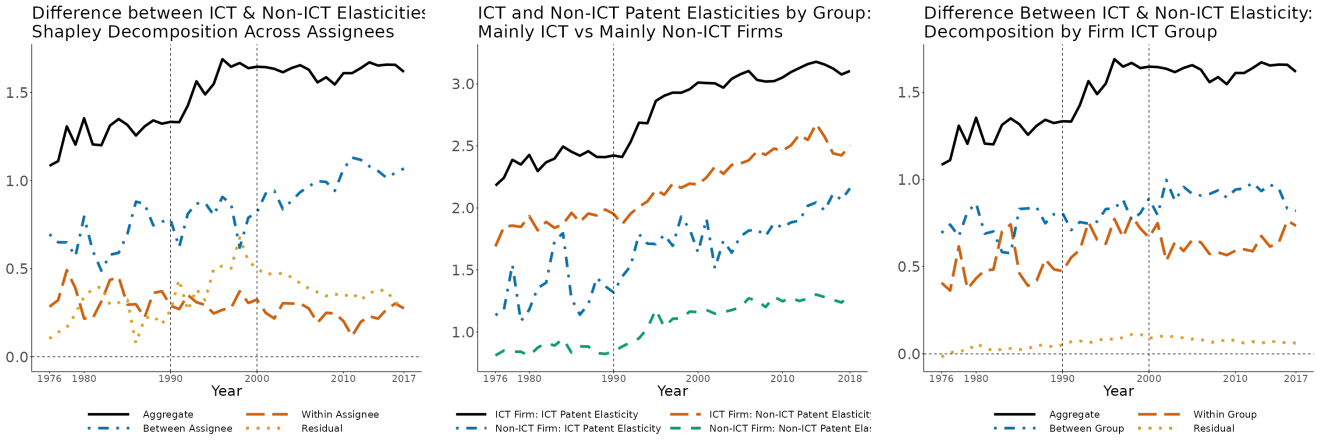


Figure 16: Decompositions of the aggregate ICT–non-ICT patent elasticity gap. Left: firm-level Shapley decomposition into within-firm and between-firm components. Middle: group-level ICT and non-ICT patent elasticities for ICT-oriented versus non-ICT-oriented firms. Right: group-level Shapley decomposition by firm ICT intensity.

5.2 Fact 7: The within-field and subfield rise in concentration (Fact 5) was primarily driven by the growth of *high-skill CZ firms*

Fact 5 showed that spatial concentration rose within fields and subfields, accounting for roughly half of the aggregate increase. I now examine whether this within-field rise reflects changes within firms or reallocation

[6] I restrict the sample to assignee-years where both ICT and non-ICT patent elasticities are well defined (i.e., the assignee has ICT and non-ICT patents in at least 2 CZs in that year), covering 52% of all patents.

across firms. Because the within-field component includes contributions from both ICT and non-ICT fields – which may be driven by different firm-level dynamics – I conduct the decomposition separately for ICT and non-ICT innovation. For each case, I implement a modified FHK decomposition, analogous to that used for the post-1990 increase in the aggregate ICT patent share.[7]

$$\alpha_{k,t^*} - \alpha_{k,1990} = \sum_{t=1991}^{t=t^*} \Delta\alpha_{k,t} = \sum_{t=1991}^{t=t^*} \Bigg[\underbrace{\sum_{f\in C} \overline{w}_{f,k,t}\Delta\alpha_{f,k,t}}_{\text{within-firm}} + \underbrace{\sum_{f\in C} \left(\overline{\alpha}_{f,k,t} - \overline{A}_{k,t}\right)\Delta w_{f,k,t}}_{\text{between-firm}} + \underbrace{\sum_{f\in E} w_{f,k,t}\left(\alpha_{f,k,t} - \overline{A}_{k,t}\right)}_{\text{entry}} + \underbrace{\sum_{f\in X} w_{f,k,t-1}\left(\alpha_{f,k,t-1} - \overline{A}_{k,t}\right)}_{\text{exit}} + \underbrace{\Delta\Big(\alpha_{k,t} - A_{k,t}\Big)}_{\text{residual}} \Bigg], \tag{7}$$

where $\alpha_{k,t}$ is the aggregate patent elasticity, $\alpha_{f,k,t}$ is firm f's patent elasticity, $A_{k,t} = \sum_f w_{f,k,t}\alpha_{f,k,t}$ is the weighted average of firm-level elasticities in sector $k \in \{\text{ICT}, \text{non-ICT}\}$ and year t. Other notation follows equation 4.

Because firm-level elasticities can be noisy, I complement the FHK decomposition with two group-level Shapley decompositions. First, I classify firms into elasticity quartiles based on their long-run average patent elasticity in sector k and implement the following group-level Shapley decomposition:

$$\alpha_{k,t^*} - \alpha_{k,1990} = \sum_{t=1991}^{t=t^*} \Delta\alpha_t = \sum_{t=1991}^{t=t^*} \Bigg[\underbrace{\sum_{g} \overline{w}_{g,k,t}\Delta\alpha_{g,k,t}}_{\text{within-group}} + \underbrace{\sum_{g} \overline{\alpha}_{g,k,t}\Delta w_{k,t}}_{\text{between-group}} + \underbrace{\Delta\Big(\alpha_{k,t} - \sum_{g} w_{g,k,t}\alpha_{g,k,t}\Big)}_{\text{residual}} \Bigg] \tag{8}$$

where g indexes the four quartiles of firms and other notation follows equation 7.[8] Second, I decompose by firm entry cohort. Using the cohort classification introduced above, I apply a similar group-level Shapley framework. Specifically, I separate the rise into three components: (i) changes among pre-1990 firms, (ii) changes among post-1990 entrants, and (iii) compositional shifts reflecting the growing patent weight of post-1990 entrants. I attribute (ii) and (iii) jointly to post-1990 entrants.

5.2.1 Firm decomposition of the rise in concentration within ICT

Figure 17 reports results for ICT. The left panel shows the FHK decomposition: between-firm reallocation explains more than two-thirds of the post-1990 rise in the aggregate ICT patent elasticity, while within-firm changes contribute little. The middle and right panels report the two group-level Shapley decompositions. The middle panel shows that most of the rise reflects the growing share of patents from higher-elasticity

[7] I restrict the sample to firm-years with defined firm-level elasticities, covering approximately 48% of ICT patents and 52% of non-ICT patents. The group-level decompositions below address this coverage limitation.

[8] In principle, there is an additional group of firms with undefined elasticities because they patent in only one CZ in each year they patent. This group accounts for a small share of patents, and dropping it does not change the decomposition results. Decomposition results excluding this group are available on request.

firms, rather than increases in elasticity within groups.[9] The right panel shows that firms patenting in ICT before 1990 account for about two-thirds of the rise in ICT elasticity during the 1990s and about half thereafter.

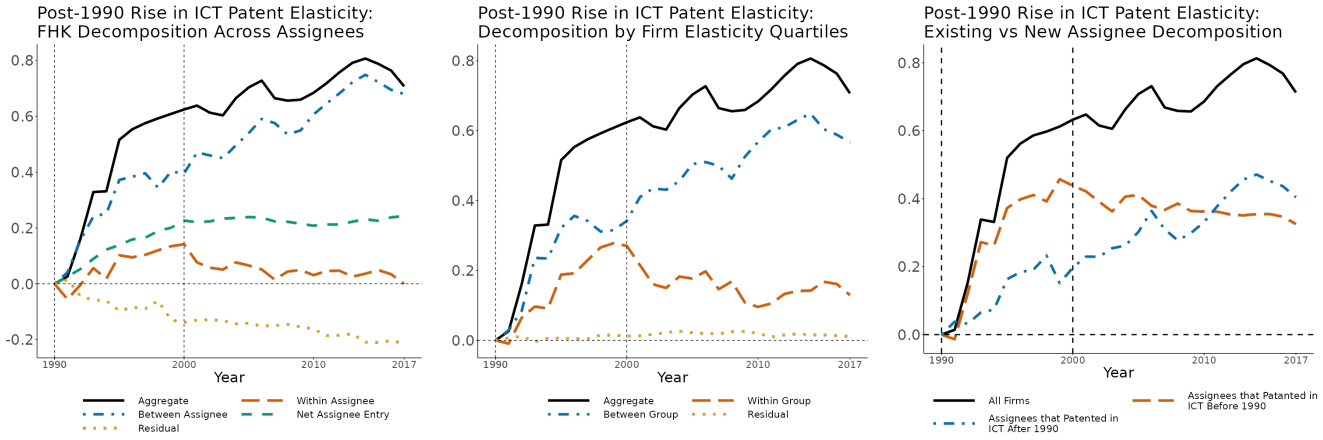


Figure 17: Decompositions of the post-1990 rise in the aggregate ICT patent elasticity. Left: FHK decomposition into within-firm, between-firm, entry, exit, and residual components. Middle: group-level Shapley decomposition by firm ICT elasticity quartiles. Right: group-level Shapley decomposition by firm entry cohort.

5.2.2 Firm decomposition of the rise in concentration within non-ICT

Figure 18 reports the analogous decompositions for non-ICT. The left panel shows that between-firm reallocation accounts for nearly all of the post-1990 rise in the aggregate non-ICT patent elasticity – an even more dominant role than for ICT. The middle panel indicates that most of the rise reflects compositional shifts toward firms in higher-elasticity groups, while the right panel shows that incumbent firms with non-ICT patents prior to 1990 account for the bulk of the increase.

Across both ICT and non-ICT, the group-level decompositions confirm that the post-1990 increase in concentration within fields is driven mainly by the rising patent shares of firms already concentrated in high-skill CZs. A key remaining question is whether this cross-firm reallocation reflects the growing importance of ICT-oriented firms (which also patent in non-ICT and have high elasticities), or instead reallocation toward higher-elasticity firms within both ICT-oriented and non-ICT-oriented groups. To distinguish between these explanations, I decompose the post-1990 rise in the aggregate ICT and non-ICT patent elasticities into the contributions of ICT-oriented versus non-ICT-oriented firms. In Figure 19, most of the post-1990 rise in the non-ICT patent elasticity is driven by non-ICT-oriented firms (left panel), while that for the ICT patent elasticity is driven by ICT-oriented firms (middle panel). The right panel reports results from an FHK decomposition of the post-1990 increase in the non-ICT patent elasticity, using a sample restricted to non-ICT-oriented firms. These results show reallocation toward higher-elasticity firms within both the ICT- and non-ICT-oriented groups.

[9]The within-group component accounts for just under half of the aggregate increase before 2000, in contrast to the much smaller within-firm component in the FHK decomposition. This reflects reallocation of patents across firms within each group.

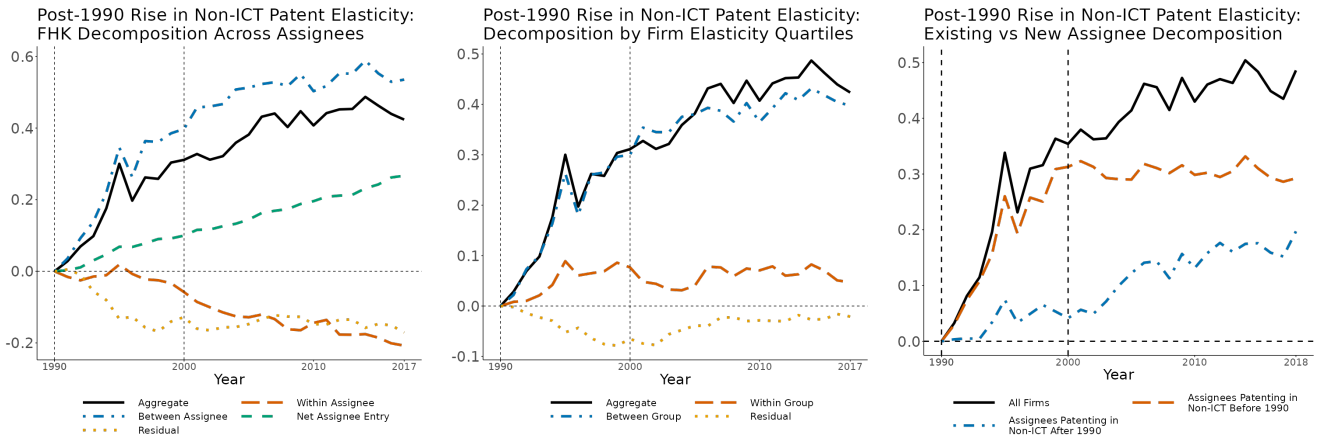


Figure 18: Decompositions of the post-1990 rise in the aggregate non-ICT patent elasticity. Left: FHK decomposition into within-firm, between-firm, entry, exit, and residual components. Middle: group-level Shapley decomposition by firm non-ICT elasticity quartiles. Right: group-level Shapley decomposition by firm entry cohort.

This finding has an important implication: the within-field rise in concentration (Fact 5) is not simply a byproduct of the ICT compositional shift documented in Fact 6. Rather, it reflects a distinct force – the reallocation of patenting toward high-skill CZ firms – that operates broadly across both ICT and non-ICT innovation.

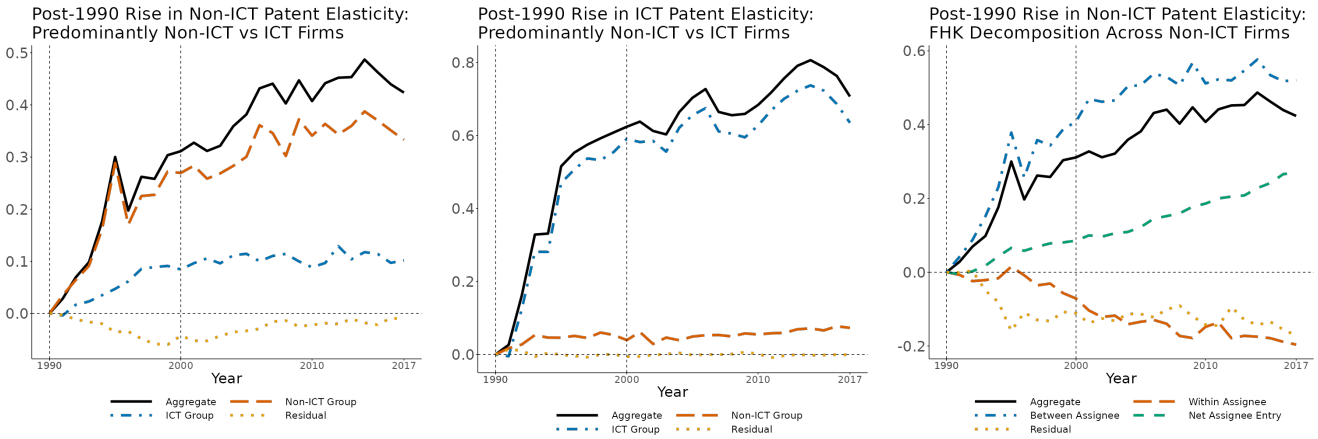


Figure 19: Left: decomposition of the post-1990 rise in non-ICT patent elasticity into contributions from ICT-oriented and non-ICT-oriented firm groups. Middle: analogous decomposition for ICT patent elasticity. Right: FHK decomposition of the post-1990 rise in non-ICT elasticity using a sample restricted to non-ICT-oriented firms.

6 Descriptive Evidence on the Underlying Mechanisms and the ICT Shock

The decompositions in Sections 4 and 5 traced the post-1990 rise in the spatial concentration of innovation to the growth of two groups of firms: *ICT-oriented firms*, whose rising patent shares drove the compositional shift toward ICT (Fact 6), and *high-skill CZ firms*, whose rising patent shares drove the within-field

concentration increase (Fact 7). This raises two questions: What drove the post-1990 expansion of these firms? And why do ICT-oriented firms concentrate their patents in high-skill CZs?

I provide preliminary descriptive evidence below that their growth was likely shaped by two components of the ICT shock: (i) the sudden emergence of influential ICT breakthroughs, which catalyzed the surge in ICT patenting and the growth of ICT-oriented firms through the *colocation* of ICT innovation with ICT production in high-skill CZs (Section 6.2), and (ii) the subsequent decline in communication costs, which enabled high-skill CZ firms to expand geographically and increase their scale and innovation output through an *asymmetric scale effect* (Section 6.3). I also briefly discuss other potential factors and argue they are unlikely to be primary drivers (Section 6.4).[10]

6.1 ICT Innovation Breakthroughs

The ICT share of all patents rose from roughly 8% in 1990 to over 32% by 2015, driving the compositional shift documented in Sections 4 and 5.1. What triggered this surge? I provide evidence that the sudden emergence of influential ICT breakthroughs played a key role. I identify breakthrough ICT patents as those in the top 1% of the *originality* distribution across all patents, where originality is measured as the ratio of forward to backward textual similarity following Kelly et al. (2021). Between 1976 and 2018, 880,927 ICT patents were produced, of which 27,419 qualify as breakthroughs. The left panel of Figure 20 traces the evolution of their influence – defined as the number of patents that cite each breakthrough as their first breakthrough, following citations up to two degrees. The figure reveals a sharp rise in the influence of breakthrough ICT patents between 1984 and 1993, coinciding with the onset of the post-1990 concentration trend documented in Section 3.1. These breakthroughs catalyzed the subsequent surge in ICT patenting. The right panel of Figure 20 displays the corresponding trend in the number of ICT patents that cite these breakthroughs (within two degrees). Across all years, 78% of all ICT patents cite a breakthrough ICT patent within two degrees.

6.2 Colocation of ICT Innovation with Production in High-Skill CZs

This wave of ICT innovation disproportionately increased patenting in *ICT-oriented firms*, rather than uniformly raising ICT patenting across all firms, as shown in Section 5.1. Most ICT-oriented firms operate plants

[10]To illustrate these mechanisms, consider *ICT-oriented firms* such as IBM and Hewlett-Packard (HP) alongside *high-skill CZ firms* such as United Technologies, Exxon Mobil, Boston Scientific, and Boeing. Before 1985, IBM concentrated both production and R&D in the highly educated Northeast corridor, with additional facilities in Boulder, CO; Rochester, MN; and Silicon Valley. HP located nearly all of its manufacturing, engineering, and managerial capacity in Silicon Valley. In both cases, innovation was colocated with production in high-skill CZs. The sudden arrival of influential ICT breakthroughs in the mid-1980s triggered sharp increases in patenting at IBM and HP, much of which remained colocated near their existing production sites. Beginning in the late 1980s – facilitated by falling communication costs – these firms, as well as high-skill CZ firms more broadly, began expanding operations across the country while simultaneously increasing their R&D employment and patenting activity. United Technologies broadened its footprint beyond its Connecticut core into engineering hubs and lower-skill regions; Exxon Mobil expanded from Houston and Northern Virginia into both high-skill and lower-skill locations; Boston Scientific and Boeing exhibit similar patterns of mixed expansion as their national networks scaled.

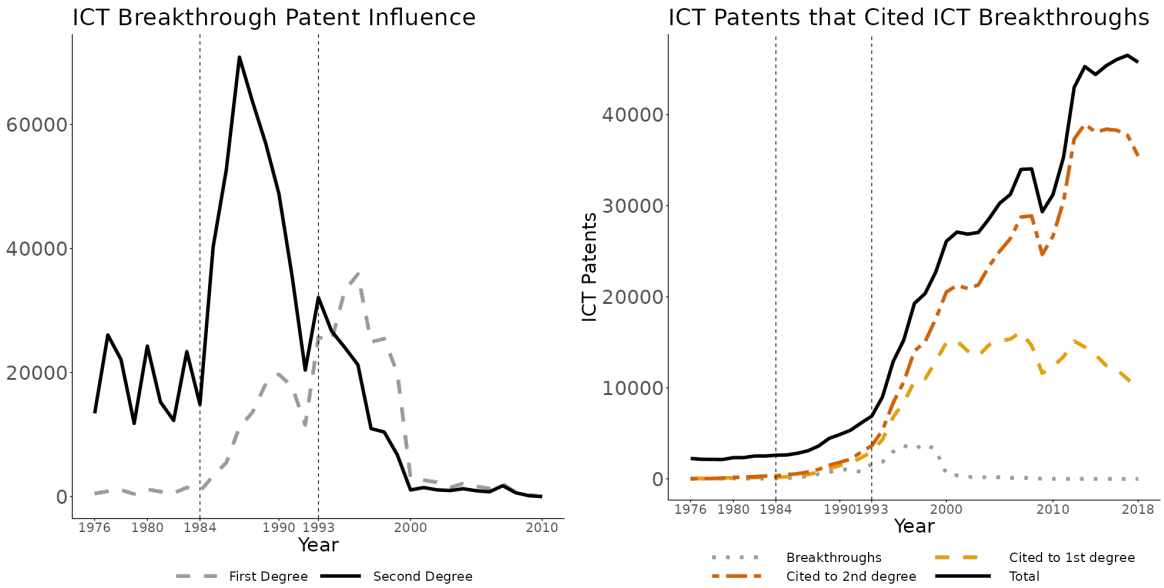


Figure 20: Origins of the ICT shock. The left panel shows the influence of these breakthrough patents, defined as the number of patents that cite each breakthrough as their first breakthrough, tracing citations up to two degrees. The right panel depicts trends in the number of ICT patents that cited these patents.

in ICT service industries – Software Publishers, Telecommunications, and Data Processing and Hosting[11] – suggesting that having an ICT service plant is a key input in the production of ICT innovation.

ICT service plants disproportionately employ high-skill workers and are therefore heavily concentrated in high-skill CZs. The left panel of Figure 21 shows that the share of college-educated workers in the ICT service industry (solid black line) is substantially higher than that in other industries (dashed blue line) across all years. The right panel of Figure 21 plots the relationship between a CZ's 1990 college ratio and its pre-1985 ICT employment share; the positive slope indicates that high-skill CZs already had disproportionately large ICT service employment before 1985.

The growth of ICT patents primarily occurred in CZs where the ICT industry was initially more concentrated. To show this, I estimate the following CZ-year specification, analogous to the first stage of a shift-share design where the "shift" is national ICT patent growth and the "share" is initial local ICT presence:

$$\text{Outcome}_{r,t} = \beta \cdot \text{ Pre-1985 ICT Employment Share}_r \times \text{Growth of National ICT Patents}_t + \gamma_r + \gamma_t + \varepsilon_{r,t} \quad (9)$$

where Growth of National ICT Patents$_t$ is the growth of national patents in year t relative to 1985, and γ_r and γ_t are CZ and year fixed effects. Specification (1) shows that CZs with a higher initial ICT employment share experienced substantially larger increases in ICT patenting. A 1 p.p. higher pre-1985 ICT employment share is associated with 8.3 more ICT patents in years when national ICT patenting is 10% above its 1985

[11] These correspond to NAICS 5112, 517, and 518 within the Information Sector (NAICS 51). In contrast to Fort et al. (2020), I categorize these industries as ICT *production*, rather than *innovation*, because establishments that focus on R&D activities are instead classified as Scientific Research and Development Services (5417) and occasionally Corporate, Subsidiary, and Regional Managing Offices (551114) under NAICS.

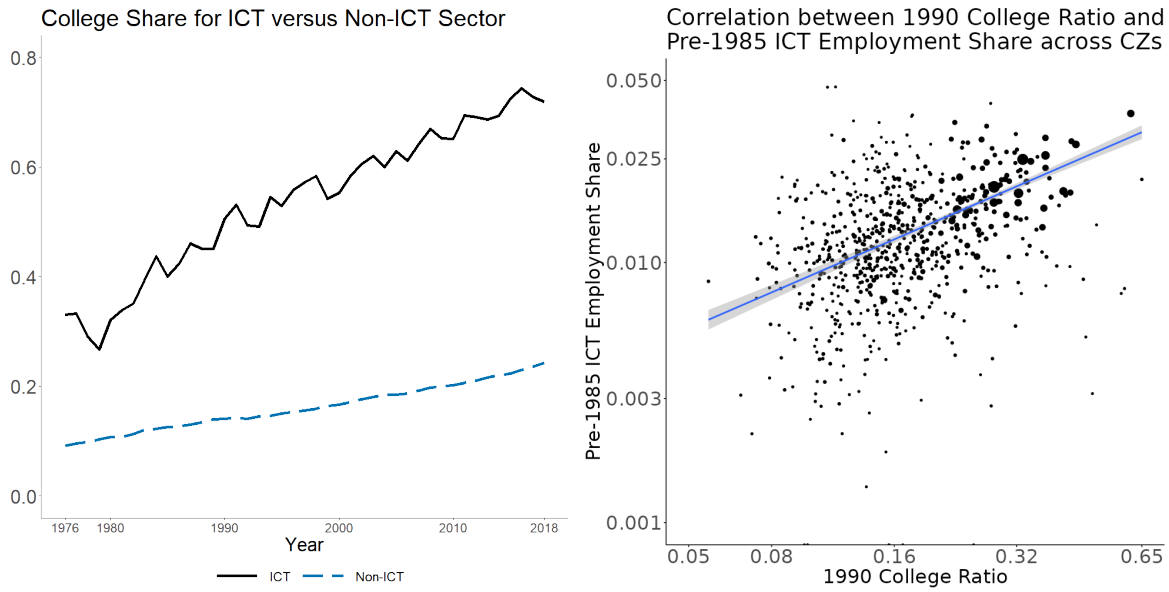


Figure 21: The ICT sector disproportionately employs college-educated workers and is concentrated in high-skill CZs. The left panel displays trends in the college share for the ICT and non-ICT sectors. The right panel shows the correlation of the pre-1985 ICT employment share against the 1990 college ratio. Each dot represents a commuting zone (CZ), where the size of the dot captures the CZ's population in 1990.

level. Specification (2) confirms that the geographical incidence of ICT patents is not driven merely by population growth: the same 1 p.p. difference predicts 1.5 additional ICT patents per thousand residents under a 10% national ICT growth shock. Turning to employment outcomes, Specification (3) shows that ICT employment also increased disproportionately in CZs where the ICT industry was initially more concentrated: a 1 p.p. higher pre-1985 ICT employment share is associated with 77 additional ICT workers when national ICT patenting rises by 10%. Taken together, the parallel rise in ICT employment and ICT patents per capita in the same CZs provides evidence of the colocation of ICT innovation and ICT production in regions with strong pre-existing ICT presence. This colocation explains why ICT-oriented firms concentrate their patents in high-skill CZs and, in turn, why the compositional shift toward ICT documented in Section 5.1 raised the aggregate spatial concentration of innovation. Finally, Specification (4) shows that these CZs did not experience larger increases in the share of local employment in ICT, likely because employment in other industries also expanded.

6.3 Firm Spatial Expansion

The colocation mechanism explains the growth of ICT-oriented firms and why they concentrate patents in high-skill CZs. But Section 5.2 also showed that *high-skill CZ firms* – including non-ICT-oriented ones – grew disproportionately after 1990. What drove their expansion?

The same ICT revolution that produced the breakthroughs documented in Section 6.1 also dramatically reduced communication costs. Table 3 lists the top five breakthrough ICT patents between 1984 and 1993 by citation influence; while most reflect advances in the application layer, patent US4644532 directly improved

Dependent Variables:	ICT Patents	ICT Patents per capita	ICT Employment	ICT Employment Share
Model:	(1)	(2)	(3)	(4)
Pre-1985 ICT Employment Share$_r$×	8,334.9***	1.474***	77,372.4***	-0.0247
Growth of National ICT Patents$_t$	(2,735.3)	(0.5416)	(28,464.3)	(0.0204)
Fixed-effects				
CZ	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
Observations	20,938	20,938	20,938	20,934
R^2	0.81282	0.81900	0.96195	0.83343
Within R^2	0.15934	0.06459	0.05407	0.01308

Clustered (CZ) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table 2: The geographical incidence of the post-1985 growth of national ICT patents.

distributed network routing, enabling more robust information systems across locations. More broadly, the transition from proprietary systems to open standards – Ethernet, TCP/IP, relational databases, and client–server architectures – occurred during the early-to-mid 1990s, when firms achieved seamless interconnection across plants, offices, and data systems (Bresnahan and Greenstein, 1996; Mowery and Simcoe, 2002). In parallel, the deployment of the National Science Foundation Network (NSFNET) between 1986 and 1995 dramatically increased the availability of high-speed Internet access (Greenstein and Prince, 2006; Greenstein, 2015).

Patent No.	**Year**	**Assignee**	**Title**	**2nd Deg.**	**1st Deg.**
4827508	1985	Personal Library	Database usage metering and protection system and device	5135	671
5172338	1990	Sandisk	Multi-state EEprom read and write circuits and methods	3771	639
4644532	**1985**	**IBM**	**Automatic update of topology in a hybrid network**	**3571**	**223**
5005122	1987	DEC/HP	Arrangement with cooperating management server node and network service node	2969	1068
4831526	1986	Chubb	Computerized insurance premium quote request and response system	2528	333

Table 3: Top 5 Breakthrough ICT Patents from 1984-1993 by Second-Degree Citation Influence

As communication costs fell, firms expanded across many regions within the United States. This pattern is well documented: Jiang (2024) shows that firm-level intranet adoption causally increased the geographic scope of multi-unit manufacturing firms; Kleinman (2025) finds a steady post-1990 rise in the number of establishments per service firm; and Hsieh and Rossi-Hansberg (2023) document a sharp increase in MSAs per firm after 1987, likely reflecting ICT diffusion and complementary management practices. Using the raw D&B files, I find closely related trends. The left panel of Figure 22 shows that the employment and sales shares of multi-CZ firms increased from about 94% to 96.5% between 1990 and 2018. The right panel shows that the employment- and sales-weighted mean number of CZs per firm rose from roughly 100 to 180 over

the same period.

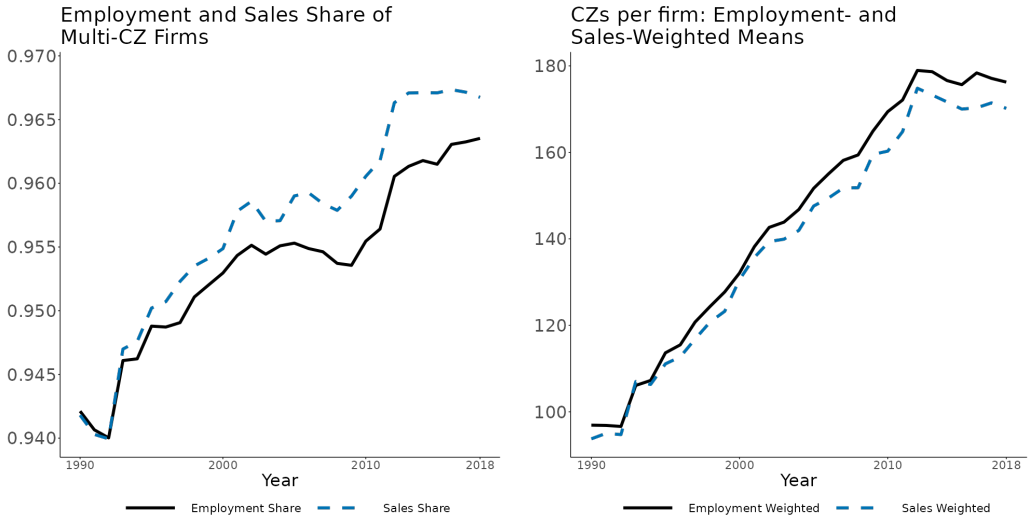


Figure 22: Trends in the *intensity* of firm spatial expansion from 1990 to 2018. The left panel displays trends in the employment and sales share of firms with establishments in multiple CZs. The right panel shows trends in the employment- and sales-weighted mean number of CZs per firm.

I next relate firm spatial expansion to the growing patent shares of high-skill CZ firms documented in Section 5.2. In Figure 23, the left panel shows that firms with initial R&D capabilities – defined as those operating a plant in Professional, Scientific, and Technical Services (NAICS 54) or Management of Companies and Enterprises (NAICS 55) between 1990 and 1992 – had a broader geographic footprint in 1990 and expanded more uniformly across high-skill and low-skill CZs thereafter than firms without such capabilities. The right panel shows that firms with R&D operations in high-skill CZs – where over 90% of all R&D-plant employment is located in every year – expanded geographically far more than firms with R&D plants in low-skill CZs between 1990 and 2005.

Importantly, firms that expanded into more CZs subsequently increased their R&D activity in high-skill CZs. Table 4 indicates that expanding a firm's geographic footprint by 10 CZs is associated with a 0.6-percentage-point increase in the likelihood of having an R&D plant (Specification 1) and a 1.1-percent increase in R&D employment (Specification 2) in the following year. Specification (3) shows that virtually all of this additional R&D employment occurred in CZs in the top skill quartile. This pattern complements the findings of Hsieh and Rossi-Hansberg (2023), who document that much of the growth in markets per firm occurred in finance, retail, and other services, where leading firms expanded into smaller MSAs and simultaneously increased employment in their R&D establishments. These results point to an *asymmetric scale effect*: high-skill CZ firms can exploit lower production and service costs in other regions, expanding their scale and innovation output, whereas low-skill CZ firms already operate in relatively low-cost environments and therefore gain less from expansion. This mechanism likely underlies the growing patent shares of high-skill CZ firms documented in Section 5.2, which drive the within-field rise in concentration.

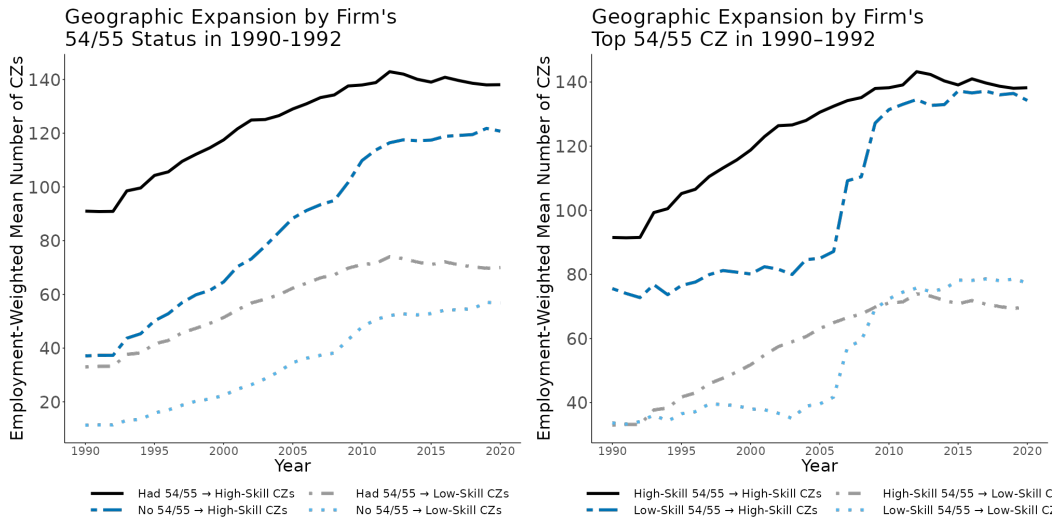


Figure 23: Trends in the *geography* of firm spatial expansion from 1990 to 2018. The left panel displays employment-weighted trends in the mean number of high-skill and low-skill CZs served by firms with an R&D plant in 1990–1992 versus those without one. The right panel shows analogous trends for firms with R&D plants in high-skill CZs versus firms with R&D plants in low-skill CZs.

Dependent Variables:	Has R&D Extensive Margin	Log R&D Employment Intensive Margin	Log R&D Employment (Q4 CZs) Intensive Margin
Model:	(1)	(2)	(3)
Lagged No. of CZs	0.0006***	0.0111***	0.0111***
	(6.22×10^{-5})	(0.0024)	(0.0024)
Fixed-effects			
Firm	Yes	Yes	Yes
Year	Yes	Yes	Yes
Fit statistics			
Observations	993,573	28,605	28,544
R^2	0.81316	0.83793	0.85087
Within R^2	0.00386	0.02642	0.02564

Clustered (firm) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table 4: Correlation of Lagged Firm Geographic Expansion on R&D Activity

6.4 Discussion of Other Factors

Worker sorting. The growth of ICT-oriented and high-skill CZ firms documented in Sections 5.1 and 5.2 was likely reinforced by the migration of workers – particularly skilled workers – toward high-skill CZs after 1990, a pattern documented in the broader literature on spatial sorting by education (Moretti, 2013; Diamond, 2016; Diamond and Gaubert, 2022). As high-skill CZs attracted more workers, firms in those regions gained access to deeper labor pools, further increasing their innovation output. Worker sorting is thus not an alternative to the colocation and firm expansion mechanisms documented in Sections 6.2 and 6.3 but rather a complementary channel that may have amplified their effects – by deepening the skilled

labor pools that attract ICT production (reinforcing colocation) and by enabling high-skill CZ firms to scale up their R&D (reinforcing the asymmetric scale effect). The decompositions in Sections 4 and 5, which are accounting identities, document *what* changed – which firms grew, in which fields, and where – regardless of whether worker sorting contributed to those firm-level dynamics. Quantifying the extent to which worker sorting contributed to these mechanisms requires a general equilibrium model, which is beyond the scope of this paper and developed in Tan (2026).

Beyond the ICT shock and worker sorting, several other factors could potentially account for part of the post-1990 rise in the spatial concentration of innovation. I briefly discuss three candidates and argue that they are unlikely to be primary drivers, though they may have played a supporting role.

Universities and academic research. Research universities are disproportionately located in high-skill CZs, and their importance for local innovation and economic growth is well documented: Hausman (2022) shows that after the Bayh-Dole Act increased the accessibility of university innovation, employment and corporate patenting grew differentially in industries related to nearby universities' technological strengths, with the effect driven by the entry and expansion of large multiunit firms. More broadly, the importance of academic science for commercial innovation has grown over time (Bikard and Marx, 2020). If the returns to proximity to university research increased after 1990, this could have reinforced the concentration of innovation in high-skill regions. However, the field decomposition in Section 4 shows that the post-1990 increase in concentration is driven primarily by the rising share of ICT patents – particularly in Digital Communication and Computer Technology – and the growth of ICT-oriented and high-skill CZ firms, rather than by fields most closely tied to academic research such as Biology and Chemistry. Moreover, the timing of the break around 1990 aligns more closely with the ICT shock than with gradual changes in university-industry linkages.

Place-based policies. Changes in state-level R&D tax credits, corporate tax rates, and other place-based policies could have differentially attracted innovative firms to certain regions. Wilson (2009) and Moretti and Wilson (2017) document that R&D tax credits affect the location of innovative activity. However, these policies varied substantially across states with no systematic bias toward high-skill CZs, and their estimated effects are modest relative to the magnitude of the aggregate increase in concentration documented in Section 3.1. Furthermore, the sharp break in innovation concentration around 1990 is difficult to reconcile with the gradual, state-by-state adoption of R&D tax credits over the 1980s and 1990s.

Changes in patent law. The creation of the Court of Appeals for the Federal Circuit in 1982 and subsequent rulings strengthened patent protection, which may have increased the propensity to patent, particularly among large firms. If these firms were disproportionately located in high-skill CZs, changes in patent law could mechanically raise measured concentration. However, the spatial concentration of patents remained approximately constant throughout the 1980s despite these legal changes and only began increasing around 1990, suggesting that patent law reforms are not a primary driver of the trends documented.

7 Conclusion

Understanding why innovation has become more spatially concentrated is essential for the design of place-based innovation policies, yet the forces driving this trend have remained elusive. This paper makes progress by combining comprehensive data on patents, firms, and inventors from 1976 to 2018 with a new empirical framework that progressively decomposes the aggregate trend – first by technology field and subfield, then by firm – to narrow down which explanations are consistent with the data.

I first document that the spatial concentration of American innovation remained roughly constant before 1990 but rose sharply thereafter, driven by a small number of high-skill regions – with Silicon Valley alone accounting for about half of the post-1990 rise. A progressive decomposition reveals that about half of the increase reflects a compositional shift toward ICT – specifically Digital Communication and Computer Technology – while the remainder reflects broad within-field concentration rises. At the firm level, both channels are driven by between-firm reallocation: the growing patent shares of ICT-oriented firms and high-skill CZ firms, rather than firm entry or within-firm changes.

Descriptive evidence points to the rapid rise of ICT as the primary driver, operating through two channels. The colocation of ICT innovation with ICT production in high-skill CZs explains why the compositional shift toward ICT raised aggregate spatial concentration. ICT-enabled reductions in communication costs then enabled high-skill CZ firms to expand production across the country while concentrating R&D further in high-skill regions – an asymmetric scale effect that drove the within-field concentration rise.

These findings have implications for the design of place-based innovation policies. The fact that rising concentration is driven largely by the location decisions of incumbent firms – rather than new entry – suggests that policies targeting startup formation may be insufficient to shift the geography of innovation. The colocation of ICT innovation with ICT production implies that attracting manufacturing or service operations without a corresponding R&D base may not generate local innovation spillovers. And the asymmetric scale effect – whereby firm spatial expansion paradoxically concentrated R&D in high-skill regions – highlights the challenge of using production subsidies to decentralize innovation. Understanding these mechanisms is a necessary first step toward designing policies that can effectively redirect or enhance innovative activity in different regions.

References

Agrawal, A. and Cockburn, I. (2003). The anchor tenant hypothesis: Exploring the role of large, local, R&D-intensive firms in regional innovation systems. *International Journal of Industrial Organization*, 21(9):1227–1253.

Agrawal, A., Cockburn, I., Galasso, A., and Oettl, A. (2014). Why are some regions more innovative than others? The role of small firms in the presence of large labs. *Journal of Urban Economics*, 81:149–165.

Andrews, M. J. and Whalley, A. (2022). 150 Years of the Geography of Innovation. *Regional Science and Urban Economics*, 94.

Audretsch, D. B. and Feldman, M. P. (1996). R&D Spillovers and the Geography of Innovation and Production. *American Economic Review*, 86(3):630–640.

Autor, D. H. and Dorn, D. (2013). The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market. *American Economic Review*, 103(5):1553–97.

Balland, P.-A., Jara-Figueroa, C., Petralia, S. G., Steijn, M. P. A., Rigby, D. L., and Hidalgo, C. A. (2020). Complex economic activities concentrate in large cities. *Nature Human Behaviour 2020 4:3*, 4(3):248–254.

Bettencourt, L. M. A., Lobo, J., and Strumsky, D. (2007). Invention in the city: Increasing returns to patenting as a scaling function of metropolitan size. *Research Policy*, 36(1):107–120.

Bikard, M. and Marx, M. (2020). Bridging Academia and Industry: How Geographic Hubs Connect University Science and Corporate Technology. *Management Science*, 66(8):3425–3443.

Bresnahan, T. and Greenstein, S. (1996). Technical Progress and Co-Invention in Computing and in the Uses of Computers. *Brookings Papers on Economic Activity. Microeconomics*, 1996:1–83.

Carlino, G. and Kerr, W. R. (2015). *Agglomeration and Innovation*, volume 5. Elsevier B.V.

Chattergoon, B. and Kerr, W. R. (2022). Winner takes all? Tech clusters, population centers, and the spatial transformation of U.S. invention. *Research Policy*, 51(2):104418.

Diamond, R. (2016). The Determinants and Welfare Implications of US Workers' Diverging Location Choices by Skill: 1980–2000. *American Economic Review*, 106(3):479–524.

Diamond, R. and Gaubert, C. (2022). Spatial Sorting and Inequality. *Annual Review of Economics*.

Eckert, F., Fort, T., Schott, P., and Yang, N. (2020). Imputing Missing Values in the US Census Bureau's County Business Patterns. Technical Report w26632, National Bureau of Economic Research, Cambridge, MA.

Ellison, G. and Glaeser, E. L. (1997). Geographic Concentration in U.S. Manufacturing Industries: A

Dartboard Approach. *Journal of Political Economy*, 105(5):889–927.

Feldman, M. P. and Florida, R. (1994). The Geographic Sources of Innovation: Technological Infrastructure and Product Innovation in the United States. *Annals of the Association of American Geographers*, 84(2):210–229.

Flood, S., King, M., Rodgers, R., Ruggles, S., Warren, J. R., Warren, D., Breton, E., Cooper, G., Drew, J. A. R., Richards, S., Van Riper, D., and Williams, K. C. (2025). IPUMS, Current Population Survey: Version 13.0.

Forman, C. and Goldfarb, A. (2020). Concentration and Agglomeration of IT Innovation and Entrepreneurship: Evidence from Patenting. Working Paper 27338, National Bureau of Economic Research.

Forman, C., Goldfarb, A., and Greenstein, S. (2016). Agglomeration of Invention in the Bay Area: Not Just ICT. *American Economic Review*, 106(5):146–151.

Fort, T. C., Keller, W., Schott, P. K., Yeaple, S., and Zolas, N. (2020). Colocation of Production and Innovation: Evidence from the United States. *Working Paper*.

Foster, L., Haltiwanger, J., and Krizan, C. (2001). Aggregate Productivity Growth: Lessons from Microeconomic Evidence. In *New Developments in Productivity Analysis*, page w6803. National Bureau of Economic Research, Cambridge, MA.

Greenstein, S. (2015). *How the Internet Became Commercial: Innovation, Privatization, and the Birth of a New Network*. Princeton University Press.

Greenstein, S. and Prince, J. (2006). The Diffusion of the Internet and the Geography of the Digital Divide in the United States. Technical Report w12182, National Bureau of Economic Research, Cambridge, MA.

Hall, B. H., Jaffe, A. B., and Trajtenberg, M. (2001). The NBER Patent Citation Data File: Lessons, Insights and Methodological Tools. *NBER Working Paper*.

Hausman, N. (2022). University Innovation and Local Economic Growth. *The Review of Economics and Statistics*, 104(4):718–735.

Hsieh, C.-T. and Rossi-Hansberg, E. (2023). The Industrial Revolution in Services. *Journal of Political Economy Macroeconomics*, 1(1):3–42.

hUallacháin, B. Ó. (2025). Organisations and the dynamics of change in the location of American invention. *Urban Studies*, 62(6):1099–1122.

Jaffe, A. B., Trajtenberg, M., and Henderson, R. (1993). Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations. Technical report.

Jiang, X. (2024). Information and Communication Technology and Firm Geographic Expansion. *Working Paper [Revise and Resubmit, Journal of Political Economy: Macroeconomics].*

Kelly, B., Papanikolaou, D., Seru, A., and Taddy, M. (2021). Measuring Technological Innovation over the Long Run. *American Economic Review: Insights*, 3(3):303–20.

Kerr, W. R. and Robert-Nicoud, F. (2020). Tech Clusters. *Journal of Economic Perspectives*, 34(3):50–76.

Kleinman, B. (2025). Wage Inequality and the Spatial Expansion of Firms. *Working Paper [Revise and Resubmit, Quarterly Journal of Economics].*

Krugman, P. R. (1992). *Geography and Trade*. MIT Press.

Li, G. C., Lai, R., D'Amour, A., Doolin, D. M., Sun, Y., Torvik, V. I., Yu, A. Z., and Lee, F. (2014). Disambiguation and co-authorship networks of the U.S. patent inventor database (1975-2010). *Research Policy*, 43(6):941–955.

Liu, J. (2024). Multinational Production and Innovation in Tandem. *Working Paper*.

Moretti, E. (2013). Real Wage Inequality. *American Economic Journal: Applied Economics*, 5(1):65–103.

Moretti, E. (2021). The Effect of High-Tech Clusters on the Productivity of Top Inventors. *American Economic Review*, 111(10):3328–3375.

Moretti, E. and Wilson, D. J. (2017). The Effect of State Taxes on the Geographical Location of Top Earners: Evidence from Star Scientists. *American Economic Review*, 107(7):1858–1903.

Mowery, D. C. and Simcoe, T. (2002). Is the Internet a US invention?—an economic and technological history of computer networking. *Research Policy*, 31(8):1369–1387.

Schmoch, U. (2008). Concept of a Technology Classification for Country Comparison. *Final report to the World Intellectual Property Organization*, (June):1–15.

Schroeder, J., Van Riper, D., Manson, S., Knowles, K., Kugler, T., Roberts, F., and Ruggles, S. (2025). National Historical Geographic Information System: Version 20.0.

Tan, W. (2026). Directed Innovation across Space. *Working Paper*.

Wilson, D. J. (2009). Beggar Thy Neighbor? The In-State, Out-of-State, and Aggregate Effects of R&D Tax Credits. *The Review of Economics and Statistics*, 91(2):431–436.

A Data Details for the Empirical Analysis

Here I provide detailed information about the data sources I used to measure and document fundamental trends in the geography of innovation in the United States, emphasizing my newly constructed six broad fields from technology classes under the Cooperative Patent Classification (CPC) scheme.

A.1 Patent Data and Cleaning Procedures

Patent data comes from bulk files from the US Patent Trademark Office (USPTO) and US PatentsView (USPV) (https://patentsview.org/). The datasets contain the universe of patents granted in the US (by the USPTO) from 1976-2022. I use only utility patents – which comprise about 95% of all granted patents from 1976 – as they are patents for inventions and provide the most appropriate measure of technological advancements, as opposed to other types of patents such as design, plant, and defensive patents. As explained by the USPTO Patent Technology Monitoring Team (PTMT)[12]:

> *Most analyses of technological activity that incorporate patent data will focus on the activity of utility patents, also known as "patents for inventions". Since design patents are granted for ornamental designs for articles of manufacture and not for inventions, they are usually perceived to have a lesser relationship with technological activity. Similarly, statutory invention registrations and defensive publications do not convey patent protection to disclosed inventions and may have a lesser relationship with technological activity. Plant patents may or may not disclose an invention resulting from technological activity; however, plant patents are numerically small relative to utility patents and are usually handled and analyzed separately.*

Inventors, assignees, and locations are disambiguated by USPV – meaning each inventor, assignee, and city-state pair has a unique ID over time – using the latest machine learning techniques, improving on the algorithms used in Li et al. (2014) and to produce the Connecting Outcome Measures in Entrepreneurship, Technology, and Science (COMETS) database. Cities and states of all inventors living in the US are provided. I use the Google Maps API to geocode all locations and assign them to counties, and standard publicly-available crosswalks to convert counties to commuting zones (CZs) [using Autor and Dorn (2013)], core-based statistical areas (CBSAs), and combined statistical areas (CSAs). The high spatial resolution of inventor locations allows me to conduct my analysis at different geographical scales.

I define the patent year as the application year since that is the closest to when the invention was produced, as there are often lags of several years between when a patent was applied and when it was granted. I assign patent shares, citations made, and citations received equally across all coinventors on a patent. I keep only patents where at least one inventor lives in the US.

[12]The USPTO PTMT explanation on the conventions of treating different types of patents can be found at https://www.uspto.gov/web/offices/ac/ido/oeip/taf/reports.htm. Should this or any of the subsequent web links become inactive, PDF copies of the contents of the archived websites that includes the date of access will be provided upon request.

A.2 Classification of Patents into Technology Fields and Subfields

Patents are assigned to 3-digit technology classes, 4-digit subclasses, and further divided into groups within subclasses under the International Patent Classification (IPC) and Cooperative Patent Classification (CPC) systems. The IPC provides a hierarchical system of language independent symbols for the classification of patents into eight sections with approximately 70,000 subdivisions. The IPC was developed under the 1971 Strasbourg Agreement and provides a common classification for patents filed in different patent offices around the world. The CPC is a unified system developed jointly by the United States Patent Trademark Office (USPTO) and European Patent Office (EPO) on 2010 to provide a common, internationally compatible classification system that provides more groups and subgroups relative to the IPC. For patents granted in the US, the CPC supercedes the US Patent Classification (USPC) developed by the USPTO and the NBER patent classification developed by Hall et al. (2001). The USPTO no longer provides the USPC and NBER patent classifications for patents granted after 2015.

Most patents have more than one field classification to facilitate easier searches to prior art. Nonetheless, under the CPC, each patent has a unique primary field classification, denoted as the "first" position. Other classifications are denoted as having a "later" position. USPTO Guideline 905.03(a)III.A.(C)[13] states:

> *There is one and only one "first" position attribute per patent family. The first attribute is associated with the invention symbol that most completely covers the technical subject matter of the disclosed invention. The first position symbol is identified as the first mandatory symbol listed on the classification form.*

The USPTO provides a bulk file of all granted patents from 1790 with their current CPC classes and position of each CPC class[14]. I use the 08/02/2022 version, downloaded on 09/09/2022[15]. I use the primary CPC class of each patent. In rare cases where there are multiple CPC classes listed as the "first" position, I use the class that is listed first. In rare cases where the CPC class listed in the "first" position is only meant as an additional classification as noted under the details of the CPC class in their classification documentation[16], I use the class in the "later" position that is listed first.

Under the IPC and CPC, the technology subclasses are grouped into 8 broad sections. These broader sections facilitate the allocation of patents to different examiner units at patent offices. **However, some sections contain patents from highly disparate economic fields. For example, Section A includes patents for medical technologies and amusement parks.** To provide more consistent categories with similar sizes, the World Intellectual Property Organization (WIPO) developed a mapping from the 4-digit technology subclasses and groups in the IPC/CPC to 35 fields (Schmoch, 2008). This WIPO report also provides 6

[13] The USPTO guidelines on the CPC scheme can be found at https://www.uspto.gov/web/offices/pac/mpep/s905.html

[14] The bulk files containing the CPC classification of every patent can be found at https://bulkdata.uspto.gov/data/patent/classification/cpc/

[15] Although USPV provides CPC classes and technology fields for most patents granted since 1976, data checks reveal that the 2022 vintages of the USPV datasets contain errors in the CPC classes and WIPO technology fields for approximately 3–5% of patents.

[16] The CPC classification documentation can be found at https://www.uspto.gov/web/patents/classification/cpc/html/cpc.html.

broad categories, but the broad category "Instruments" is difficult to interpret and the "Electrical Engineering" category includes fields in information technology as well as older electrical machinery. Thus, I provide an alternative mapping of these 35 WIPO fields into 6 broad categories: (1) Physics, Electrical Engineering & Electronics; (2) Information Technology; (3) Chemistry; (4) Biology & Medicine; (5) Mechanical Engineering; (6) Civil Engineering & Consumer Goods. These categories are similar in scope to the older field classifications in the NBER and COMETS datasets and may be seen as an updated version of them.

The IPC/CPC class to WIPO field concordance is primarily at the level of 4-digit technology subclasses, but there are 6 subclasses with different groups that map to different WIPO fields – A61K, B01D, C13B, E01F, G01N, H04N – due to significantly different subject matter across groups within these subclasses. For example, within the subclass A61K, there are patents for cosmetics and patents for medical technologies. Thus, I provide a further decomposition for these 6 subclasses that corresponds to the WIPO field they are assigned to. My added subclasses are: A61K-14; A61K-16; B01D-23; B01D-24; C13B-18; C13B-29; E01F-24; E01F-35; G01N-10; G01N-11; H04N-2; H04N-3; H04N-4. The numbers after the hyphen refer to the WIPO fields that groups within each subclass are assigned to. Thus, my data has 586 technology classes, comprising the 579 CPC/IPC subclasses along with this decomposition of 6 subclasses. In this paper, I refer to the 35 World Intellectural Property Organization (WIPO) fields as "subfields", my 6 broad categories as "fields", and the 579 CPC/IPC subclasses and 7 additional group categories as "classes".

A.3 Data Sources for Regional Outcomes

I use data on county-level educational attainment for each decade from the Economic Research Service (ERS) of the US Department of Agriculture (USDA) to construct the college ratio and percent of workers that are college educated for each CZ in 1970, 1980, 1990, 2000, 2010 and 2017. The 1970, 1980, 1990 and 2000 data are constructed from the Decennial Censuses of Population and equivalent to the 5% samples from the IPUMS used by Moretti (2013); Diamond (2016) to construct college percent and ratio respectively for metropolitan statistical areas (MSAs) in 1980 and 2000. The 2010 and 2017 data are constructed from 5-year averages from the annual American Community Survey (ACS). Data from the IPUMS NHGIS (Schroeder et al., 2025) are identical.

I obtain annual data on the college-educated and non-college educated workers in each IND1990 industry from IPUMS CPS. I define the ICT sector to include

I obtain annual county-level employment by NAICS industry for 1975-2018 from the harmonized County Business Patterns (CBP) Database produced by Eckert et al. (2020). They develop a linear programming method to impute suppressed county-industry-year cells in raw CBP files released by the US Census Bureau and show that total non-agricultural employment from their dataset is highly correlated with the panel from the Longitudinal Business Database (LBD). I use their county-industry employment panel to construct my primary measures for the annual industrial composition of each CZ and the United States, such as the employment share in the ICT service sector. Building on Fort et al. (2020), I define the ICT service sector to include the following industries in the Information Sector (NAICS 51): Software Publishers (5112);

Telecommunications (517); Data Processing, Hosting, and Related Services (518).

I also obtain annual county-level data on income and resident population for 1969-2018 from the Bureau of Economic Analysis to construct basic annual measures of each CZ such as population, population density, and income per capita.

B Additional Empirical Results

Here I present supplementary empirical results on the main facts and likely drivers of the post-1990 rise in the spatial concentration of US innovation. Appendix B.1 presents robustness checks on the aggregate trends (Section 3.1), Appendix B.2 presents additional figures on the regional drivers of rising concentration (Section 3.2), Appendix B.3 presents additional results on the concentration of innovation in high-skill CZs (Section 3.3), Appendix B.4 presents additional results and robustness checks on the field-level decomposition (Section 4), and Appendix B.5 presents details of the NSFNET (Section 5).

B.1 Fact 1: Robustness on Rising Spatial Concentration of Innovation after 1990

Figure 24 presents robustness checks on the trend in the aggregate spatial concentration of innovation in Figure 1. Innovation became more concentrated in space only after 1990, whether I use alternative measures of spatial concentration such as the Herfindahl index, coefficient of variation, and simplified Ellison-Glaeser measure of patent shares (left graph), or the annual share of patents produced by the top 10 CZs (right graph).

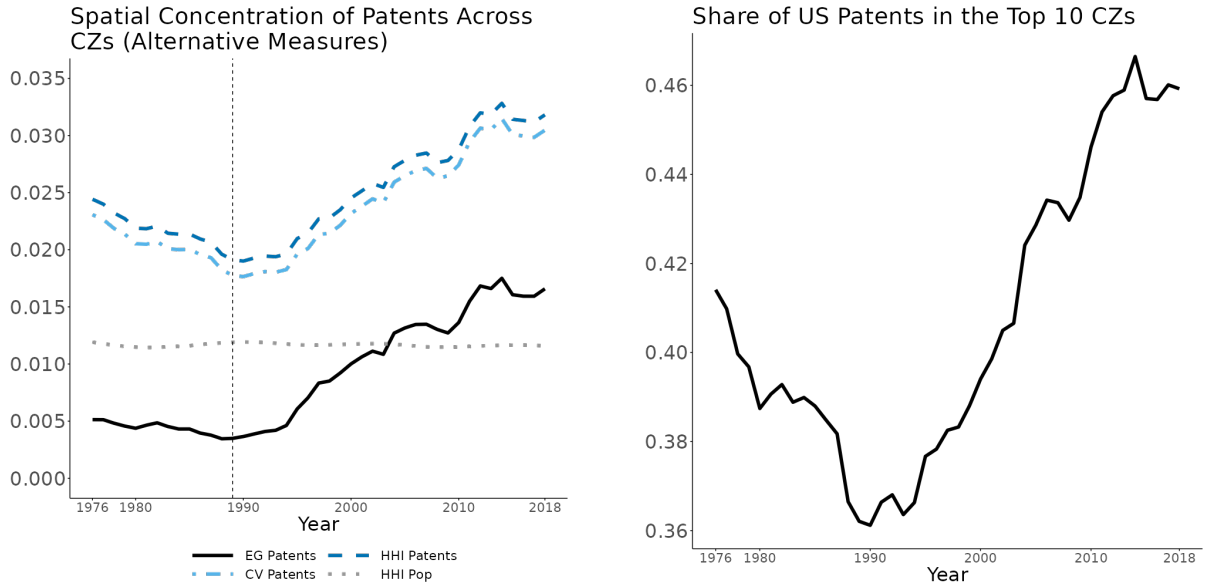


Figure 24: Robustness of trends in the spatial concentration of patents across CZs. The left graph plots trends in alternative measures of spatial concentration, such as the Herfindahl index (blue dashed line), coefficient of variation (light blue dotdashed line), and simplified Ellison-Glaeser measure (black solid line) of patent shares. The right graph plots trends in the annual share of patents produced by the top 10 CZs.

B.2 Fact 2: Additional Figures on Regional Drivers of Rising Concentration

Figure 25 plots the change in patent share against the change in innovation location quotient for all CZs, revealing four distinct groups. CZs in the upper-right quadrant — including San Jose, San Francisco, Seattle, and San Diego — experienced rising innovation on both dimensions: they captured a growing share of national patents while also becoming more patent-intensive relative to their population. CZs in the lower-left quadrant, such as Detroit and Rochester, declined on both measures. The off-diagonal quadrants highlight where the two measures diverge: CZs in the upper-left saw rising innovation intensity despite a falling national patent share, typically because their population shrank faster than their patenting, while CZs in the lower-right experienced the reverse.

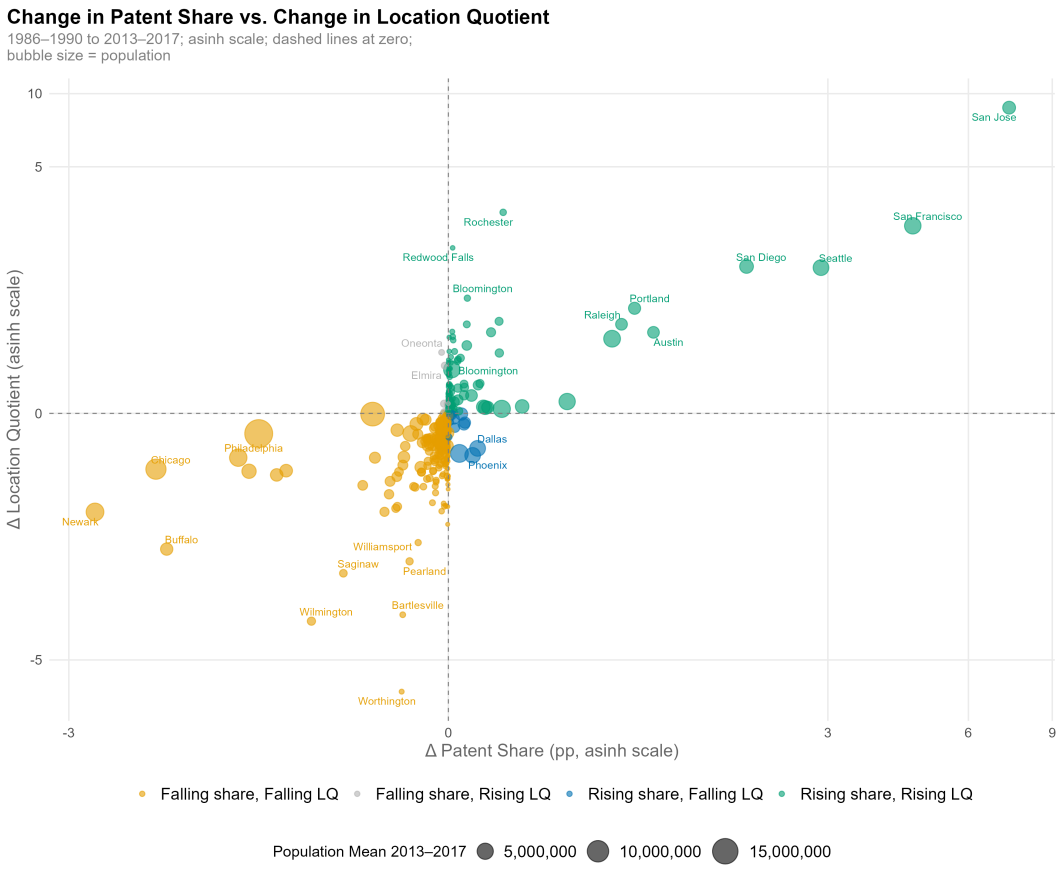


Figure 25: Change in patent share vs. change in innovation location quotient for all CZs, from 1986–1990 to 2013–2017 averages. Axes are on an inverse hyperbolic sine (asinh) scale to spread the distribution while preserving the zero crossing. Bubble size reflects CZ population (2013–2017 average). Dashed lines mark zero change on each axis, dividing CZs into four quadrants: rising share and rising intensity (upper right), rising share but falling intensity (lower right), falling share but rising intensity (upper left), and falling on both dimensions (lower left).

Table 5 lists the top 15 CZs by changes in patent-to-population share between 5-year averages around 1990 and 2015. Apart from well-known superstar cities such as San Jose, San Francisco, San Diego, Seattle and Boston, CZs like Portland, Boise, Wayne, Provo and Fort Collins also became some of the most innovative

regions in the US from 1990-2015. Despite narratives of the apparent collapse of innovation in Rochester in several papers and the popular press, the region actually experienced the second greatest increase in patenting intensity from 1990-2015.

	CZ	Change in patent-to-population share 1990-2015	Patent-to-population share	Rank of patent-to-population share	Patent Share	Patents	Population
1	San Jose	9.494	12.421	1	0.104	14581.5	2684061
2	Rochester	3.608	4.618	2	0.004	531.4	263093
3	San Francisco	3.029	4.440	3	0.073	10153.8	5229363
4	San Diego	2.094	3.430	5	0.035	4882.0	3254818
5	Seattle	2.063	2.955	8	0.043	5974.0	4622131
6	Portland	1.618	2.523	11	0.018	2521.7	2285666
7	Boise	1.411	1.962	18	0.004	604.7	704711
8	Raleigh	1.398	2.282	13	0.015	2064.7	2068936
9	Wayne	1.396	2.303	12	0.001	132.5	131541
10	Austin	1.258	2.959	7	0.019	2660.2	2055406
11	Provo	1.232	1.763	19	0.003	467.1	605648
12	Bloomington	1.186	1.587	25	0.001	157.2	226557
13	Burlington	1.167	2.224	14	0.002	333.9	343199
14	Boston	0.885	2.730	9	0.046	6488.4	5434761
15	Fort Collins	0.871	1.756	20	0.004	495.9	645656
	Oneonta	-0.944	4.414	4	0.002	300.3	155554
	Albany	0.466	2.959	6	0.010	1461.1	1128888
	Elmira	0.341	2.627	10	0.003	401.9	349814
	Poughkeepsie	0.554	2.185	15	0.006	897.9	939410
	Minneapolis	0.192	2.021	16	0.021	2989.4	3381353
	Detroit	0.588	1.967	17	0.032	4510.1	5243719

Table 5: The top 15 CZs by change in ratio of patent-to-population share from 1990-2015 with a population of at least 100,000 in 2015. To minimize the effect of idiosyncratic fluctuations, I take five-year averages around 1990 and 2015. I append this list with the 6 CZs that are in the top 20 by ratio of patent-to-population share in 2015 but did not experience the greatest increase from 1990 to 2015.

B.3 Fact 3: Concentration in High-Skill CZs, Not Large or Dense CZs

Figure 26 plots trends in the annual elasticity of CZ patents per capita with respect to the 1990 population (left graph) and population density (right graph). The left graph shows that while the annual elasticity of patents per capita with respect to the 1990 population increased moderately from 1990-2018, these increases did not overcome the decrease in annual elasticity from 1976-1990. The right graph shows that annual elasticity of patents per capita with respect to the 1990 population density did not increase after 1990.

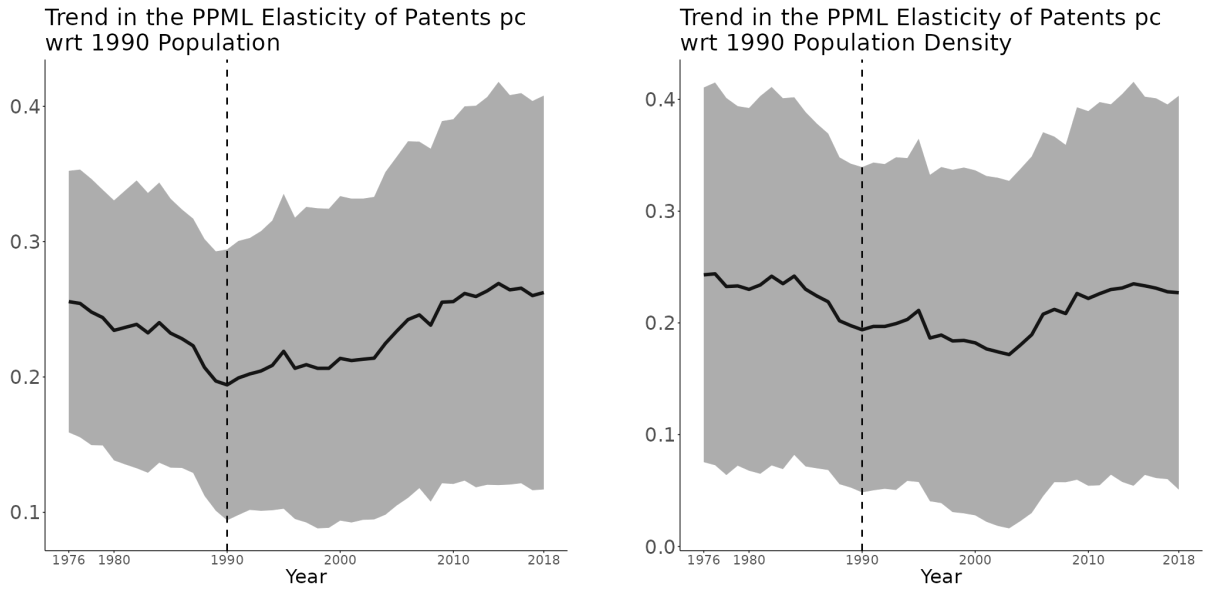


Figure 26: Absence of evidence on patent sorting to large or dense CZs after 1990. The left panel displays the trend in the annual elasticity of CZ patents per capita with respect to the 1990 population while the right panel displays an analogous trend with respect to the 1990 population density. The confidence intervals reflect the statistical uncertainty of the slope estimates for each year given heteroskedastic-robust standard errors.

B.4 Additional Results from the Field Decompositions

B.4.1 Trends on ICT versus non-ICT Patents

In the main text, Figure 13 presents trends in the Gini coefficient and patent elasticity by technology field. To better illustrate how ICT differs from other fields, I now aggregate all non-ICT fields into a single category. The left panel of Figure 27 plots the 1990 patent elasticity for ICT vs non-ICT, i.e. the correlation between patents per capita and the college ratio in 1990, where each dot represents a CZ and the size of the dot reflects the CZ's population. The graph shows that in high-skill CZs, ICT and non-ICT patents per capita are similar, whereas in lower-skill CZs, non-ICT patents far exceeds that of ICT. This highlights the stronger concentration of ICT innovation in high-skill CZs. The middle panel shows that the patent elasticity for ICT has been consistently higher than for non-ICT, and the right panel confirms that ICT patents have consistently been more spatially concentrated than non-ICT patents as measured by the Gini coefficient, with spatial concentration increasing for both groups after 1990.

B.4.2 Trends and Decomposition by Field without Top Patenting CZs

Section 3.2 shows that Silicon Valley accounts for about half of the post-1990 rise in the locational Gini. A natural concern is whether the field-level trends and decompositions in Section 4 are also driven primarily by Silicon Valley. This appendix shows that while excluding Silicon Valley attenuates the Gini-based results, the key qualitative findings are robust.

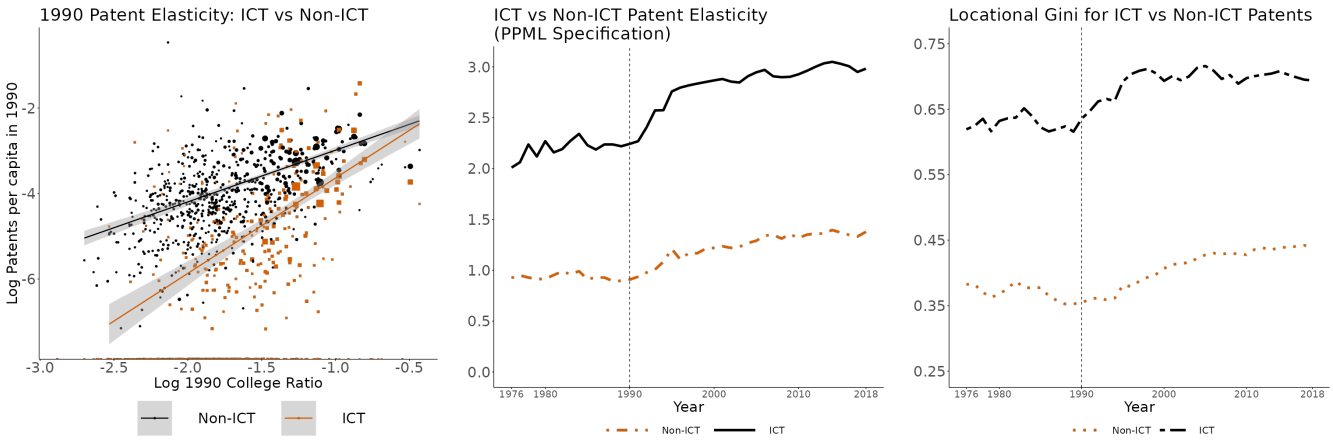


Figure 27: Concentration of ICT versus non-ICT patents. The left panel displays the correlation between patents per capita and the college ratio in 1990, where each dot represents a CZ and the size of the dot reflects the CZ's population. The gradients represent the ICT and non-ICT patent elasticity respectively in 1990. The middle panel shows trends in the patent elasticity for ICT versus non-ICT. The right panel shows trends in the Gini coefficient of patents per capita by ICT versus non-ICT.

Figure 28 replicates the field-level trends from Section 4.1 after dropping San Jose and San Francisco. The left panel shows that the compositional shift toward ICT is virtually unchanged, as expected since it reflects national patent counts. The middle panel confirms that ICT patents remain more spatially concentrated than other fields and became more concentrated between 1990 and 2000, though excluding Silicon Valley leads to a modest decline in ICT spatial concentration after 2000. The right panel shows that trends in the patent elasticity by field are largely unaffected, confirming that the increasing concentration of innovation in high-skill CZs is a broad national pattern rather than a Silicon Valley artifact.

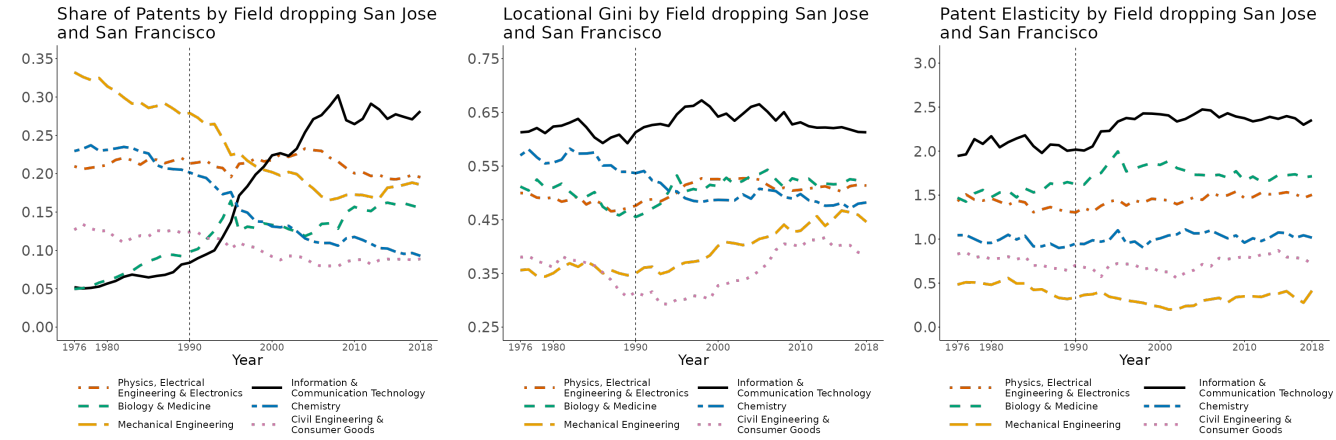


Figure 28: Trends in the share of annual patents (left panel), locational Gini coefficient of patents per capita (middle panel), and patent elasticity (right panel) by technology field dropping San Jose and San Francisco.

Figure 29 replicates the Gini decomposition from Section 4.2 after excluding Silicon Valley. The left panel shows that the post-1990 rise in the locational Gini is attenuated by 44%, and the cross-field component now accounts for 67% of the remaining increase (compared to 52% in the full sample). The middle panel

confirms that the cross-field component is still driven almost entirely by the rising share of ICT patents. The right panel reveals an important difference: the within-field contribution from ICT becomes smaller over time and turns negative after 2008, reflecting the decline in ICT spatial concentration in the middle panel once Silicon Valley is excluded. This implies that much of the within-field Gini rise in ICT is attributable to Silicon Valley, while the cross-field compositional shift is not. Since trends in the patent elasticity by field are largely unaffected by excluding Silicon Valley (right panel of Figure 28), the elasticity-based field decomposition is also robust.[17]

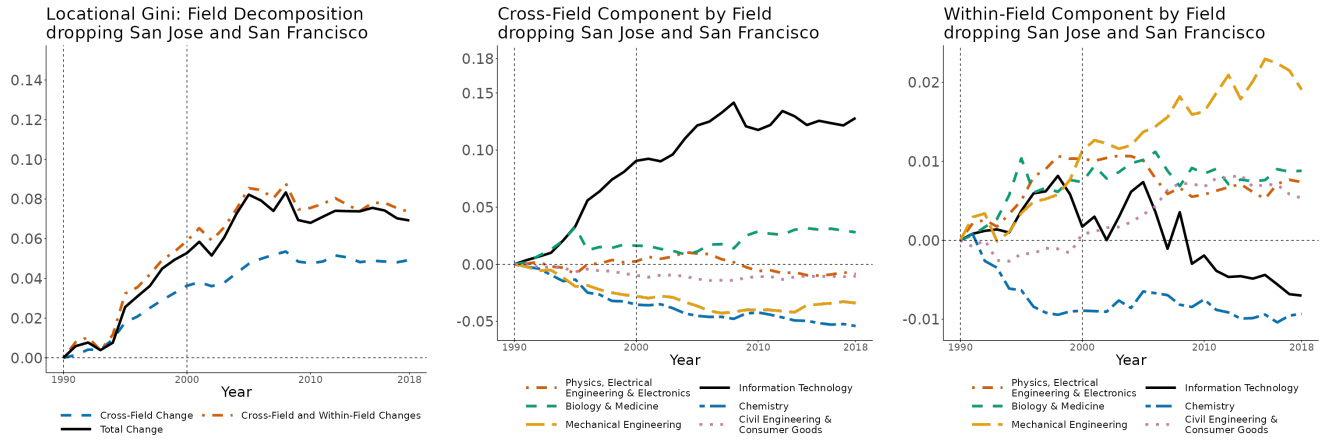


Figure 29: Trends in the decomposition of the locational Gini coefficient of patents per capita into cross-field and within-field components (left panel) and a further decomposition of the cross-field and within-field components into each individual field (middle and right panels respectively) after dropping San Jose and San Francisco.

B.5 Details of the NSFNET

The National Science Foundation Network (NSFNET), established by the National Science Foundation and developed between 1986 and 1995, was created to facilitate collaboration among researchers at universities and military bases across the United States. It eventually became a crucial bridge between ARPANET – the first public packet-switched computer network operated by the Defense Advanced Research Projects Agency from 1969 to 1989 – and the commercial networks that began providing internet access to the public, particularly during the internet boom starting in 1995.

The history of the NSFNET is defined by five key events:

- Initial establishment: In 1985, the NSF funded the creation of five supercomputing centers. In 1986, it established a long-haul backbone network with a data speed of 56 Kbps, connecting these new centers to the existing supercomputing facility at the National Center for Atmospheric Research.

- First round of upgrading: On June 15, 1987, the NSF issued a solicitation to upgrade and expand the backbone network, addressing the overwhelming demand that had saturated the existing infrastructure. On November 24, 1987, a contract was awarded to a team comprising IBM, MCI, Merit, and the

[17] I omit the elasticity decomposition here, but results are available upon request.

University of Michigan. By July 1988, the upgraded T1 backbone was completed, increasing the number of backbone sites from 6 to 13 and raising network speeds to 1.5 Mbps. The upgraded backbone also enabled connections to regional and campus networks. Each partner played a specific role: Merit developed user support and information services, IBM provided hardware, software, and network management tools, and MCI supplied transmission circuits with reduced tariffs.

- Second round of upgrading: In 1991, the NSF completed the upgrade from the T1 network to a T3 network, increasing the broadband speed to 45 Mbps and adding three new backbone sites.

- **Privatization**: Discussion about privatization began quietly in 1989 and become public by 1990. **In March 1991, the Acceptable Use Policy – which previously required the network to be used solely for research and education – was revised to allow private users**. In 1992, the NSF announced plans to decommission the NSFNET by 1995. By May 1993, it issued a solicitation to encourage more companies to contribute to the development of the Internet's privatized structure[18]

- Decommission: 1995

The most significant event in the history of the NSFNET for U.S. patenting – predominantly conducted by private firms – was the modification of the Acceptable Use Policy in March 1991. This change granted firms access to the NSFNET, significantly reducing their communication costs. Table 6 lists the NSFNET nodes, including the year each was established and the commuting (CZ) in which it is located. Note that although the NSFNET backbone consisted of just 11 nodes in 1991 and 15 in its full version in 1993, many of these nodes connected to regional networks, thereby extending high-speed internet access to universities and firms across most US regions.

[18]Taken from Kesan, Jay P., and Rajiv C. Shah. 2001. "Fool Us Once, Shame on You—Fool Us Twice, Shame on Us: What We Can Learn from the Privatizations of the Internet Backbone Network and the Domain Name System." Washington University Law Quarterly 79: 89–220.

S/N	NSFNET Node Location	CZ	CZ Name	Year	Selected Universities
1	John von Neumann Center in Princeton, NJ	19600	Newark	1986	Princeton University
2	Cornell Theory Center in Ithaca, NY	18100	Elmira	1986	Cornell University
3	Pittsburgh Supercomputing Center in Pittsburgh, PA	16300	Pittsburgh	1986	Carnegie Mellon University, University of Pittsburgh
4	San Diego Supercomputer Center in San Diego, CA	38000	San Diego	1986	UC San Diego
5	National Center for Supercomputing Applications in Urbana, IL	23500	Decatur	1986	University of Illinois, Urbana-Champaign
6	National Center for Atmospheric Research in Boulder,CO	28900	Denver	1986	University of Colorado Boulder, Colorado School of Mines
7	Palo Alto, CA	37500	San Jose	1988	Stanford University
8	Houston, TX	32000	Houston	1988	Rice University
9	Ann Arbor, MI	11600	Detroit	1988	University of Michigan at Ann Arbor
10	College Park, Maryland	11304	Washington DC	1988	Georgetown University, University of Maryland
11	Salt Lake City, UT	36100	Salt Lake City	1988	
12	Seattle, WA	39400	Seattle	1988	
13	Lincoln, NE	28101	Lincoln	1988	
14	Cambridge, MA	20500	Boston	1991	Harvard University, MIT, Tufts University
15	Argonne National Laboratory in Lemont, IL	24300	Chicago	1991	University of Chicago, Northwestern University
16	Atlanta, Georgia	9100	Altanta	1991	Emory University

Table 6: List of NSFNET backbone nodes