Talent as a Bottleneck to the Spatial Diffusion of Digital Innovations

August 31, 2025

This work was led by doctoral students.

Abstract

Although skilled technical talent has been recognized as a critical constraint when adopting new information technologies, little attention has been paid to how labor markets adapt to meet these shifting needs. Using a large-scale dataset that tracks *applications* to job listings, this study uncovers three empirical regularities in the spatial diffusion of IT skills. First, for emerging IT innovations, the pool of applicants for new roles is geographically concentrated, often limited to those workers who are willing to relocate significant distances. Second, as the technology matures, their diffusion is associated with broader participation from workers across diverse regions, reducing spatial constraints. However, the diffusion of worker participation consistently lags behind the geographic diffusion of job postings. Finally, regions that pioneered the technology attract a disproportionate share of applicant interest, and the gap is widening over time.

1 Introduction

Recent scholarship has highlighted that artificial intelligence (AI) activity is highly concentrated in superstar cities (McElheran et al., 2024). This phenomenon aligns with a long standing body of literature examining the role of geographic agglomeration in IT adoption, particularly noting the pivotal roles of regions such as San Francisco and Silicon Valley in the diffusion of information technology (IT) innovations (Saxenian, 1996; Fallick et al., 2006; Forman et al., 2012). Despite considerable IT advancement, the diffusion of IT innovations remains a time-consuming and challenging process (Brynjolfsson et al., 2021), constrained not only by technological barriers but also by the availability of factors like human capital necessary to apply these technologies effectively (Attewell, 1992; Fichman, 2000).

Previous research has explored the geographic dimensions of IT innovation diffusion both within the United States and globally, underscoring the critical role of localized factors like skilled engineers in overcoming initial adoption barriers and accelerating diffusion (Caselli and Coleman, 2001; Benhabib and Spiegel, 2005; Kraemer et al., 2005; Forman et al., 2005). This work has consistently identified spillover effects—beneficial impacts stemming from geographic concentration and proximity—as facilitators of early diffusion (Forman et al., 2012; Tambe, 2014). While such spillovers may diminish as technologies mature and spread to broader markets (Desmet and Rossi-Hansberg, 2009), initial concentration patterns can create advantages for pioneering regions.

Building on this understanding of geographically constrained spillovers, the study of "new work" created by new technologies (Autor et al., 2024) has primarily been at the city level (Lin, 2011) or from the employer (demand side) perspective in the literature. For example, Kalyani et al. (2023) examine how new technologies spread and impact labor markets, skill distribution, and regional US economies by employing job posting data to infer diffusion. Their findings reveal evidence of a slow geographic spread of jobs (on average over 50 years) with persistent pioneering region advantage in jobs. This is particularly true for skill-biased technological change—innovations which disproportionately advantage high skilled talent—where the diffusion process into jobs is longer (Lin, 2011; Kalyani et al., 2023), suggesting significant barriers that prevent labor markets from quickly equilibrating (training new talent to meet demand).

Under these circumstances, understanding the behavior and preferences of the *supply* side of the labor market becomes critical, since the "war for talent" is likely to be an apt description in the face of augmenting innovations and their geographically concentrated diffusion patterns. Nonetheless, relatively little attention has been devoted to the role of labor market dynamics in shaping the diffusion trajectories of IT innovations. Specifically, the implications of firms' IT investment decisions on the labor pool they face remains understudied.

Using data from a leading job posting platform specializing in high growth startups, Wellfound (formerly AngelList Talent), we examine the application behavior of high skill talent to job postings, paying particular attention to the geographic and requisite skill "maturity" dimensions. Although

the study of job postings (or in some cases, resume data) are relatively widespread, our analyses are uniquely enabled by a database of applications to job postings. We are not aware of other work that uses both supply and demand activity for full-time job openings to evaluate adjustment in markets for new technological innovations, including digital technologies. We find that applicants express disproportionate interest in jobs in pioneering regions, and that this gap grows over time.

Our research bridges the IS and strategy literatures to provide insights on how geographic talent mobility shapes the spread of digital innovation. The IS field has studied technology diffusion and value, emphasizing human capital alongside organizational capabilities and complementary assets as critical adoption factors (Melville et al., 2004; Tambe, 2014; Brynjolfsson et al., 2021). While the role of geography on IT use has been less central in IS research[1], it represents a cornerstone of the economics and strategy literature (Marshall, 1920), particularly through theories of agglomeration economies and spatial clustering that extend to digitization (Saxenian, 1996). The intersection of these fields becomes increasingly relevant as recent scholarship documents geographic concentration in frontier technology diffusion patterns (Forman et al., 2012) and most recently AI (McElheran et al., 2024), and as political winds bring changes to firms' abilities to draw on diverse (global) labor pools, with implications for competitive advantage and market power. Given these trends, understanding the mechanisms that either reinforce regional divergence or promote convergence across geographies has significant implications for theory and practice in the digital economy.

By addressing this gap, we uniquely examine how the pool of available talent responds to employers' IT investment choices. Our analysis reveals the perhaps surprising finding that initial talent advantages in a region are self-reinforcing, further deepening regional disparities in technological innovation and adoption. We conclude by discussing the implications of these geospatial skill diffusion patterns for managerial practice and policy.

[1] There has been considerable work on how information systems affect geographic constraints through remote collaboration, virtual worlds, digital platforms, etc. but less on how geographic factors affect IT adoption and use, with some exceptions. For instance, see Forman et al. (2005); Kraemer et al. (2005).

2 Related Literature & Hypotheses

2.1 What attracts technical talent?

Most research on IT workers has focused on the role of wage compensation (Ang et al., 2002; Mithas and Krishnan, 2008; Whitaker et al., 2019). However, in addition to financial incentives such as salary and potential equity options, the literature has suggested non-pecuniary job aspects which may attract talent. We categorize these as features of the job and features of the labor market; we discuss these in turn.

Non-monetary dimensions for attracting talent include preferences for autonomy, meaningfulness of work, and opportunities for upskilling. Studies by Sauermann (2018) and Roach and Sauermann (2023) show graduate students self-select into startups or established firms on these dimensions. Research by Stern (2004) and Tambe et al. (2020) further indicates that opportunities to pursue research and learn technical skills represent important non-pecuniary benefits for high-skill talent, with these studies quantifying these features through compensating differentials analyses.

Research has examined the career effects of working in technology clusters. Freedman (2008) shows software workers in clusters job hop more frequently, accepting lower initial wages but experiencing steeper wage growth. Bleakley and Lin (2012) identify wage premiums in thicker markets due to both ex-ante effects (greater investment in industry-specific skills) and ex-post effects (better skill matching). They find job hopping patterns vary by experience, with experienced workers churning less and inexperienced workers more in thick markets.[2]

Employees working in pioneering regions may signal to future employers that they possess frontier skills, regardless of other considerations like compensation or skill fit. Future employers may use location as a proxy for skill level, especially when evaluating talent from startups lacking established reputations. Firms also strategically locate where they can access necessary talent. This is particularly true for startups, though all firms consider human capital when selecting locations.

[2] Regarding career dynamics, Sorenson et al. (2021) find startup employees generally don't earn more long-term than those at established firms, while Kim (2018) shows startup employees experience faster wage growth despite starting lower. These studies suggest startups attract talent through non-monetary benefits rather than compensation.

Tambe (2014) demonstrates that firms investing in big data technologies only see benefits when located in labor markets with sufficient technical expertise.

2.2 Knowledge spillovers and geographic agglomeration

The dynamics of labor market influences on IT innovation diffusion may be better understood by examining similar patterns in other sectors. For example, in the early days of the biotechnology industry, a high percentage of biotechnology startup firms aiming to leverage new technology were founded in close geographic proximity to the star scientists themselves (Zucker and Darby, 1996). This is consistent with the proposition that the tacit knowledge these star scientists had at the birth of the industry shaped the early geographic locus of the industry.[3]

Jaffe et al. (1993) established that patent citations, which proxy for knowledge spillovers, occur disproportionately locally due to agglomerated social processes and institutions. This localization affects technology diffusion, as pioneering clusters experience concentration on both demand and supply sides when skills differ sufficiently from prior innovations. Labor markets and knowledge sharing represent the two primary Marshallian agglomeration forces that maintain geographic concentration of industrial activity (Marshall, 1920).

Specialized labor serves as a key Marshallian agglomerating force as workers seek to hedge against firm-specific shocks by locating where multiple firms can compete for their talents. Similarly, firms locate where specialized skills are available, leading to industry agglomeration. Fallick et al. (2006) find greater talent circulation within rather than across technology clusters (particularly in Silicon Valley's computer industry), supporting the idea that knowledge spillovers function as a geographically self-reinforcing social process.

Marshall (1920) also discussed firm clustering due to capturing knowledge spillovers. Larger-scale empirical support arrived decades later (Jaffe et al., 1993), showing that knowledge flows (proxied by patent citations) are a local process even after controlling for the existing geography of

[3] Such geospatial patterns intersect with the finding in the knowledge spillovers literature that such spillovers tend to occur *locally* (Jaffe et al., 1993), even taking into account the pre-existing geographic industrial location of economic activity (Audretsch and Feldman, 1996). Consequently, localized knowledge spillovers can be a powerful contributor to industry agglomeration.

production (Audretsch and Feldman, 1996). Zucker and Darby (1996) found that biotechnology venture founding locations were often co-located with scientific stars at the birth of that industry, likely due to concentrated tacit knowledge. Gambardella and Giarratana (2010) find that localized knowledge spillovers become a positive self-reinforcing cycle with knowledge workers gaining higher productivity. Owen-Smith and Powell (2004), within the Boston biotechnology cluster, show that geographic propinquity and organizational form link spillovers to increased innovation.

The literature has paid less attention to cluster collapse or de-agglomeration. Four main factors have been identified: (1) global competition and adverse macroeconomic events, as seen in Detroit's auto industry (Klier and Rubenstein, 2008) and northern England's textile cluster (Toms and Zhang, 2016); (2) technological changes, exemplified by typewriter clusters' decline after PCs emerged (Grabher, 1993), Akron's tire industry following radial tire innovations (Sull, 1999), and digital platforms reducing proximity advantages (Bathelt and Glückler, 2011); (3) institutional changes, such as losing anchor institutions like Pittsburgh's steel mills (Krugman, 1991; Henderson, 1974); and (4) social changes, including eroded trust leading to skilled labor exodus (Maskell and Malmberg, 2007; Glaeser, 2011). However, few studies specifically isolate labor market and skill-based explanations for de-agglomeration.

2.3 Hypotheses

Guided by mechanisms highlighted in the agglomeration and skill-biased technical-change literature, we formulate four hypotheses that relate the spatial evolution of labor demand to the evolution of labor supply.

Hypothesis 1. *The spatial diffusion of a technology's labor supply is slower than labor demand.*

Although job listings for new technologies diffuse gradually (Kalyani et al., 2023), labor demand still spreads faster than the supply of specialized workers, whose relocation and up-skilling may be slowed by training pipelines, migration costs, and local knowledge spillovers. These frictions produce a persistent spatial mismatch that grows over time.

Hypothesis 2. *The imbalance of spatial diffusion between labor supply and demand is more severe for high-skilled / specialized jobs.*

High-skill and specialized fields require steep learning curves and tacit expertise clustered around top universities or incumbent leaders, amplifying the shortage of qualified workers outside pioneering regions. Lower-skill technologies, by contrast, can draw on a broader, more mobile workforce and therefore exhibit a smaller gap.

Hypothesis 3. *Technology pioneering regions attract a disproportionate share of applicant interest compared to the non-pioneering regions.*

Regions that spawn new digital innovations typically house flagship firms, dense professional networks, and universities or research centers that specialize in the focal technology. Together, these institutions offer stronger career prospects, richer learning opportunities, and clear signaling benefits, drawing applicants nationwide. As a result, the share of applications flowing to the pioneer remains higher than its share of job postings even after similar vacancies appear elsewhere.

Hypothesis 4. *The advantage of pioneering regions in attracting applicant interest is increasing over time.*

This hypothesis may appear counter-intuitive – one might expect the pioneer's edge to fade as other regions catch up. Indeed, some diffusion theories (and regional convergence models) predict that early advantages dissipate once knowledge spreads (Desmet and Rossi-Hansberg, 2009). Our argument, however, is that in skill-centric innovations the pioneer region can enter a self-reinforcing cycle, ultimately widening its lead in attracting talent. Several mechanisms could drive this. As the technology matures, the pioneer region accumulates experienced specialists and complementary industries, bolstering its attractiveness to workers. Knowledge spillovers and professional networks in the region make it ever more valuable for experts to congregate there. Empirically, this would manifest as an increasing share of applications (supply) flowing to pioneer-region jobs, even as jobs themselves diffuse widely.

3 Data and Key Measure Construction

3.1 AngelList Data on Job Postings and Applications

We leverage hundreds of millions of job applications submitted to thousands of startup job postings on AngelList Talent (now branded as "Wellfound"), a platform for job listings in the high-tech startup ecosystem, from 2012 to 2023. AngelList Talent has since its inception in 2011 attracted some 62,000 companies to post over 215,000 jobs. The site has had approximately 3.6 million unique job candidates upload profiles and has connected about 1 million applicants to jobs. We are unaware of other papers investigating both sides of the technical job market at this scale.

We have job listing data on all job postings employers made to the platform including a description of the firm, open position titles, remote work designation, and compensation details (wages and equity). From the applicant data, we have information on demographics, some information on prior experience, and job listings that they applied to (including details of timing). To engage on the platform, job candidates fill in their profile information (such as resume data, including educational history, skills, prior employers, and LinkedIn page) and this is the source of our demographic information. They are then able to search for positions (and can filter their searches by role, compensation, location, skills required, startup size, etc.) and apply for a position from within the platform.

We also observe the matching process between the applicants and jobs. The dataset includes *pairings* between job listings and applicants. A pairing is generated when the applicant views a job listing. Conditional on viewing, applicants decide whether they want to apply to the listing. These pairings provide information on the applicants' consideration sets, allowing us to empirically examine applicants' preference for different job characteristics. Figure 1 plots the geographic distribution of job listings and applicants, demonstrating that our sample has broad geographic coverage of major population centers in the U.S.

3.2 Matching Job Listings to Influential Technologies

We use data from Kalyani et al. (2023), which collects pioneering locations for influential technologies. The dataset includes 1,286 unique technologies (e.g., cloud computing, autonomous ve-

hicles) mapped to the Core-Based Statistical Areas (CBSA) where the technologies were first developed.[4] For each technology, an "emergence year" is identified by marking the year in which the technology reaches 100 cite-weighted patents and has at least 10% annual growth rate in patent filings in the next five years. Within the 10 years following the emergence year of a technology, Kalyani et al. (2023) compute the total number of patents filed in each CBSA. CBSAs are sorted based on the number of patents filed. The top CBSAs, collectively accounting for at least 50% of patents filed, are denoted as *pioneering locations* for the technology.[5]

Job listings are matched to technologies using a text embeddings-based approach. We concatenate skills required by a job listing into a single skill string. We use the OpenAI text embeddings API to produce the embeddings for (1) job listings' skill strings and (2) technologies in the Kalyani et al. (2023) dataset. For each job listing, we identify the set of technologies where the cosine similarity of their text embeddings exceeds some threshold and count them as matches. With this approach, a listing may be assigned to zero or multiple technologies. Our results are robust to different thresholds for the cosine similarity. Table 1 shows the top 10 CBSA by their number of technologies pioneered in our sample.

3.3 Measurement of Labor Market Outcomes

We construct labor market outcomes at the CBSA-Technology-Year level to examine the geographic diffusion of technologies. *Labor demand* at the CBSA-Technology-Year level is measured as the total number of active job listings that mention the focal technology during a given year. *Labor supply* measures the number of unique applicants who apply to job listings mentioning the technology. To measure diffusion, we follow Kalyani et al. (2023) to construct the *coefficient of variation* for demand and supply. For CBSA r, technology k, and year t, the regional share of demand

[4] Their data construction process is roughly as follows. New technologies are identified by first obtaining bi-grams appearing in US patent filings since 1976. Within these bi-grams, only those appearing in patents with at least 1,000 cites are kept. Technology-related bi-grams are filtered through matching with Wikipedia pages.

[5] For example, if San Francisco and Seattle take up 40% and 15% of the initial patents for cloud computing, they collectively produce 55% of the patents which is larger than 50%. Then both are taken as the pioneering locations for cloud computing.

can be computed as:

$$\text{Regional Share (Demand)}_{r,k,t} = \frac{\text{Num. Listings}_{r,k,t}}{\text{Num. Listings}_{k,t}}$$

A higher value indicates that a greater share of listings requiring technology k are from CBSA r. We rely on this measure for our analyses on the pioneering regions in section 4.2. We further perform the following normalization on the regional shares:

$$\text{Normalized Share (Demand)}_{r,k,t} = \text{Regional Share (Demand)}_{r,k,t} \Big/ \left(\frac{\text{Num. Listings}_{r,t}}{\text{Num. Listings}_{t}}\right)$$

The regional demand share is normalized by the CBSA's share of all job listings to ensure that a large population CBSA does not always have a large regional share. A normalized share greater than one indicates that region r accounts for a larger fraction of job listings for a technology k, compared to r's overall share of job listings across all technologies. The coefficient of variation of the normalized share is then computed as follows.

$$\text{CV(Demand)}_{k,t} = \frac{Std.Dev_r\left(\text{Normalized Share(Demand)}_{r,k,t}\right)}{Mean_r\left(\text{Normalized Share(Demand)}_{r,k,t}\right)}$$

If job listings for the technology are distributed uniformly across CBSAs, the normalized share of demand would equal one for all CBSAs, so the standard deviation (i.e. the numerator of $\text{CV(Demand)}_{k,t}$) would equal zero. The coefficient of variation would therefore be zero. A higher value of the coefficient of variation indicates a *more* concentrated distribution of the technology. The same measure is calculated for the labor supply. To compare the diffusion of supply versus demand, we calculate the following CV ratio:

$$\text{CV Ratio}_{k,t} = \frac{CV(Supply)_{k,t}}{CV(Demand)_{k,t}}$$

A CV ratio greater than one indicates that labor supply is more concentrated than labor demand.

4 Analyses

4.1 Aggregate Diffusion

Our first analyses take an aggregate view of how each technology spreads across space. To quantify diffusion, we model the geographic concentration of each technology in each year using the coefficient of variation (CV) measures. We estimate separate panel regressions for labor demand (CV of job postings) and labor supply (CV of applicants), and then for their ratio. All specifications include technology fixed-effects and cluster standard errors at the technology level. Table 2 and Table 4 fit variants of the following specification. For technology k in year t:

$$y_{kt} = \alpha_k + \beta\, Years\ Since\ Emergence_{kt} + \mathbf{X}'_{kt}\gamma + (Years\ Since\ Emergence_{kt} \times \mathbf{X}'_{kt})'\,\delta + \varepsilon_{kt}$$

where $y_{kt} \in \{CV(Supply)_{k,t}, CV(Demand)_{k,t}, log(CV\ Ratio_{k,t})\}$, $Years\ Since\ Emergence_{kt}$ represents the number of years since technology k emerged, and $\mathbf{X}'$ represent select covariates.

Table 2 shows that, conditional on technology fixed-effects, the CV for job postings declines sharply with time: each additional year since a technology's emergence reduces the demand-side CV by 3–10 percentage points (Models 1–2, $p < 0.001$). In contrast, the supply-side CV is statistically flat over the same horizon (Models 3–4), indicating that the geographic dispersion of applicants does not keep pace with the dispersion of postings. This divergence is precisely what *Hypothesis 1* predicts: postings spread, but talent remains tied to early hubs, consistent with the view that human-capital mobility rather than information flow constrains adoption. Because the two CVs evolve differently, the CV ratio (supply divided by demand) rises with time. Across the three specifications in Table 4, the coefficient on $Years\ Since\ Emergence$ ranges from 0.018 to 0.124 and is always significant at conventional levels. What begins as a modest mismatch becomes a pronounced bottleneck as technologies mature: by year ten, supply is appreciably more concentrated than demand, and the gap continues to widen thereafter.

Hypothesis 2 posits that high-skill technologies should experience the greatest supply-demand

mismatch. Table 4 confirms this intuition. The share of technology-related jobs requiring a PhD—a proxy for skill intensity[6]—enters positively and significantly (coefficient approximately 0.08) in all CV-ratio regressions. Although the negative interaction with *Years Since Emergence* suggests a modest narrowing over time (–0.004, marginally significant), the gap remains large throughout the diffusion window. These results align with the skill-biased technical-change literature: specialized expertise diffuses slowly, likely because it depends on local training pipelines, tacit knowledge, and agglomeration economies (Marshall, 1920; Saxenian, 1996).

Taken together, the empirical results provide formal support for *Hypothesis 1* and *Hypothesis 2*. Job postings for new technologies become less geographically concentrated as the technology ages, but the pool of qualified applicants remains relatively more concentrated, so the spatial gap between demand and supply widens—especially when the skills are advanced.

4.2 Pioneer Region Analysis

We now shift from national aggregates to a region-level lens, asking whether and how the metropolitan area in which a technology first emerges (the pioneer CBSA) maintains an advantage over later-adopting regions. Table 3 isolates the labor-market consequences of being the first metropolitan area to host a new technology. The specification regresses the log regional share of job postings (demand) or job applications (supply) on a *Pioneer Region* indicator and its interaction with the number of years since the technology's emergence, while absorbing CBSA, technology, and year fixed effects:

$$y_{rkt} = \beta_1 \, Pioneer\ Region_{rk} + \beta_2 \left(Pioneer\ Region_{rk} \times Years\ Since\ Emergence_{kt}\right) + \alpha_k + \lambda_r + \gamma_t + \varepsilon_{rkt}$$

where y_{rkt} is, in turn, the log of (i) the regional share of job listings (demand), or (ii) the regional share of applications (supply).

In the year of emergence the pioneer city commands a large share of activity. On the demand

[6] Following Kalyani et al. (2023), skill intensity for each technology is constructed at the six-digit SOC level using ACS 2015 data by computing, for each relevant occupation, the share of workers with a doctoral degree.

side the coefficient on *Pioneer Region* is 0.652 log points, implying roughly a 92 percent larger vacancy share than a comparable non-pioneer CBSA. On the supply side the pioneer attracts 0.586 log points more applications—about 80 percent above the benchmark non-pioneering region. Both advantages erode as the technology diffuses, yet at different rates. The interaction term between the pioneer dummy and *Years Since Emergence* is –0.029 log points per year for postings but only –0.025 for applications. This asymmetric convergence implies that the ratio of supply to demand in pioneer regions rises over time, an empirical pattern that sets the stage for the dynamic test of *Hypothesis 4* in the event-study figures. More broadly, the slower decay of applicant concentration echoes *Hypothesis 1* by underscoring that talent disperses more sluggishly than job opportunities.

To assess how pioneering digital regions evolve relative to followers, we implement an event-study regression at the region–tech–year level. We estimate a dynamic specification of the form:

$$y_{rkt} = \sum_{s \neq 0} \beta_s \mathbf{1}\{Years\ Since\ Emergence = s\}_{kt} \times Pioneer\ Region_{rk} + \alpha_k + \lambda_r + \gamma_t + \varepsilon_{rkt}$$

where y_{rkt} is, alternatively, i) the log share of job postings (labor demand) in region r for technology k at time t, ii) the log share of job applications (labor supply), or iii) the log ratio of the two. The coefficients β_s thus trace out the trajectory of the "pioneer effect" over time, relative to the emergence year. All specifications include region (α_r), technology (γ_k), and year fixed effects (δ_t), ensuring that estimation comes from within-region, within-technology comparisons over time. The results of this analysis are presented in Figure 2.

Consistent with Table 3, the pioneer advantage decays over time as the technology diffuses geographically. On the demand side, the pioneer region's share of job postings declines steadily over time. On the supply side, the decline is notably more gradual; the pioneer's share of applications also trends downward, but at a slower rate than on the demand side. These patterns suggest that while new job opportunities eventually spread out across regions, worker interest remains relatively concentrated in the original hub of the innovation.

An important consequence of these diverging paths is that pioneer regions increasingly exhibit

an excess of labor supply relative to demand as diffusion unfolds. We see this in the trajectory of the applications-to-postings ratio. In the emergence year, the pioneer's share of applications is roughly in line with its share of job postings – both are extremely high – so the relative supply-demand ratio shows little difference between pioneer and follower regions. However, as postings diffuse outward faster than applications, the pioneer region begins to attract more applications per job than elsewhere. Figure 2 plots the event-time estimates for the log supply/demand ratio: the coefficients for pioneer regions rise over time, indicating an increasing gap in favor of the pioneer. In practical terms, as the technology matures, a job opening in the pioneer city draws a larger pool of applicants (from across the country) than an equivalent opening in a non-pioneering region. By about 10–15 years after emergence, pioneer regions receive substantially higher relative application volumes per job than non-pioneers, reflecting a persistent labor supply advantage. This result aligns with the idea that even when job vacancies for the new technology become ubiquitous, workers continue to preferentially target opportunities in the pioneering hub.

5 Discussion

Our study highlights the role of talent in technology adoption across geographic regions. While prior research has documented how technologies spread, we focus on the human capital that enables this diffusion. Our analyses reveal that the geographic distribution of skilled workers—not just information flow or technology access—often determines the pace and pattern of technology adoption across regions. As new technologies emerge and job opportunities spread geographically, the supply of qualified workers remains disproportionately concentrated in pioneering regions. This talent-demand mismatch grows *more* pronounced over time, particularly for technologies requiring advanced skills. The persistence of this pattern suggests that the geographic mobility of skilled labor is a mechanism in technology adoption that has been underexplored in previous research.

Our contributions are two-fold. First, we provide an empirically-founded analysis of how talent responds to job opportunities across geographic space and time following a technology's emergence. While existing research has primarily focused on the demand side of technology diffusion, our study incorporates the crucial supply-side dynamics of the labor market. Second, we demonstrate that

technology adoption depends significantly on the availability of skilled workers who can effectively implement these technologies. Unlike technologies that can be adopted through simple purchase decisions, the technologies we study require human application of specialized skills, making the geographic distribution of talent a critical factor in their diffusion patterns.

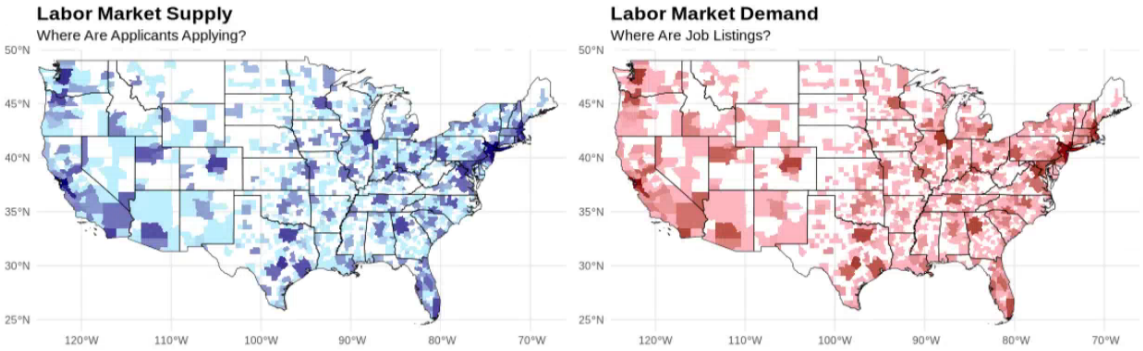
These findings have implications for firms adopting emerging technologies. Organizations outside pioneering regions face a persistent talent disadvantage that could constrain their ability to implement new technologies. Managers should recognize that simply creating job opportunities for new technologies is insufficient—they must also develop strategies to overcome the concentration of skilled talent. This might include offering remote work options, investing in internal training programs, establishing satellite offices in talent hubs, or creating compelling relocation packages.

Table 1: Top 10 Core-Based Statistical Area (CBSA) by Number of Technologies Pioneered

CBSA	Num. Tech. Pioneered	Share of Tech.
San Francisco-Oakland-Fremont, CA	367	0.64
San Jose-Sunnyvale-Santa Clara, CA	273	0.48
New York-Newark-Jersey City, NY-NJ	245	0.43
Boston-Cambridge-Newton, MA-NH	150	0.26
Los Angeles-Long Beach-Anaheim, CA	132	0.23
Seattle-Tacoma-Bellevue, WA	66	0.12
Washington-Arlington-Alexandria, DC-VA-MD-WV	50	0.09
Chicago-Naperville-Elgin, IL-IN	49	0.09
San Diego-Chula Vista-Carlsbad, CA	31	0.05
Austin-Round Rock-San Marcos, TX	27	0.05

Notes: This table shows the top 10 CBSAs by number of technologies pioneered in the region. Share of technologies is the number of technologies pioneered in the CBSA divided by the total number of technologies. There are 572 technologies in total.

Figure 1: Geography of Job Listings and Job Application



Notes: The left-hand plot shades every U.S. metropolitan statistical area (CBSA) in blue according to the number of applications we observe, while the right-hand plot shades the same CBSAs in red according to the number of job postings. Nearly all major labor markets appear in both maps, illustrating that our dataset has broad geographic coverage. This provides the spatial variation necessary to study regional differences in technology diffusion and talent flows.

Table 2: Regression Models of CV Demand and Supply

	CV Demand		CV Supply	
	Model 1	Model 2	Model 3	Model 4
Years Since Emergence	-0.033***	-0.101***	-0.011	-0.039
	(0.006)	(0.022)	(0.007)	(0.026)
Log(Cited Patents)		-0.553***		-0.420***
		(0.085)		(0.103)
Years Since Emergence		0.013***		0.006
× Log(Cited Patents)		(0.004)		(0.005)
Num.Obs.	2131	2130	2131	2130
R2 Adj.	0.192	0.232	0.215	0.250
Std.Errors	by: Technology	by: Technology	by: Technology	by: Technology
FE: Technology	X	X	X	X

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Notes: The models are technology–year level regressions testing how the coefficient of variation (CV) changes with years since the emergence of the technology. A lower CV indicates lower geographic concentration. Demand refers to job listings. Supply refers to number of applicants. $Cited\ Patents$ is the citation-weighted count of USPTO patents that mention the focal technology: each patent's contribution is scaled by its patent-class- and year-normalized forward-citation score, and the weighted counts are then summed for the technology-year in question, which we obtained directly from Kalyani et al. (2023).

Table 3: Pioneer Region Analysis

	Regional Share Demand		Regional Share Supply	
	Model 1	Model 2	Model 3	Model 4
Pioneer Region	0.652**	0.559**	0.586**	0.450*
	(0.209)	(0.192)	(0.192)	(0.179)
Pioneer x Years Since Emergence	-0.029**	-0.026**	-0.025**	-0.021*
	(0.011)	(0.010)	(0.010)	(0.009)
Log(Cited Patents)		0.020*		0.033*
		(0.010)		(0.017)
Num.Obs.	19241	17183	19241	17183
R2 Adj.	0.856	0.845	0.701	0.687
Std.Errors	by: Technology	by: Technology	by: Technology	by: Technology
FE: CBSA	X	X	X	X
FE: Technology	X	X	X	X
FE: Year	X	X	X	X

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Notes: The models are CBSA-technology–year level regressions testing how the regional share of labor demand and supply changes with years since the emergence of the technology and the pioneering status of the region. Pioneering region = 1 if the CBSA is where the technology was first pioneered. Regional supply share measures CBSAs' share of job applicants applying to listings containing the technology over all applicants applying to listings containing the technology. Demand share is analogous but focusing on the share of job listings. $Cited\ Patents$ is the citation-weighted count of USPTO patents that mention the focal technology: each patent's contribution is scaled by its patent-class- and year-normalized forward-citation score, and the weighted counts are then summed for the technology-year in question, which we obtained directly from Kalyani et al. (2023).

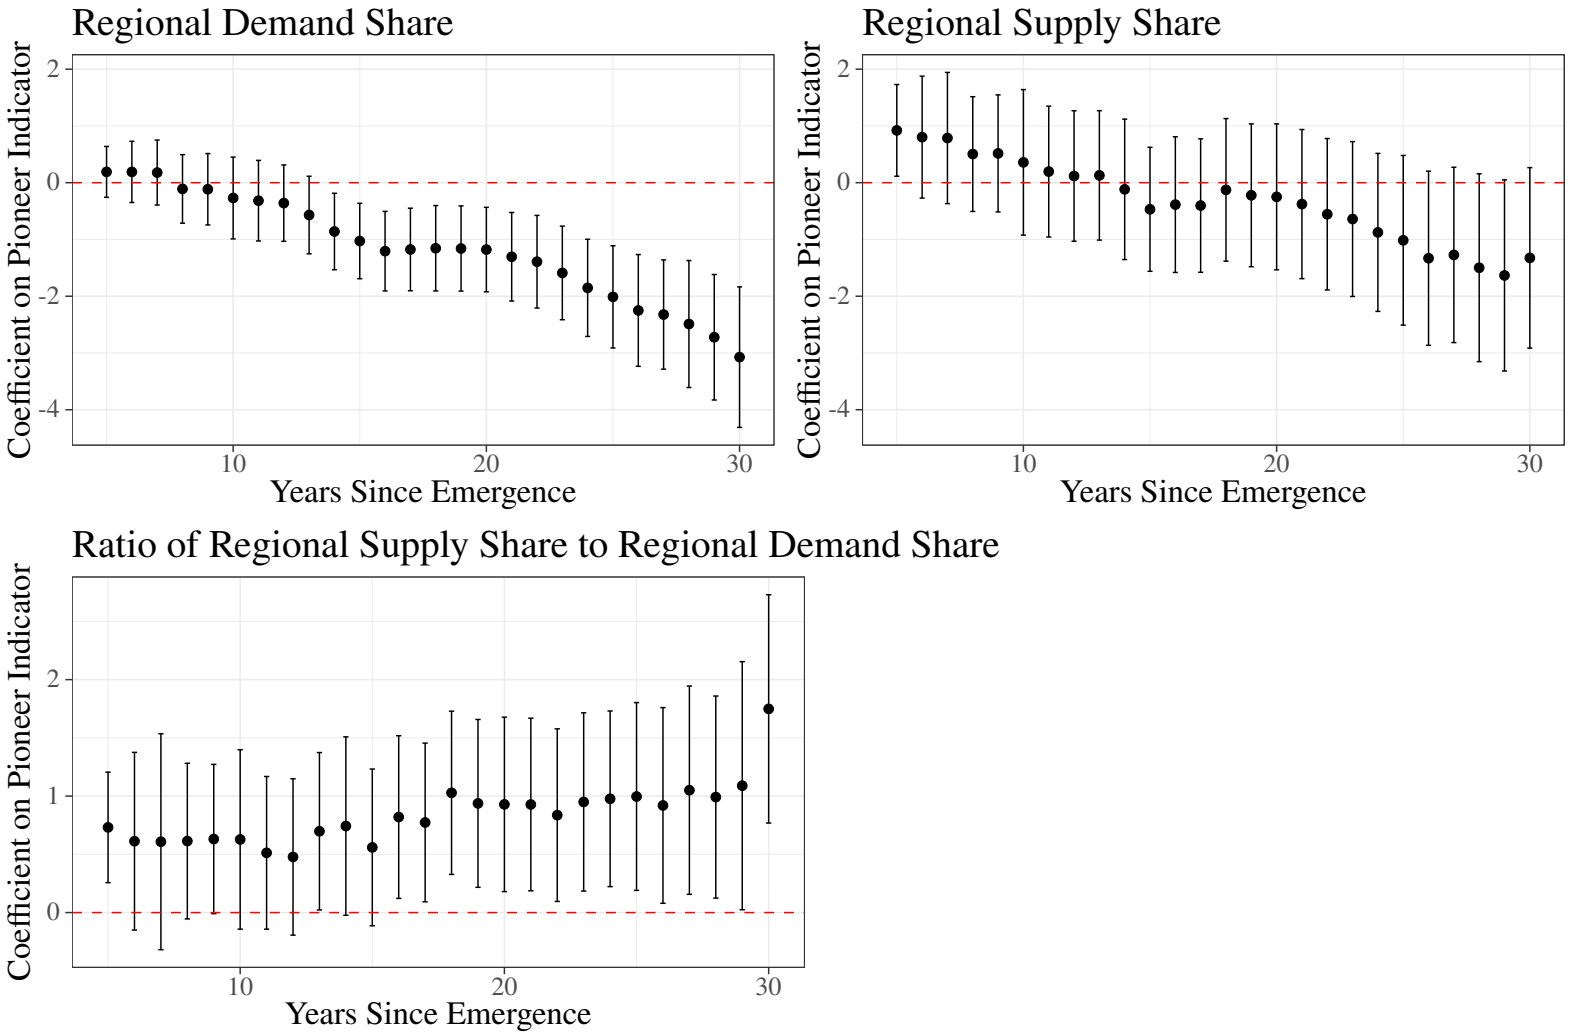
Table 4: CV Ratio over Time

	Model 1	Model 2	Model 3
Years Since Emergence	0.018***	0.049**	0.124**
	(0.004)	(0.018)	(0.045)
Years Since Emergence × Log(Cited Patents)		-0.006+	-0.018*
		(0.003)	(0.008)
Log(Cited Patents)		0.158*	0.448*
		(0.072)	(0.182)
PhD Percentage			0.079*
			(0.038)
Years Since Emergence × PhD Percentage			-0.004+
			(0.002)
Log(Cited Patents) × PhD Percentage			-0.014+
			(0.008)
Years × PhD Percentage × Log(Cited Patents)			0.001
			(0.000)
Num.Obs.	2129	2128	2117
R2 Adj.	0.076	0.079	0.081
Std.Errors	by: Technology	by: Technology	by: Technology
FE: Technology	X	X	X

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Notes: CV ratio is the ratio of CV(Supply) over CV(Demand). A higher value indicates that for a focal technology, the distribution of labor supply (i.e., applicants) is more concentrated than the distribution of labor demand (i.e., job listings). A positive coefficient on year since emergence suggests that over time, applicants are getting relatively more concentrated than job listings. In other words, the labor supply diffuses slower geographically than labor demand. $PhD\ Percentage$ is the share of technology-related jobs requiring a PhD. $Cited\ Patents$ is the citation-weighted count of USPTO patents that mention the focal technology: each patent's contribution is scaled by its patent-class- and year-normalized forward-citation score, and the weighted counts are then summed for the technology-year in question. These measures were obtained directly from Kalyani et al. (2023).

Figure 2: Event Study Plots - Pioneer Region Advantage



Notes: The coefficients are extracted from the event study models. It represents the relative difference in outcome between the pioneering regions and non-pioneering ones. Pioneering region = 1 if the CBSA is where the technology was first pioneered. Regional supply share measures CBSAs' share of job applicants applying to listings containing the technology over all applicants applying to listings containing the technology. Demand share is analogous but focusing on the share of job listings. The last plot measures the ratio of supply share to demand share. A higher value indicates that the region has a disproportionate higher share of job applicants over its job listings.

References

Soon Ang, Sandra Slaughter, and Kok Yee Ng. Human capital and institutional determinants of information technology compensation: Modeling multilevel and cross-level interactions. *Management Science*, 48(11):1427–1445, 2002.

Paul Attewell. Technology diffusion and organizational learning: The case of business computing. *Organization science*, 3(1):1–19, 1992.

David B Audretsch and Maryann P Feldman. R&d spillovers and the geography of innovation and production. *The American economic review*, 86(3):630–640, 1996.

David Autor, Caroline Chin, Anna Salomons, and Bryan Seegmiller. New frontiers: The origins and content of new work, 1940–2018. *The Quarterly Journal of Economics*, page qjae008, 2024.

Harald Bathelt and Johannes Glückler. *The Relational Economy: Geographies of Knowing and Learning*. Oxford University Press, Oxford, 2011. doi: 10.1093/acprof:oso/9780199587384.001.0001.

Jess Benhabib and Mark M Spiegel. Human capital and technology diffusion. *Handbook of economic growth*, 1: 935–966, 2005.

Hoyt Bleakley and Jeffrey Lin. Thick-market effects and churning in the labor market: Evidence from us cities. *Journal of urban economics*, 72(2-3):87–103, 2012.

Erik Brynjolfsson, Daniel Rock, and Chad Syverson. The productivity j-curve: How intangibles complement general purpose technologies. *American Economic Journal: Macroeconomics*, 13(1):333–372, 2021.

Francesco Caselli and Wilbur John Coleman. Cross-country technology diffusion: The case of computers. *American Economic Review*, 91(2):328–335, 2001.

Klaus Desmet and Esteban Rossi-Hansberg. Spatial growth and industry age. *Journal of Economic Theory*, 144(6): 2477–2502, 2009.

Bruce Fallick, Charles A Fleischman, and James B Rebitzer. Job-hopping in silicon valley: some evidence concerning the microfoundations of a high-technology cluster. *The review of economics and statistics*, 88(3):472–481, 2006.

Robert G Fichman. The diffusion and assimilation of information technology innovations. *Framing the domains of IT management: Projecting the future through the past*, 105127:105–128, 2000.

Chris Forman, Avi Goldfarb, and Shane Greenstein. How did location affect adoption of the commercial internet? global village vs. urban leadership. *Journal of urban Economics*, 58(3):389–420, 2005.

Chris Forman, Avi Goldfarb, and Shane Greenstein. The internet and local wages: A puzzle. *American Economic Review*, 102(1):556–575, 2012.

Matthew L Freedman. Job hopping, earnings dynamics, and industrial agglomeration in the software publishing industry. *Journal of Urban Economics*, 64(3):590–600, 2008.

Alfonso Gambardella and Marco S. Giarratana. Localized knowledge spillovers and skill-biased performance. *Strategic Entrepreneurship Journal*, 4(4):323–339, 2010. doi: 10.1002/sej.96.

Edward Glaeser. *Triumph of the City*. Penguin Books, New York, 2011.

Gernot Grabher. The weakness of strong ties: The lock-in of regional development in the ruhr area. In Gernot Grabher, editor, *The Embedded Firm: On the Socioeconomics of Industrial Networks*, pages 255–277. Routledge, London, 1993.

J. Vernon Henderson. The sizes and types of cities. *American Economic Review*, 64(4):640–656, 1974.

Adam B Jaffe, Manuel Trajtenberg, and Rebecca Henderson. Geographic localization of knowledge spillovers as evidenced by patent citations. *the Quarterly journal of Economics*, 108(3):577–598, 1993.

Aakash Kalyani, Nicholas Bloom, Marcela Carvalho, Tarek Hassan, Josh Lerner, and Ahmed Tahoun. The diffusion of new technologies. 2023.

J Daniel Kim. Is there a startup wage premium? Evidence from MIT graduates. *Research Policy*, 47(3):637–649, 2018.

Thomas Klier and James M. Rubenstein. *Who Really Made Your Car?* W.E. Upjohn Institute for Employment Research, Kalamazoo, MI, 2008. doi: 10.17848/9781449632317.

Kenneth L Kraemer, Dale Ganley, and Sanjeev Dewan. Across the digital divide: A cross-country multi-technology analysis of the determinants of it penetration. *Journal of the Association for Information Systems*, 6(12):10, 2005.

Paul Krugman. Increasing returns and economic geography. *Journal of Political Economy*, 99(3):483–499, 1991. doi: 10.1086/261763.

Jeffrey Lin. Technological adaptation, cities, and new work. *Review of Economics and Statistics*, 93(2):554–574, 2011.

Alfred Marshall. *Principles of Economics*. Macmillan, London, 8th edition, 1920.

Peter Maskell and Anders Malmberg. Myopia, knowledge development and cluster evolution. *Journal of Economic Geography*, 7(5):603–618, 2007. doi: 10.1093/jeg/lbm020.

Kristina McElheran, J Frank Li, Erik Brynjolfsson, Zachary Kroff, Emin Dinlersoz, Lucia Foster, and Nikolas Zolas. Ai adoption in america: Who, what, and where. *Journal of Economics & Management Strategy*, 33(2):375–415, 2024.

Nigel Melville, Kenneth Kraemer, and Vijay Gurbaxani. Information technology and organizational performance: An integrative model of it business value. *MIS quarterly*, pages 283–322, 2004.

Sunil Mithas and Mayuram S Krishnan. Human capital and institutional effects in the compensation of information technology professionals in the united states. *Management Science*, 54(3):415–428, 2008.

Jason Owen-Smith and Walter W. Powell. Knowledge networks as channels and conduits: The effects of spillovers in the boston biotechnology community. *Organization Science*, 15(1):5–21, 2004. doi: 10.1287/orsc.1030.0054.

Michael Roach and Henry Sauermann. Can technology startups hire talented early employees? ability, preferences, and employee first job choice. *Management Science*, 2023.

Henry Sauermann. Fire in the belly? Employee motives and innovative performance in start-ups versus established firms. *Strategic Entrepreneurship Journal*, 12(4):423–454, 2018. ISSN 1932-443X. doi: 10.1002/sej.1267. URL https://onlinelibrary.wiley.com/doi/abs/10.1002/sej.1267.

AnnaLee Saxenian. *Regional advantage: Culture and competition in silicon valley and route 128, with a new preface by the author*. Harvard University Press, 1996.

Olav Sorenson, Michael S Dahl, Rodrigo Canales, and M Diane Burton. Do startup employees earn more in the long run? *Organization Science*, 32(3):587–604, 2021.

Scott Stern. Do scientists pay to be scientists? *Management Science*, 50(6):835–853, 2004.

Donald N. Sull. The dynamics of standing still: Firestone tire & rubber and the radial revolution. *Business History Review*, 73(3):430–464, October 1999.

Prasanna Tambe. Big data investment, skills, and firm value. *Management science*, 60(6):1452–1469, 2014.

Prasanna Tambe, Xuan Ye, and Peter Cappelli. Paying to program? engineering brand and high-tech wages. *Management Science*, 66(7):3010–3028, 2020.

Steven Toms and Qinglu Zhang. The rise and decline of the lancashire textile industry. *Business History*, 58(5):646–664, 2016. doi: 10.1080/00076791.2015.1122712.

Jonathan Whitaker, Sunil Mithas, and Che-Wei Liu. Beauty is in the eye of the beholder: Toward a contextual understanding of compensation of information technology professionals within and across geographies. *Information Systems Research*, 30(3):892–911, 2019.

Lynne G Zucker and Michael R Darby. Star scientists and institutional transformation: Patterns of invention and innovation in the formation of the biotechnology industry. *Proceedings of the National Academy of Sciences*, 93(23): 12709–12716, 1996.

A Model

A.1 Setup

Consider two regions, indexed by $r \in \{A, B\}$, and two types of skills: frontier skills (F) and mature skills (M). Let L_{rF} and L_{rM} denote the number of workers in region r with frontier and mature skills, respectively.

Skill Acquisition and Knowledge Spillovers. The acquisition of each skill in each region exhibits skill-specific knowledge spillovers. The rate at which workers in region r acquire skill $s \in \{F, M\}$ is given by:

$$\lambda_{rs} = \alpha_s + \gamma_s L_{rs} \tag{1}$$

where $\alpha_s > 0$ is the baseline acquisition rate and $\gamma_s \geq 0$ captures the intensity of within-skill knowledge spillovers. Frontier skills have stronger spillovers, i.e., $\gamma_F > \gamma_M$, and spillovers do not cross skill categories.

Production and Complementarity. Output in each region is determined by a CES production function with strong complementarity between skills:

$$Y_r = \left[\delta_F L_{rF}^{\rho} + \delta_M L_{rM}^{\rho}\right]^{1/\rho} \tag{2}$$

where $\delta_F, \delta_M > 0$ are productivity parameters and $0 < \rho < 1$ controls the elasticity of substitution between skills. For strong complementarity, ρ is chosen close to zero (i.e., $\sigma = \frac{1}{1-\rho}$ is low).

Labor Markets and Migration. Wages for each skill are determined by marginal products:

$$w_{rs} = \frac{\partial Y_r}{\partial L_{rs}} = \left[\delta_F L_{rF}^{\rho} + \delta_M L_{rM}^{\rho}\right]^{\frac{1}{\rho}-1} \cdot \delta_s \cdot \rho L_{rs}^{\rho-1} \tag{3}$$

Workers are mobile and can migrate between regions to maximize their utility. Migration occurs until wage differences, net of migration costs $c > 0$, are arbitraged away.

A.2 Dynamics and Steady State

Suppose region A has an initial advantage in frontier skills ($L_{AF,0} > L_{BF,0}$), resulting in a higher skill acquisition rate for frontier skills in A due to stronger local spillovers. Over time, this advantage amplifies as additional frontier-skilled workers are attracted and trained in A, reinforcing its lead.

Because production requires both skills as complements, the rising stock of frontier skills in A increases the marginal product and thus the wage for mature-skilled workers in that region. This creates a wage premium for mature skills in A, inducing migration of mature-skilled workers from B to A, even though A lacks strong mature-skill spillovers. In equilibrium, both L_{AF} and L_{AM} are higher in region A than in B.