New Technologies and the Skill Premium*

Tarek Hassan†, Aakash Kalyani‡, and Pascual Restrepo§

December 22, 2025

Abstract

This paper shows that the pace of technology creation is a key driver of wage inequality. It develops a model in which skilled workers learn to use new technologies faster than other workers. This advantage dissipates over time as technologies become standardized and easier to use. The skill premium is determined by the interplay between the pace of technology creation and standardization. A rapid pace of technology creation causes a long-lasting increase in the skill premium. We calibrate the model using novel text-based data on new technologies and their changing demand for skills as they age. These data show that new technologies initially require more skilled workers but see a reversal as they age. The data also point to an increased rate of new technology creation starting in the 1980s and tapering off in the 2000s. In response to this measured acceleration in the pace of technology creation, the model generates a 25 log-point increase in the college premium that begins to revert in the 2010s. In extensions, we allow new technologies to diffuse from dense to lower-density cities, and younger workers to have a comparative advantage in new technologies. These explain why the college premium was higher in dense cities and its increase was mainly an urban phenomenon, and why it had a marked age pattern, rising first for young workers and then for older workers.

*The views expressed here are solely those of the authors and do not necessarily represent those of the Federal Reserve Bank of St Louis or the Federal Reserve System. We thank Erik Hurst, Loukas Karabarbounis, and Larry Schmidt for helpful comments.

†**Boston University**, 270 Bay State Rd, Boston, MA 02215; thassan@bu.edu.

‡**Federal Reserve Bank of St Louis**, Federal Reserve Bank Plaza, St. Louis, MO 63102; aakash.kalyani@stls.frb.org.

§**Yale University**, 28 Hillhouse Ave, New Haven, CT 06511; pascual.restrepo@yale.edu.

Introduction

In 1965, College-educated Americans earned 1.5 times as much as those without a college degree. Between 1980 and 2010, this college wage premium doubled, resulting in college graduates earning twice the wages of non-graduates. In densely populated areas, such as New York City, this premium rose even further, approaching 2.5 times. Since 2010, the college premium has plateaued and experienced a modest reversal. Figure 1 documents these well-known trends using data from the CPS and Census.[1]

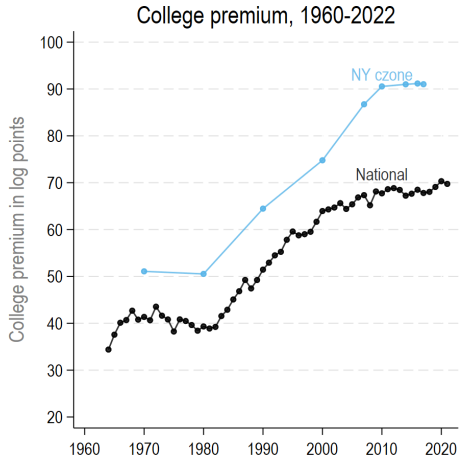


Figure 1: The figure shows the wage college premium for the US at the national level (in black, from the Current Population Survey) and for the New York commuting zone (in blue, from the Decennial Census and the American Community Survey). The college wage premium is given in log points.

Existing theories attribute the rising college premium to structural changes in the production process. Katz and Murphy (1992) emphasize the role of skill-biased technical change—shifts in technology that enhanced the productivity of college workers relative to others. Krusell et al. (2000) emphasize the rapid decline in equipment prices starting in the 1970s. The intensive use of capital goods favors college workers, who have an advantage in using them. Acemoglu and Restrepo (2022) argue that advances in automation reduced the demand for non-college workers by displacing them from tasks they performed. Others highlight structural changes in production

[1] For work documenting these trends see Autor et al. (2008), Acemoglu and Autor (2011), Beaudry et al. (2016), and Autor (2019).

due to information technology (see Beaudry et al., 2016; Burstein et al., 2019).

In this paper, we offer a complementary explanation emphasizing the pace of technology creation rather than structural changes in the production process. In our account, what distinguishes college-educated workers is their ability to learn to operate the newest technologies.[2] This conception of skill goes back to Schultz (1975) and has the clear implication that the college premium depends on the pace of technology creation, even if there are no underlying structural changes in production. This paper demonstrates this implication and quantifies the contribution of changes in the pace of technology creation to rising wage inequality in the US by employing novel text-based data on the emergence and use of new technologies in the workplace.

We develop a model where new technologies are introduced at an exogenous rate. College-educated workers learn to use new technologies faster than others, but their advantage diminishes as the technology ages, becomes standardized, and the knowledge on how to use it spreads. The equilibrium college premium then depends on the interplay between the pace of technology creation, the initial advantage of college workers in these technologies, and how fast it diminishes. An acceleration in the pace of new technology creation causes a long-lasting increase in the college premium. This is true even if newly introduced technologies are not systematically different from previous waves, so that there are no structural changes. Moreover, the increase in the college premium is temporary and reverses as recently introduced technologies are standardized or as all workers learn to use them.

We then show how the model can be quantified using data on the pace of creation of new technologies, their use in the workplace over time, and their changing demand for skills as they age. We use and extend the data from Kalyani et al. (2025), which leverages the text of patents and job postings to trace new technologies and their diffusion. These data show that new technologies initially require more college-educated workers but that their skill demands reverse as they age.

[2]The presumption that college-educated workers have a comparative advantage at adapting to new technologies makes intuitive sense. As Schultz (1975) puts it, "The presumption is that education—-even primary schooling—enhances the ability of students to perceive new classes of problems, to clarify such problems, and to learn ways of solving them." Even if the ability to adapt and learn is innate and not enhanced by education, one would expect fast learners to become highly educated on average. For previous work adopting this assumption, see Greenwood and Yorukoglu (1997), Caselli (1999), Galor and Moav (2000), Mukoyama (2004), and Acemoglu and Restrepo (2018). For empirical support, see Bartel and Lichtenberg (1987) and Doms et al. (1997).

This fact disciplines college workers' comparative advantage in using new technology and how quickly it wanes. The data also reveal a significant acceleration in the pace of technology creation beginning in the 1970s, with a subsequent slowdown in the 2000s,

In response to the increased pace of technology creation during this period, the model generates a 25 log-point increase in the college premium between 1980 and 2010 that begins to flatten and revert afterward, precisely as in the data. Using the model, we decompose the data into the contribution of changes in the supply of college-educated workers, changes in the pace of technology creation (our mechanism), and residual structural changes in the production process (as in Katz and Murphy, 1992). We conclude that the total demand for college-educated workers has increased by 105 log points since 1980. Changes in the pace of technology creation account for a fourth of the shift, while the remainder is attributed to residual structural changes.

After presenting these main results, we introduce two extensions of our theory. The extensions show how our framework offers a unified explanation of for two otherwise puzzling dimensions of wage inequality.

The first extension explores the divergent trends in inequality across US regions. As documented by Autor (2019), Rubinton (2020), and Eckert et al. (2022), and seen in Figure 1, the college premium is generally higher in more densely populated places, and also rose the most in these regions. We extend the model to include locations of different population densities, with new technology diffusing over time from high to low-density places (as in Glaeser, 1999). High-density locations will then use younger technologies and have greater demand for college workers, explaining why the college premium is higher in New York than in the rest of the country. Moreover, a heightened pace of technology creation initially generates a larger increase in the college premium for high-density places.

We estimate the rate of technology diffusion by density using the data from Kalyani et al. (2025). We find that newly invented technologies arrive first in densely populated areas and diffuse to others over time. The modal technology mentioned in job postings for the top 1% densest locations is 34 years old, while the modal technology in the bottom 50% lowest density locations is 52 years old, pointing to sizable diffusion gaps. The estimated model matches these gaps and generates

large differences in the steady-state level of inequality across space, with the 1970 college premium ranging from 40 log points in mid-density regions to 52 log points in high-density ones. Moreover, the heightened pace of technology creation over 1970–2000 explains 30%-50% of the larger increase in inequality in dense regions relative to the rest of the country. Our results thus help explain why the increase in the college premium has been mainly an urban phenomenon and suggest that wages for jobs in low-density regions may catch up as the diffusion process continues.[3]

Our second extension explores the role of worker age. As shown by Card and Lemieux (2001), the college premium rose first for young workers, while older ones saw a smaller and protracted increase. Card and Lemieux (2001) attribute this pattern to supply forces. College completion rates increased continuously since 1910 and plateaued for cohorts entering the labor market after 1970. This slows down the supply of young college workers on impact and that of older college workers later on as the post-1970 cohorts age.

The extension shows that the age-specific profile of the college premium can also be explained by demand-side forces. This is the case if young college-educated workers are the ones who excel at mastering newly arrived technologies ahead of others.[4] We are currently working on calibrating this version of the model using data on technology use by worker age and education. Preliminary findings suggest that the extended model can generate age-specific trends in inequality that qualitatively align with the data.

Relationship to the literature and contribution: This paper contributes to the literature exploring how technology affects the college premium, including work by Katz and Murphy (1992) and Krusell et al. (2000). Our paper identifies a novel and specific driving force behind the increase in inequality: the high pace of new technology creation from 1980 to 2000. This driving force can be measured, which allows us to quantify its contribution to trends in the college premium.

Our approach differs from Katz and Murphy (1992), who infer the extent of skill-biased technical change needed to match wage data but do not identify its determinants. Our micro-founded

[3]These findings complement Rubinton (2020) and Eckert et al. (2022). These papers also propose mechanisms that make the effects of skill-biased technical change stronger in urban areas.

[4]The idea that young workers are better at learning aligns with our experience, evidence from cognitive and educational psychology (see Horn and Cattell, 1967; Salthouse, 1996; Lindenberger and Baltes, 1997; Burke and Barnes, 2006) and reduced-form evidence in economics (see Lehr, 2023).

approach also allows us to predict the future trajectory of the skill premium. Absent any further changes in the pace of new technology creation, the model predicts the college premium will revert half of the way to its steady state level by 2080.

Our approach is closer in spirit to Krusell et al. (2000), who identify capital prices as a distinct driving force behind the increase in inequality that can be quantified from available data. The key difference with their work is that we emphasize college-educated workers' ability to use new technologies, while Krusell et al. (2000) emphasize their ability to use equipment, old or new. This is why, in our theory, the relevant driving force is the pace of new technology creation, and in theirs' is the price of capital goods. To see the difference, consider computers. In our theory, college-educated workers use computers more frequently at first, but this difference diminishes as everyone learns to use computers. Instead, in their framework, skilled workers have a permanent advantage in using computers. This is also why their theory predicts permanent effects from a decline in capital prices, while ours predicts temporary (though persistent) effects.[5]

The idea that workers differ in their ability to learn new technologies and the implications of this mechanism for inequality were explored in previous work by Greenwood and Yorukoglu (1997) and Caselli (1999).[6] Relative to these papers, our value added is in quantifying the contribution of this mechanism to observed trends in US inequality over time, space, and age groups. Greenwood and Yorukoglu (1997) propose a model where skilled workers learn how to use new technologies first and are then followed by other workers. They present comparative statics showing that technological breakthroughs raise inequality temporarily. Caselli (1999) emphasizes the role of heterogeneity in learning costs. When new technologies have high learning costs, their introduction brings inequality. When new technologies have low learning costs, their introduction reduces inequality. Caselli (1999) interprets the increase in inequality since the 1970s as being due to the arrival of a wave of technologies with high learning costs. Instead, previous waves had low learning costs and generated wage compression in the 1910s. This is different from our work,

[5] Similar distinctions apply to the literature that focuses on the role of computer prices and the IT revolution, including Autor et al. (1998), Beaudry et al. (2016), and Burstein et al. (2019), among others. Our model pinpoints the pace of new technology creation as a key driver of inequality, above and beyond the specific nature of computers and the "IT revolution."

[6] A separate branch of this literature focused on implications for technology adoption decisions (see Mukoyama, 2004).

which emphasizes the effects of changes in the pace of new technology creation, keeping learning costs unchanged.

Finally, our empirical work draws on text-based approaches to measuring technical change pioneered by Kalyani et al. (2025), Kogan et al. (2023), Kalyani (2024), and Autor et al. (2023). These works identify new technologies and patterns of adoption using textual analysis of patents, job postings, and corporate disclosures. These data provide a deeper micro-founded account of how, when, and where new technologies increase the return to education.

Organization: Section 1 introduces our baseline model and derives its properties. Section 2 explains how we use the data from Kalyani et al. (2025) to estimate the primitives of the model and quantify the effects of changes in the pace of technology creation on the college premium. Section 3 provides our main quantitative results and decompositions. Section 4 explores the interplay of our mechanism with geography, while Section 5 explores the interplay with worker age.

1 A Model of Technology Creation and the Skill Premium

We first consider an economy where only the pace of new technology creation changes over time. We discuss the role of residual structural changes and shifts in supply in Section 2.

We consider an economy in which new technologies arrive exogenously over time and are subsequently used to produce goods and services. Let $m(b)$ denote the mass of technologies introduced at date (cohort) b. At any time t, the mass of technologies of age u (i.e., born at $b = t - u$) is:

$$m_t(u) \;=\; m\big(t-u\big).$$

Once introduced, technologies of age u produce output at time t according to

$$y_t(u) \;=\; A\big(t-u\big)\cdot z(u)\cdot\left[\alpha(u)^{\frac{1}{\gamma}}\, h_t(u)^{\frac{\gamma-1}{\gamma}} \;+\; \big(1-\alpha(u)\big)^{\frac{1}{\gamma}}\, \ell_t(u)^{\frac{\gamma-1}{\gamma}}\right]^{\frac{\gamma}{\gamma-1}},$$

where $h_t(u)$ and $\ell_t(u)$ denote the inputs of skilled (college) and unskilled (non-college) labor, respectively, and γ is their elasticity of substitution.

The terms $A\big(t-u\big)\cdot z(u)$ characterize technologies' productivity. The first term is cohort-specific and assumed to grow exponentially in birth date

$$A(b) \;=\; A_0\, e^{gb}, \text{ with } g > 0,$$

so that newer cohorts are more productive than previous ones. These "quality" improvements are the source of growth in the model.

The second term $z(u)$ captures the productivity life-cycle of technologies as they age. This term reflects the idea that technologies may take some time to reach their potential so that, as technologies age, they become more productive.

Throughout the paper, we assume

$$g > \lim_{u\to\infty} \frac{\dot{z}(u)}{z(u)},$$

so that old technologies eventually become less productive than new ones.

The function $\alpha(u)$ captures the skill intensity of technologies of age u. The function $\alpha(u)$ decreases with age u, reflecting the idea that college-educated workers initially have a comparative advantage at new technologies, and this advantage wanes over time. The level of $\alpha(u)$ represents the extent of this advantage and its slope how quickly the technology is standardized or the knowledge on how to use it disseminates.

Note that both $\alpha(u)$ and $z(u)$ are functions of technology age that do not change with calendar time. That is, all technology waves are identical, with the only force changing over time being the number of technologies in each cohort, $m(b)$. This choice ensures that there are no structural changes over time shifting the demand for college workers.

Aggregate output at time t is a CES aggregate (with elasticity $\sigma > 1$) of all technologies:

$$Y_t \;=\; \Big(\int_0^\infty m_t(u)\, y_t(u)^{\frac{\sigma-1}{\sigma}}\, du\Big)^{\frac{\sigma}{\sigma-1}}.$$

The total supply of non-college and college labor are fixed at ℓ and h, respectively. Labor-market

clearing requires that workers are allocated across all vintages:

$$\int_0^\infty m_t(u)\,\ell_t(u)\,du \;=\; \ell, \quad \int_0^\infty m_t(u)\,h_t(u)\,du \;=\; h.$$

Equilibrium: We treat the final good as the numeraire and normalize its price to 1. Given the path of technology vintages $\{m_t(u)\}$, an equilibrium is a sequence for output Y_t and real wages $\{W_{h,t},\, W_{\ell,t}\}$, so that labor markets clear and factor payments add up to the value of total production

$$W_{h,t}h + W_{l,t}l = Y_t.$$

This yields three equilibrium conditions at each point in time:

i. Income adds up,

$$1 \;=\; \int_0^\infty m_t(u)\,\left(\frac{c\big(\alpha(u),W_{h,t},W_{\ell,t}\big)}{A(t-u)\,z(u)}\right)^{1-\sigma} du, \tag{1}$$

where

$$c\big(\alpha(u),\,W_{h,t},W_{\ell,t}\big) \;=\; \big(\,\alpha(u)\,W_{h,t}^{1-\gamma} \;+\; \big(1-\alpha(u)\big)\,W_{\ell,t}^{1-\gamma}\big)^{\frac{1}{1-\gamma}}$$

is the unit cost of producing one unit of technology vintage u's output (normalizing productivity to 1). This cost depends on $\alpha(u)$ which determines the relative weights of the wage rates.

ii. The labor market for college workers clears. This requires total wage payments to college labor, $W_{h,t}\,h$ to match the revenue it generates across all vintages:

$$W_{h,t}\,h \;=\; Y_t \int_0^\infty m_t(u)\,\left(\frac{c\big(\alpha(u),W_{h,t},W_{\ell,t}\big)}{A\big(t-u\big)\,z(u)}\right)^{1-\sigma} \alpha(u)\,\left(\frac{W_{h,t}}{c\big(\alpha(u),W_{h,t},W_{\ell,t}\big)}\right)^{1-\gamma} du. \tag{2}$$

iii. The labor market for non-college workers clears:

$$W_{\ell,t}\,\ell \;=\; Y_t \int_0^\infty m_t(u)\,\left(\frac{c\big(\alpha(u),W_{h,t},W_{\ell,t}\big)}{A\big(t-u\big)\,z(u)}\right)^{1-\sigma} \big[\,1-\alpha(u)\big]\,\left(\frac{W_{\ell,t}}{c\big(\alpha(u),W_{h,t},W_{\ell,t}\big)}\right)^{1-\gamma} du. \tag{3}$$

The three equations determine unique equilibrium values for the endogenous variables Y_t, $W_{h,t}$, and $W_{\ell,t}$

at each point in time. These depend on the pace of technology introduction, summarized by $m_t(u)$, which is the key driving force in the model.

Our first proposition shows that the model admits a balanced growth path when the pace of technology creation is constant.

Proposition 1 (Balanced Growth Path)**.** *Suppose the pace of technology creation is constant, i.e.* $m(b) = m$*, so that* $m_t(u) = m$*. There exists a unique balanced growth path along which real wages and output grow at rate* g*, the college premium remains constant and is independent of* m*.*

All proofs are presented in the appendix.

To illustrate the proposition, consider the special case with $\gamma = \sigma$. When $m_t(u) = m$, the economy aggregates to a CES in h and ℓ:

$$Y_t = A(t)\, m^{\frac{1}{\sigma-1}} \left(\underbrace{\left(\int_0^\infty e^{-(\sigma-1)gu}\, z(u)^{\sigma-1}\, \alpha(u)\, du \right)^{\frac{1}{\sigma}}}_{\equiv \alpha_h} h^{\frac{\sigma-1}{\sigma}} + \underbrace{\left(\int_0^\infty e^{-(\sigma-1)gu}\, z(u)^{\sigma-1}\, [1-\alpha(u)]\, du \right)^{\frac{1}{\sigma}}}_{\equiv \alpha_\ell} \ell^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}.$$

The weights α_h and α_ℓ give the overall skill intensity of the economy and determine the college premium along the BGP:

$$\frac{W_{h,t}}{W_{\ell,t}} = \left(\frac{\alpha_h}{\alpha_\ell} \right)^{\frac{1}{\sigma}} \left(\frac{h}{\ell} \right)^{-\frac{1}{\sigma}}. \tag{4}$$

This is reminiscent of the model of Katz and Murphy (1992), with the key difference that the weights α_h and α_ℓ are micro-founded and depend on the skill requirements and productivity of individual technologies.

This special case illustrates several properties of the balanced growth path:

- Quality improvements, subsumed in the $A(t)$ term, drive growth over time.
- Along the BGP, old and new technologies coexist. Technologies of age u account for constant shares of employment and output. The life cycle of market shares depends on the growth rate of $z(u)$—how fast incumbent technologies improve—relative to g—competition from new entrants. Old technologies eventually lose market share at a rate $(\sigma-1)(g - \lim_{u\to\infty} \frac{\dot{z}(u)}{z(u)}) > 0$, due to competition from new entrants.

- The aggregate skill intensity of the economy is a weighted average of individual skill intensities of technology of age u, $\alpha(u)$. The weights $e^{-(\sigma-1)gu}\ z(u)^{\sigma-1}$ capture the dynamics of market shares by age.

- The BGP college premium then depends on the speed at which technologies become standardized (governed by $\alpha(u)$) and the evolution of their market shares as they age (governed by $z(u)$).

- The BGP college premium is independent of m. Increases in m shift the level of wages and GDP but do not affect the college premium in the long run. The mass of technologies created in each period does not alter the steady-state distribution of technology by age, which is what matters for inequality.

- As usual, the ratio h/ℓ of labor endowments reduces the college premium.

Changes in the pace of technology creation. Although the college premium in (4) is independent of m in the long run, time variation in $m(b)$ can generate temporary increases in inequality.

Proposition 2. *Consider an economy in its BGP at time t_0. A permanent increase in $m(b)$ at t_0 from m to $m' > m$ generates a permanent shift in the level of output and wages and a transitory increase in the college premium. A temporary increase in m generates a transitory increase in the college premium, output, and wage levels.*

Let's illustrate the proposition in the example with $\gamma = \sigma$. The economy continues to aggregate to a CES

$$Y_t = A(t) \Big(\underbrace{\Big(\int_0^\infty m_t(u)\, e^{-(\sigma-1)gu}\, z(u)^{\sigma-1}\, \alpha(u)\, du \Big)^{\frac{1}{\sigma}}}_{\equiv \alpha_{h,t}} h^{\frac{\sigma-1}{\sigma}} + \underbrace{\Big(\int_0^\infty m_t(u)\, e^{-(\sigma-1)gu}\, z(u)^{\sigma-1}\, [1-\alpha(u)]\, du \Big)^{\frac{1}{\sigma}}}_{\equiv \alpha_{\ell,t}} \ell^{\frac{\sigma-1}{\sigma}} \Big)^{\frac{\sigma}{\sigma-1}},$$

where the weights $\alpha_{h,t}$ and $\alpha_{\ell,t}$ are now time varying. This variation is driven by changes in the pace of technology creation, summarized by $m_t(u)$, which determines the distribution of technologies

by age. Moreover, the college premium along the transition is now

$$\frac{W_{h,t}}{W_{\ell,t}} = \left(\frac{\alpha_{h,t}}{\alpha_{\ell,t}}\right)^{\frac{1}{\sigma}} \left(\frac{h}{\ell}\right)^{-\frac{1}{\sigma}}. \tag{5}$$

An increase in $m(b)$ starting at t_0 shifts the distribution of technology by age toward newer technologies in the short run. Because college workers have a comparative advantage in these technologies (i.e., $\alpha(u)$ is decreasing in u), $\alpha_{h,t}$ rises relative to $\alpha_{\ell,t}$, creating an increase in the college premium.

The increase in inequality is transitory. Over time, the distribution of technology by age returns to its initial steady-state level. This is true for both temporary and permanent increases in m. In the latter case, new technologies eventually become old technologies, and the distribution of technology by age reverts back to its initial steady state. This aligns with Proposition 1, which shows that a permanent increase in m raises the level of GDP and wages in the long run proportionally without altering inequality.

Propositions 1 and 2 clarify the link between the pace of new technology creation, the age distribution of technology, and the college premium. In the long run, the equilibrium skill bias of the economy depends only on how quickly technology is standardized ($\alpha(u)$) and how rapidly older vintages fall behind ($z(u)$), as well as on the exogenous supply ratio h/ℓ. In the short run, a rise in the pace of new technology creation, $m(b)$, places greater weight on newer, more skill-intensive technologies, raising the college premium. We now explore these implications in the data and quantify the contribution of changes in the pace of technology creation to observed trends in US inequality.

2 Data and Measurement

This section describes the construction of our dataset, highlighting how we identify the birth of new technologies from patent text data and track their diffusion in job postings. We then use these data to measure the rate of technology creation over time, $m(b)$, and to estimate the functions $\alpha(u)$ and $z(u)$, which mediate the impact of a boom in new technology on the college premium.

2.1 Measuring the Emergence of New Technology and its Demand for Skills

We use and extend the data from Kalyani et al. (2025), who trace the origin of disruptive technologies from 1976 to 2007. Relative to their work, we extend the data back to 1950 and include a wider range of technologies (not only those considered disruptive). This extension captures earlier innovation waves and yields a longer time series essential for our application. All the steps described below follow Kalyani et al. (2025) with minor modifications to accommodate our longer sample.

In the first step, we isolate a list of technical bigrams (two-word combinations) that appear in U.S. patents but were not in common use in English prior to 1930. We combine the full text of patents from the U.S. Patent and Trademark Office—-covering all utility patents filed between 1930 and 2014-—with a historical corpus of English usage–the Corpus of Historical American English (Davies, 2010). Bigrams that appear in patents but not in pre-1930 English text are deemed "technical." We then keep those bigrams that are both "technical" and "influential," defining "influential" to mean that a bigram has been mentioned in at least 100 patents since 1930.

A challenge here is that these bigrams do not necessarily correspond to new technologies. To address it, we filter the technical and influential bigrams using the set of Wikipedia technology entries. These are entries in Wikipedia which have sections on 'application', 'use' or 'features' of the technology. This is done by entering each bigram into the Wikipedia search engine and keeping it if any one of the top-5 results of the search returns a technology entry (e.g., "semiconductor laser" or "fiber optic"). This filter yields a final list of 8,533 technology bigrams associated with influential technological innovations, which we group by the 6,072 technology titles from their corresponding Wikipedia pages. This gives the list of technologies used in our analysis.

Next, to isolate the emergence year for each technology, we track the year in which each bigram crosses a threshold of accelerated patenting. We define a technology's emergence year as the first year when it meets two criteria: (1) at least 100 patents mention the bigram, and (2) the number of patents referencing the bigram doubles within five years thereafter.[7]

[7]The Appendix provides additional details on these filters. The main difference with Kalyani et al. (2025) is that this previous study focused on a subset of 276 disruptive technologies, where this group includes technologies that are extensively discussed in conference calls by listed firms. We also modify their procedure in two minor ways. First,

After identifying this list of new technologies along with their emergence year (b), we trace their deployment in the labor market by cross-referencing with the raw text of online job postings from Lightcast for 2010–2023. Lightcast collects and postings from online job boards and employer websites. They provide data on the job's location and detailed occupational codes, and the full, unstructured, text of the posting. We flag a posting as involving a given technology if the text contains an associated technology bigram. Validation exercises conducted by Kalyani et al. (2025) show that 91% of these mentions reflect a task to be performed with or for that technology.

In the final step, we measure the educational requirements for each job posting. Although Lightcast flags job postings that require a college degree, this variable is sparsely populated, precluding its use. We follow Kalyani et al. (2025) and measure the educational requirements of a job from its detailed occupation, which provides information on the qualifications needed for a job. We assign to each job posting the average share of college educated workers in its target occupation, computed from the 2010 American Communities Survey. For example, a job posting for a Judge is assumed to have a college intensity of 98%. A job posting for a paralegal, on the other hand, has a college intensity of 45%.

The resulting dataset spans 300 million job posts for 2010–2023, allowing us to observe how technologies of varying ages—from those that emerged in the 1950s to those just a few years old in 2010—change their demand for skills and overall employment levels as they age.

2.2 The Pace of Technology Creation 1930-2007 ($m(b)$)

The key driving force in our model is changes in the pace of technology creation over time, $m(b)$. Figure 2 plots the empirical counterpart of $m(b)$, computed as the number of technologies that emerged each year. The dashed line depicts the raw data. The solid line reports a smoothed version, computed by fitting a local polynomial in time, reflecting low-frequency changes in the pace of new technology creation.

The figure shows a temporary increase in the pace of technology creation, with a heightened

because patent citations were sparse prior to 1976, we do not weigh patents by citations when defining influential bigrams and determining emergence years. Second, because some of the older technologies we identify might be less well known, we check the top-5 results from the Wikipedia search engine, as opposed to only the top result. All the remaining steps in the analysis are identical to those in the prior work.

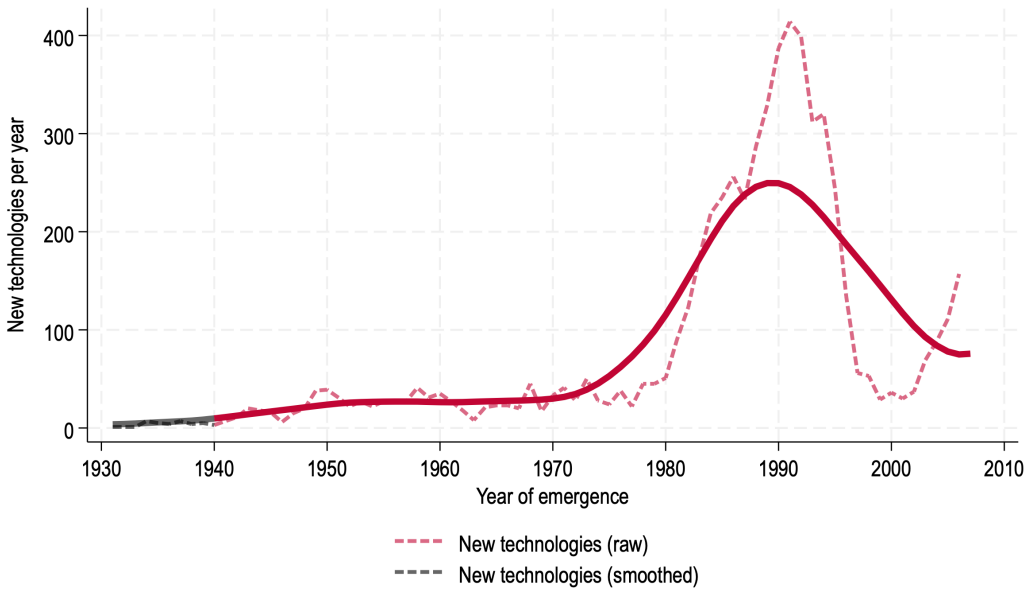


Figure 2: **Time series for arrival of new technologies, $m(b)$.** The figure plots the measured number of new technologies created each year. The dashed line reports the raw data and the solid line is a smoothed version isolating low-frequency variation in the pace of technology creation.

pace in the 1980s and 1990s. Prior to 1970, the pace of technology creation was stable, with 25-30 new influential technologies introduced per year. The pace then accelerates between 1980 and 1990, reaching a peak of 250 technologies per year before 1990. The data also point to a slowdown in the pace of technology creation relative to this peak from 1990 to 2003, with a modest uptick from 2003–2007.[8]

Table 1 lists the top technological innovation in each year since 1930, defined as the technology with more associated patenting activity. These range from computer systems in the 1950s, nucleic acid in the 1970s, and eventually to mobile devices in the 1990s.

We conducted various checks to establish the robustness of the pattern found in Figure 2. These alternative series are shown in Figure 3 Panel A. The college premium in our model depends on changes in the pace of technology creation and not its level. To facilitate comparisons between different measurement approaches, we therefore normalize $m(b)$ to one in 1970 (so that it shows variation in growth rates), and smooth the series to focus on low-frequency changes (short run

[8]We are currently working on extending these series post 2007.

Table 1: Top technologies by emergence year

Title	Top Bigram	Year	Title	Top Bigram	Year
pulse generator	pulse_generator	1951	nucleic acid	nucleic_acid	1979
polymer solution	polymer_solution	1952	targeted drug delivery	drug_delivery	1980
level shifter	voltage_level	1953	solar cell	solar_cell	1981
electrical element	circuit_element	1954	fiber-optic cable	optic_cable	1982
pull-up resistor	resistor_connected	1955	operational data store	datum_stored	1983
shift register	shift_register	1956	holographic optical element	optical_element	1984
epoxy	epoxy_resin	1957	personal computer	personal_computer	1985
manual transmission	output_shaft	1958	communications system	communication_system	1986
comparator	digital_output	1959	user interface	user_interface	1987
flowchart	flow_chart	1960	polymerase chain reaction	chain_reaction	1988
ac power plugs and sockets	ac_power	1961	southern blot	dna_fragment	1989
pulse-width modulation	pulse_width	1962	universal remote	remote_control	1990
antioxidant	free_radical	1963	electrical connector	electrically_connected	1991
laser beam machining	laser_beam	1964	internet protocol suite	internet_protocol	1992
duty cycle	duty_cycle	1965	cable & wireless communications	wireless_communication	1993
solar cell	organic_polymer	1966	stem-cell therapy	stem_cell	1994
particle-size distribution	particle_size	1967	real-time polymerase chain reaction	pcr_product	1995
printed circuit board	printed_circuit	1968	web query	search_query	1996
tracheal tube	plastic_tube	1969	machine learning	machine_learning	1997
multiplexer	input_signal	1970	mobile device	mobile_device	1998
pressure measurement	pressure_sensor	1971	carbon nanotube	carbon_nanotube	1999
release agent	release_agent	1972	marker-assisted selection	desired_trait	2000
size-exclusion chromatography	pore_size	1973	multi-core processor	multi_core	2001
processing units	processing_unit	1974	silo	bulk_storage	2002
analog-to-digital converter	digital_converter	1975	wi-fi	wi_fi	2003
immortalised cell line	cell_line	1976	coaxial cable	coaxial_cable	2004
lateral flow test	reaction_zone	1977	combined charging system	charging_system	2005
growth factor	growth_factor	1978	tablet computers	tablet_computer	2006

Notes: The table reports the top emerging technology by year, defined as the one with the most patents associated to it.

fluctuations do not have sizable effects on the college premium). With this normalization, we re-calculate our estimates of $m(b)$ in a number of plausible ways and compare results.

- *Series (2)* increases the patent-mention threshold from 100 to 1,000, requiring technologies to appear in ten times as many patents to be considered "influential."
- *Series (3)* plots the number of technological bigrams appearing each year independently of whether these are matched to a Wikipedia technology entry, removing the Wikipedia validation filter entirely.
- *Series (4)* normalizes patent counts by the total number of patents filed each year to account for the overall growth in patenting activity over time.
- *Series(5)* removes technologies with early patents in ICT technology classes (patent classes G06, H04, H01, H03).

All series point to a temporary increase in the pace of new technology creation in the 1980s and 1990s of similar magnitude.

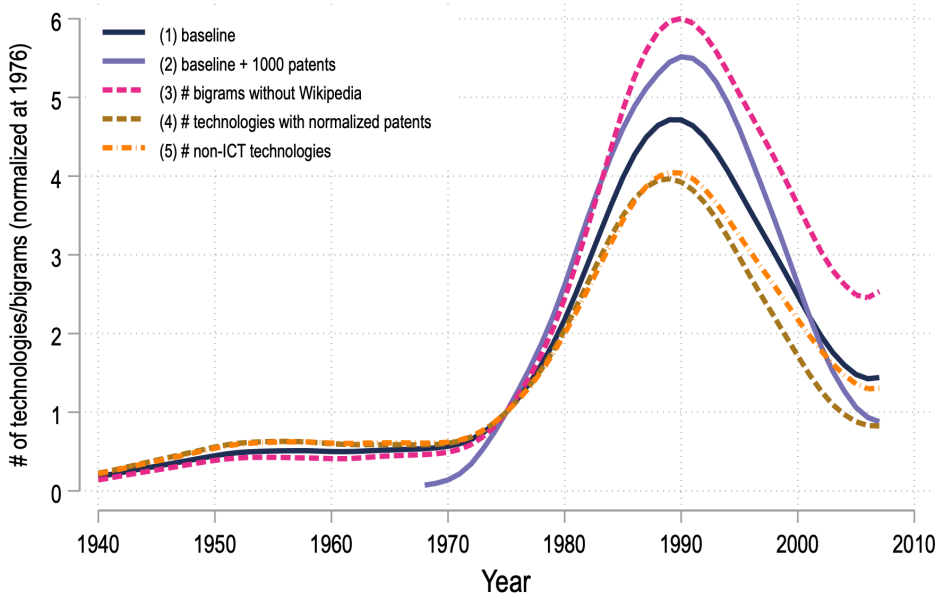


Figure 3: **Robustness checks on the time series for arrival of new technologies,** $m(b)$**.** This figure plots alternative measures of the number of new technologies created each year. Series (1) is the baseline measure. Series (2) increases the patent threshold from 100 to 1,000. Series (3) counts all technological bigrams, including those without a Wikipedia technology page. Series (4) normalizes patent counts by total patents filed each year to define an adjusted threshold. Series (5) removes technologies in ICT technology classes (defined as those with patents in classes G06, H04, H01, H03). All series are smoothed using local polynomial regression and normalized to 1 in 1976 for comparability.

2.3 Estimating $\alpha(u)$ and $z(u)$

We now explain how we estimate the primitives of the model. Our approach is to set the elasticities γ and σ externally and then estimate the functions $\alpha(u)$ and $z(u)$ using the job-posting by technology data. For the elasticities of substitution, we set $\gamma = 1.4$, as in Katz and Murphy (1992). We also set the elasticity of substitution between technologies to $\sigma = 3.5$. This aligns with the median

estimate of elasticities of substitution across varieties in Broda and Weinstein (2006).[9] Finally, we set $g = 2\%$ per year to match the growth rate of the US economy.

Estimating $\alpha(u)$: The function $\alpha(u)$ controls the extent of college workers' comparative advantage at new technologies and how rapidly this diminishes in time.

We estimate $\alpha(u)$ to match the changing demand for college workers in a given technology as it ages. As documented first in Kalyani et al. (2025), there is a clear and robust pattern of skill broadening in the data: new technologies first demand highly educated workers and, over time, increase their demand for non-college workers. Panel A of Figure 4 replicates this finding using our expanded dataset. The figure plots the average college intensity of all job postings associated with technologies of age u. The figure spans ages ranging from $u = 3$ for technologies emerging in 2007 (and seen in Lighcast postings in 2010) to $u = 80$ for technologies emerging in 1943 (and seen in Lighcast postings in 2023).

The figure reveals a clear decline in the college intensity of technologies as they age. In their year of emergence, 56% of jobs involving new technology require a college degree. As technologies age, they become standardized or the knowledge of how to use them diffuses, resulting in a consistent decline in their college demand. For example, 35% of jobs involving 80-year-old technologies require a college degree. This figure supports the presumption that college-educated workers are better equipped to learn to use new technologies but lose this advantage as technology ages.

To capture this pattern in the model, we parametrize $\alpha(u)$ using a logistic function of the form

$$\alpha(u) = \frac{exp(\theta_0 - \theta_1 u)}{1 + exp(\theta_0 - \theta_1 u)}.$$

Here, θ_0 gives the initial advantage of college labor at new technologies and θ_1 the rate at which the advantage diminishes (either because of standardization or knowledge diffusion to non-college workers).

[9] The aggregate elasticity of substitution between college and non-college workers in our model is higher than γ, due to substitution across technologies. Quantitatively, this effect is not strong enough, which justifies equating γ to the aggregate elasticity of substitution in Katz and Murphy (1992).

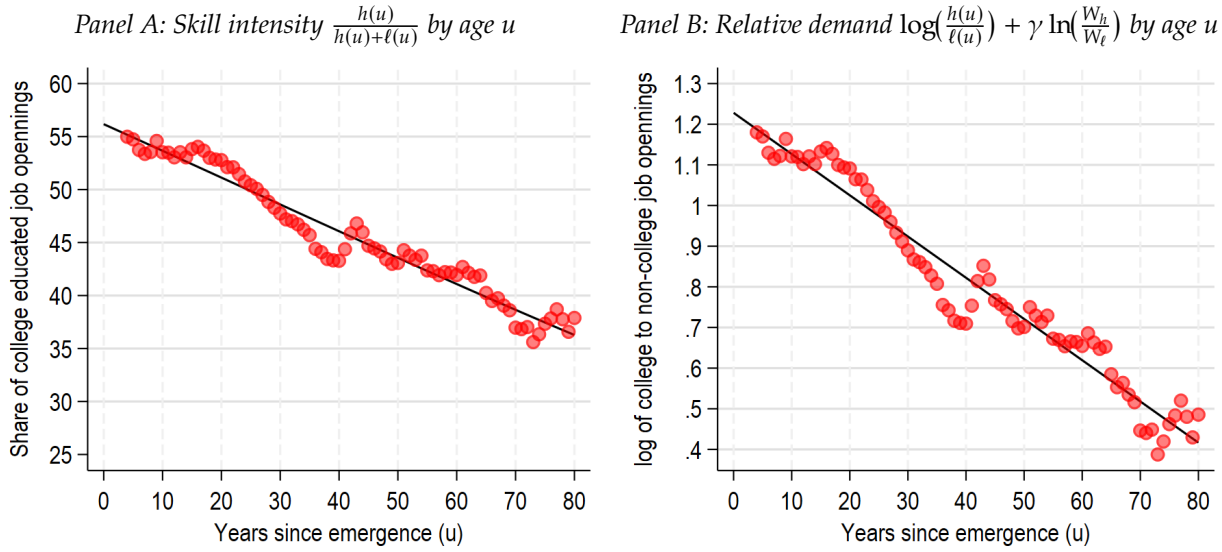


Figure 4: **Standardization and the reduced skill intensity of technologies as they age.** Panel A plots the average skill intensity of job postings by technology age, computed as the share of college-educated workers $\frac{h(u)}{h(u)+\ell(u)}$ for technologies of age u. Panel B plots the relative demand for college labor, computed as $\log(\frac{h(u)}{\ell(u)}) + \gamma \ln(\frac{W_h}{W_\ell})$, where $\gamma = 1.4$ and wages are averaged over 2010–2023. Red markers represent data points aggregated across all technologies at each age. Solid black lines show the fitted values from estimating equation (6). The data span technology ages from $u = 3$ (technologies emerging in 2007) to $u = 80$ (technologies emerging in 1943), all observed in Lightcast job postings from 2010–2023.

We then exploit the fact that the relative demand for college workers by technology age is

$$\frac{h(u)}{l(u)} = \frac{\alpha(u)}{1-\alpha(u)} \left(\frac{W_h}{W_\ell}\right)^{-\gamma}.$$

Taking logs, using the assumed functional form for $\alpha(u)$, and rearranging yields the estimating equation

$$\ln\left(\frac{h(u)}{l(u)}\right) + \gamma \ln\left(\frac{W_h}{W_\ell}\right) = \theta_0 - \theta_1 u + \epsilon_u, \tag{6}$$

where the error term is interpreted as a measurement error. The left side is observed in the data and is shown in Panel B of Figure 4 plotted against technology age. To compute it, we take college requirements by technology age from the job posting data in Panel B, set $\gamma = 1.4$ as discussed above, and use the average college premium $\frac{W_h}{W_\ell}$ over 2010–2023 (matching the period in which job postings are observed).

Estimating equation (6) via OLS yields $\theta_0 = 1.22$ (s.e.=0.01) and $\theta_1 = 0.01$ (s.e.=0.0002), with an R^2 of 0.95. This shows that our functional form provides a good fit to the underlying data. The fitted line and predicted values for college intensity by technology age from our model are shown by the solid black lines in both panels.

The Appendix provides a series of robustness checks for these estimates. For these, we estimate (6) aggregating the data at the cohort× calendar year level. This specification allow us to control for calendar-year, which account for national trends affecting labor markets as well as trends in coverage in the Lightcast data. Doing so yields an estimate for the standardization rate of $\theta_1 = 0.01$. This specification also allow us to control for cohort fixed effects, either by using 20-year cohort dummies (for 1940–1960, 1960–1980, 1980–2000, and 2000–2020) or a quadratic polynomial in cohort. In both cases, we continue to find evidence of standardization as technologies age, with an estimated rate of $\theta_1 = 0.024$. This suggests that the pattern shown in Figure 4 is not merely the result of new technology vintages being more skill intensive than previous ones. We also present estimates removing ICT technologies and using only the early (2010–2015) and late years (2016–2023) in our sample. We continue to find similar point estimates for θ_1, suggesting that the declining demand for skills in a technology is stable in time and not exclusive to ICT technologies.

Estimating $z(u)$: The function $z(u)$ captures the changing productivity of a technology as it ages. This function controls the behavior of the market share of technologies of age u and determines how their influence on aggregates changes as they age.

We estimate $z(u)$ to match the share of workers using technologies of each age u. Panel A of Figure 5 plots the share of employment (regardless of education level) per technology of age u, computed from the job posting data for 2010 to 2023. The figure spans ages ranging from $u = 3$ for technologies emerging in 2007 (and seen in Lighcast postings in 2010) to $u = 80$ for technologies emerging in 1943 (and seen in Lighcast postings in 2023). The figure shows a clear life-cycle for market shares: the share of employment per technology rises from 0.2% to 1.8% 35 years after emerging and then declines slowly as they age.[10]

[10]The shares are computed as a fraction of all job postings associated with technologies 80 years old or younger, as these are the ones in our sample.



Figure 5: **Market share life cycle by technology age.** Panel A plots the average market share in employment per technology of age u, computed as $\frac{h(u)+\ell(u)}{h+\ell}$ where $h(u)$ and $\ell(u)$ are college and non-college workers employed in technologies of age u, and h and ℓ are total college and non-college workers. Panel B plots the transformed market share corresponding to the left-hand side of equation (8), which adjusts for productivity growth, cost differences, and the elasticity of substitution. Red markers represent data points aggregated across all technologies at each age. Solid black lines show the fitted values from the non-linear least squares estimation of equation (8). Market shares are computed from Lightcast job postings data for 2010–2023, covering technology ages from $u = 3$ to $u = 80$.

The market share (in employment) per technology of age u in the model is proportional to

$$\text{Market share}(u) \propto e^{-(\sigma-1)gu}\; z(u)^{\sigma-1} \underbrace{\frac{\alpha(u)W_h^{-\gamma} + (1-\alpha(u))W_\ell^{-\gamma}}{c(\alpha(u), W_h, W_\ell)^{\sigma-\gamma}}}_{\equiv \kappa(\alpha(u), W_h, W_\ell)}. \tag{7}$$

The first term on the right captures the increased competition of new frontier technologies, which reduces the market share of incumbents at a rate that depends on g. The second term captures the dynamics of $z(u)$, reflecting the relationship between the technology's life cycle and productivity. The last term adjusts for differences in cost structures across technologies, which depend on $\alpha(u)$ and the prevailing factor prices in our post-2010 sample, W_h and W_ℓ.

To match the pattern shown in Panel A of Figure 5, we parametrize $z(u)$ as

$$\ln z(u) = g_m u + \frac{1}{\lambda}(g_M - g_m)\,(e^{-\lambda u} - 1).$$

This specification normalizes $z(0)$ to 1 and allows its growth rate to change with age:

$$\frac{\dot{z}(u)}{z(u)} = g_m + (g_M - g_m)e^{-\lambda u}.$$

g_M gives the initial growth rate of $z(u)$. This growth rate decelerates at a rate λ and converges to g_m. The specification generates the observed market share dynamics if $g_M > g$—new technologies first grow fast and gain market share—and $g > g_m$—old technologies lose market share to new, more productive ones.

Taking logs in (7), rearranging, and bringing in the specification for $z(u)$ yields the estimating equation

$$\frac{1}{\sigma - 1} \ln \text{Market share}(u) + gu - \frac{1}{\sigma - 1}\kappa(\alpha(u), W_h, W_\ell) = \text{constant} + g_m u + \frac{1}{\lambda}(g_M - g_m)e^{-\lambda u} + \varepsilon_u, \tag{8}$$

where we added a measurement error ε_u.[11] This equation shows that $z(u)$ can be recovered from the log of market shares, adjusting for the growth of frontier technologies and differences in skill intensity.[12]

The right side in (8) is a transformation of market shares that can be measured from the job posting data, market wages, our previous estimates for $\alpha(u)$ and the assumed values for γ, σ, and g. Panel B of Figure 5 plots the behavior of these transformed market shares.

Fitting (8) via non-linear least squares yields $g_m = 0.54\%$ (se= 0.002), $g_M = 9.52\%$ (se= 0.013), and $\lambda = 0.057$ (se= 0.01). The model fits the data well, with an R^2 of 97%. The fit is shown in Panel B of Figure 5 by the solid line. Panel A also shows the fit of our estimated model in terms of predicted market shares.

The Appendix provides a series of robustness checks for these estimates. As before, we estimate (8) aggregating the data at the cohort× calendar year level. This specification allow us to control for

[11] The constant term in the model reflects the proportionality constant in (7). Because of this, z is only identified up to a multiplicative constant. Our parametrization normalizes $z(0) = 1$. This is without loss of generality since multiplying $z(u)$ by a constant scales all aggregates by the same amount and leaves the college premium unchanged.

[12] This is in analogy to a large literature in industrial organization and international trade that measures productivity by inverting market shares (Syverson, 2004; Foster et al., 2008; De Loecker, 2011)

either calendar-year effects, cohort dummies (or a cohort polynomial), remove ICT technologies, and estimate the model separately for the early (2010–2015) and late years (2016–2023) in our sample. In all cases, we obtain a similar pattern, with $z(u)$ growing at a rate between 8%–26% early on, and switching to a long-term growth rate close to zero. This shows that the pattern shown above, where technologies initially gain market share and then slowly are pushed out of the market by new entrants is a robust finding and is not driven by differences across cohorts, ICT technologies, or other national trends.

Matching the BGP college premium: The data allows us to trace the behavior of $z(u)$ for technologies of up to 80 years of age. However, we do not know how fast these technologies continue to lose market share thereafter. In our model, old technologies lose market share at a rate that depends on $g - \lim_{u\to\infty} \frac{\dot{z}(u)}{z(u)} > 0$. This limit rate of obsolescence is important because it determines how much of employment and economic activity is accounted by old, highly standardized technologies.

To discipline this limit behavior, we assume that after age 80, the obsolescence rate $g - \lim_{u\to\infty} \frac{\dot{z}(u)}{z(u)}$ is constant. We calibrate this rate to ensure that the BGP college premium in our model is 40 log points, matching the inequality level seen before 1980. This ensures that, along the BGP, old technologies generate sufficient demand for non-college labor to match the pre-1980 data. This approach yields a limit obsolescence rate of X% per year. This implies that along the BGP, X% of output is generated by technologies 80 years and older.[13]

3 The College Premium Over Time.

This section uses the estimated model to quantify how measured shifts in the pace of technology creation affected the college wage premium in the US in the last 50 years.

[13] One could alternatively play with the limit properties of $\alpha(u)$ to ensure that the BGP level of inequality in the model aligns with the pre-1980 data. Note also that the limit behavior of $\alpha(u)$ and $z(u)$ is important for getting levels right but has little influence on the estimated changes in inequality reported below.

3.1 Effects of Changes in the Pace of Technology Creation:

We assume the economy was in its BGP before 1970, reflecting the stable rate of new technology creation per year before 1970 shown in Figure 2. We then feed in the changes in the pace of technology creation measured over 1970–2007. We use the smoothed series for $m(b)$ shown by the solid line in Figure 2 to isolate its low-frequency variation. We keep all other determinants of the college premium, including the supply of college-educated workers, fixed at their current levels.

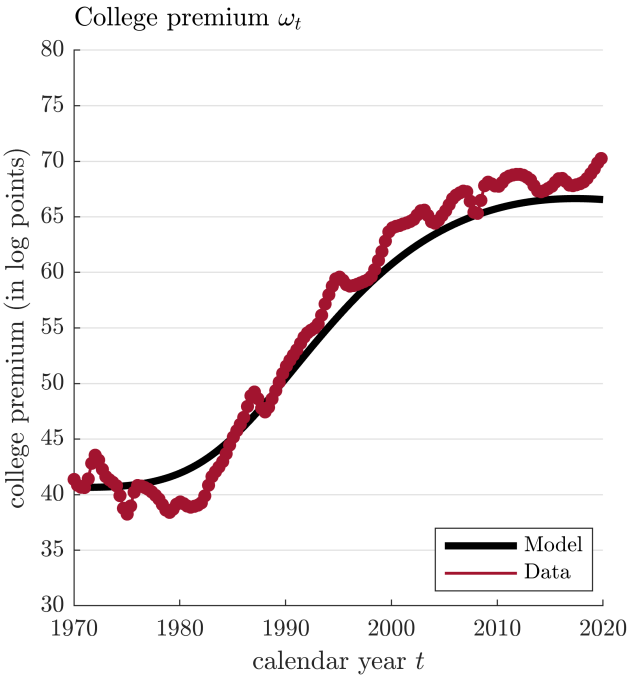


Figure 6: **College premium over time: Model vs. data.** The figure compares the college wage premium from the model (black line) with the data (red line). The model series shows the predicted path of the college premium in response to measured changes in the pace of technology creation $m(b)$ between 1970 and 2007, holding all other factors constant. The data series is computed from the Current Population Survey and shows the log wage differential between college and non-college workers. The model is calibrated to match the 1970 level of the college premium at 40 log points, with all subsequent dynamics driven by changes in the pace of technology creation shown in Figure 3. Both series are measured in log points.

Figure 6 reports the estimated effects on the college premium (in black) and compares these effects to the US data (in red). Recall that the model is calibrated to match the BGP level of the college premium to 1970, while the subsequent increase is untargeted. In response to the heightened rate of technology creation during 1970–2000, the model generates a 27 log point

increase in the college premium. The timing and size of the increase align with the data, where the college premium increased by 28 log points. The model also captures the more recent flattening of the college wage premium. As technologies mature and become standardized, they draw in more low-skill workers, contributing to the slowdown of the wage premium's growth.

The reason why the model matches the timing of inequality so well has to do with the estimated shapes of the $z(u)$ and $\alpha(u)$ functions. The estimated $z(u)$ function implies that technologies have a sizable impact 20–35 years after being introduced. The estimated $\alpha(u)$ function implies that technologies are standardized slowly. Both effects combined imply that a temporary increase in the pace of new technology creation on the college premium is less pronounced, more protracted, and persistent and tapers off much more slowly than the series for $m(b)$ alone would suggest.

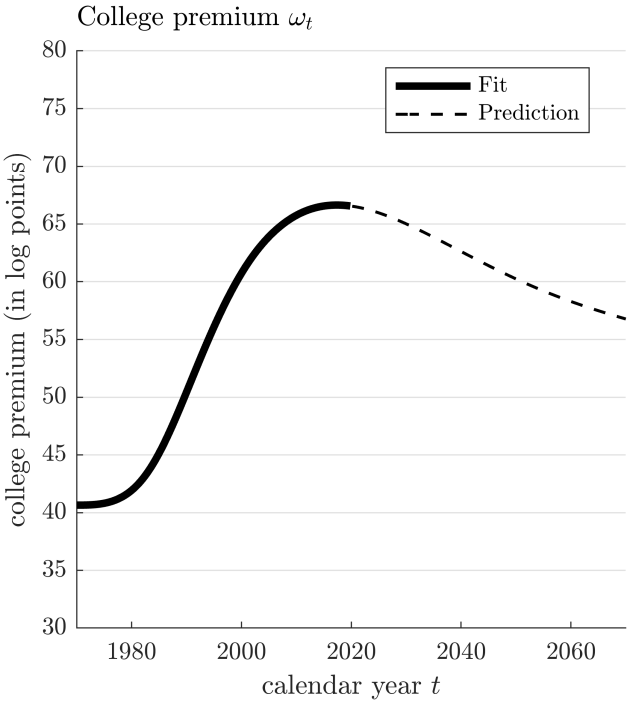


Figure 7: **College premium: Historical fit and future prediction** The figure shows the model-predicted path of the college wage premium. The solid black line shows the model fit from 1970 to 2020, driven by measured changes in the pace of technology creation $m(b)$ between 1970 and 2007 (see Figure 3). The dotted line shows the predicted future trajectory of the college premium from 2020 onward, assuming no further changes in $m(b)$ after 2007. All other model parameters (supply of college workers, structural factors) are held constant to isolate the effect of changes in the pace of technology creation.

The model also predicts the future path of the college premium under the assumption that

there are no subsequent changes in the pace of technology creation after 2007. Figure 7 reports the model prediction. Absent further changes in $m(b)$, the model predicts the college premium to slowly return to half of its BGP level by 2080. The slow transitional dynamics show that the model generates highly persistent effects and that high levels of inequality due to temporary changes in the pace of technology creation can last for generations.

Panel A in Figure 8 reports the estimated effects of the changing pace of technology creation $m(b)$ on wage and output levels. Our estimates show that the upsurge of new technology had a minor impact on wage levels for the first 10 years. It then increased wages for college-educated workers first, between 10 and 20 years after impact. Wage growth for non-college workers lags behind it, as new technologies benefit these workers later in time as they age and are standardized.

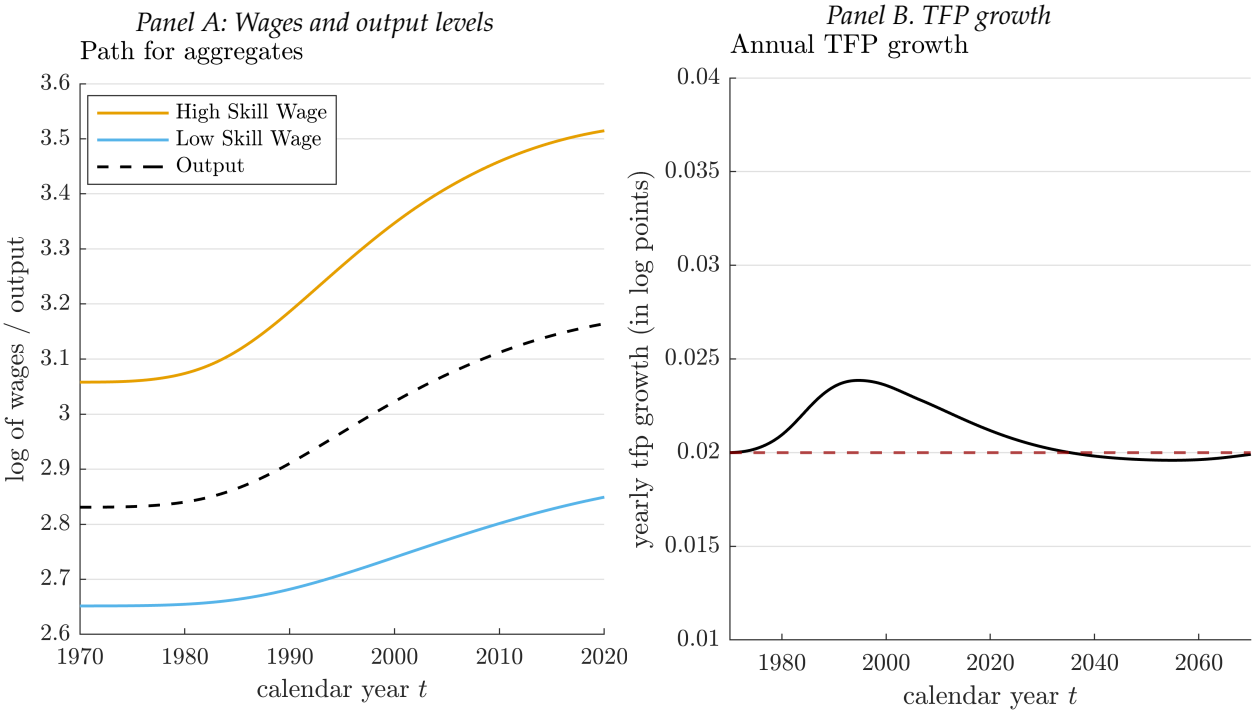


Figure 8: **Model-implied effects on wages, output, and productivity growth.** The figure shows the model-predicted responses of aggregate variables to measured changes in the pace of technology creation $m(b)$ between 1970 and 2007. Panel A plots the log deviations from trend for GDP (dashed line), wages of college-educated workers (orange line), and wages of non-college workers (blue line). The deviation from trend isolates the effect of changes in $m(b)$ on levels, holding the growth rate g constant. Panel B plots annual TFP growth rates. All other model parameters are held constant to isolate the effects of changes in the pace of technology creation.

Panel B in Figure 8 reports the implications of the model for annual TFP growth. Here too, we present the predictions of our model under the assumption that there are no further changes in $m(b)$ after 2007. Our model generates a boom-bust cycle in in line with US data. The effects of the increased pace of technology creation on TFP reach a peak in 1995 when TFP growth goes from its steady-state level of 2% to 2.4%. TFP growth then reverts to its initial BGP level with some minor overshooting around 2040, as the pace of technology creation slows down. Naturally, the model cannot generate a slowdown in the long-run growth rate of the economy, as seen in the US data, since g is taken as exogenous and held constant in these exercises.

3.2 Decomposing Changes in the College Premium Over Time

The previous section assumed the supply of college workers and structural determinants of production remained fixed in (calendar) time. We now account for changes in these determinants of the college premium.

We let the supply h_t and ℓ_t to vary as in the data. Our series for the relative supply of college-educated labor, $\frac{h_t}{\ell_t}$, is computed using hours worked from the Current Population Survey and is shown in Figure 9. The ratio of college to non-college workers rose from .14 in 1970 to .75 in 2010.

We also allow the skill intensity of technology to vary over time, and not just with technology age:

$$\alpha_t(u) \;=\; \frac{exp(\theta_t - \theta_1 u)}{1 + exp(\theta_t - \theta_1 u)}.$$

In this specification, θ_t captures structural changes in production that increase the demand for college-educated workers in all technologies. These shifters capture in a reduced-form way the role of capital-skill complementarity and declining capital prices, automation, and other drivers of the college premium.

We follow Katz and Murphy (1992) and recover θ_t as a residual, computed so that the model matches the college-premium data when we feed in changes in (i) the supply of college-educated labor h_t/ℓ_t, (ii) the residual shifters θ_t, and (iii) the pace of technology creation $m(b)$,

Panel A in Figure 10 reports the recovered series for θ_t in black. The series shows that θ_t increased 3.1% per year, pointing to large residual changes in demand for college workers. For

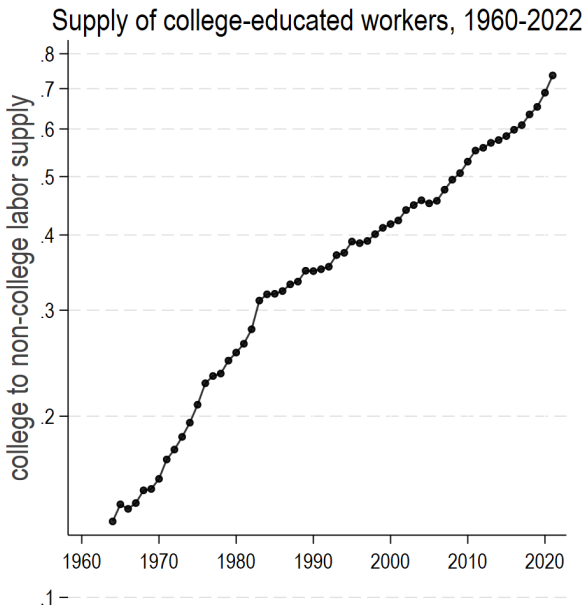


Figure 9: **Supply of college-educated workers, 1960–2022.** The figure plots the ratio of hours worked by college-educated workers relative to non-college-educated workers in the United States from 1960 to 2022, computed using data from the Current Population Survey (CPS). College-educated workers are defined as those with a bachelor's degree or higher. The vertical axis uses a logarithmic scale to facilitate interpretation of proportional changes.

reference, the figure shows in red estimates for θ_t that do not account for the pace of technology creation, $m(b)$. These estimates attribute the full increase in demand for college workers to structural changes in technology, as in Katz and Murphy (1992). The red series shows that one would need a much larger residual, growing at a rate of 4.6% per year, to match the data if we had not accounted for our mechanism. The comparison shows that changes in the pace of technology are half as important as the residual structural changes previously studied in the literature.

Panel B in Figure 10 reports the contribution of (i), (ii), and (iii) to the college premium. The black line gives the contribution of the change in the supply of college educated workers. The vast increase in the share of college-educated workers in the US above would have reduced the college premium by 80 log points. The red dotted line factors in the role of residual structural shifts in technology, which raised the college premium by 80 log points. The blue line factors in the effects of changes in the pace of technology $m(b)$. By construction, this line matches the data and shows that our mechanism adds 25 log points to the increase in the college premium. In this

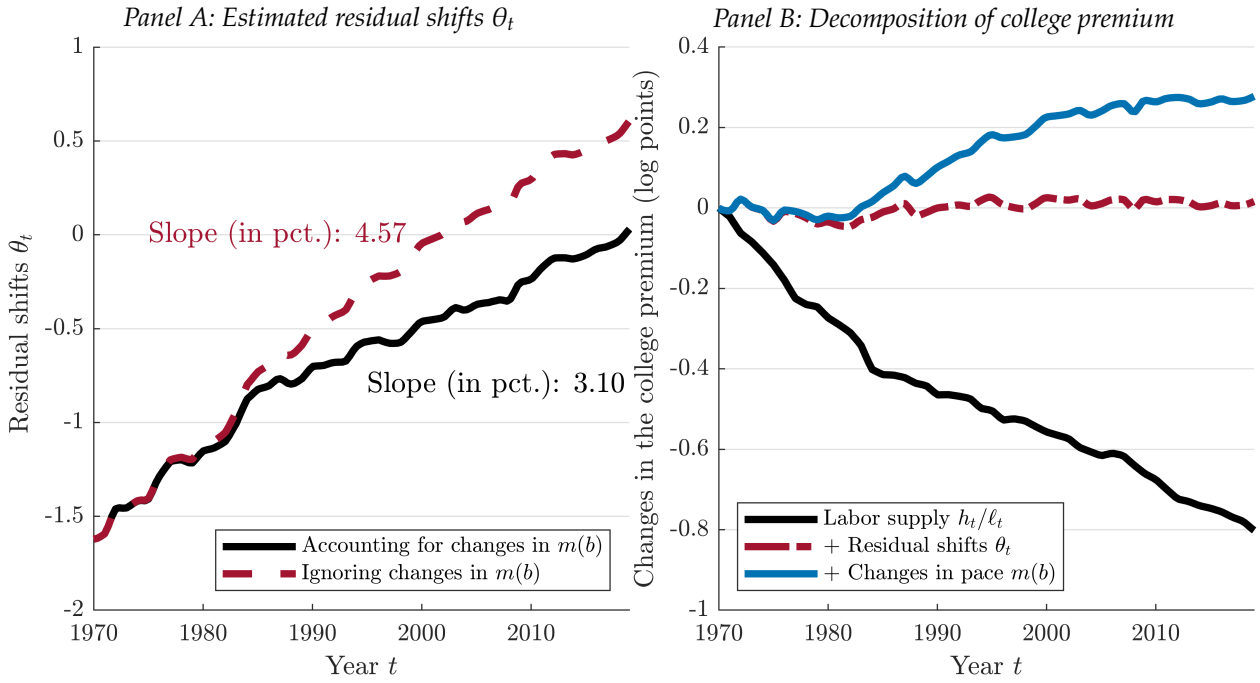


Figure 10: **Decomposition of college premium changes.** Panel A plots estimated residual shifts in technology θ_t that capture structural changes in skill demand beyond those explained by the pace of technology creation. The black line shows θ_t estimated while accounting for changes in $m(b)$, growing at 3.1% per year. The red line shows θ_t estimated without accounting for changes in $m(b)$, as in Katz and Murphy (1992), growing at 4.6% per year. Panel B decomposes the observed log change in the college premium since 1970 into three components: (1) the effect of changes in the relative supply of college workers h_t/ℓ_t (black line), which would have reduced the premium by 80 log points in isolation; (2) residual structural shifts in technology θ_t (red dotted line), which raised the premium by 80 log points; and (3) changes in the pace of technology creation $m(b)$ (blue line), which added 25 log points. By construction, the blue line matches the data.

decomposition, the effect of measured changes in the pace of technology is a third as important as residual structural shifts in technology favoring college workers.

4 The College Premium in Space

This section explores how the pace of technology creation affects the college premium across US regions. Figure 11 portrays the differences in the college premium across commuting zones of different densities and how these have evolved over time. The figure echoes the findings from Autor (2019), Rubinton (2020), and Eckert et al. (2022). It shows the college premium was slightly higher in dense areas by 1980. Since then, its increase has been a marked urban phenomenon. The

college premium rose by 30 log points in dense cities and only 16 log points in less dense areas.

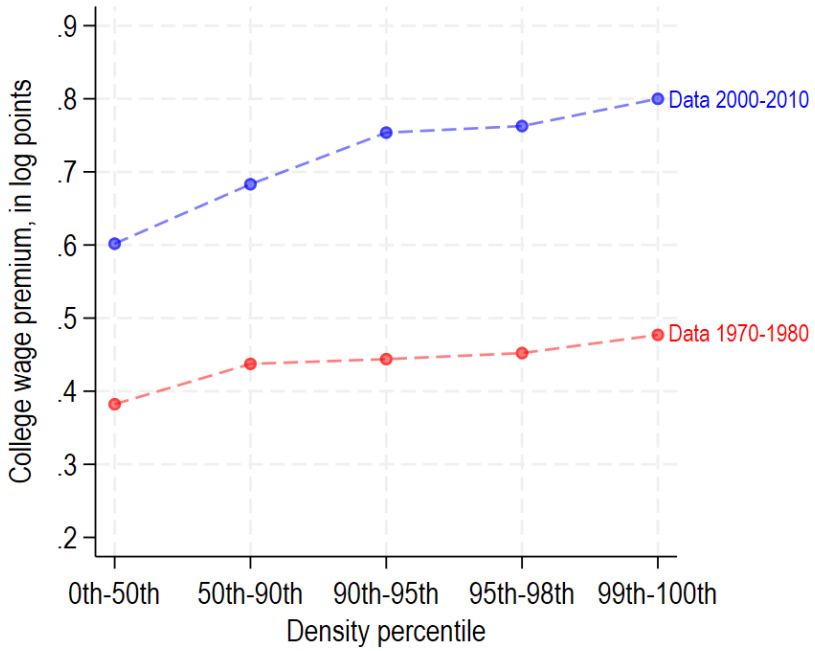


Figure 11: **The college premium in space.** The figure shows the wage college premium plotted against the 1970 density of US commuting zones. The figure sorts commuting zones by density and then estimates the average college premium across different bins. The curves are for the two periods 1970-80 and 2000-10.

The reason why our theory can shed light on these patterns is that there are sizable differences in the rate of diffusion of new technology from high to low-density places. Figure 12 uses our data and plots the modal age of technology used across US regions of different densities. The modal technology used in the bottom 50% lowest-density regions of the US is 52 years old. Instead, the modal technology used in the top 1% highest-density regions in the US is 34 years old. This shows that technologies reach a peak use rate about 20 years early on in the top 1% cities relative to the bottom 50%.

We show that an extension of our model that accounts for this diffusion across space can explain why the college premium was higher in more urban places and rose the most in these regions.

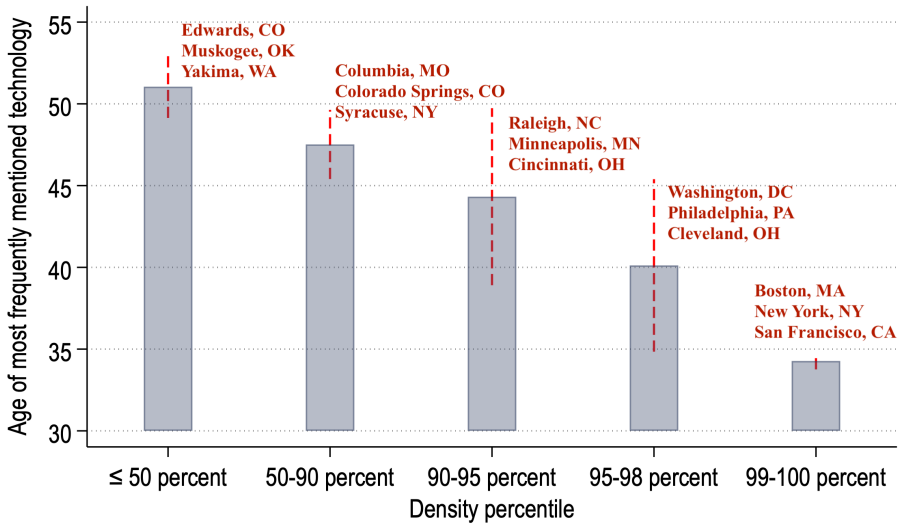


Figure 12: **Technology diffusion from dense to less dense places.** The figure reports the modal age (i.e., the one with the highest employment share) of technologies used in US regions of different densities.

4.1 Extending the model

We consider an economy with multiple locations. Each location is indexed by its density d. Locations do not trade and there is no migration. Locations only interact via technology diffusion.

New technologies arrive exogenously over time. Let $m(b)$ denote the mass of technologies introduced at date (cohort) b. At any time t, the mass of technologies of age u (i.e., born at $b = t - u$) is:

$$m_t(u) \;=\; m\big(t-u\big).$$

Once created, a technology of age u diffuses to a location of density d with probability $p(u,d)$. We assume $p(u,d)$ is increasing and log-submodular in $\langle u,d\rangle$. This implies technology arrives first at high-density places and low-density regions caught up in time, matching the evidence in Figure 12.

Once introduced, technologies of age u produce output at time t according to

$$y_t(u,d) \;=\; A\big(t-u\big)\cdot z(u)\cdot\left[\alpha(u)^{\frac{1}{\gamma}}\,h_t(u,d)^{\frac{\gamma-1}{\gamma}} \;+\; \big(1-\alpha(u)\big)^{\frac{1}{\gamma}}\,\ell_t(u,d)^{\frac{\gamma-1}{\gamma}}\right]^{\frac{\gamma}{\gamma-1}},$$

where $h_t(u,d)$ and $\ell_t(u,d)$ denote labor inputs in location d. As before, $A(b) = A_0 e^{gb}$, and $g > \lim_{u\to\infty} \frac{\dot{z}(u)}{z(u)}$.

Location output at time t is a CES aggregate (with elasticity $\sigma > 1$) of all technologies:

$$Y_t(d) = \left(\int_0^\infty p(u,d)\, m_t(u)\, y_t(u,d)^{\frac{\sigma-1}{\sigma}}\, du \right)^{\frac{\sigma}{\sigma-1}}.$$

Labor-market clearing requires that workers in each location are allocated across all vintages:

$$\int_0^\infty p(u,d)\, m_t(u)\, \ell_t(u,d)\, du = \ell(d), \quad \int_0^\infty p(u,d)\, m_t(u)\, h_t(u,d)\, du = h(d).$$

Here $h(d)$ and $\ell(d)$ are exogenous labor endowments by region. In the analysis, we keep the relative supply $h(d)/\ell(d)$ constant cross regions to emphasize the role of demand-side forces.

Note that locations in our model only differ in that they use a different set of technologies to produce, and these differences are entirely subsumed in the diffusion probability $p(u,d)$. There are no other structural differences across regions affecting inequality.

The equilibrium is defined as before. Given the path of technology vintages $\{m_t(u)\}$, an equilibrium is a sequence for output $Y_t(d)$ and real wages $\{W_{h,t}(d), W_{\ell,t}(d)\}$, that now vary over time—due to changes in the pace of technology creation—and regions—due to differences in diffusion. The equilibrium conditions are shown in the appendix, as these are small modifications of the ones in Section 1.

Proposition 3 (Balanced Growth Path in Space)**.** *Suppose the pace of technology creation is constant, i.e. $m(b) = m$, so that $m_t(u) = m$. There exists a unique balanced growth path along which real wages and output grow at rate g, the college premium $W_{h,t}(d)/W_{\ell,t}(d)$ is constant in time, is independent of m, and increases in density d.*

The proposition shows that the model admits a balanced growth path when the pace of technology creation is constant. The new result is that the BGP level of the college premium rises with density. This is because, relative to low-density regions, the distribution of technology by age in high-density locations is shifted towards young, more skill-intensive technologies (in the first-order stochastic sense).

Proposition 4. *Consider an economy in its BGP at time t_0. A permanent (or temporary) increase in $m(b)$ at t_0 from m to $m' > m$ generates a transitory increase in the college premium in all locations. The increase is more front-loaded in high-density locations. That is, there exists some $\bar{t} > t_0$, such that for $d' > d$ and $t < \bar{t}$:*

$$\ln \frac{W_{h,t}(d')}{W_{\ell,t}(d')} - \ln \frac{W_{h,t_0}(d')}{W_{\ell,t_0}(d')} > \ln \frac{W_{h,t}(d)}{W_{\ell,t}(d)} - \ln \frac{W_{h,t_0}(d)}{W_{\ell,t_0}(d)}.$$

The proposition shows that a heightened pace of technology creation generates an increase in the college premium that, at least for some years, is more pronounced in high-density regions. This is because the newly created technologies reach the high-density first and the low-density regions later on.

4.2 Estimating the diffusion process across space $p(u, d)$

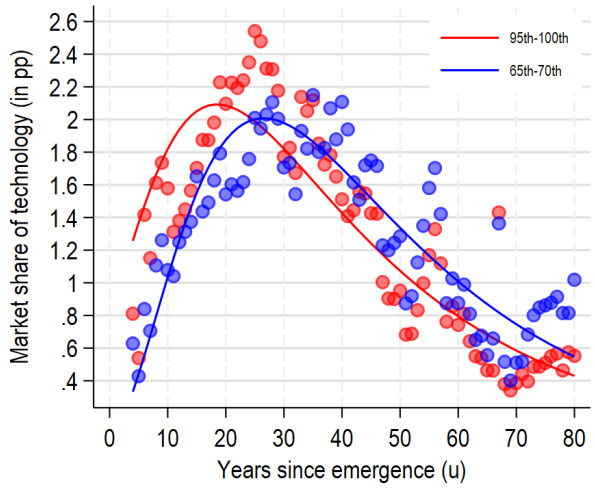
To explore the quantitative implications of Propositions 3 and 4, we extend our estimation procedure to account for differences in technology diffusion across space.

As above, we set $\gamma = 1.4$, $\sigma = 3.5$ and $g = 2\%$ per year. We keep the estimation of $\alpha(u)$ unchanged since the patterns shown in Figure 4 are visible and similar in magnitude within locations. We also calibrate the limit obsolescence rate to generate a BGP college premium of 46 log points in cities with density $d = 0.93$, as in the data.

The difference is that now, the market share (in employment) per technology of age u varies across locations, as implied from the differences in modal ages in Figure 12 and the slow diffusion of new technology to low-density places. We jointly estimate $z(u)$ and $p(u, d)$ to match the share of workers using technologies of age u in regions of different densities. Panel A of Figure 13 plots the share of employment (regardless of education level) per technology of age u, computed from the job posting data for 2010 to 2023. The figure shows this separately for high (in blue, given by the top 5% densest regions) and lower-density (in red, given by places between the 65th and 70th density percentiles). The figure confirms that the distribution of technology use by age in dense regions is shifted towards young technologies relative to that of lower-density regions.[14]

[14]The shares are computed as a fraction of all job postings associated with technologies 80 years old or younger, as these are the ones in our sample. Note also that the figure does not condition on technologies arriving at a place.

Panel A: Market share(u) (in employment) by technology age u.



Panel B: Transformed market share by technology age u (left side of equation 10).

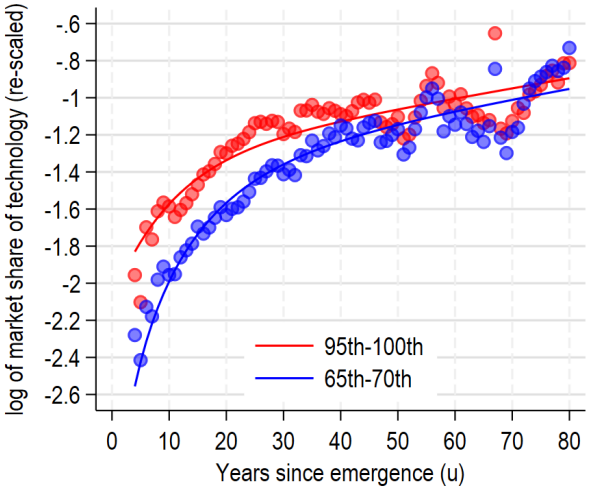


Figure 13: The figures plot the market share in employment for technologies as a function of their age and separately for low and high-density regions. Panel A plots average market share per technology $\frac{h(u,d)+\ell(u,d)}{h(d)+\ell(d)}$. Panel B plots the transformed market share in the right side of equation (10). The red markers are data and the solid black line the model fit.

The market share (in employment) per technology of age u in the model is proportional to

$$\text{Market share}(u,d) \propto p(u,d)\; e^{-(\sigma-1)gu}\; z(u)^{\sigma-1}\; \kappa(\alpha(u), W_h(d), W_\ell(d)). \tag{9}$$

The first term on the right captures the arrival of technology. The remaining terms as the same as before. The last term adjusts for differences in cost structures across technologies, which now depend on prevailing factor prices in our post-2010 sample in each location, $W_h(d)$ and $W_\ell(d)$.

To match the pattern in Panel A of Figure 13, we parametrize $z(u)$ as before and $p(u,d)$ as

$$p(u,d) = 1 - e^{-(\psi_0+\psi_1\frac{d}{1-d})u}.$$

This specification assumes new technology reaches locations with density d at a constant Poisson rate $\psi_0 + \psi_1\frac{d}{1-d}$ per year. ψ_0 gives the arrival rate for low-density locations. ψ_1 parametrizes the higher arrival rate in denser locations. This is normalized to infinity at the highest-density places

with $d = 1$.[15]

Taking logs in (9), rearranging, and bringing in the specification for $z(u)$ and $p(u, d)$ yields the estimating equation

$$\frac{1}{\sigma - 1} \ln \text{Market share}(u, d) + gu - \frac{1}{\sigma - 1}\kappa(\alpha(u), W_h(d), W_\ell(d))$$
$$= \text{constant}(d) + g_m u + \frac{1}{\lambda}(g_M - g_m)e^{-\lambda u} + \frac{1}{\sigma - 1} \ln\left(1 - e^{-(\psi_0 + \psi_1 \frac{1}{1-d})u}\right) + \varepsilon_{u,d}, \quad (10)$$

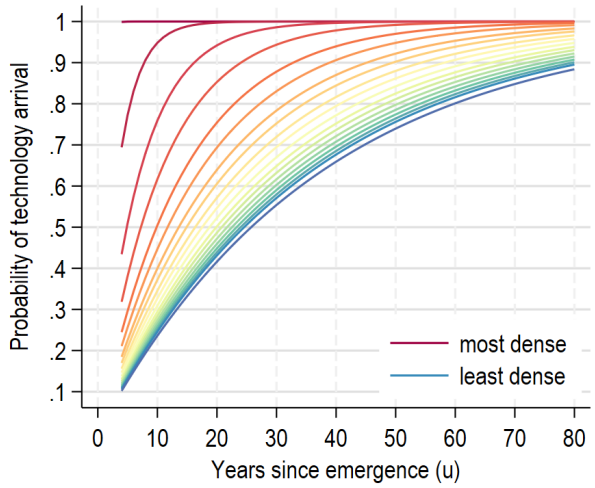
where we added a measurement error $\varepsilon_{u,d}$ and the location-specific proportionality constant in (9). This equation shows that $z(u)$ and $p(u, d)$ can be recovered from the log of market shares, adjusting for the growth of frontier technologies and differences in skill intensity. Market shares in high-density places identify $z(u)$. Differences in market shares by age across densities identify $p(u, d)$.

The right side in (8) is shown in Panel B of Figure 5, for both high and lower-density places. Fitting (10) via non-linear least squares yields $g_m = 0.005$ (se= 0.0003), $g_M = 0.076$ (se= 0.008), and $\lambda = 0.093$ (se= 0.008). These estimates differ from those above because they separate the role of diffusion across space. We also obtain $\psi_0 = 0.026$ (se=0.003) and $\psi_1 = 0.01$ (se=0.001). These estimates imply an arrival rate of 2.6% per year in the lowest density places and of 20% per year in a place with $d = 0.95$. Panel B of Figure 13 plots the fitted curve with a solid line, and Panel A shows the fit of our estimated model in terms of predicted market shares.

Figure 14 shows that the estimated function $p(u, d)$ generates sizable differences in the distribution of technology age across space. The left panel plots the estimated probabilities of diffusion $p(u, d)$ as we move gradually from high-density places ($d = 1$ in red) to low-density ones ($d = 0.075$ in blue). The right panel plots the model-implied market shares by technology age along the BGP, using the same color scheme.

[15] This normalization is inconsequential since we can pool $p(u, 1)$ with $z(u)$ and re-interpret our estimates as capturing the probability of arrival relative to high-density regions.

Panel A: probability of diffusion $p(u, d)$ across densities.



Panel B: Market share(u, d) across densities.

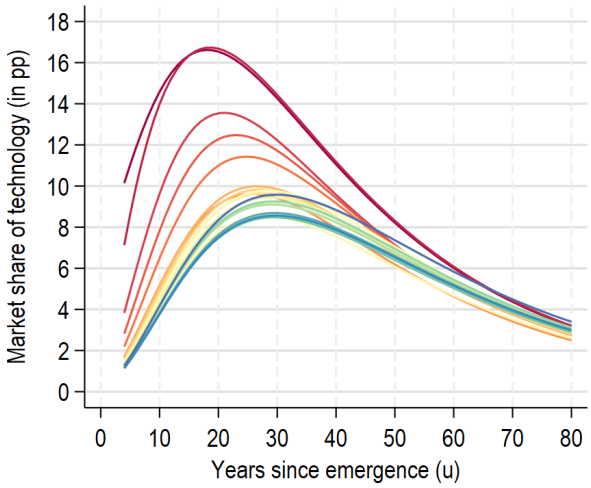


Figure 14: Estimated diffusion probability $p(u, d)$ and market shares Market share(u, d) for different city densities. Blue tones indicate higher densities and red tones lower densities, ranging from $d = 1$ to $d = 0.075$.

4.3 Quantitative results

As above, we assume the economy was in its BGP before 1970. We then feed in the changes in the pace of technology creation measured over 1970–2007. We use the smoothed series for $m(b)$ shown by the solid line in Figure 2 to isolate its low-frequency variation. All other determinants of the college premium, including the supply of college-educated workers, are kept fixed in time and across locations to isolate the role of demand.

Figure 6 reports the estimated effects on the college premium (using a solid line) and compares these effects to the US data (in the dashed line with markers). Recall that the model is calibrated to match a 46 log point college premium in BGP in locations with $d = 0.93$ by 1970. All subsequent increases and estimates of the college premium across locations are untargeted.

The comparison for 1970–1980 shows that the model generates sizable differences in inequality across regions along its BGP. These differences align with what we see in the data. The model matches the 47 log point college premium seen in the densest cities during this period by construction, and predicts a lower college premium of 35 log points in the bottom half cities (vs a 37 log point college premium in the data). The model can, therefore, account for the full 12 log-point gap

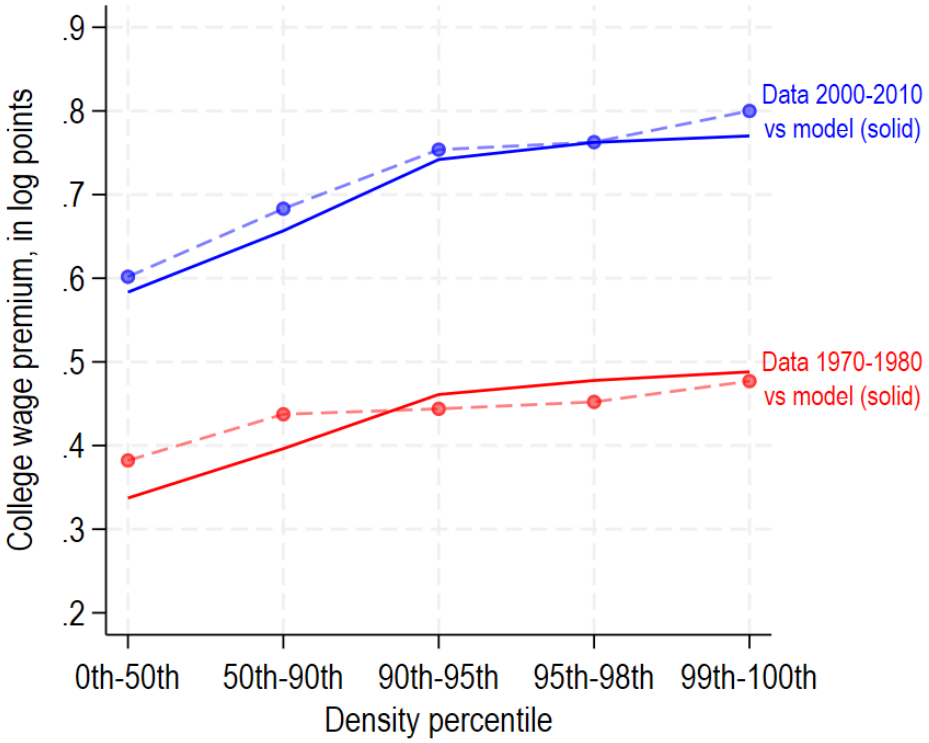


Figure 15: The figure reports the college premium generated by the model across locations of different densities. The panels report data for 1970-1980 (red) and 2000-10 (green). The data are shown in colors, and the model predicted levels in black.

in inequality between the densest and lowest density regions in the US in the 1970–1980 data.

In response to the heightened rate of technology creation during 1970–2000, the model generates a more pronounced increase in the college premium in high-density regions. By 2000-2010, the model generates an increase in the college premium from 47 to 77 log points in the densest locations and a more modest increase from 35 to 57 log points in the lowest-density locations. Both estimates match the data closely. The model can thus account for the 20 log-point gap in inequality between the densest and lowest density regions in the US.

The results show that changes in the pace of technology creation coupled with lags in diffusion can account for some of the spatial patterns in the inequality data.

5 The College Premium by Worker Age

This section explores how the pace of technology creation affects the college premium across worker ages. Figure 16 shows trends in the college premium by worker age or experience groups, measured in years since they finished their schooling. As in Card and Lemieux (2001), we see the college premium increased and plateaued first for young, less experienced workers. Older workers saw a protracted and continuous increase that has not yet plateaued.

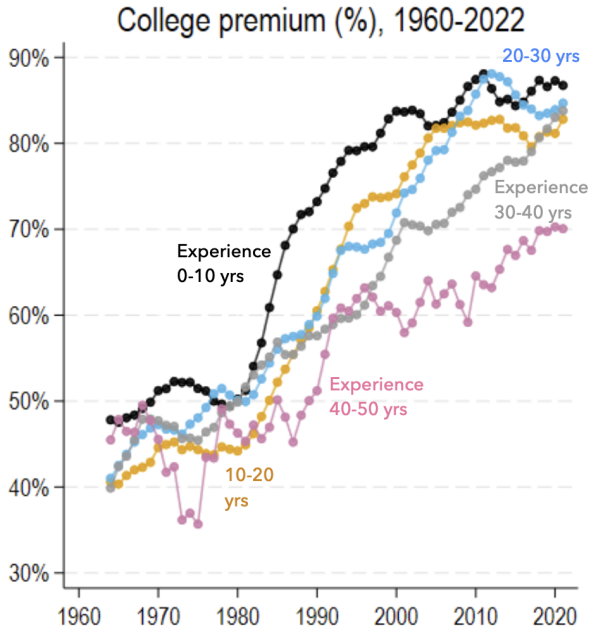


Figure 16: The figure shows the wage college premium by worker age, defined as years since they completed their schooling. Data from the Current Population Survey.

The reason why our theory can shed light on these patterns is the presumption that young college workers are better suited to learn new technologies. If so, an increase in the pace of technology creation should raise the demand for college workers first among the young and less so among older workers who do not use frontier technologies.

5.1 Extending the model

We return to the baseline economy with a single location but now account for worker demographics. Population grows exponentially at a rate n. At each point in time b a cohort of people of size

$$N(b) = e^{n\,b}$$

is born. People are born of age $u_p = 0$, age, and retire at age $\bar{u}$. The population of age u_p at time t is then

$$N_t(u_p) = e^{n\,(t-u_p)}.$$

A fraction h of newborns attends college, and a fraction ℓ does not.

The production side is the same as in the baseline economy. The only difference is that now, the college-labor input used per technology of age u is

$$h(u) = \int_0^{\bar{u}} a(u,u_p)\; h(u,u_p)\; du_p$$

and the non-college input is

$$\ell(u) = \int_0^{\bar{u}} a(u,u_p)\; \ell(u,u_p)\; du_p.$$

Here $h(u,u_p)$ and $\ell(u,u_p)$ are the masses of workers of age u_p employed per technology of age u.

Labor-market clearing requires that workers in each age group $u_p \in [0,\bar{u}]$ are allocated across all vintages:

$$\int_0^{\infty} m_t(u)\,\ell_t(u,u_p)\,du \;=\; \ell\; e^{n\,(t-u_p)}, \qquad \int_0^{\infty} m_t(u)\,h_t(u,u_p)\,du \;=\; h\; e^{n\,(t-u_p)}.$$

The function $a(u,u_p)$ gives the productivity of workers depending on their and technology age. These functions are log super modular in $\langle u,u_p\rangle$, conferring younger workers a comparative advantage at new technology.

As in Costinot and Vogel (2010), The equilibrium now features specialization by age. Young

workers specialize in using new technology and older workers specialize in older ones.

The equilibrium definition parallels the previous one but must now specify the assignment of workers and technologies by age. Given the path of technology vintages $\{m_t(u)\}$, an equilibrium is given by strictly increasing assignment functions $\mu_{h,t}(u_p)$ and $\mu_{\ell,t}(u_p)$ given the age of the technology used by workers of age u, a sequence for output Y_t and real wages $\{W_{h,t}(u_p),\ W_{\ell,t}(u_p)\}$, that now vary over time—due to changes in the pace of technology creation—and workers' age—due to differences in comparative advantage.

The equilibrium conditions are shown in the appendix.

Proposition 5 (Balanced Growth Path and Demographics)**.** *Suppose the pace of technology creation is constant, i.e. $m(b) = m$, so that $m_t(u) = m$. There exists a unique balanced growth path along which real wages and output per worker grow at rate g, the college premium $W_{h,t}(u_p)/W_{\ell,t}(u_p)$ and the matching functions $\mu_h(u_p)$ and $\mu_\ell(u_p)$ are constant in time and independent of m.*

The proposition shows that the model admits a balanced growth path when the pace of technology creation is constant. As in the data, younger workers specialize in using new technologies, while older workers use older ones. Specialization patterns and relative wages are independent of the mass of technologies m, since this leaves the distribution of technology by age unchanged.

Proposition 6. *Consider an economy in its BGP at time t_0. A permanent (or temporary) increase in $m(b)$ at t_0 from m to $m' > m$ generates a transitory increase in the college premium across all worker age groups. The increase is more front-loaded among young workers. That is, there exists some $\bar{t} > t_0$, such that for $u_p < u_p'$ and $t < \bar{t}$:*

$$\ln \frac{W_{h,t}(u_p)}{W_{\ell,t}(u_p)} - \ln \frac{W_{h,t_0}(u_p)}{W_{\ell,t_0}(u_p)} > \ln \frac{W_{h,t}(u_p')}{W_{\ell,t}(u_p')} - \ln \frac{W_{h,t_0}(u_p')}{W_{\ell,t_0}(u_p')}.$$

We are working on the proof of this proposition, which, at this point, is a conjecture based on numerical results.

The proposition shows that a heightened pace of technology creation generates an increase in the college premium that, at least for some years, is more pronounced for young workers. This is because these are the workers who specialize in this segment of the market.

5.2 Quantitative results

We now explore whether the model can generate the age-specific trends in Figure 16.

The key object governing the strength of comparative advantage of young workers in new technologies is the function $a(u, u_p)$. We calibrate this function to match data on Computer use at work by worker age, obtained from the CPS. This dataset is ideal because it includes information on computer use at different points in time, allowing us to see how computer use varies with age as the technology becomes older and more commonplace.

Panel A in Figure 17 plots the data on computer use by worker age across different CPS waves. To facilitate the interpretation, we plot the share of workers using computers at work in each age group relative to the share of 30 year old users. We can see that in early CPS waves, usage was heavily tilted towards young workers. Over time, usage is more commonplace and has risen also for older workers. The same data can be visualized by plotting relative use by worker age group as computers aged, as done in Panel B. Here the horizontal axis gives the years since computers were introduced and the vertical axis plots relative use for 70, 60, 50, and 40 year olds relative to 30 year olds. In all cases we see that older workers use the technology less when it is introduced but catch up over time.

To connect with these facts, we assume that reported use rates in the CPS are a proxy for how proficient people are at using the technology. That is, we assume that

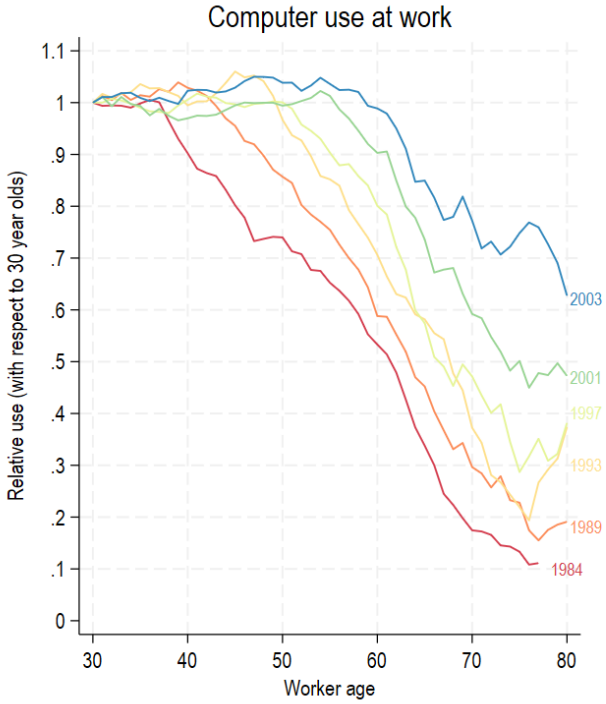
$$a(u, u_p) \;=\; M \text{ share users}(u, u_p)\; e^{\varepsilon_{u,u_p}},$$

where M is a positive constant and ε_{u,u_p} an error term orthogonal to worker or technology age. We also parametrize

$$\ln a(u, u_p) \;=\; -\rho_0\; u_p^{\rho_1}\; e^{-\rho_2 u},$$

and map workers of age 30 to $u_p = 0$ in the model. In this specification, workers of age 30 have a productivity of 1 at new technologies. Older workers have an initial productivity gap $\rho_0\; u_p^{\rho_1}$, which then vanishes as the technology ages at a rate ρ_2.

Panel A: Computer use (by worker age)



Panel B: Computer use (by technology age)

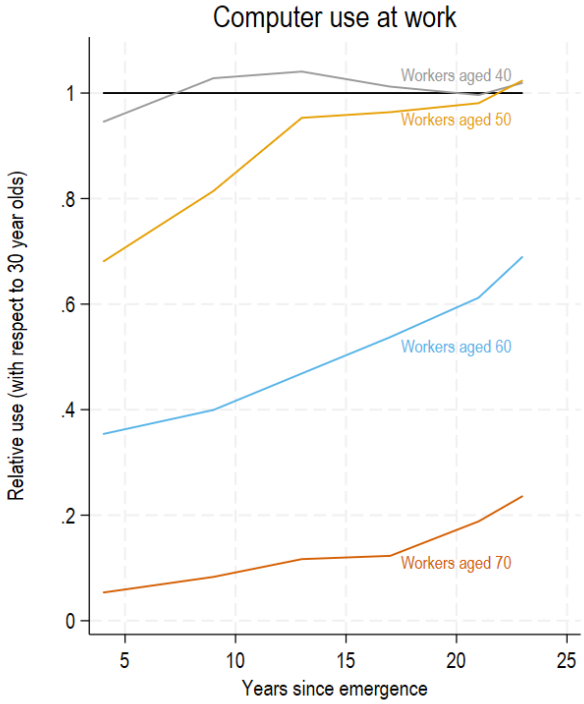


Figure 17: **Computer use at work among workers of different ages**. The figure plots data on computer use by worker (Panel A) and technology age (Panel B) using data from the CPS.

With these assumptions in place, we can estimate $a(u, u_p)$ from the model

$$\ln \frac{\text{share users}(u, u_p)}{\text{share 30-year old users}(u)} = -\rho_0 \, u_p^{\rho_1} \, e^{-\rho_2 u} + \underbrace{\varepsilon_{u,u_p} - \varepsilon_{u,0}}_{v_{u,u_p}}. \tag{11}$$

Estimating this equation via NLLS yields $\rho_0 = 3.45$ (se=0.1), $\rho_1 = 5.01$ (se=0.19), and a low catching-up rate of $\rho_3 = 0.035$ (se=0.002). Our parametrization fits the data on technology use well, yielding an R^2 of 99%.

For the remaining parameters, we use the same values as in the calibration of our model from Section 2. We set $\gamma = 1.4$, $\sigma = 3.5$, and $g = 2\%$ per year, and we use the estimates for $\alpha(u)$ and $z(u)$ from this section.

Figure 18 reports the changes in the college premium generated by the model for workers of different ages. The figure shows that the model is capable of generating the age-specific patterns

seen in the data, with the college premium rising first for young workers and only slowly for older ones. In total, the model explains about 70% of the cohort effects seen in the data.

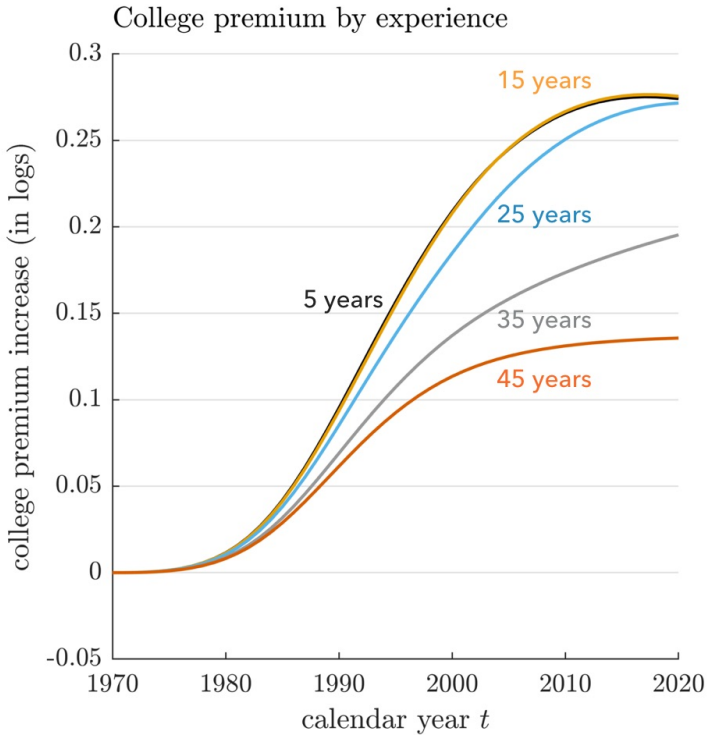


Figure 18: The figure reports the changes in the college premium for workers of different ages.

6 Concluding Remarks

The increase in the college wage premium is a major driver of growing inequality across households and regions in the United States and other developed economies (Autor et al., 2008; Glaeser and Saiz, 2004). It raises critical policy questions, fueling debates about falling social mobility, widening regional disparities in economic opportunities (Chetty et al., 2014), political instability (Autor et al., 2020), and societies' capacity for long-run economic growth (Galor and Zeira, 1993).

This paper argues that changes in the pace of technology creation are a key driver of the increased wage inequality over time, across space, and among worker age groups.

In our model, changes in the pace of new technology affect inequality along these dimensions because: (i) College-educated workers have a temporary advantage in using new technology

(Schultz, 1975), (ii) young technologies arrive first in dense cities, and (iii) young workers are the earliest adopters of young technologies.

Uniquely, due to recent advances in text-to-data techniques, we are able to evaluate each of these forces quantitatively, using data on the development and diffusion of new technologies in the workplace (Kalyani et al., 2025).

In the time series, measured changes in the pace of technology creation generate a 25 log point increase in the college premium and account for a fourth of measured changes in the demand for college labor since 1970. Across regions, our mechanism generates 30% of the urban bias of rising inequality. Our numerical examples suggest that the model is also capable of generating age-specific patterns for inequality seen in the data.

References

Acemoglu, D. and D. H. Autor (2011). Skills, tasks and technologies: Implications for employment and earnings. In O. Ashenfelter and D. Card (Eds.), Handbook of Labor Economics, Volume 4B, pp. 1043–1171. Elsevier.

Acemoglu, D. and P. Restrepo (2018). The race between man and machine: Implications of technology for growth, factor shares, and employment. American Economic Review 108(6), 1488–1542.

Acemoglu, D. and P. Restrepo (2022). Tasks, automation, and the rise in u.s. wage inequality. Econometrica 90(5), 1973–2016.

Autor, D. H. (2019, May). Work of the past, work of the future. AEA Papers and Proceedings 109, 1–32.

Autor, D. H., C. Chin, A. Salomons, and B. Seegmiller (2023). New frontiers: The origins and content of new work, 1940–2018. Technical report, Unpublished Manuscript. Working Paper or Forthcoming.

Autor, D. H., L. F. Katz, and M. S. Kearney (2008). Trends in u.s. wage inequality: Revising the revisionists. The Review of Economics and Statistics 90(2), 300–323.

Autor, D. H., L. F. Katz, and A. B. Krueger (1998). Computing inequality: Have computers changed the labor market? The Quarterly Journal of Economics 113(4), 1169–1213.

Autor, D. H., D. A. Mindell, and E. B. Reynolds (2020). The work of the future: Building better jobs in an age of intelligent machines. MIT Task Force on the Work of the Future.

Bartel, A. P. and F. R. Lichtenberg (1987). The comparative advantage of educated workers in implementing new technology. The Review of Economics and statistics, 1–11.

Beaudry, P., D. A. Green, and B. M. Sand (2016). The great reversal in the demand for skill and cognitive tasks. Journal of Labor Economics 34(S1), S199–S247.

Broda, C. and D. E. Weinstein (2006, 05). Globalization and the gains from variety. The Quarterly Journal of Economics 121(2), 541–585.

Burke, S. N. and C. A. Barnes (2006). Neural plasticity in the ageing brain. Nature Reviews Neuroscience 7(1), 30–40.

Burstein, A., E. Morales, and J. Vogel (2019). Changes in between-group inequality: Computers, occupations, and international trade. American Economic Journal: Macroeconomics 11(2), 289–326.

Card, D. and T. Lemieux (2001). Can falling supply explain the rising return to college for younger men? a cohort-based analysis. The Quarterly Journal of Economics 116(2), 705–746.

Caselli, F. (1999, March). Technological revolutions. American Economic Review 89(1), 78–102.

Chetty, R., N. Hendren, P. Kline, and E. Saez (2014). Where is the land of opportunity? the geography of intergenerational mobility in the united states. The Quarterly Journal of Economics 129(4), 1553–1623.

Costinot, A. and J. Vogel (2010). Matching and inequality in the world economy. Journal of Political Economy 118(4), 747–786.

Davies, M. (2010). The corpus of historical american english (coha). Online database, https://www.english-corpora.org/coha/. Accessed YYYY-MM-DD.

De Loecker, J. (2011). Product differentiation, multi-product firms and estimating the impact of trade liberalization on productivity. Econometrica 79(5), 1407–1451.

Doms, M., T. Dunne, and K. R. Troske (1997, 02). Workers, wages, and technology. The Quarterly Journal of Economics 112(1), 253–290.

Eckert, F., S. Ganapati, and C. Walsh (2022). Urban-biased growth: a macroeconomic analysis. Technical report, National Bureau of Economic Research.

Foster, L., J. Haltiwanger, and C. Syverson (2008). Reallocation, firm turnover, and efficiency: Selection on productivity or profitability? American Economic Review 98(1), 394–425.

Galor, O. and O. Moav (2000, 05). Ability-biased technological transition, wage inequality, and economic growth. The Quarterly Journal of Economics 115(2), 469–497.

Galor, O. and J. Zeira (1993). Income distribution and macroeconomics. The Review of Economic Studies 60(1), 35–52.

Glaeser, E. L. (1999). Learning in cities. Journal of Urban Economics 46(2), 254–277.

Glaeser, E. L. and A. Saiz (2004). The rise of the skilled city. In Brookings-Wharton Papers on Urban Affairs, pp. 47–94.

Greenwood, J. and M. Yorukoglu (1997). 1974. Carnegie-Rochester Conference Series on Public Policy 46, 49–95.

Horn, J. L. and R. B. Cattell (1967). Age differences in fluid and crystallized intelligence. Acta Psychologica 26, 107–129.

Kalyani, A. (2024). The Creativity Decline: Evidence from US Patents. Federal Reserve Bank of St. Louis, Research Division.

Kalyani, A., N. Bloom, M. Carvalho, T. A. Hassan, J. Lerner, and A. Tahoun (2025). The diffusion of new technologies. Technical report, National Bureau of Economic Research.

Katz, L. F. and K. M. Murphy (1992). Changes in relative wages, 1963–1987: Supply and demand factors. The Quarterly Journal of Economics 107(1), 35–78.

Kogan, L., L. Schmidt, and B. Seegmiller (2023). Technology and labor displacement: Evidence from linking patents with worker-level data. Technical report, Unpublished Manuscript. Working Paper or Forthcoming.

Krusell, P., L. E. Ohanian, J.-V. Ríos-Rull, and G. L. Violante (2000). Capital-skill complementarity and inequality: A macroeconomic analysis. Econometrica 68(5), 1029–1053.

Lehr, N. H. (2023). Innovation in an Aging Economy.

Lindenberger, U. and P. B. Baltes (1997). Intellectual functioning in old and very old age: Cross-sectional results from the berlin aging study. Psychology and Aging 12(3), 410–432.

Mukoyama, T. (2004). Diffusion and innovation of new technologies under skill heterogeneity. Journal of Economic Growth 9(4), 451–479.

Rubinton, H. (2020). The geography of business dynamism and skill biased technical change. FRB St. Louis Working Paper (2020-20).

Salthouse, T. A. (1996). The processing-speed theory of adult age differences in cognition. Psychological Review 103(3), 403–428.

Schultz, T. W. (1975, September). The Value of the Ability to Deal with Disequilibria. Journal of Economic Literature 13(3), 827–846.

Syverson, C. (2004). Market structure and productivity: A concrete example. Journal of Political Economy 112(6), 1181–1222.