Shifts in the Composition of College Workers: Implications for US Inequality and Labor Demand

Bharat Chandar *

Link to Latest Version

January 10, 2025

Abstract

I study how changes in the skill composition of workers affect the college wage gap in the US. In a range of wage determination models I show that skill may be measured as log wages subtracting any factors that shift wages for all workers within a firm or over time, such as differences in technology or labor demand. Using LEHD data I find that the average skill of the college workforce declined by 16% between 2002 and 2019, consistent with lower skill for marginal college graduates over time. Average skill for non-college workers changed little. I estimate that growth in the college wage gap would have been four times greater absent changes in skill composition. Through the lens of a supply and demand model, the college wage gap continues to rise despite narrowing skill differences because college workers sort to firms with much faster relative growth in labor demand. **JEL:** J2

*Stanford University; chandarb@stanford.edu. I thank my advisors Rebecca Diamond, Chris Tonetti, and Nick Bloom for invaluable guidance. I also thank Lanwei Yang and Joseph Ballegeer at the US Census and colleagues at Stanford. Any views expressed are those of the authors and not those of the U.S. Census Bureau. The Census Bureau has reviewed this data product to ensure appropriate access, use, and disclosure avoidance protection of the confidential source data used to produce this product. This research was performed at a Federal Statistical Research Data Center under FSRDC Project Number 3042. This research uses data from the Census Bureau's Longitudinal Employer-Household Dynamics Program, which was partially supported by the following National Science Foundation Grants SES-9978093, SES-0339191 and ITR-0427889; National Institute on Aging Grant AG018854; and grants from the Alfred P. Sloan Foundation.

1 Introduction

The gap in average real hourly wages between workers with and without a college degree has grown substantially since 1980, reaching over 60% by the late 2010s (Valletta, 2018). Growing divides by education account for a majority of the increase in income inequality in recent decades (Hoffmann et al., 2020). At the same time, the share of college graduates in the workforce has shot up. From 2002 to 2019 alone, it grew from 32% to 43%.[1] Prior work reconciles growing relative supply of college workers with a higher college wage gap via relative growth in labor demand for educated workers (Katz and Murphy, 1992). Much less empirical research studies how increases in the college share affect the *composition* of college graduates.

Nearly all expansion in US postsecondary education has come within less-selective institutions (Deming, 2022). Growth also predominantly occured among students from disadvantaged backgrounds with worse opportunities even before college (Cilluffo, 2019; Chetty et al., 2020). Average SAT scores for college-bound seniors declined over recent years, particularly verbal scores.[2] These patterns indicate changes in who gets a college degree and which universities they attend. In turn they may affect the college wage gap, a key quantity for understanding growth in inequality and relative labor demand by education.

In this paper, I bring together a model of wage setting with large-scale administrative data to measure changes in the composition of college and non-college workers in the US. I first develop a method for measuring changes in average skill using matched worker-firm panel earnings data. I characterize a class of models in which wages are log additive in worker skill and time-varying firm fixed effects, a prediction with empirical support in the literature that uses the Abowd, Kramarz, and Margolis (1999) (AKM) decomposition (Abowd et al., 1999). Under these conditions, I define average skill for a group of workers as the mean of

[1] Growth in the college share comes from CPS MORG data, described in Appendix C.2. These numbers restrict to individuals with positive earnings, following Valletta (2018).

[2] Data come from College Board (2016). See Figure C1 in Appendix C. Similar declines are evident in individual states. Other standardized tests from the US and other rich countries similarly find declining average scores in recent decades (Dworak et al., 2023; OECD, 2024; Ichino et al., 2024).

log wages after subtracting time-varying firm-level wage shifters. The average skill measure isolates variation in wages that comes from differences across workers, removing variation that arises from changes in technology, labor demand, labor market power, or any other factors that shift wages for all workers within a firm or over time.

I compute the new skill measures using Longitudinal Employer-Household Dynamics (LEHD) data from the US Census. I find that the average skill of the college workforce declined by 15 log points (16%) between 2002 and 2019; the average skill for non-college workers changed little. Overall, I find that the college wage gap would have grown by 15 log points (16%) absent changes in skill composition compared to 4 log points (4%) in the raw wage data, showing that shifts in skill composition substantially reduce growth in wage inequality by education. Through the lens of a Katz and Murphy (1992) supply and demand model, the results suggest significantly faster growth in relative labor demand for educated workers than implied by the increase in the college wage gap alone.

Method. In the data, log wages increased by 7% for college graduates compared to 3% for non-college graduates from 2002 to 2019, but wage changes need not correspond to changes in skill. Recent work shows that, conditional on worker effects, firm-level differences explain a substantial portion of wage inequality (Abowd et al., 1999; Card et al., 2013; Song et al., 2019). These findings suggest that wages may vary either because of skill differences or because of sorting to different firms. As an example, baristas in San Francisco get paid more than baristas in rural parts of California, but these wage discrepencies likely reflect cost of living or labor market differences as much as skill variation. Using wage differences alone may thus lead to biased skill measures.

To measure changes in skill, I regress log wages on worker effects, time-varying firm effects, and controls, subtracting the firm-time effects from log wages to isolate skill. I microfound the approach by showing conditions under which wages are log additive in worker and firm characteristics. The key theoretical assumption is additively substitutable labor. I estimate

time-varying firm effects to allow for changes in firm wage-setting behavior (TV-AKM), in contrast to a standard AKM model that assumes time-invariant firm effects (Lachowska et al., 2023).[3] The firm-time effects are identified via workers who move between firms and correspond to firm pay premia documented in prior literature; they may reflect differences in labor demand, rent sharing, compensating differentials, labor market frictions, adjustment costs, technology, or other factors that may vary across firms and over time.

I compute changes in average skill by education group by averaging the worker component of wages in each year, subtracting firm-time effects from log wages for each worker. I find quantitatively large differences in the growth of pay premia across firms that employ college and non-college workers, a new fact in the AKM literature which demonstrates the importance of allowing for time variation in the firm effects in this setting.

The approach has several nuances. First, the time-varying firm effects absorb any changes that increase pay for all workers at a firm or over time, such as growing labor demand or improved productivity and technology. The precise interpretation of the skill measure supposes that a worker counterfactually employed at another firm and time period would experience the same firm-time level shock to productivity and wages as other workers at that firm. Second, while the approach offers several advantages in my setting compared to a standard AKM model, it inherits restrictions. To identify parameters I must rule out differential rates of persistent skill accumulation, dependence of wages on past and future firm indicators, and endogeneous sorting based on match value (see Kline (2024) for a recent review of the AKM assumptions). In accordance with the prior AKM literature, I find symmetric wages when workers move between firms with high and low pay premia, a result consistent with these assumptions (Card et al., 2013).[4] Third, I impose a normalization that

[3]Other papers that estimate AKM-style models with time-varying firm coefficients include Bagger et al. (2014); Chan et al. (2023b); Engbom et al. (2023). Relative to this prior work, I document substantial heterogeneity in changes in the firm-time effects across industries and counties, and I analyze the implications for wage inequality by education.

[4]I find quantiatively robust results by gender, race, and age; across alternative ways of constructing the sample; when allowing for flexible controls for experience; and when running separate TV-AKM estimations for college and non-college workers, suggesting that while the assumptions may not hold exactly the model provides a useful approximation to the data.

50-year old workers have zero marginal returns to experience, an approach common in the literature that lets me separately identify skill changes due to aging from growth in firm-time effects. (Heckman et al., 1998; Card et al., 2018).

Results. Applying the approach to the data, I find that the average firm pay premium increased by 22 log points (25%) for college workers compared to 7 log points (7%) for non-college workers. After adjusting for pay premia, I find that the average skill of college graduates declined over time by 15 log points (16%), the main result of the paper. The average skill for non-college grads changed little. Overall average skill across all workers also changed little, as the growing share of college workers offset their declining average skill. Absent changes in skill composition, the college wage gap would have increased by 15 log points, similar in magnitude to growth from 1980-1990 when the wage gap increased most rapidly in recent decades (Valletta, 2018).

While the population of college graduates increased across all skill deciles between 2002 and 2019, it increased much faster at the bottom of the college skill distribution. For non-college workers, I find that population declined across the non-college skill distribution as more workers went to college. Together, these results provide evidence of marginal college grads being drawn from across the non-college skill distribution but primarily falling at the bottom of the college skill distribution.

How could the college wage gap continue to increase even as the gap in average skill declined? The reason is that college educated workers sorted to counties, industries, and firms with rapidly growing firm pay premia, while non-college workers sorted to firms with slower growth in pay premia. For example, average pay premia increased by 29 log points in high tech industries that disproportionately employ college-educated workers, compared to only 12 log points in other industries. Did college workers increasingly sort to high-paying firms, or did they already work for firms at baseline that subsequently experienced fast growth in pay premia? I find that over 90% of the change in pay premia for college and non-

college workers can be explained by baseline sorting across counties and industries. These results are consistent with spikes in labor demand in sectors of the economy that already employed high shares of college workers, such as high tech industries or big coastal cities. Marginal college graduates sort to similar sectors over time, benefiting from growth in pay premia that boosts their wages. However, their wages do not keep pace with inframarginal, older cohorts of college graduates, who on average have higher average skill.

These findings complement prior research such as Katz and Murphy (1992) that explains the growing college wage gap via growth in relative labor demand, interpreting higher labor demand as growth in firm pay premia. I revisit the canonical supply and demand model described in Katz and Murphy (1992) and Acemoglu and Autor (2011), showing that an extension of my model nests that framework. I adjust estimates of the growth in relative labor demand for college workers using my estimated changes in skill composition. I find this adjustment increases estimates of growth in relative labor demand by 15-46% compared to standard approaches in the literature. In turn my estimates imply larger growth in potential causes of relative demand shifts such as skill-biased technical change (Acemoglu, 2002).

This study thus highlights new ways to interpret recent growth in inequality across education groups. Declining average skill compresses the college wage gap by drawing down average wage growth for college workers, partially offsetting growing labor demand for their services. Failing to account for skill compositional changes leads to understating the growth in labor demand for educated workers by a substantial amount.

Finally, my findings have implications for measuring aggregate labor demand and TFP growth. Recent work notes slowing growth in college attainment rates, warning that these trends could harm long-term growth (Jones, 2023). The results suggest that changes in skill composition mean that the growth in the college share may overstate growth in human capital, leading to underestimates of TFP growth.

Related Literature. My work relates to literature that studies the college wage gap through the lens of supply and demand for college graduates, such as Tinbergen (1974, 1975); Katz and Murphy (1992); Juhn et al. (1993); Autor et al. (1998); Krusell et al. (2000); Card and Lemieux (2001); Acemoglu (2002); Acemoglu and Autor (2011); Diamond (2016); Valletta (2018); Autor et al. (2020, 2023); and others.[5] My study contributes to this research by measuring changes in average skill by education and studying its implications. The most closely related papers within this group attempt to empirically measure how shifts in the supply of college workers affect skill composition (Juhn et al., 1993; Acemoglu, 2002; Carneiro and Lee, 2011; Bowlus and Robinson, 2012).[6] The primary approach in past work compares earnings growth in different cohorts to separate changes in the relative price of education from changes in relative skill.[7] For example, Bowlus and Robinson (2012) estimate changes in the relative price for different eduction groups using earnings growth for older workers, under the assumption that people over 50 have minimal marginal gains in skill with experience (Heckman et al., 1998). They then subtract these estimated prices from wages to measure skill.

While an advantage of these prior methods is that they can be used with public microdata such as the Current Population Survey, they must assume common changes in relative prices across cohorts.[8] In the case of Bowlus and Robinson (2012), changes in human capital

[5] Valletta (2018) and Autor et al. (2023) find evidence for deceleration or even reversal in the growth of the college wage gap in recent years. Autor et al. (2023) find that much of the wage compression comes from non-college workers sorting to higher paying firms after the Covid-19 pandemic, highlighting the importance of separately measuring skill from firm pay premia and sorting.

[6] See also Juhn et al. (2005); Carneiro and Lee (2009); Bowlus et al. (2023).

[7] Carneiro and Lee (2011) follow a different approach, studying the effect of college enrollment rates in a person's Census region of birth on their earnings, controlling for a rich set of fixed effects. The intuition is that higher college enrollment rates draw in marginal graduates, who may have different skill compared to inframarginal graduates. Hendricks and Schoellman (2014) and Ichino et al. (2024) measure skill changes of college graduatesusing standardized test scores.

[8] Using the CPS allows prior work to study compositional changes over longer time horizons. Skill changes may have differed in earlier decades. As one bit of evidence, Katz and Murphy (1992) find steeper growth in the college wage gap for workers just entering the labor force from 1979 to 1987. In contrast I find older cohorts have steeper growth in the wage gap, suggesting potentially different relationships between marginal and inframarginal college graduates in the two periods. These could arise from changes in selection into college or changes in the average treatment effect of college. Hendricks et al. (2021) show that family characteristics were better predictors of college attendance before World War II, but academic ability was afterwards. While having a college degree became more meritocratic over the middle of the 20th century,

prices for older workers must equal changes for younger workers within an education group. My model shows price changes, which I measure as growth in average firm pay premia acoss workers, may vary because of changes in the level of pay premia or differential sorting patterns across firms with different pay premia. Thus, for different cohorts to have the same price changes they must have the same sorting patterns across firms and over time.[9] I relax this assumption by directly separating the role of sorting in price changes using matched worker-firm earnings data.[10] Relative to prior work I also use the data to delve more into the mechanisms for relative price changes, studying the role of county and industry in particular.

My work also connects research on causes of the college wage gap with the literature that uses the AKM method to study pay differences across firms.[11] My primary point of departure from the AKM literature is aggregating the estimated parameters for college and non-college workers and tracking changes in the aggregates over time. The most closely related paper among this group is Card et al. (2013). They find that relative sorting to high wage establishments explains 70% of the rise in wages for University graduates compared to apprentices in Germany. Distinct from Card et al. (2013), by estimating a single TV-AKM model over the panel I measure changes in pay premia and skill over time within groups, not just differences-in-differences across groups, resolving a missing intercept problem. This contrast lets me determine that average skill declined for college workers both in a relative sense compared to non-college workers and also in an absolute sense. I then further decompose the effects of sorting patterns by county and industry and measure implications for estimating changes in relative labor demand.

Conceptually, I characterize a class of models that generate log additivity in wages,

the evidence suggests further expansion in the college share in recent decades reversed this effect.

[9]Different age groups differentially sort to urban, suburban, and rural areas, showing one potential threat to an assumption of common sorting patterns. See Parker et al. (2018).

[10]The data also let me avoid complications with top coded earnings. Bowlus and Robinson (2012) must assume that changes in prices for median workers equal changes for average workers. This assumption would not hold if growth in pay premia have long tails, as I show in my results.

[11]See Abowd et al. (1999); Card et al. (2013, 2016); Barth et al. (2016); Abowd et al. (2018); Sorkin (2018); Song et al. (2019); Haltiwanger et al. (2022); and others.

extending results in Yeh et al. (2022) and Chan et al. (2023a). This persepective connects the methodology and microdata underpinning the AKM model to the literature on supply and demand determinants of the college earnings gap.

In Section 2 I describe the conceptual framework, define the skill measures, and characterize the assumptions underlying the approach. In Section 3 I describe the data sources. In Section 4 I discuss the implementation of the conceptual approach in the data. In Section 5 I present the main results. In Section 6 I explain mechanisms for growth in the college wage gap. In Section 7 I discuss implications for growth in relative labor demand. Section 8 concludes.

2 Conceptual Framework: Defining Skill

2.1 Baseline Model

I first show a simpler version of the model before generalizing the results. Consider the behavior of a firm j seeking to hire workers for production at time t. Let $R_{jt}(L_{jt})$ be the firm's revenue function when it hires L_{jt} units of labor. R_{jt} evaluates any inputs other than labor at their optimal levels and subtracts their costs.

Firm j faces total labor cost $w_{jt} \times L_{jt}$ when it hires L_{jt}, where the unit cost of labor w_{jt} may vary across firms and over time. The firm solves its profit maximization problem,

$$\max_{L_{jt}} R_{jt}(L_{jt}) - w_{jt}L_{jt} \tag{2.1}$$

Assuming the revenue function is differentiable, the firm's optimality condition can be rearranged as

$$w_{jt} = \frac{\partial R_{jt}}{\partial L_{jt}} \tag{2.2}$$

I now assume that workers are additively subsitutable within a firm, an assumption I relax in Section 2.2.2. First, let h_{it} be the number of hours and z_{it} be the number of efficiency units per hour that a worker i supplies at time t. I refer to z_{it} as i's skill. Let $l_{it} \equiv h_{it} \times z_{it}$

be the total number of efficiency units supplied by i at t.

ASSUMPTION **1.** *The labor aggregate for a firm equals*

$$L_{jt} = \sum_{i \in \mathcal{N}_{jt}} l_{it} \tag{2.3}$$

where $\mathcal{N}_{jt}$ is the set of workers employed at firm j at time t.

Assume also that workers are roughly infinitesimal, so that for all workers and firms l_{it} is small relative to L_{jt}. Then applying the chain rule to Equation 2.2 and taking a Taylor approximation, Assumption 1 implies

$$w_{it} \approx \frac{\partial R_{jt}}{\partial L_{jt}} \times z_{it} \tag{2.4}$$

with w_{it} the hourly wage received by worker i. The left side of Equation 2.4 is the marginal cost from hiring i, while the right side is the marginal revenue. The firm equates these quantities on the margin. Marginal revenue is multiplicative in skill under the the assumption of additive labor inputs.

Ignoring approximation error and taking logs,

$$\log w_{it} = \xi_{jt} + \log z_{it} \tag{2.5}$$

with $\xi_{jt} \equiv \log\left(\frac{\partial R_{jt}}{\partial L_{jt}}\right)$. Equation 2.5 shows log additivity in wages in a firm-time effect and skill. Both of these may vary over time. To separately identify time variation in the two components, I impose a particular parametric assumption for skill motivated by the AKM model,

$$\log z_{it} = \zeta_i + \Gamma^T X_{it} + \epsilon_{it} \tag{2.6}$$

where X_{it} a vector of time-varying covariates. Combining Equations 2.5 and 2.6 gives

$$\log w_{it} = \xi_{jt} + \zeta_i + \Gamma^T X_{it} + \epsilon_{it} \tag{2.7}$$

This equation adapts a standard AKM model to allow for time-varying firm effects, and I refer to it as the TV-AKM model following Lachowska et al. (2023). It decomposes log wages

w_{it} between a time-varying firm effect, ξ_{jt}; a time-invariant person effect, ζ_i; time-varying covariates, X_{it}; and a residual term, ϵ_{it}. Following the literature I refer to the firm-time effects, ξ_{jt}, as pay premia.

The model shows that wages vary with z_{it}. However, they also vary with firm characteristics ξ_{jt}. In this simple case, ξ_{jt} reflects variation in hiring costs across firms and over time. These may result from compensating differentials, for example (Sorkin, 2018; Lamadon et al., 2022). Below I will discuss an array of other reasons why ξ_{jt} may vary, several of which are plausibly time-varying. What matters is that if I can estimate ξ_{jt}, I can recover a worker's log skill $\log z_{it}$ by subtracting ξ_{jt} from the log wage, $\log z_{it} = \log w_{it} - \xi_{jt}$.[12]

2.1.1 Defining Average Skill

Before discussing estimation of Equation 2.7, I define the average skill measures. These are the key quantities I compute in the empirical analysis. I define changes in *average total skill* for college and non-college workers relative to a base year as

$$\Delta \bar{z}_{C,t} \equiv \frac{1}{N_{C,t}} \sum_{i \in \mathcal{N}_{C,t}} \log(z_{it}) - \frac{1}{N_{C,0}} \sum_{i \in \mathcal{N}_{C,0}} \log(z_{i0}) \tag{2.8}$$

$$\Delta \bar{z}_{NC,t} \equiv \frac{1}{N_{NC,t}} \sum_{i \in \mathcal{N}_{NC,t}} \log(z_{it}) - \frac{1}{N_{NC,0}} \sum_{i \in \mathcal{N}_{NC,0}} \log(z_{i0}) \tag{2.9}$$

where $N_{C,t}$ is the number of college workers at time t, $N_{NC,t}$ is the number of non-college workers, $\mathcal{N}_{C,t}$ is the set of college workers, and $\mathcal{N}_{NC,t}$ is the set of non-college workers. The changes for all workers pooled together equals

$$\Delta \bar{z}_{t} \equiv \frac{1}{N_{t}} \sum_{i \in \mathcal{N}_{t}} \log(z_{it}) - \frac{1}{N_{0}} \sum_{i \in \mathcal{N}_{0}} \log(z_{i0}) \tag{2.10}$$

with N_t the total number workers at time t and $\mathcal{N}_t$ the set of all workers.

Changes in average total skill incorporate shifts in the mean covariates, X_{it}, and mean

[12] A threat to interpreting this quantity as skill is discrimination in wage setting based on factors such as gender or race. Discrimination means pay variation within a firm need not equate to skill variation. Card et al. (2016) shows that men and women face different firm effects in Portugal. I discuss this issue further in Section 5.3, showing that the empirical findings do not seem to be driven by discrimination.

residuals, ϵ_{it}, over time. In the model I interpret these quantities as time variation in worker skill. I also consider skill measures that remove shifts in the covariates and residuals, isolating only changes in the person effects, ζ_i.

$$\Delta\bar{\zeta}_{C,t} \equiv \frac{1}{N_{C,t}} \sum_{i \in \mathcal{N}_{C,t}} \zeta_i - \frac{1}{N_{C,0}} \sum_{i \in \mathcal{N}_{C,0}} \zeta_i \tag{2.11}$$

$$\Delta\bar{\zeta}_{NC,t} \equiv \frac{1}{N_{NC,t}} \sum_{i \in \mathcal{N}_{NC,t}} \zeta_i - \frac{1}{N_{NC,0}} \sum_{i \in \mathcal{N}_{NC,0}} \zeta_i \tag{2.12}$$

$$\Delta\bar{\zeta}_{t} \equiv \frac{1}{N_{t}} \sum_{i \in \mathcal{N}_{t}} \zeta_i - \frac{1}{N_{0}} \sum_{i \in \mathcal{N}_{0}} \zeta_i \tag{2.13}$$

I refer to these as *average person skill*. Because the person effects ζ_i do not vary over time, changes in average person skill arise only because of compositional changes across cohorts. Changes in average total skill additionally incorporate changes due to the age structure of the economy and trends in residuals. I plot the time series of the average skill measures in the empirical application.[13]

2.1.2 Estimating the TV-AKM Model

To compute the average skill measures I need to estimate the TV-AKM model in Equation 2.7. Suppose there are N total workers in the sample and J total firms. I estimate the parameters using movement between firms under standard AKM assumptions (see, for example, Abowd et al. (1999); Bonhomme et al. (2023)).

Assumption **2.**

$$\mathbb{E}[\epsilon_{it} | X_{11}, \ldots, X_{NT}, j(1,1), \ldots, j(N,T), \zeta_1, \ldots, \zeta_N, \xi_{11}, \ldots, \xi_{JT}] = 0 \tag{2.14}$$

Assumption 2 embeds the requirements that worker mobility between firms is not driven by idiosyncratic time-varying wage fluctuations and that expected log wages are additively

[13]Because my interest in this study is in evolution of the means over time, under Assumption 2 defined below my estimates of these series are unbiased. I thus do not require bias adjustments needed for estimating quadratic forms of estimated parameters as in Kline et al. (2020), though such adjustments would be needed for stastical inference.

separable. It does not restrict in any way the joint distribution of worker and firm effects, with workers allowed to sort to firms based on any function of their own ζ_i and the vector of firm effects ξ (Card et al., 2013; Kline, 2024).

To elucidate the requirements of Assumption 2, I show some of its implications. The following conditions follow from an extension of results in Kline (2024). Let $\log w_{it}(j_1, \ldots, j_T)$ denote potential log wages for worker i at time t who works at firm j_1 in period 1, j_2 in period 2, etc. In Appendix A.1 I prove that under Assumption 2 potential outcomes satisfy

1. Exclusion:

$$E[\log w_{i,t}(j_1, \ldots, j_T)] = E[\log w_{i,t}(j_t)] \tag{2.15}$$

2. Parallel trends:

$$E[\log w_{i,t+1}(j) - \log w_{i,t}(j) | j(i,t) = j, j(i,t+1) = k] =$$

$$\xi_{jt+1} - \xi_{jt} + \Gamma^T(X_{it+1} - X_{it}) \tag{2.16}$$

3. No selection on treatment effects:

$$E[\log w_{i,t+1}(k) - \log w_{i,t}(j) | j(i,t) = j, j(i,t+1) = k]$$

$$= E[\log w_{i,t+1}(k) - \log w_{i,t}(j)] \tag{2.17}$$

I relax the stationarity condition from Kline (2024) by allowing for time-varying firm effects. In Appendix A.1 I show under these conditions that worker mobility identifies the unconditional treatment effect on wages when moving between firm-time pairs.

Since Card et al. (2013), prior work tests whether the data accord with Equations 2.15 through 2.17 using event study plots of wage changes around job moves. Exclusion requires that wages for a given worker do not depend on past or future firm indicators.[14] If Exclusion

[14]This rules out models of on-the-job search with sequential bargaining, as in Postel–Vinay and Robin (2002), or heterogeneous persistent skill accumulation across employers, as in de la Roca and Puga (2017). In some sequential bargaining models such as Bagger and Lentz (2019), wages within a firm need not even be monotone in skill as more productive workers may have high option value that leads them to accept lower

holds, then wages at the destination firm should not depend on the origin firm.[15] Parallel trends is consistent with no systematic changes in wages before or after a move. Prior studies find evidence in concert with these predictions. (Card et al., 2013, 2018; Song et al., 2019).[16] No selection on treatment effects requires that workers more likely to move to a firm do not have systematically different wage changes from moving.[17] I test this by comparing wage changes for different types of workers, as I discuss in Section 4.3.

2.1.3 Relation to Prior Work

My approach primarily differs from prior literature in its focus on aggregating the estimated parameters and tracking trends in their magnitudes over time. To suit my application I estimate a single TV-AKM model over the entire panel. This serves two benefits. First, it circumvents missing intercept problems and lets me compare magnitudes of the parameters across time. The parameters for the model are identified only up to a normalization of the worker and firm-time effects within a connected set of the worker-firm-time graph (Abowd et al., 2002). One of the person or firm-time effects equals 0, and the other worker and firm-time effects are estimated relative to that normalization. By estimating a single model over the entire sample I introduce only a single normalization that falls out of Equations 2.8 through 2.13 since it gets subtracted at both time t and the base period. If instead I estimated separate AKM equations in distinct periods, as in Card et al. (2013) or Song et al. (2019), I would have a distinct normalization in the base period and time t, introducing a missing intercept preventing me from comparing magnitudes of the average skill measures over time.

Second, I allow for time variation in the firm effects, as in recent work by Lachowska

wages.

[15] Di Addario et al. (2023) show that origin effects explain a relatively inconsequential share of earnings in an augmented AKM model that includes fixed effects for the origin firm.

[16] Bonhomme et al. (2019) shows that the model provides a good approximation in Swedish data even in a model that allows for flexible complementarities, sorting, and dynamics.

[17] This rules out endogeneous mobility based on ϵ_{it}, as in Eeckhout and Kircher (2011) or Borovickova and Shimer (2024). It also rules out differential trends in residuals for different types of workers.

et al. (2023), Engbom et al. (2023), and Chan et al. (2023b). The TV-AKM model allows for heterogeneous growth rates in ξ_{jt} across different firms. If growth rates do indeed vary across firms, assuming time-invariant pay premia could bias parameter estimates and skill measures. In Section 5 I indeed find highly heterogeneous growth rates in pay premia across different counties and industries. Forcing equal growth rates in pay premia would load heterogeneous changes in the firm-time effects onto the person components of wages, biasing average skill measures. This issue interacts with estimating the model on a long panel. While in a short time period biases from assuming time-invariant firm effects may be small, over a long panel the bias may accumulate.[18]

If workers do have systematic differences in persistent skill accumulation, my model requires these are captured by the covariates. Estimating an AKM in separate short panels offers more flexibility in this regard because it allows an individual worker's person effect to vary in each estimation. I forego this advantage so that I can track changes in the magnitudes of the parameters over time.

2.2 General Case and Discussion

In Section 2.1 I showed that under Assumption 1 wages are log additive in a firm effect and skill when hiring costs exogenously vary across firms and over time. I now extend the simple case to allow for features such as market power, rent sharing, and adjustment costs while maintaining log additivity. I also clarify the interpretation of the skill measures.

Suppose now that the firm's profit maximization problem is

$$\max_{L_{jt}} \underbrace{R_{jt}(\chi_{jt}L_{jt})}_{\text{Revenue}} - \overbrace{w_{jt}(L_{jt})L_{jt}}^{\text{Wages}} - \overbrace{\Phi_{jt}(\chi_{jt}L_{jt}, \chi_{jt}L_{jt-1})}^{\text{Adjustment Costs}} \tag{2.18}$$

Equation 2.18 generalizes Equation 2.1 in several dimensions. First, the Revenue term includes a firm-specific multiplicative shifter χ_{jt}. This term could just absorbed within the

[18] Engbom et al. (2023) show that the difference in variance between firm effects and firm-time effects increases with the length of the panel.

revenue function, but I write it explicitly to draw out intuition for the approach below. It represents variation across firms and over time in how effectively they utilize some amount of labor L_{jt}. When χ_{jt} is high, all workers at the firm work more productively than they would elsewhere.

The Wages term is similar to before, except now the unit wage rate $w_{jt}(L_{jt})$ is a function of the amount of labor employed, allowing for market power. It may also incorporate a rent sharing term,

$$w_{jt}(L_{jt})L_{jt} = \underbrace{\nu_{jt}R(\chi_{jt}L_{jt})}_{\text{Rent Sharing}} + \overbrace{w^*_{jt}(L_{jt})L_{jt}}^{\text{Other Wages}} \tag{2.19}$$

Allowing for rent sharing captures the possibility that firms may share a fraction of their revenues with workers, a feature of some standard labor market models.

Finally, the Adjustment Costs term captures costs of changing the amount of labor employed over time. Here I assume for simplicity that the firm is myopic and does not consider the effects of current period choices on future adjustment costs. In Appendix A.2 I show that the results below generalize with dynamic adjusment costs.[19]

Assuming differentiability of revenue, the wage schedule, and adjustment costs, the firm's optimality condition can be rearranged as

$$w_{jt}(L_{jt}) = \chi_{jt}\left(\frac{\partial R_{jt}}{\partial L_{jt}} - \frac{\partial \Phi_{jt}(L_{jt}, L_{jt-1})}{\partial L_{jt}}\right) \times \left[e_{jt}^{-1} + 1\right]^{-1} \tag{2.20}$$

with $e_{jt}^{-1} \equiv \frac{w'_{jt}(L_{jt})L_{jt}}{w_{jt}(L_{jt})}$ the firm's inverse labor supply elasticity. Under Assumption 1, the wage for an individual worker again satisfies Equation 2.5,

$$\log w_{it} = \xi_{jt} + \log z_{it}$$

except now

$$\xi_{jt} = \log\left[\chi_{jt}\left(\frac{\partial R_{jt}}{\partial L_{jt}} - \frac{\partial \Phi_{jt}(L_{jt}, L_{jt-1})}{\partial L_{jt}}\right) \times \left[e_{jt}^{-1} + 1\right]^{-1}\right]$$

Applying the parametric form in Equation 2.6 gives the TV-AKM model. Consequently,

[19] Yeh et al. (2022) show that this setup nests a range of models in the literature, including wage-posting models, additive random utility models, and monopsonistic competition models.

under Assumption 2 I can compute unbiased estimates of Equation 2.7 and estimate the quantities in Equations 2.8 through 2.13. The methodology can thus accomodate market power, rent sharing, and adjustment costs. ξ_{jt} may also reflect compensating differentials, shocks to aggregate labor demand that make hiring more expensive, technological shifts, or any other factors that change hiring costs at the firm-time level.

2.2.1 Interpreting the Average Skill Measures

I now offer clarifications about what the skill measures identify and how to interpret them. First, note that χ_{jt} appears in the expression for ξ_{jt}. This term captures variation in how effectively a firm makes use of workers' skills. For example, a firm with a superior technology or work environment may get more out of its workers, holding z_{it} fixed. This sort of variation would get absorbed by ξ_{jt} and not get measured in $\log(z_{it})$.

The appropriate way to interpret variation in total average skill is thus in the following counterfactual sense. Consider a worker i with person effect ζ_i and covariates X_{it}. Consider exogeneously moving i to another firm and time period, j', t'. i's expected log hourly wage would then by

$$E[\log w_{it'}(j')] = \xi_{j't'} + \zeta_i + \Gamma^T X_{it}$$

In particular, i would receive the same multiplicative firm shifter $\xi_{j't'}$ as other workers at j', t'. When I compare workers between 2002 and 2019, I am thus supposing that a worker i from 2002 counterfactually would receive the same technological improvements enjoyed by a worker i' from 2019 if i instead were employed in the future. i would use these technologies just as effectively, in expectation.

While the approach allows for heterogeneous growth rates in ξ_{jt}, it does not allow for systematic heterogeneous growth rates in persistent skill. Suppose that workers employed at firm j' at time t' receive a persistent skill shock η. Expected log wages at some time $t > t'$ equal

$$\log(w_{it}) = \xi_{jt} + \zeta_i + \Gamma^T X_{it} + \eta \mathbf{1}\{j(i, t') = j'\} + \epsilon_{it} \tag{2.21}$$

This represents a violation of Assumption 2 since $\eta \mathbf{1}\{j(i,t') = j'\}$ gets absorbed in the error term and depends on past firm indicators.

2.2.2 Additivity and Skill Comparisions

I show how the result generalizes under a relaxation of the labor additivity assumed in Assumption 1. Suppose that the firm's decision problem depends on quantities of different types of labor,

$$\max_{L_{j1t},\ldots,L_{jKt}} R_{jt}(L_{j1t},\ldots,L_{jKt}) - W_{jt}(L_{j1t},\ldots,L_{jKt}) \tag{2.22}$$

with L_{jkt} the quantity of type k labor employed at the firm and $W_{jt}(L_{j1t},\ldots,L_{jKt})$ the firm's total wage bill. I abstract from adjustment costs for simplicity. This set up allows for complementarity or substitutability across types in the revenue function, with additive substitutability within types. It nests models such as CES substitution between worker types. Then log additivity holds *within* types,

$$\log(w_{it}) = \xi_{jt}^{k(i)} + \log(z_{it}) \tag{2.23}$$

with $k(i)$ worker i's type and $\xi_{jt}^{k(i)}$ a type-specific firm-time shifter.

This result suggests estimating separate type-specific firm-time shifters in the TV-AKM model. Estimating separate models allows for differences-in-differences estimates of relative skill changes across groups, as in Card et al. (2013). In addition it allows for skill comparisons within a group over time because I estimate each model on a single long panel. However, it precludes estimates of skill differences across groups if workers have immutable types. In the case of skill comparisons by education, with separate AKM models I can determine growth in skill for college and for non-college workers and differences between these growth rates, but I cannot identify differences in the level of skill between the two groups. In turn this means I cannot compute changes in average skill overall across all workers.[20]

[20]Due to limitations with the education data in the LEHD I do not use Equation 2.23 as the baseline estimating equation. The LEHD data I use impute education information for 85% of workers, so for the main results I instead use education data from the ACS. I discuss this issue more in Section 4.2. Nevertheless,

One intuition for this is that allowing for type-specific firm-time effects equates to estimating the model on a worker-type-firm graph instead of worker-firm graph. The model can estimate skill differences only within connected components of the graph. Workers have no mobility between immutable types, preventing skill comparisons across types.

Why does additivity matter for skill comparisons? To measure average skill I require estimates of cardinal skill differences between worker-time observations. With additivity of labor in the wage function I can measure skill differences within a firm-time via differences in wages. I can then compare skills across firms and over time by studying wage changes when workers switch employers. In the more general form from Equation 2.22 wage differences within a firm confound skill differences with variation in type-specific shifters, so I cannot compute cardinal skill differences for workers with different types, though I can still compare skills within types.

3 Data

3.1 LEHD Data

I use the U.S. Census Bureau's Longitudinal Employer-Household Dynamics (LEHD) data, which is a quarterly matched employer-employee data set constructed from state unemployment insurance records.[21] Every quarter, employers subject to state UI laws—approximately 95% of all private sector employers plus state and local governments—are required to submit total earnings information for their employees. I observe data in 30 states for each year from 2002 to 2019.[22]

An employer in the data set is a state-level unemployment insurance account (SEIN).

Appendix C.8 shows the main results are robust to estimating Equation 2.23 separately by imputed education type from the LEHD.

[21] See Abowd et al. (2009) for details.

[22] Alabama, Arizona, California, Colorado, Connecticut, Delaware, Indiana, Iowa, Kansas, Maine, Maryland, Massachusetts, Minnesota, Montana, Nebraska, Nevada, New Jersey, New Mexico, North Dakota, Ohio, Oklahoma, Pennsylvania, South Carolina, South Dakota, Tennessee, Utah, Virginia, Washington, Wisconsin, Wyoming.

For single establishment firms, this is equivalent to a firm. For firms operating in multiple states or that have multiple UI accounts in a state, an SEIN is smaller than a firm.[23] I match SEINs to their national firm identifier using the Employer Characteristics File (T26), which maps each SEIN to its parent firm in the US. For the purposes of my analysis, I refer to an "firm" as a national parent firm × state and treat all workers at the same national parent firm × state as working for the same firm. This definition of a firm allows parent firms to have different pay premia in different states in the estimation.

The LEHD Individual Characteristics File (ICF) additionally provides limited demographic data on workers, such as their age and sex. It also has information on the worker's education level, though only for about 15% of workers. The remaining education data in the ICF file are imputed by the Census based on information such as gender, age, race/ethnicity, earnings, industry, and characteristics of co-workers and co-residents.

I do not observe hours data in the LEHD. In Section 4 I discuss how I address hours variation in total earnings.

3.2 American Community Survey (ACS)

I also use restricted-access microdata from the American Community Survey from 2002 to 2019. The ACS has detailed information on education, gender, age, race, county, industry, and occupation for a random sample of people in the US. I use 3-digit Census Industry and Occupation codes.

The ACS also has information on the number of hours worked. In each year the ACS reports the typical number of hours worked per week in weeks in which the person worked. Between 2002 and 2007 and in 2019, the ACS also reports the number of weeks worked in the past year. Between 2008 and 2018, the ACS reports the number of weeks worked in bins.[24] For consistency, I convert the number of weeks worked to these bins in all years. I

[23] Sorkin (2018) notes that few firms have multiple SEINs within a state.

[24] These bins are 13 weeks or less, 14 to 26 weeks, 27 to 39 weeks, 40 to 47 weeks, 48 to 49 weeks, and 50 to 52 weeks.

then impute the number of weeks worked to the midpoint of the bin and multiply by the typical number of hours worked per week to impute the total number of hours worked in the past year.

3.3 Current Population Survey

As a benchmark I compare my wage data with Current Population Survey Merged Outgoing Rotation Group Earnings (CPS MORG) data downloaded from the NBER. I clean the data following procedures in Valletta (2018). I provide more details in Appendix C.2.

In Section 7 I use data from the CPS Annual Social and Economic Supplement (CPS ASEC) to measure changes in labor demand, following Autor et al. (2008) and Acemoglu and Autor (2011). I discuss the data further in Appendix D.

4 Implementation

I now discuss how I map the data to the conceptual framework discussed in Section 2. I first estimate the TV-AKM model using the LEHD to measure parameters of the model. I then join the estimates to the ACS to study changes in skill for different education and demographic groups.

4.1 Estimating the TV-AKM Model

4.1.1 Baseline LEHD Sample

I first estimate the TV-AKM model using data from the LEHD. I make several restrictions to the data for the baseline estimation. First, I keep workers between the ages of 20 and 60, following Song et al. (2019). Because I do not observe hours data in the LEHD, I then make several restrictions that aim to keep only full time workers. I keep only worker-years in which the worker worked at only one firm for the whole year. This restriction is distinct from the approach taken in other work such as Sorkin (2018) and Song et al. (2019) that links workers

to the firm providing their largest source of earnings in a year. Keeping workers who work at only one firm in a year associates earnings for each worker with only one firm-time effect, which maps more closely to my model.[25]

Next, I aggregate the data from the worker-quarter to the worker-year level. I then restrict to worker-year observations in which the worker earned at least the federal minimum wage for a full year, 52 weeks times 40 hours a week.[26] I then adjust wages using the CPI.

I then split the sample into two, with each sample containing a random 50% sample of workers. I call the first random sample Sample A and the second sample Sample B. I split the sample in this way to test for errors in non-linear forms of model parameters, as in the leave-one-out approach from Kline et al. (2020). Estimating the model in two different samples is computationally much simpler than leave-one-out approaches.[27]

Next, I take Sample A and compute the largest connected component of the resulting worker-firm graph. I estimate the TV-AKM model on this subset of the data. Before discussing alternative ways of constructing the sample, I explain precisely how I implement the TV-AKM model,

$$\log(w_{it}) = \xi_{jt} + \zeta_i + \Gamma^T X_{it} + \epsilon_{it}$$

4.1.2 Controlling for Age

As in Song et al. (2019), I control for a polynomial in age in X. At baseline, I set $X_{it} = ((\text{age}_{it} - 50)^2, (\text{age}_{it} - 50)^3)$. I omit a linear term to avoid the classic identification problem that age can be computed as calendar year minus birth year, which get absorbed by ξ_{jt} and ζ_i, respectively (Card et al., 2018).[28] This form for X restricts the age profile to be flat at

[25] If I observed hours worked for each firm, I could control for hours in the TV-AKM estimation and include separate observations for each firm a worker worked at in a year.

[26] This is a stricter earnings restriction than in prior work such as Song et al. (2019) because I aim to keep only full-time workers. Song et al. (2019) show that their estimated AKM parameters seem largely insensitive to the earnings threshold.

[27] Sorkin (2018), Drenik et al. (2023), and Card et al. (2024) perform similar exercises.

[28] In simulation exercises I verify that this functional form cannot identify the effect of age if it has a completely linear effect on the outcome. Identification requires that the effect of age has some non-linear component, with a better fit when the model has more curvature. In practice I find a non-linear age-earnings profile, so the functional form provides a reasonable fit to the data.

50, in accordance with the age-earnings profile documented in prior work (Heckman et al., 1998; Bowlus and Robinson, 2012; Lagakos et al., 2018). In Appendix C.4.1 plot expected wages by age, showing that they peak between 50 and 55.[29]

4.1.3 Controlling for Hours

I do not observe hours worked in the LEHD. As discussed, I make several restrictions on the data that aim to keep only full-time workers. In addition, I control for mean log hours by demographic group in the ACS. Since the ACS surveys a random set of people each year, I cannot control for individual log hours, but I can control for group means. In each year I compute the mean hours worked for each age bin $\times$ sex $\times$ education group $\times$ Census region, restricting to people who earn at least the federal minimum wage times 40 hours a week times 52 weeks. I create 5-year age bins and define education group as dropped out of high school, high school graduate, some college, and college graduate using the imputed education data in the LEHD. I then merge the mean hours data to the LEHD sample and control for mean log hours in the TV-AKM model.[30]

4.1.4 Alternative Ways of Constructing the Sample

In Appendix C.6 I discuss robustness to several other ways of constructing the AKM sample and estimating the model. These include keeping all workers who meet the sample criteria instead of taking a random sample; estimating the model on a "sandwich" sample that drops workers if they do not work at the same firm in Q4 of the prior year or Q1 of the following year; estimating a separate age-earnings profile by sex and education; estimating a fifth order polynomial in age instead of a third order polynomial; and constructing alternative controls for log hours.

[29] Card et al. (2013) restrict the age profile to be flat at age 40, an approach adopted by several other papers in the literature. Empirically, earnings continue to increase in the US until between the ages of 50 and 55, suggesting 40 may be too early.

[30] In Section C.6 I show robustness to controlling for mean log hours in alternative groupings that exclude imputed education.

4.2 Computing the Average Skill Measures

I now explain how I compute the skill average measures using the estimated parameters from the TV-AKM model. A straightforward approach would be to use the imputed education data in the LEHD to measure average skill by college status. However, imputation error for the education variables may propogate to the average skill measures. I instead use the education data from the ACS by merging it to the LEHD sample by person and year. I thus compute the average skill measures for workers I observe in both the LEHD and the ACS, weighting by ACS sample weights.

I make several additional restrictions to the merged LEHD-ACS data before computing the average skill measures. I subset to people between the ages of 25 and 60 who report positive hours worked in the ACS.I also make geographic restrictions. From 2002 to 2005, the ACS surveyed individuals from a subset of counties in the US. It sampled the same subset in each year. In subsequent years the survey expanded to full national coverege and added additional counties to the sample. I restrict to observations from the original set of counties that get sampled over the entire period from 2002 to 2019. This restriction helps me study how geography explains the growth of pay premia and skill over time in Section 6.

For certain robustness exercises, I also require two distinct estimates of the firm-time effects ξ. As discussed in Section 4.1.1 I construct two random samples of workers in the LEHD, Sample A and Sample B. I estimate the TV-AKM model again on the largest component of Sample B and join the estimated firm-time effects to the merged ACS-LEHD data. I keep only observations with two distinct estimates of the firm-time effects. This drops firm-years that fail to appear in the largest connected component of both Sample A and Sample B, in practice only small firms that account for a small fraction of employment.

I use this final sample to compute the average skill measures in Equations 2.8 through 2.13.

4.3 Validating the Model

Before presenting the empirical results, I first discuss tests for whether the model provides a reasonable fit to the data. If the model is correctly specified, then on average when workers move from one firm to another they should experience earnings changes corresponding to the difference in firm-time effects. Prior studies find empirical support for this using event study analyses (Card et al., 2013, 2018; Song et al., 2019; Kline, 2024).

Similar to past research, I create event study plots to evaluate the predictions from Equations 2.15 through 2.17. I consider workers who work at one firm from 2012 to 2013, then switch firms in 2014, and then stay at the subsequent firm from 2015 to 2017. For the pre-move period, I divide firms based on quartiles of their 2013 firm-time effects, and for the post-move period I use quartiles of 2015 effects. I then consider changes in firm-time effects and residuals for workers who move between pay premium quartiles.

Results in Appendix C.1 show large changes in firm-time effects relative to changes in mean residuals following moves between firms. This implies earnings at the destination firm do not seem to depend on the origin firm, providing empirical support for Exclusion in Equation 2.15. I find relatively little sign of systematic changes in residuals before or after the move, consistent with parallel trends (Equation 2.16). I repeat the exercise for college and non-college workers separately, finding similar results. Finally, I show event study plots for workers employed in protective services, foor preparation and serving, or building and grounds cleaning and maintenance.[31] These professions likely have less firm-specific indiosyncratic match value. Changes in firm-time effects and residuals show similar patterns for these workers, providing support for no selection on treatment effects (Equation 2.17). The findings complement a large and growing body of work that shows AKM models provide an imperfect but reasonable approximation to the data (Card et al., 2018; Kline, 2024).

[31] Census occuption codes 3700-3950, 4000-4150, and 4200-4250. I code workers as working in these occupations if they ever report working in them in some ACS survey year.

As an additional exercise, in Appendix B I conduct extensive simulations across five Midwestern states in the LEHD to test for potential biases in the model.[32] I simulate counterfactual outcomes by permuting the parameter estimates I get from actually estimating the TV-AKM model on the real earnings. I then run the TV-AKM model on the simulated data. The goal of the simulation exercise is to test if I recover the parameters under various threats to the model while preserving the mobility structure and characteristics of the data as closely as possible.

In one exercise I simulate faster skill growth for college workers compared to non-college workers. I find that this change can lead to biases in the estimated firm-time effects, with corresponding trends in average residuals over time. In practice I find modest trends in residuals, implying such forces do not drive the results. Second, I test whether more flexible age-earnings profiles or different profiles across demographic groups bias the model. I find these factors may introduce modest biases in parameter estimates but these biases may be mitigated by allowing for a more flexible age-earnings profile in estimation. In Appendix C.6 I replicate the main results in Section 5 allowing for a separate fifth order polynomial by education and sex in the TV-AKM model, suggesting that the way I model returns to experience likely does not drive my findings. Overall the model seems to be a useful tool for decomposing trends in wages and skill.

5 Explaining Changes in Wages

Having described the conceptual framework and data, I now present descriptive statistics about earnings and wages for college and non-college workers. In Figure 1, I plot mean log earnings and hourly wages for college and non-college workers in the matched LEHD-ACS data, relative to non-college workers in 2002. I also plot the series for all workers pooled together. I find about a 50-60 log point difference in earnings between college and non-college workers across years. Log hourly wages differ by about 45-55 log points. In Appendix C.2 I

[32] Indiana, Iowa, Ohio, Kansas, and Wisconson.

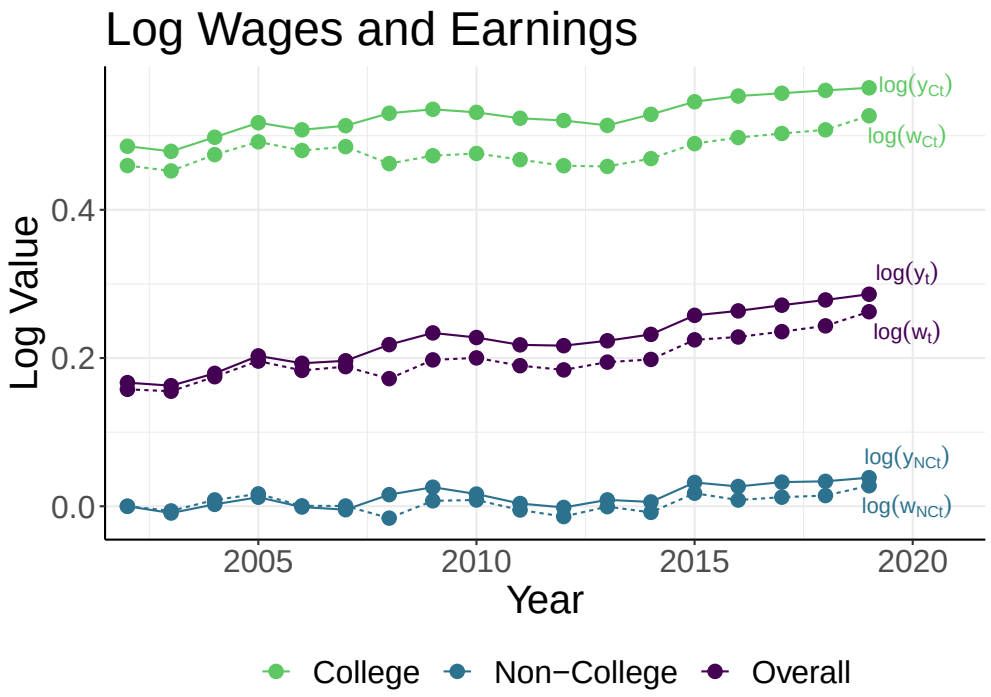


Figure 1: Mean log real earnings and hourly wages for college and non-college workers in the LEHD-ACS data, relative to non-college workers in 2002. The Overall lines shows values for both groups pooled together. y denotes earnings, and w denotes hourly wages.

show similar differences in log hourly wages in the CPS MORG data.

To highlight the growth over time, in Figure 2 I plot each log hourly wage series but normalize them to start at 0 in 2002. The results show that between 2002 and 2019 log hourly wages for college workers increased by 6.7 log points compared to 2.8 log points for non-college workers, implying about a 4 log point increase in the college wage gap. Wage growth for all workers pooled together equaled about 10.5 log points. Changes in the share of college graduates over time explain higher overall growth in wages.

The conceptual framework in Section 2 shows that changes in hourly wages can reflect both skill changes and shifts in firm pay premia, with $\log(w_{it}) = \xi_{jt} + \log(z_{it})$. Figure 3 shows changes in average estimated firm pay premia ξ over time for each group. Average firm pay premia increased by 22 log points for college workers compared to only 7 log points for non-college workers. The growth in firm pay premia far exceeds growth in overall earnings for college workers in particular.

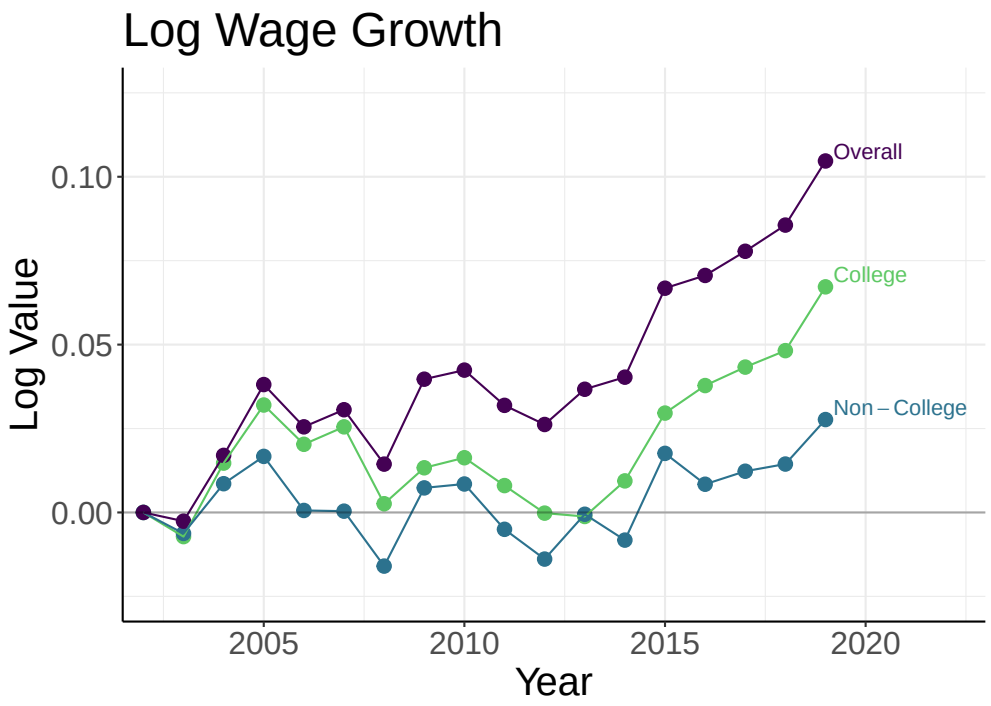


Figure 2: Growth in mean log hourly wages by education in the LEHD-ACS data, starting from 2002.

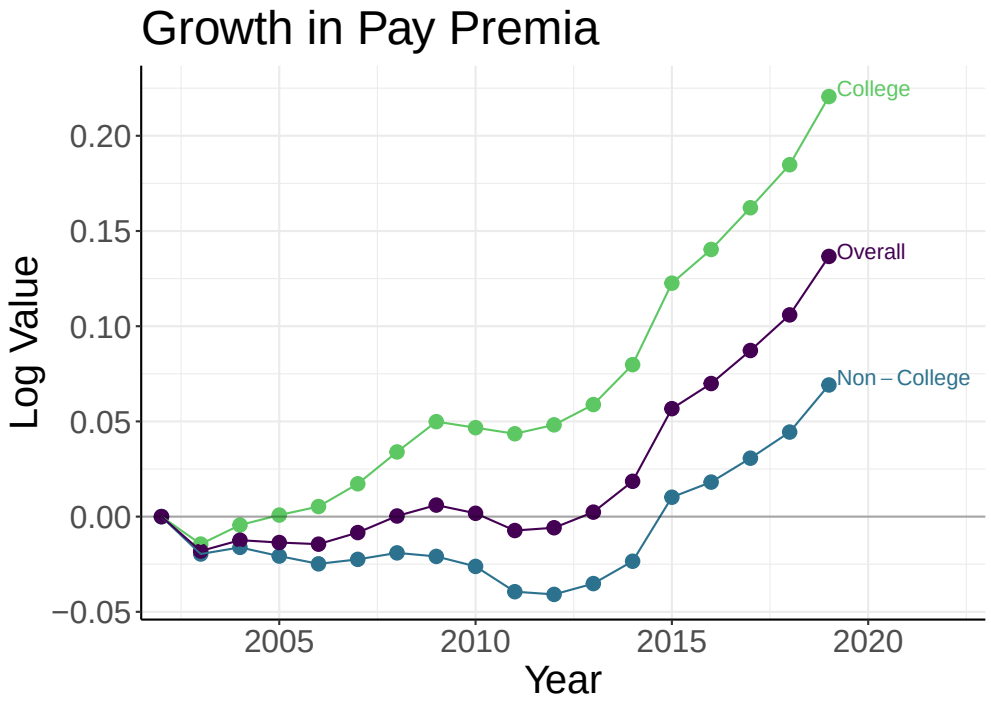


Figure 3: Growth in mean estimated pay premia for the LEHD-ACS data by education. Pay premia equal the firm-time fixed effects from the TV-AKM estimation. I average them across workers in a given year.

Figure 4 shows the main result of the paper, plotting changes in the average skill measures defined by Equations 2.8 through 2.13. The first panel shows the average total skill measures in Equation 2.8 through 2.10, which subtract firm pay premia from log hourly wages. The second panel shows the average person skill measures in Equations 2.11 and 2.13 which average the person fixed effects, removing the effect of age and residuals. Both series suggest steeply declining average skill for college workers, with a 15 log point decline in average total skill and 17 log point decline in average person skill. I find smaller changes for non-college workers and overall, with higher values in the middle of the sample and lower values at the beginning and end.[33] Despite large declines in average skill for college workers, widening differences in pay premia lead to an increase in the college wage gap across years.

5.1 Implications for Inequality

Changes in skill composition have important implications for growth in inequality. The model implies the following simple decomposition for growth in log wages for group k,

$$\overbrace{\left(\overline{\log(w_{k,t})} - \overline{\log(w_{k,0})}\right)}^{\Delta \log(w_{k,t})} = \overbrace{\left(\overline{\log(z_{k,t})} - \overline{\log(z_{k,0})}\right)}^{\Delta \log(z_{k,t})} + \overbrace{\left(\overline{\xi_{k,t}} - \overline{\xi_{k,0}}\right)}^{\Delta \xi_{k,t}} \tag{5.1}$$

with bars denoting average values across workers in a given year. The change in average log wages over time equals the change in average total skill plus the change in average pay premia. Taking differences across college and non-college workers gives a differences-in-differences decomposition of the growth in the college wage gap,

$$\begin{aligned} &\underbrace{[\Delta \log(w_{C,t}) - \Delta \log(w_{NC,t})]}_{\text{Growth in College Wage Gap}} \\ &= \underbrace{[\Delta \log(z_{C,t}) - \Delta \log(z_{NC,t})]}_{\text{Difference in Average Skill Growth}} + \underbrace{[\Delta \xi_{C,t} - \Delta \xi_{NC,t}]}_{\text{Difference in Pay Premium Growth}} \end{aligned} \tag{5.2}$$

[33] In Appendix C.2 I show that the average age of the workforce in the CPS MORG data increased slightly from 41.5 to 42 from 2002 to 2019. However, the share of workers under the age of 30 also increased over the sample period from 16% to 17.5%. Overall I find a small role for the age controls in the changes in average total skill. Differences between average person skill and average total skill primarily come from modest trends in the residual, plotted in Figure C7.

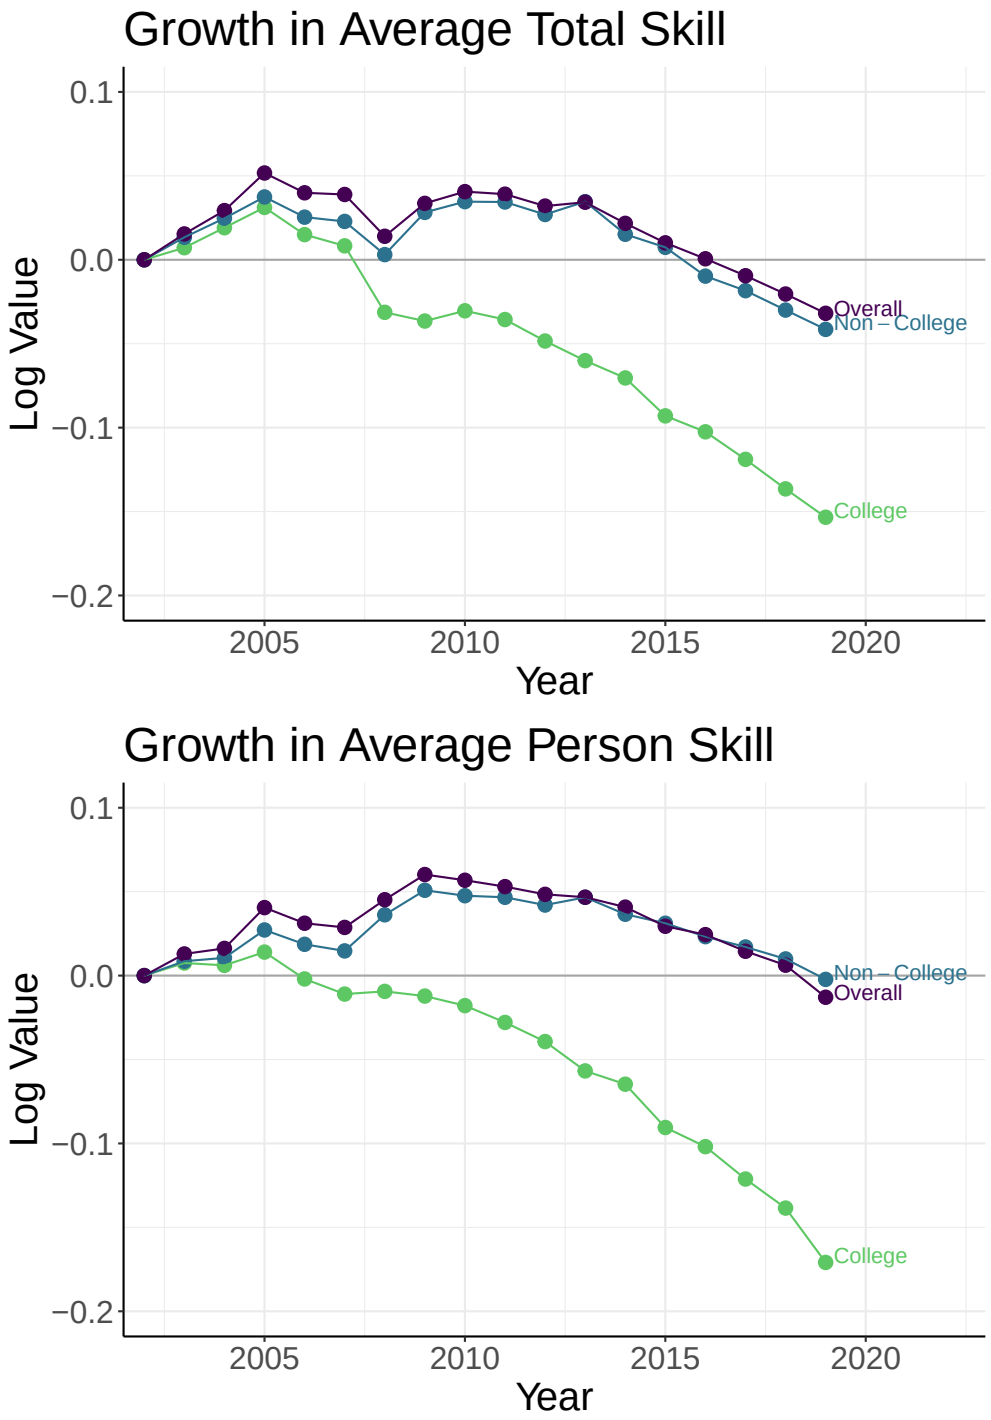


Figure 4: Growth in average total skill and average person skill by education. Total skill subtracts firm-time fixed effects from log hourly wages. Person skill equals the estimated person fixed effects from the TV-AKM equation.

Holding fixed average skill at baseline levels, the model suggests growth in the college wage gap would equal the difference in growth in the average pay premium. Results in Figure 3 imply this difference equals 15 log points (16%) compared to a 4 log point (4%) change in the college wage gap in the raw data, so changes in skill composition reduced wage inequality between college and non-college workers. As a point of comparison, Valletta (2018) finds that the wage gap between college and high school graduates increased by about 15 log points between 1980 and 1990, the period with the most rapid growth in the wage gap in recent decades (Katz and Murphy, 1992; Goldin and Katz, 2008; Acemoglu and Autor, 2011). Recent work notes an unexpected flattening in the college wage gap over the early 21st century (Beaudry et al., 2014, 2016; Valletta, 2018; Autor, 2018; Autor et al., 2023). My results suggest that changes in average skill for college workers could be a key driver of these trends.

5.2 Distributional Changes in Skill

I find that average skill for college graduates declined by 15-17 log points, while for non-college workers it remained mostly flat. I now explore shifts in the population that underlie these trends.

In the top panel of Figure 5, I show growth in the number of college workers from 2002 to 2019 by decile of total skill in 2002. The x-axis show deciles of estimated $\widehat{\log(z_{ij2002})}$ for college workers in 2002. The y-axis shows the population of these workers in 2002 and 2019.[34] The results show that the number of college workers grew over 95% below the median of the 2002 college skill distribution, compared to 21% above the median. The college population grew faster at the lower end of its skill distribution. For comparison the total population of the US grew by 14.6% over the same period.

Unlike means of the skill measures, quantiles of skill suffer from bias due to estimation

[34]To measure population I first scale the ACS sampling weights so that the sum of the weights across all people in a given year matches the total US population measured by the Census Bureau. I then sum the rescaled weights.

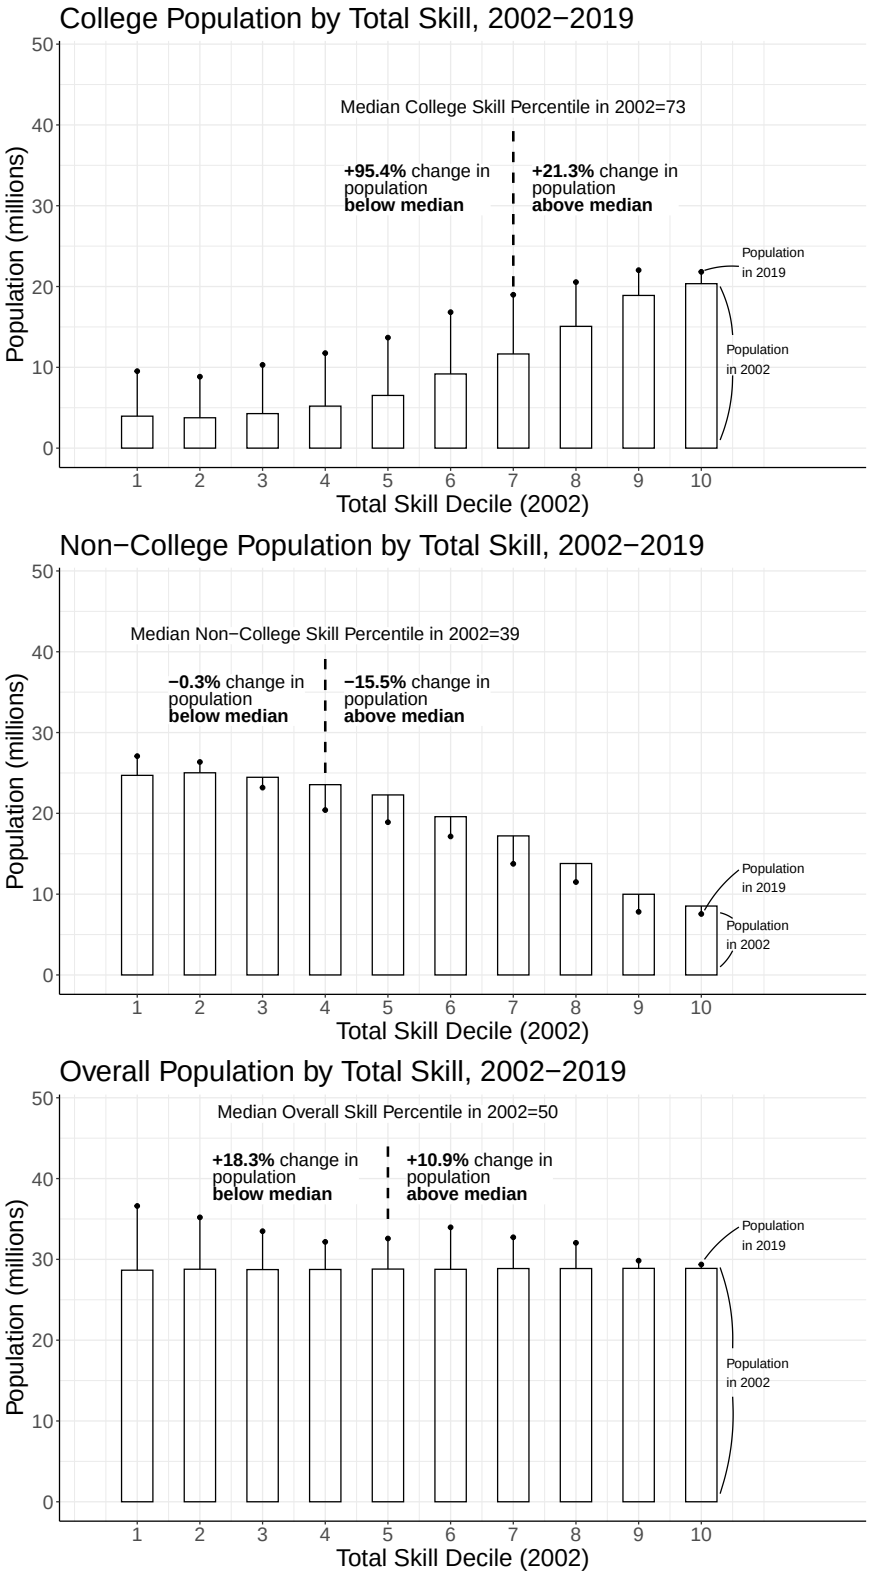


Figure 5: Changes in population by decile of total skill in 2002 for college workers, non-college workers, and all workers. Total skill equals log wages minus firm-time fixed effects.

error, similar to the issues with quadratic forms discussed in Kline et al. (2020). This merits caution in interpreting shifts in deciles of the skill measures if the distribution of estimation error changes over time. In Appendix C.12 I plot changes in the distribution of estimation error for the firm-time effects using differences between my two independent estimates from Sample A and Sample B. I find little shift in the distribution of estimation error between 2002 and 2019, suggesting changes do not drive the results in Figure 5.

The middle panel of Figure 5 shows the equivalent figure for non-college workers. The population of non-college workers declined, in particular at higher deciles of the skill distribution. This occured even as the total population increased by close to 15%. Together the top and middle panels coincide with a story that increases in college attainment primarily come at the bottom of the college skill distribution. The bottom panel shows growing population across each decile overall, with smaller changes at the very top of the distribution.

These results indicate that the growing share of college graduates has primarily come at the bottom of the college skill distribution. I find much smaller population growth at the top of the college skill distribution. The increase in college graduates at the bottom of the skill distribution coincides with decreases in the non-college population. These patterns are consistent with marginal college graduates being less skilled than the average, weighing down growth in average college wages.[35]

5.3 Robustness

I consider several robustness exercises for the main results. In Appendix C.4 I consider including part time workers and people who work at multiple firms in a year. While I still estimate the TV-AKM model on the baseline sample as in Section 4.1.1, for the average skill mesaures I include any worker in the ACS who works at a firm within the connected set.

[35]Section C.5 shows population growth by decile of the person fixed effects, ζ_i. These results warrant particular caution because of potential changes in estimation error for the person effect, a possibility I cannot test for using my split sample approach since each worker only appears in one sample. Nonetheless I find much larger growth in the college population at the bottom of the ζ distribution, along with declines in the non-college population.

This allows me to avoid restricting the skill measures to workers satisfying the minimum earnings threshold and who only work at one firm in a given year. I find larger earnings gaps between college and non-college earnings within this sample because non-college workers are more likely to fall below the earnings threshold. Otherwise, I find similar changes in pay premia and average skill as in the main results.

In Appendix C.6 I consider several alternative strategies for estimation. I compute the pay premia and average skill measures when averaging the estimated pay premia from Samples A and B; exclude imputed education when constructing hours controls for the TV-AKM estimation; estimate a richer age-earnings profile, allowing for a fifth order polynomial in age; and allow for a fifth order polynomial in age separately by sex and education. Across all specifications I find quantitatively very similar results.

In Appendix C.7 I show results in the sample of 5 Midwestern states. I find similar results when estimating the model on all workers rather than a random sample of them. I also show similar findings in a sandwich sample that excludes people who work at a different firm in Q4 of the previous year or Q1 of the following year.

In Appendix C.8 I show results within the Midwest Sample when estimating separate TV-AKM models for college and non-college workers, in line with the discussion in Section 2.2.2. I find quantitatively quite similar results. This suggests that the baseline model that pools all workers in the TV-AKM estimation provides a good approximation. It also suggests divergent trends in firm pay premia *withn a firm* between college and non-college workers do not drive the results.

In Appendix C.9 I show changes in average skill separately for non-college workers, workers with a college degree only, and workers with post-college schooling. I find that both workers with a college degree and workers with post-college experience have declining average skill over time, with post-college workers experiencing the steepest drops in skill.

In Appendix C.10 I show results separately for men and women and for White and Black workers. Patterns are broadly similar across groups. These findings give evidence that

discrimination does not drive the results.

In Appendix C.11 I show steeper declines in skill for college workers below the age of 40, with more modest declines for workers above the age of 40. These results suggest a larger role for compositional changes in recent cohorts. In addition, I find a 12 log point increase in wages for older college workers over the sample compared to a 3 log point increase for young college workers. The gap in wage growth between older and newer college cohorts provides decriptive evidence in the raw data of declining average skill in recent cohorts of college workers.

6 What Explains Growth in the Wage Gap?

The results show that changes in average skill alone would have reduced the college wage gap by 11 log points. How did it then continue to rise by 4 log points? My framework explains this via differential growth in firm pay premia. I now explore what factors seem to drive growth in pay premia, finding a central role for location and industry.

Figure 6 shows average pay premia over time by location and industry. In the top panel I show average pay premia for counties above and below the median college share in 2000.[36] Gaps in pay premia grew from about 3 log points in 2002 to 11 log points in 2019. I find even greater differences by industry. Differences in average pay premia grew from 11 log points to 28 log points between high tech and other industries, with a 29 log point increase in pay premia for high tech industries between 2002 and 2019.[37] These results complement findings from Abowd et al. (2018) and Haltiwanger et al. (2022) that a small number of industries

[36] College shares come from the 2000 Decennial Census.

[37] I use the definition of high-tech industries from the Census Bureau's BDS-High Tech Data, coded to match Census industry codes. High tech Census industry codes are: 336, Computer and Peripheral Equipment Manufacturing; 337, Communications Equipment Manufacturing; 338, Navigational, Measuring, Electromedical, and Control Instruments Manufacturing; 339, Semiconductor and Other Electronic Component Manufacturing; 359, Aerospace Product and Parts Manufacturing; 649, Software Publishers; 678, Other Information Services; 679, Data Processing, Hosting, and Related Services; 729, Architectural, Engineering, and Related Services; 738, Computer Systems Design and Related Services; 746, Scientific Research and Development Services.

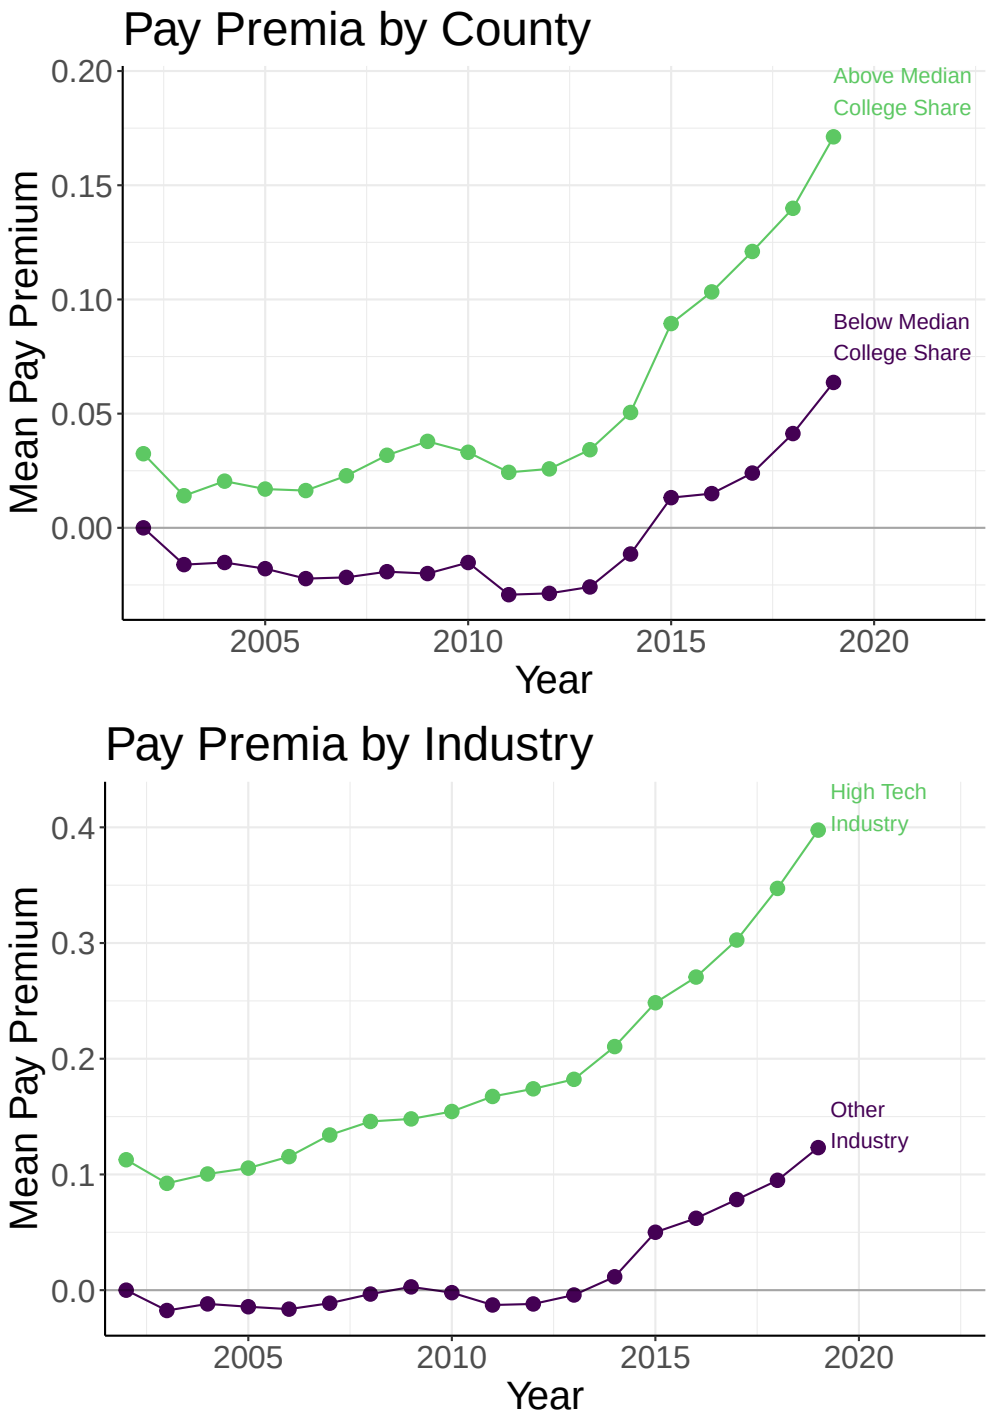


Figure 6: Average pay premia for different sectors of the economy. Pay premia equal the firm-time fixed effects in the TV-AKM model. The top panel shows changes in average pay premia in counties above and below the median college share in 2000. I normalize the series to 0 for below median counties in 2002. The bottom panel shows changes for high-tech industries compared to other industries. I normalize the series to 0 for other industries in 2002.

account for a large share of the growth in earnings inequality.[38] Relative to prior work I show that a substantial portion of divergence in earnings across industries comes from rising pay premia, which I can document because of the time variation in the firm effects.

What do these trends imply about the role of geography and industry in wage growth by education? I show the contribution of mean county and industry effects on wage growth. Let

$$\bar{\xi}_{lt} = \sum_{\mathcal{N}_{lt}} \xi_{jt}/N_{lt}$$

be the average firm-time effect in county l at time t and

$$\bar{\xi}_{It} = \sum_{\mathcal{N}_{It}} \xi_{jt}/N_{It}$$

be the mean firm-time effect in industry I, with $\mathcal{N}_{lt}$ the set of workers in county l at time t and $\mathcal{N}_{It}$ the set of workers in industry I at time t. N_{lt} and N_{It} denote the cardinality of these sets. Then for education group k define

$$\Delta \left(\log(w) - \xi\right)_{kt}^{COU} \equiv \frac{\sum_{i \in \mathcal{N}_{k,t}} \left(\log(w_{it}) - \bar{\xi}_{l(i,t)t}\right)}{N_{k,t}} - \frac{\sum_{i \in \mathcal{N}_{k,0}} \left(\log(w_{ij0}) - \bar{\xi}_{l(i,0)0}\right)}{N_{k,0}} \tag{6.1}$$

and

$$\Delta \left(\log(w) - \xi\right)_{kt}^{IND} \equiv \frac{\sum_{i \in \mathcal{N}_{k,t}} \left(\log(w_{it}) - \bar{\xi}_{I(i,t)t}\right)}{N_{k,t}} - \frac{\sum_{i \in \mathcal{N}_{k,0}} \left(\log(w_{ij0}) - \bar{\xi}_{I(i,0)0}\right)}{N_{k,0}} \tag{6.2}$$

These measures decompose how much of the growth in wages is explained by sorting across counties and industries with different average growth in pay premia. I compare them to the series for $\Delta \log(z_{k,t})$, which subtracts workers' actual pay premia instead of the average for

[38] These findings indicate that time variation in the firm effects plays a key role in estimation since growth rates in pay premia are highly heterogeneous. Estimating AKM models with time-invariant firm effects could lead to biases in parameter estimates over long panels. The issue may be less severe over shorter panels of 7 years, as implemented by Card et al. (2013) and Song et al. (2019).

all workers in a county or industry,

$$\Delta \left(\log(w) - \xi\right)_{kt}^{FIRM} \equiv$$
$$\Delta \log(z_{k,t}) = \frac{\sum_{i \in \mathcal{N}_{k,t}} \left(\log(w_{it}) - \xi_{j(i,t),t}\right)}{N_{k,t}} - \frac{\sum_{i \in \mathcal{N}_{k,0}} \left(\log(w_{ij0}) - \xi_{j(i,t),0}\right)}{N_{k,0}} \tag{6.3}$$

Figure 7 shows that county and industry explain much of the change in wages, especially for college graduates. For college workers, the series that subtracts average industry pay premia nearly equals the series that subtracts the firm pay premia, indicating that industry-level changes explain most of the growth in average pay premia. The series that subtracts firm pay premia lies below the county and industry lines, implying that within counties and industries college workers sort to firms with higher growth in pay premia. The opposite pattern holds for non-college workers, implying that within counties and industries non-college workers sort to firms with lower growth in pay premia.

Figure 7 also shows a 19 log point gap between $\Delta \log(w_{Ct})$ and $\Delta \left(\log(w) - \xi\right)_{Ct}^{IND}$ by the end of the sample, compared to a 10 log point gap between $\Delta \log(w_{Nt})$ and $\Delta \left(\log(w) - \xi\right)_{Nt}^{IND}$. This implies that college workers sort to industries with faster growth in pay premia. Interestingly the gap between $\Delta \log(w_{kt})$ and $\Delta \left(\log(w) - \xi\right)_{kt}^{COU}$ is 15 log points for $k =$ college workers and 12 log points for $k =$ non-college workers, implying comparable growth in average pay premia for both groups by county. Within counties college workers sort to firms with fast growing pay premia, while non-college workers sort to firms with lower growth.

Overall Figure 7 shows that county and industry play a key role in explaining changes in wages for college and non-college workers. One interpretation is that certain industries and counties experienced large labor demand shocks over the time period, with college workers more likely to sort to those sectors of the economy.

6.1 Baseline Sorting or Sorting Over Time?

Did college workers increasingly sort to high-paying industries over time, or did they already work in those industries at baseline? The top panel of Figure 8 plots the share of workers in

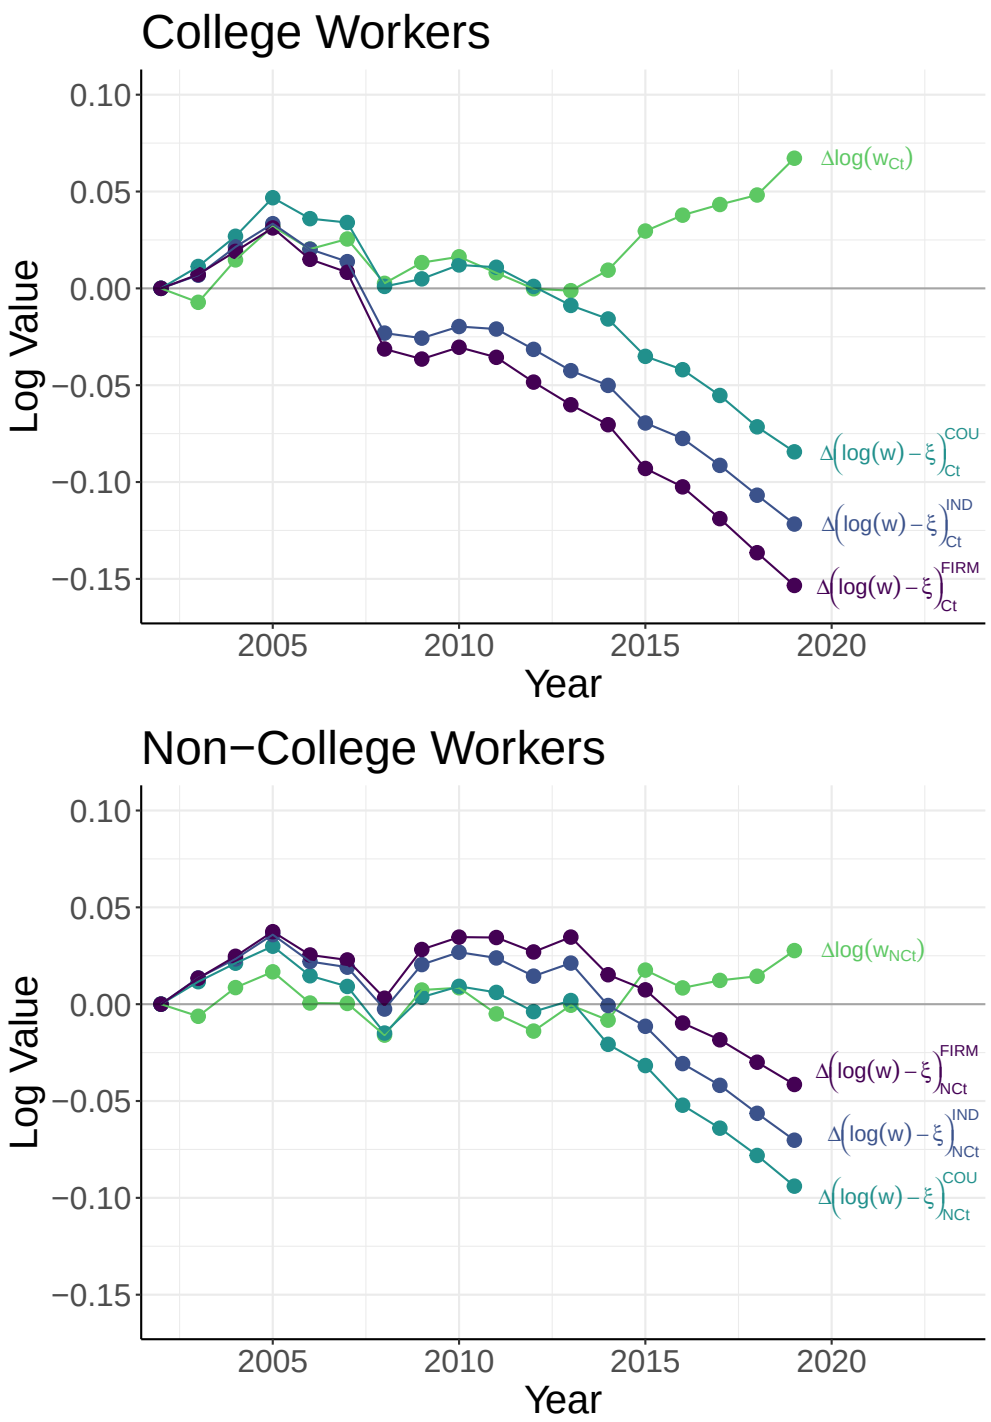


Figure 7: Decomposing amount of earnings growth explained by county and industry. The top panel shows changes for college workers. The bottom panel shows changes for non-college workers. $\Delta\left(\log(w) - \xi\right)_{kt}^{COU}$ takes the average of log wages subtracting the mean county pay premium. $\Delta\left(\log(w) - \xi\right)_{kt}^{IND}$ takes the average of log wages subtracting the mean 3-digit industry pay premium. $\Delta\left(\log(w) - \xi\right)_{kt}^{FIRM}$ equals the average total skill measure. See Equations 6.1 through 6.3.

2002 and 2019 in each quintile of the distribution of $\bar{\xi}_{l2019} - \bar{\xi}_{l2002}$. This splits counties into quintiles of growth in their average pay premia and then shows what share of workers fall in each quintile at baseline and at the end of the sample. At baseline college workers sorted into higher pay growth counties than non-college workers.[39] They particularly have higher representation in the top quintile of county pay growth. These patterns do not change much between 2002 and 2019.

The bottom panel of Figure 8 shows sharper baseline sorting patterns by industry, with the x-axis showing quintiles of $\bar{\xi}_{I2019} - \bar{\xi}_{I2002}$. The majority of college workers already worked at baseline in the top two quintiles for pay premia growth. Non-college workers showed surprisingly little sorting at baseline, with nearly equal proportions in all bins. The results again show relatively small changes over time in the shares.

I decompose the role of baseline sorting and changes in sorting over time in explaining growth in pay premia as follows:

$$\overbrace{\frac{\sum_{\mathcal{N}_{k2019}} \xi_{j(i)2019}}{N_{k2019}} - \frac{\sum_{\mathcal{N}_{k2002}} \xi_{j(i)2002}}{N_{k2002}}}^{\text{Total Change}} = \underbrace{\sum_g \left[\frac{\sum_{\mathcal{N}^g_{k2019}} \xi_{j(i)2019}}{N^g_{k2019}} - \frac{\sum_{\mathcal{N}^g_{k2002}} \xi_{j(i)2002}}{N^g_{k2002}}\right] \left(\frac{N^g_{2002}}{N_{2002}}\right)}_{\text{Level Change}} + \underbrace{\sum_g \left[\left(\frac{N^g_{2019}}{N_{2019}}\right) - \left(\frac{N^g_{2002}}{N_{2002}}\right)\right] \frac{\sum_{\mathcal{N}^g_{k2019}} \xi_{j(i)2019}}{N^g_{k2019}}}_{\text{Share Change}} \tag{6.4}$$

Equation 6.4 decomposes the change in pay premia between changes in share across groups, holding the level of pay premia constant at 2019 levels, and changes in levels across groups, holding constant the share of workers in each group at 2002 levels.

In Table 1 I show these decompositions separately for industries and counties, with groups defined as quintiles of $\bar{\xi}_{l2019} - \bar{\xi}_{l2002}$ and $\bar{\xi}_{I2019} - \bar{\xi}_{I2002}$. The level change explains over 90% of

[39] Card et al. (2023) show sorting of highly skilled workers to counties with high firm pay premia. My results show that educated workers sort to counties with high *growth* in firm pay premia. My results are closer to findings from Diamond (2016) that document increased spatial sorting of college workers to counties with growing labor demand.

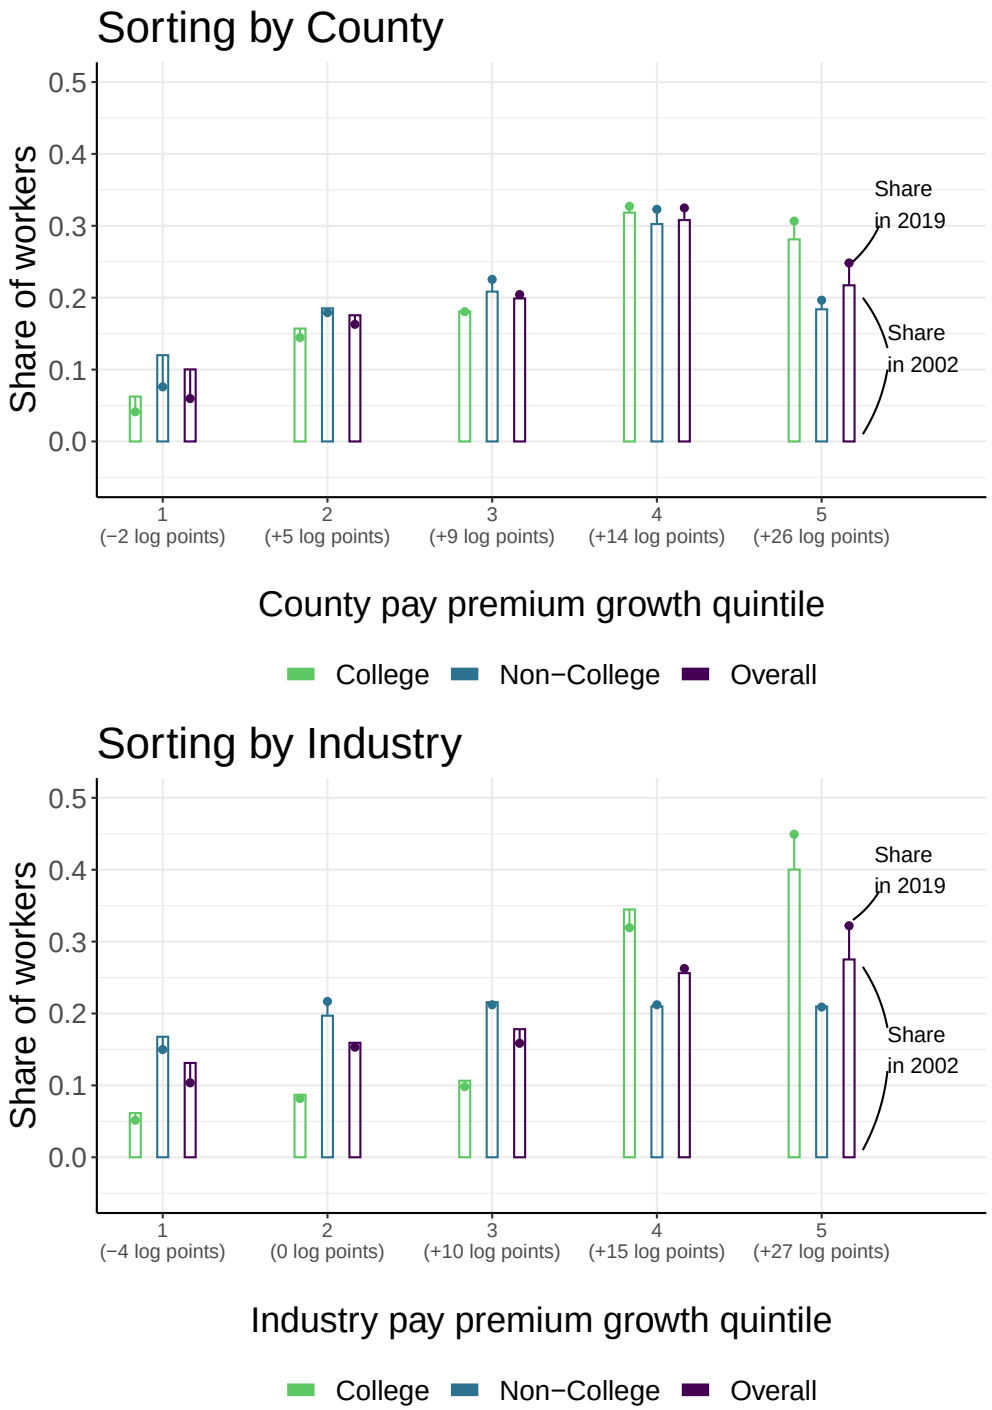


Figure 8: The top panel shows the share of workers across different groups of counties. The x-axis shows quintiles of $\bar{\xi}_{l2019} - \bar{\xi}_{l2002}$, the growth in the average firm effect, across counties. I display the change in the average firm effect within the quintile in parentheses. Bars represent shares in 2002 and points the shares in 2019. Colors represent different education groups. The bottom panels shows a similar figure by quintiles of $\bar{\xi}_{I2019} - \bar{\xi}_{I2002}$, the growth by industry.

County			
	Total Change	% Explained by Level Change	% Explained by Share Change
College	22 log points	95.7%	4.3%
Non-College	7 log points	91.6%	8.4%
Overall	14 log points	91.7%	8.3%
Industry			
	Total Change	% Explained by Level Change	% Explained by Share Change
College	22 log points	95.5%	4.5%
Non-College	7 log points	100.0%	0.0%
Overall	14 log points	92.1%	7.9%

Table 1: Decomposing Changes in Firm Pay Premia. Total Change shows the growth in average firm pay premia from 2002 to 2019. In the top panel, % Explained by Level Change shows the percent of the Total Change explained by changes in pay premia holding fixed the shares of workers in each quintile of county pay premia growth at 2002 levels. % Explained by Share Change holds fixed the levels of pay premia at 2019 and allows only shares of workers in each quintile to change. The bottom panel shows equivalent results by industry.

the growth in pay premia across each cell of the table, demonstrating that even at baseline college and non-college workers had predominantly sorted based on subsequent growth in pay premia. These baseline sorting patterns largely persisted even as the share of workers with a college degree increased by close to 40% over the same time period.[40]

Together, results in this section show that county and industry explain quite a meaningful share of the growth in the college earnings gap. College workers work in counties and industries that experienced large growth in average pay premia. They worked in these sectors of the economy even at baseline, with changes in sorting patterns explaining little of the growth in pay premia.

7 Implications for Labor Demand

How do skill composition changes impact estimates of the relative supply and demand for college educated workers? To study this I follow an exercise from Katz and Murphy (1992) and Acemoglu and Autor (2011) that decomposes the college wage gap between relative supply shifts and relative demand shifts over time.[41] Consider the canonical supply and demand model as discussed in Acemoglu and Autor (2011). Suppose each worker supplies one unit of labor inelastically. The total supplies for college and non-college labor equal

$$\mathbb{C} \equiv \sum_{\mathcal{N}_C} z_i, \quad \mathbb{NC} \equiv \sum_{\mathcal{N}_{NC}} z_i \tag{7.1}$$

with $\mathcal{N}_C$ the set of college workers, $\mathcal{N}_{NC}$ the set of non-college workers, and z_i the skill or efficiency units of worker i. Assume an aggregate production function for the economy with constant elasticity of substitution (CES) form,

$$Y = \left[(A_{NC} \times \mathbb{NC})^{\frac{\sigma-1}{\sigma}} + (A_C \times \mathbb{C})^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}} \tag{7.2}$$

[40] Diamond and Gaubert (2022) similarly find a greater role for level changes compared to share changes in explaining spatial differences in the college wages from 1980 to 2000.

[41] See also Acemoglu (2002); Autor et al. (2008); Goldin and Katz (2008); Autor (2014, 2018).

where $\sigma \in [0, \infty]$ is the elasticity of substitution between college and non-college labor, and A_C and A_{NC} are factor-augmenting technology terms. The extension of the model in Section 2.2.2 nests this framework.[42]

With competitive labor markets, the wage per efficiency unit for college and non-college workers equal

$$w_C = \frac{\partial Y}{\partial \mathbb{C}} = A_C^{\frac{\sigma-1}{\sigma}} \left[A_{NC}^{\frac{\sigma-1}{\sigma}} (\mathbb{C}/\mathbb{NC})^{-\frac{\sigma-1}{\sigma}} + A_C^{\frac{\sigma-1}{\sigma}} \right]^{\frac{1}{\sigma-1}} \tag{7.3}$$

$$w_{NC} = \frac{\partial Y}{\partial \mathbb{NC}} = A_{NC}^{\frac{\sigma-1}{\sigma}} \left[A_{NC}^{\frac{\sigma-1}{\sigma}} + A_C^{\frac{\sigma-1}{\sigma}} (\mathbb{C}/\mathbb{NC})^{\frac{\sigma-1}{\sigma}} \right]^{\frac{1}{\sigma-1}} \tag{7.4}$$

Because of additivity within education groups, log wages for individual workers equal

$$\begin{aligned} \log(w_i) &= \log(w_C) + \log(z_i) \quad \text{if} \quad i \in \mathcal{N}_C, \\ \log(w_i) &= \log(w_{NC}) + \log(z_i) \quad \text{if} \quad i \in \mathcal{N}_{NC} \end{aligned} \tag{7.5}$$

Let

$$\overline{\log(w_C)} = \frac{\sum_{i \in \mathcal{N}_C} \log(w_i)}{N_C}, \quad \overline{\log(w_{NC})} = \frac{\sum_{i \in \mathcal{N}_{NC}} \log(w_i)}{N_{NC}} \tag{7.6}$$

denote average log wages for each group, with N_C and N_{NC} the number of college and non-college workers. Combining the above equations, adding time subscripts, and taking differences implies

$$\begin{aligned} \Delta \log \left(\frac{A_C}{A_{NC}} \right)_t = &\left(\frac{\sigma}{\sigma - 1} \right) \Delta \left[\overline{\log(w_C)} - \overline{\log(w_{NC})} \right]_t \\ &+ \left(\frac{1}{\sigma - 1} \right) \Delta \log \left(\frac{N_C}{N_{NC}} \right)_t - \Delta \left[\overline{\log(z_C)} - \overline{\log(z_{NC})} \right]_t \end{aligned} \tag{7.7}$$

with $\overline{\log(z_C)} = \sum_{i \in \mathcal{N}_C} \log(z_i)/N_C$ and $\overline{\log(z_{NC})} = \sum_{i \in \mathcal{N}_{NC}} \log(z_i)/N_{NC}$ the average log skill for each group.

Following Katz and Murphy (1992), the left-hand side of Equation 7.7 has typically been interpreted as growth in relative labor demand for college workers. The expression shows

[42] Evidence from Appendix C.8 indicates that estimating separate TV-AKM models by education gives comparable results to estimating one joint model. In this section I thus use the main estimates from Section 5 to measure skill changes for college and non-college workers.

Table 2: Changes in relative demand from 2002 to 2019 for different values of elasticity of substitution σ between college and non-college workers. Demo Adj. adjusts relative wages and relative labor supply for demographic changes as in Autor et al. (2008) and Stephenson (2024). Skill Adj. adjusts for changes in average skill as in Equation 7.7. % Diff. shows the percentage difference in relative demand growth between the Skill Adj. and Demo Adj. columns. Sample years are the years over which the source study estimated σ. Goldin and Katz (2008) estimate σ using data from 1915, 1940, 1950, 1960, and 1963-2005.

Changes in relative demand, 2002-2019.

σ	Source, sample years	Standard Adj.	Skill Adj.	**% Diff.**
1.40	Katz and Murphy (1992), 1963-1987	1.10	1.27	**+15%**
1.80	Acemoglu (2002), 1939-1996	0.56	0.68	**+20%**
1.64	Goldin and Katz (2008), 1915-2005	0.70	0.82	**+18%**
2.90	Acemoglu and Autor (2011), 1963-2008	0.26	0.33	**+31%**
3.30	Autor (2018), 1963-2012	0.22	0.29	**+35%**
4.70	Own Estimate, 2002-2019	0.15	0.21	**+46%**

that growth in relative labor demand varies with growth in average skill. For any given value of σ, estimating relative growth in labor demand without correcting for skill would lead to bias. My empirical estimates of mean relative growth in average skill indicate that measures of labor demand growth from 2002 to 2019 would be downward biased by 11 log points absent skill adjustments.

Prior literature uses alternative approaches to correct for skill changes. Rather than using raw counts of workers, past work weights different demographic groups based on their relative wages across the sample period; see Autor et al. (2008) and Acemoglu and Autor (2011) for examples. These papers also adjust the wage series for composition by taking a weighted average of mean log wages for different demographic groups in each year, with fixed weights

equal to the average hours share over the full sample. In Table 2 I follow the methodology of these papers for different values of σ using data from the March CPS.[43] I find about 6-17 log point faster growth in relative labor demand for college workers when using the decomposition in Equation 7.7 compared to results that use composition adjustments from prior literature.[44] My results suggest that commonly used composition adjustments may underestimate changes in skill composition, in turn downward biasing estimates of growth in relative labor demand for college workers by 15-46%. This may consequently downward bias estimates of the potential causes of relative demand shifts such as skill-biased technical change.

8 Conclusion

In this work I develop a measure of average worker skill. The method subtracts time-varying firm effects from log wages, isolating wage variation that comes from differences across workers. I apply this approach to study changes in the average skill for college and non-college workers. I find that average declined by 15 log points for college workers, with smaller changes for non-college workers. Growth in the college wage gap in recent decades thus came from differential sorting to firms with high pay growth, with much of the effect explained by baseline sorting across counties and industries. The results have important implications for understanding sources of inequality and for measuring growth in relative labor demand for college workers.

Future work may apply the methodology to other contexts to measure how changes in

[43] I follow the replication from Stephenson (2024) of Autor et al. (2008). I discuss more details of the data construction in Appendix D. A difference in my work compared to prior papers is that I define worker types as college graduates and people without college degrees, in line with my empirical analysis in Section 5. Previous papers define worker types as college graduates and high school graduates, splitting people with some college equally between these two groups.

[44] Standard approaches can also lead to biases in estimates for σ. In Appendix D I estimate σ to equal 4.7, corresponding to the final row in Table 2. Estimates in prior literature range from 1.4 to 3.3 (see also Heckman et al. (1998) and Krusell et al. (2000)). Table 2 shows higher estimates of σ in studies using data from more recent years. Correspondingly I estimate a higher value of σ compared to the literature over 2002 to 2019.

average skill contribute to wage growth in different sectors of the economy. One fruitful direction for research is studying how skill changes shape wage growth across industries, occupations, and geographies. Another application is studying how skill changes contribute to aggregate GDP and TFP growth in the economy.

References

Abowd, J. M., R. H. Creecy, and F. Kramarz (2002): "Computing Person and Firm Effects Using Linked Longitudinal Employer-Employee Data," Longitudinal Employer-Household Dynamics Technical Papers 2002-06, Center for Economic Studies, U.S. Census Bureau. 2.1.3

Abowd, J. M., F. Kramarz, and D. N. Margolis (1999): "High Wage Workers and High Wage Firms," *Econometrica*, 67, 251–333, _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/1468-0262.00020. 1, 1, 11, 2.1.2

Abowd, J. M., K. L. McKinney, and N. L. Zhao (2018): "Earnings Inequality and Mobility Trends in the United States: Nationally Representative Estimates from Longitudinally Linked Employer-Employee Data," *Journal of Labor Economics*, 36, S183–S300, publisher: The University of Chicago Press. 11, 6

Abowd, J. M., B. E. Stephens, L. Vilhuber, F. Andersson, K. L. McKinney, M. Roemer, and S. Woodcock (2009): "The LEHD Infrastructure Files and the Creation of the Quarterly Workforce Indicators," in *Producer Dynamics: New Evidence from Micro Data*, ed. by T. Dunne, J. B. Jensen, and M. J. Roberts, University of Chicago Press, 149–234. 21

Acemoglu, D. (2002): "Technical Change, Inequality, and the Labor Market," *Journal of Economic Literature*, 40, 7–72. 1, 1, 41

Acemoglu, D. and D. Autor (2011): "Skills, Tasks and Technologies: Implications for Employment and Earnings*," in *Handbook of Labor Economics*, ed. by D. Card and O. Ashenfelter, Elsevier, vol. 4, 1043–1171. 1, 1, 3.3, 5.1, 7, 7, D1

Autor, D. (2018): "How Long Has This Been Going On? A Discussion of "Recent Flattening in the Higher Education Wage Premium: Polarization, Skill Downgrading, or Both?" by Robert G. Valletta," in *Education, Skills, and Technical Change: Implications for Future US GDP Growth*, University of Chicago Press, 342–353. 5.1, 41

Autor, D., A. Dube, and A. McGrew (2023): "The Unexpected Compression: Competition at Work in the Low Wage Labor Market," . 1, 5, 5.1

Autor, D., C. Goldin, and L. F. Katz (2020): "Extending the Race between Education and Technology," *AEA Papers and Proceedings*, 110, 347–351. 1

Autor, D. H. (2014): "Skills, education, and the rise of earnings inequality among the "other 99 percent"," *Science*, 344, 843–851, publisher: American Association for the Advancement of Science. 41

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, 300–323. 3.3, 41, 2, 7, 43, D, D1

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, 1169–1213. 1

Bagger, J., B. Christensen, and D. Mortensen (2014): "Wage and Labor Productivity Dispersion: The Roles of Total Factor Productivity, Labor Quality, Capital Intensity, and Rent Sharing," . 3

Bagger, J. and R. Lentz (2019): "An Empirical Model of Wage Dispersion with Sorting," *The Review of Economic Studies*, 86, 153–190. 14

Barth, E., A. Bryson, J. C. Davis, and R. Freeman (2016): "It's Where You Work: Increases in the Dispersion of Earnings across Establishments and Individuals in the United States," *Journal of Labor Economics*, 34, S67–S97, publisher: The University of Chicago Press. 11

Beaudry, P., D. A. Green, and B. M. Sand (2014): "The Declining Fortunes of the Young since 2000," *American Economic Review*, 104, 381–386. 5.1

——— (2016): "The Great Reversal in the Demand for Skill and Cognitive Tasks," *Journal of Labor Economics*, 34, S199–S247, publisher: The University of Chicago Press. 5.1

Bonhomme, S., K. Holzheu, T. Lamadon, E. Manresa, M. Mogstad, and B. Setzler (2023): "How Much Should We Trust Estimates of Firm Effects and Worker Sorting?" *Journal of Labor Economics*, 41, 291–322, publisher: The University of Chicago Press. 2.1.2

Bonhomme, S., T. Lamadon, and E. Manresa (2019): "A Distributional Framework for Matched Employer Employee Data," *Econometrica*, 87, 699–739, _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA15722. 16

Borovickova, K. and R. Shimer (2024): "Assortative Matching and Wages: The Role of Selection," Working Paper. 17

Bowlus, A., L. Lochner, C. Robinson, and E. Suleymanoglu (2023): "Wages, Skills, and Skill-Biased Technical Change: The Canonical Model Revisited," *Journal of Human Resources*, 58, 1783–1819, publisher: University of Wisconsin Press Section: Articles. 6

Bowlus, A. J. and C. Robinson (2012): "Human Capital Prices, Productivity, and Growth," *American Economic Review*, 102, 3483–3515. 1, 10, 4.1.2

Card, D., A. R. Cardoso, J. Heining, and P. Kline (2018): "Firms and Labor Market Inequality: Evidence and Some Theory," *Journal of Labor Economics*, 36, S13–S70, publisher: The University of Chicago Press. 1, 2.1.2, 4.1.2, 4.3

Card, D., A. R. Cardoso, and P. Kline (2016): "Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women *," *The Quarterly Journal of Economics*, 131, 633–686. 11, 12

Card, D., J. Heining, and P. Kline (2013): "Workplace Heterogeneity and the Rise of West German Wage Inequality*," *The Quarterly Journal of Economics*, 128, 967–1015. 1, 1, 11, 2.1.2, 2.1.2, 2.1.3, 2.2.2, 29, 4.3, 38

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, 705–746. 1

Card, D., J. Rothstein, and M. Yi (2023): "Location, Location, Location," . 39

——— (2024): "Industry Wage Differentials: A Firm-Based Approach," *Journal of Labor Economics*, 42, S11–S59, publisher: The University of Chicago Press. 27

Carneiro, P. and S. Lee (2009): "Estimating distributions of potential outcomes using local instrumental variables with an application to changes in college enrollment and wage inequality," *Journal of Econometrics*, 149, 191–208. 6

——— (2011): "Trends in Quality-Adjusted Skill Premia in the United States, 1960-2000," *American Economic Review*, 101, 2309–2349. 1, 7

Chan, M., E. Mattana, S. Salgado, and M. Xu (2023a): "Wage Setting and

Passthrough: The Role of Market Power, Production Technology and Adjustment Costs," Working Paper. 1

Chan, M., S. Salgado, and M. Xu (2023b): "Heterogeneous Passthrough from TFP to Wages," Working Paper. 3, 2.1.3

Chetty, R., J. N. Friedman, E. Saez, N. Turner, and D. Yagan (2020): "Income Segregation and Intergenerational Mobility Across Colleges in the United States*," *The Quarterly Journal of Economics*, 135, 1567–1633. 1

Cilluffo, R. F. a. A. (2019): "A Rising Share of Undergraduates Are From Poor Families, Especially at Less Selective Colleges," . 1

College Board (2016): "Total Group Profile Report," Tech. rep. 2, C1

de la Roca, J. and D. Puga (2017): "Learning by Working in Big Cities," *The Review of Economic Studies*, 84, 106–142. 14

Deming, D. J. (2022): "Four Facts about Human Capital," *Journal of Economic Perspectives*, 36, 75–102. 1

Di Addario, S., P. Kline, R. Saggio, and M. Sølvsten (2023): "It ain't where you're from, it's where you're at: Hiring origins, firm heterogeneity, and wages," *Journal of Econometrics*, 233, 340–374. 15

Diamond, R. (2016): "The Determinants and Welfare Implications of US Workers' Diverging Location Choices by Skill: 1980-2000," *American Economic Review*, 106, 479–524. 1, 39

Diamond, R. and C. Gaubert (2022): "Spatial Sorting and Inequality," *Annual Review of Economics*, 14, 795–819, publisher: Annual Reviews. 40

Drenik, A., S. Jäger, P. Plotkin, and B. Schoefer (2023): "Paying Outsourced Labor: Direct Evidence from Linked Temp Agency-Worker-Client Data," *The Review of Economics and Statistics*, 105, 206–216. 27

Dworak, E. M., W. Revelle, and D. M. Condon (2023): "Looking for Flynn effects in a recent online U.S. adult sample: Examining shifts within the SAPA Project," *Intelligence*, 98, 101734. 2

Eeckhout, J. and P. Kircher (2011): "Identifying Sorting—In Theory," *The Review of Economic Studies*, 78, 872–906. 17

Engbom, N., C. Moser, and J. Sauermann (2023): "Firm pay dynamics," *Journal of Econometrics*, 233, 396–423. 3, 2.1.3, 18

Goldin, C. and L. F. Katz (2008): *The Race between Education and Technology*, Harvard University Press. 5.1, 41, 2

Haltiwanger, J., H. R. Hyatt, and J. R. Spletzer (2022): "Rising Top, Falling Bottom: Industries and Rising Wage Inequality," *American Economic Review*. 11, 6

Heckman, J. J., L. Lochner, and C. Taber (1998): "Explaining Rising Wage Inequality: Explorations with a Dynamic General Equilibrium Model of Labor Earnings with Heterogeneous Agents," *Review of Economic Dynamics*, 1, 1–58. 1, 1, 4.1.2, 44

Hendricks, L., C. Herrington, and T. Schoellman (2021): "College Quality and Attendance Patterns: A Long-Run View," *American Economic Journal: Macroeconomics*, 13, 184–215. 8

Hendricks, L. and T. Schoellman (2014): "Student abilities during the expansion of US education," *Journal of Monetary Economics*, 63, 19–36. 7

Hoffmann, F., D. S. Lee, and T. Lemieux (2020): "Growing Income Inequality in the

United States and Other Advanced Economies," *Journal of Economic Perspectives*, 34, 52–78. 1

Ichino, A., A. Rustichini, and G. Zanella (2024): "College, cognitive ability, and socioeconomic disadvantage: policy lessons from the UK in 1960-2004," . 2, 7

Jones, C. I. (2023): "The Outlook for Long-Term Economic Growth," . 1

Juhn, C., D. I. Kim, and F. Vella (2005): "The Expansion of College Education in the United States: Is There Evidence of Declining Cohort Quality?" *Economic Inquiry*, 43, 303–315, _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1093/ei/cbi020. 6

Juhn, C., K. M. Murphy, and B. Pierce (1993): "Wage Inequality and the Rise in Returns to Skill," *Journal of Political Economy*, 101, 410–442, publisher: University of Chicago Press. 1

Katz, L. F. and K. M. Murphy (1992): "Changes in Relative Wages, 1963–1987: Supply and Demand Factors*," *The Quarterly Journal of Economics*, 107, 35–78. 1, 1, 1, 8, 5.1, 7, 7, D

Kline, P., R. Saggio, and M. Sølvsten (2020): "Leave-Out Estimation of Variance Components," *Econometrica*, 88, 1859–1898, _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA16410. 13, 4.1.1, 5.2, C.5

Kline, P. M. (2024): "Firm Wage Effects," . 1, 2.1.2, 2.1.2, 4.3, C.5

Krusell, P., L. E. Ohanian, J.-V. Ríos-Rull, and G. L. Violante (2000): "Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis," *Econometrica*, 68, 1029–1053, publisher: [Wiley, Econometric Society]. 1, 44

Lachowska, M., A. Mas, R. Saggio, and S. A. Woodbury (2023): "Do firm effects drift? Evidence from Washington administrative data," *Journal of Econometrics*, 233, 375–395. 1, 2.1, 2.1.3

Lagakos, D., B. Moll, T. Porzio, N. Qian, and T. Schoellman (2018): "Life Cycle Wage Growth across Countries," *Journal of Political Economy*, 126, 797–849, publisher: The University of Chicago Press. 4.1.2

Lamadon, T., M. Mogstad, and B. Setzler (2022): "Imperfect Competition, Compensating Differentials, and Rent Sharing in the US Labor Market," *American Economic Review*, 112, 169–212. 2.1

OECD (2024): "Do Adults Have the Skills They Need to Thrive in a Changing World?: Survey of Adult Skills 2023," Tech. rep., OECD Publishing, Paris. 2

Parker, K., J. M. Horowitz, A. Brown, R. Fry, D. Cohn, and R. Igielnik (2018): "Demographic and economic trends in urban, suburban and rural communities," . 9

Postel–Vinay, F. and J. Robin (2002): "Equilibrium Wage Dispersion with Worker and Employer Heterogeneity," *Econometrica*, 70, 2295–2350, _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1468-0262.2002.00441.x. 14

Song, J., D. J. Price, F. Guvenen, N. Bloom, and T. von Wachter (2019): "Firming Up Inequality*," *The Quarterly Journal of Economics*, 134, 1–50. 1, 11, 2.1.2, 2.1.3, 4.1.1, 4.1.2, 26, 4.3, 38

Sorkin, I. (2018): "Ranking Firms Using Revealed Preference*," *The Quarterly Journal of Economics*, 133, 1331–1393. 11, 2.1, 23, 4.1.1, 27

Stephenson, C. (2024): "Trends in U.S. Wage Inequality: Revising the Revisionists. A Replication Study of Autor, Katz, and Kearney (The Review of Economics and Statistics, 2008)," *Journal of Comments and Replication in Economics (JCRE)*, 3, 1–26, publisher: Kiel, Hamburg: ZBW - Leibniz-Informationszentrum Wirtschaft. 2, 43, D

Tinbergen, J. (1974): "Substitution of Graduate by Other Labour," *Kyklos*, 27, 217–226, _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1467-6435.1974.tb01903.x. 1

——— (1975): *Income Differences: Recent Research*, North-Holland Publishing Company, google-Books-ID: 2hguAAAAMAAJ. 1

Valletta, R. G. (2018): "Recent Flattening in the Higher Education Wage Premium Polarization, Skill Downgrading, or Both?" in *Education, Skills, and Technical Change: Implications for Future US GDP Growth*, ed. by C. R. Hulten and V. A. Ramey, University of Chicago Press, 313–354. 1, 1, 1, 1, 5, 3.3, 5.1

Yeh, C., C. Macaluso, and B. Hershbein (2022): "Monopsony in the US Labor Market," *American Economic Review*, 112, 2099–2138. 1, 19, A.2

Online Appendix (Not for Publication)

A Model Appendix

A.1 Interpreting Assumption 2

First I prove each of 2.15 through 2.17. Each follows simply from Assumption 2.

Exclusion The TV-AKM model specifies

$$\log w_{it}(j_1, \ldots, j_T) = \xi_{jt} + \zeta_i + \Gamma^T X_{it} + \epsilon_{it}$$

Since $E[\epsilon_{it}|j_1, \ldots, j_T] = 0$, $E[\log w_{it}(j_1, \ldots, j_T)] = E[\log w_{it}(j_t)]$.

Parallel Trends

$$E[\log w_{i,t+1}(j) - \log w_{i,t}(j)|j(i,t) = j, j(i,t+1) = k] =$$

$$E[\xi_{jt+1} + \zeta_i + \Gamma^T X_{it+1} + \epsilon_{it+1}|j(i,t) = j, j(i,t+1) = k]$$

$$-E[\xi_{jt} + \zeta_i + \Gamma^T X_{it} + \epsilon_{it}|j(i,t) = j, j(i,t+1) = k]$$

$$= \xi_{jt+1} - \xi_{jt} + \Gamma^T(X_{it+1} - X_{it})$$

No Selection

$$E[\log w_{i,t+1}(k) - \log w_{i,t}(j) | j(i,t) = j, j(i,t+1) = k]$$

$$= \xi_{kt+1} - \xi_{jt} + \Gamma^T(X_{it+1} - X_{it})$$

$$= E[\log w_{i,t+1}(k) - \log w_{i,t}(j)]$$

Given these three conditions, wage changes when moving between firms or over time identify unconditional average treatment effects of transitions across firm-times, after adjusting for covariates.

$$E[\log w_{i,t+1} - \log w_{i,t} | j(i,t) = j, j(i,t+1) = k]$$

$$= E[\log w_{i,t+1}(k) - \log w_{i,t}(j) | j(i,t) = j, j(i,t+1) = k]$$

$$= E[\log w_{i,t+1}(k) - \log w_{i,t}(j)]$$

$$= \xi_{kt+1} - \xi_{jt} + \Gamma^T(X_{it+1} - X_{it})$$

A.2 Dynamic Labor Adjustment Costs

In the main text I show that with myopic firms log additivity of wages holds even with adjustment costs. I now extend the result to allow for dynamic decision-making, following results from Yeh et al. (2022). Assume continuously differentiable revenue, wages, and adjustment costs in labor. As in Section 2.2 I allow adjustment costs to depend on prior period labor. Consider the firm's dynamic profit maximization problem,

$$V_{jt}(L_{jt-1}) = \max_{L_{jt}} R_{jt}(L_{jt}) - w_{jt}(L_{jt})L_{jt} - \Phi_{jt}(L_{jt}, L_{jt-1}) + \beta E[V_{j,t+1}(L_{jt})]$$

with $\beta \in [0,1]$ the firm's discount factor. I allow for stochasticity in future value. Differentiating,

$$R'_{jt}(L_{jt}) = w'_{jt}(L_{jt})L_{jt} + w_{jt}(L_{jt}) + \frac{\partial \Phi_{jt}(L_{jt}, L_{jt-1})}{\partial L_{jt}} - \beta E[V'_{j,t+1}(L_{jt})]$$

$$= w'_{jt}(L_{jt})L_{jt} + w_{jt}(L_{jt}) + \frac{\partial \Phi_{jt}(L_{jt}, L_{jt-1})}{\partial L_{jt}} - \beta E\left[\frac{\partial \Phi_{jt+1}(L^*, L_{jt})}{\partial L_{jt}}\right]$$

where the second equality uses the Envelope Theorem, and L^* denotes the optimal choice of labor next period. Then

$$R'_{jt}(L_{jt}) = \left[e_{jt}^{-1} + 1\right] w_{jt}(L_{jt}) + \frac{\partial \Phi_{jt}(L_{jt}, L_{jt-1})}{\partial L_{jt}} - \beta E\left[\frac{\partial \Phi_{jt+1}(L^*, L_{jt})}{\partial L_{jt}}\right]$$

so

$$w_{jt}(L_{jt}) = \left(R'_{jt}(L_{jt}) - \frac{\partial \Phi_{jt}(L_{jt}, L_{jt-1})}{\partial L_{jt}} + \beta E\left[\frac{\partial \Phi_{jt+1}(L^*, L_{jt})}{\partial L_{jt}}\right]\right)\left[e_{jt}^{-1} + 1\right]^{-1}$$

Under additive labor inputs, the chain rule again implies log additive wages.

B Simulation

To test for potential biases in the model, I conduct simulation exercises. I aim to mimic the mobility and characteristics of the data as closely as possible.

I subset to Indiana, Iowa, Ohio, Kansas, and Wisconson. Within these states I take a 1/3 random sample of individuals. I also restrict to people who work for at least two firms over the sample period. I estimate the TV-AKM model on this sample, getting an adjusted R^2 of 0.84. I use the imputed data in the LEHD for education.

B.1 Baseline Simulation

I test if the model recovers the true parameters in simulation. In the baseline simulation, I model the simulated outcome as

$$\log(\tilde{w}_{ijt}) = \tilde{\xi}_{jt} + \log(\tilde{h}_{ijt}) + \tilde{\zeta}_i + \tilde{\Gamma}^T X_{it} + \tilde{\epsilon}_{ijt} \tag{B.1}$$

I sample $\tilde{\xi}_{jt}$ from the model's estimated firm-time effects, conditioning on state-year; measure $\log(\tilde{h}_{ijt})$ as mean hours for by education x sex x age group x Census region x year in the ACS; sample $\tilde{\zeta}_i$ from the model's estimated person effects, conditioning on decade of birth,

education, and sex; set the age effect $\tilde{\Gamma}^T X_{it}$ to equal model's estimated value; and sample $\tilde{\epsilon}_{ijt}$ randomly from the model's estimated residuals. I then estimate the TV-AKM model on the simulated data and test for biases in the parameters.

As a first test of the models' fit, I regress the estimated firm-time effects, $\widehat{\xi}_{jt}$, on the true simulated firm-time effects, $\tilde{\xi}_{jt}$. I fail to reject the null hypothesis that estimated coefficient equals 1, with a standard error of 0.0003. The R^2 for the regression equals 0.80. To verify that this result is not some artifice of the simulation protocol, I also regress the estimated firm-time effects in the simulation on the estimated firm-time effects on the real data. These should have no relationship with each other since I randomly sample the firm-time effects in the simulation. I fail to reject that the coefficient equals 0.

Even if the firm-time effects are unbiased on average pooling across years, they may have biased trends. In Figure B1 I plot the average difference between $\widehat{\xi}_{jt}$ and $\tilde{\xi}_{jt}$ across years. I show this separately for firms where more than 50% of workers have college degrees and firms where less than 50% do. I find no trends in the differences. In Figure B2 I plot the average residual for college and non-college workers across years, again finding no trends or deviations from 0.

B.2 Differences in Skill Accumulation

I next probe various threats to the model. I start with differences in skill accumulation between college and non-college workers. First, I simulate faster skill growth for college workers compared to non-college workers. I modify the baseline simulation so that college workers have 1% growth in skill accumulation per year:

$$\log(\tilde{w}_{ijt}) = \tilde{\xi}_{jt} + \log(\tilde{h}_{ijt}) + \tilde{\zeta}_i + \tilde{\Gamma}^T X_{it} + College_i \times .01 \times (t - 2002) + \tilde{\epsilon}_{ijt}$$

I use the same simulated values for $\tilde{\xi}_{jt}, \log(\tilde{h}_{ijt}), \tilde{\zeta}_i, \tilde{\Gamma}^T X_{it}$, and $\tilde{\epsilon}_{ijt}$ as in the baseline model.

I again regress the model's estimate firm-time effects on the true simulated firm-time effects, finding no bias on average. However, I do find biased trends in Figure B3. Firms with

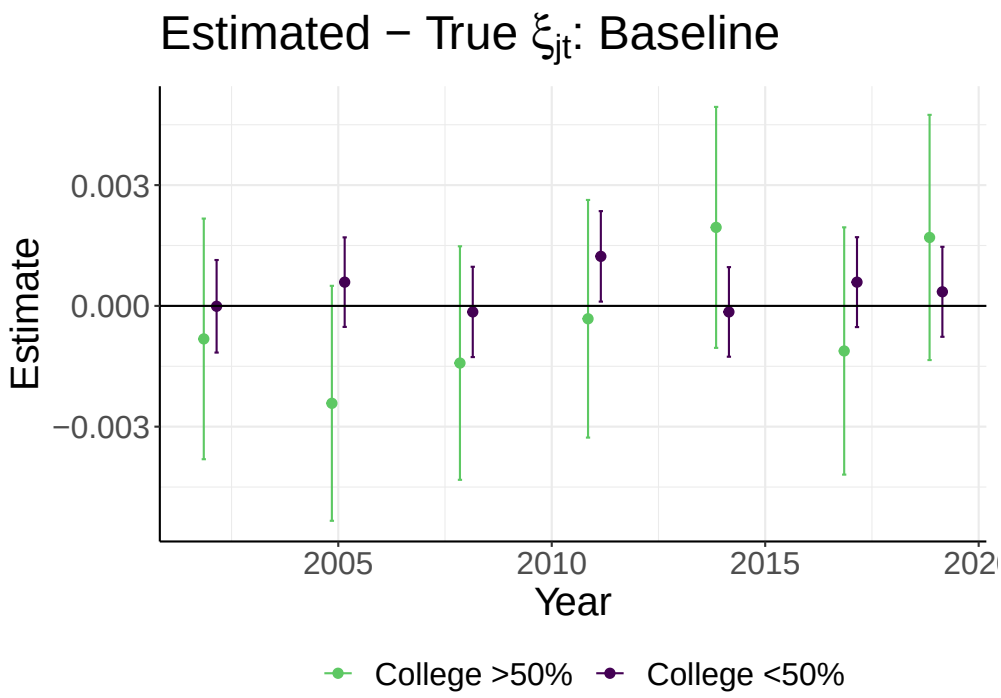


Figure B1: Average value of $\widehat{\xi}_{jt} - \tilde{\xi}_{jt}$ by year in the baseline simulation. Colors denote whether more than 50% of the firm's workers have college degrees.

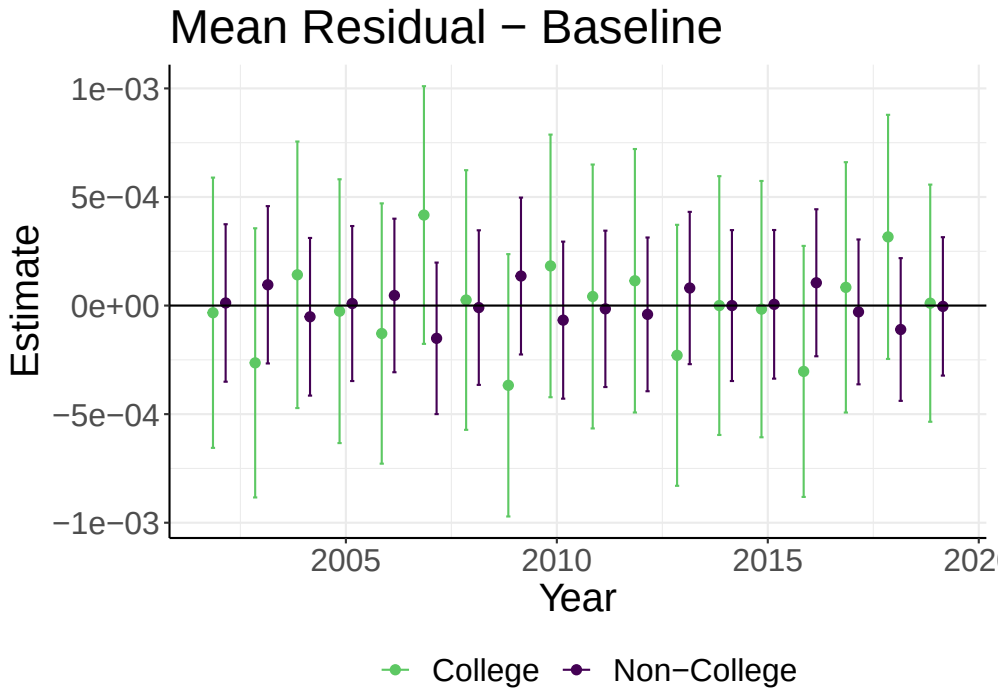


Figure B2: Average residual by year for college and non-college workers in the baseline simulation.

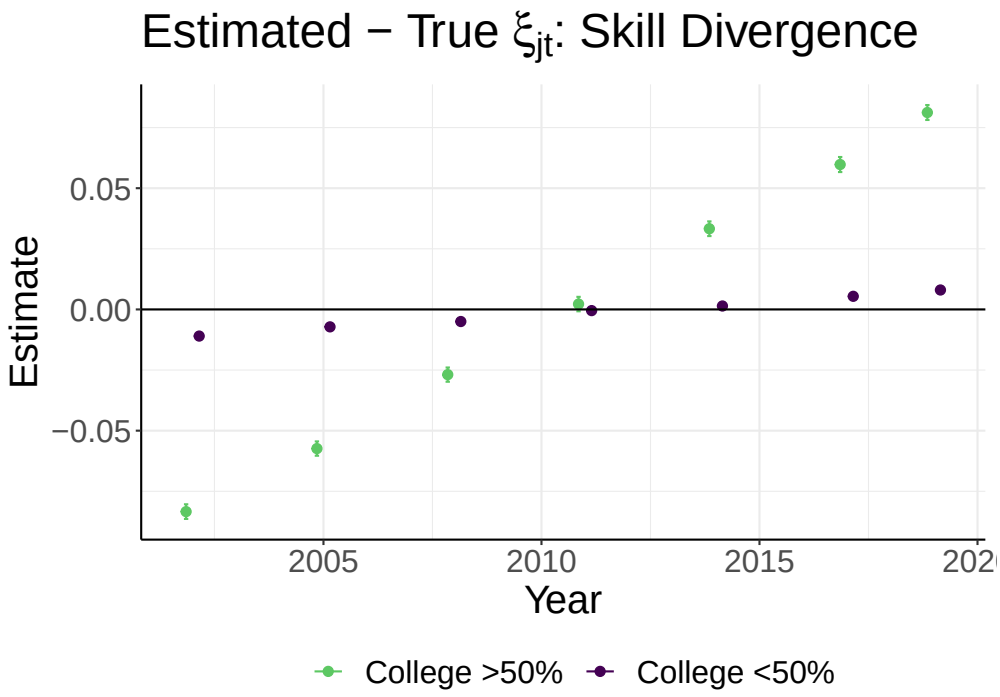


Figure B3: Average value of $\widehat{\xi}_{jt} - \tilde{\xi}_{jt}$ by year in the model with 1% growth in skill per year for college workers.

a high share of college workers have growing differences between the estimated and true firm-time effects. Further, Figure B4 shows that college workers have increasing residuals over time, and non-college workers have decreasing residuals. Differences in skill accumulation across groups can lead to biases in the parameters of the model and subsequently the average skill measures.

B.3 Modeling the Age Term

I next test whether the model can recover the true parameters under alternative specifications of the age term. I cannot separately identify firm-time effects, worker effects, and a linear term in age because of collinearity. Instead for my main results I control for a quadratic and cubic in age-50.

Figures B5 through B7 plot estimated model fits of earnings in age for a non-parametric estimator, a cubic esimator with a linear term, and a cubic estimator with no linear term.

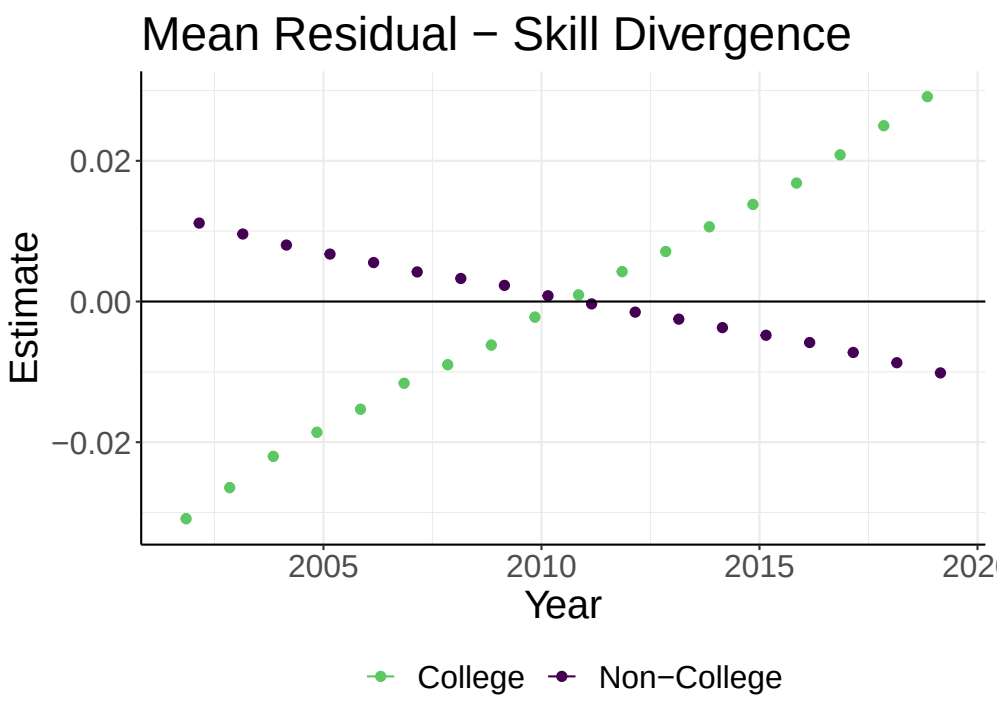


Figure B4: Average residual by year for college and non-college workers in the model with 1% growth in skill per year for college workers.

Each specification controls for year fixed effects. The non-parametric estimator estimates a dummy for each year of age,

$$\log(w_{it}) = \sum_{a=20}^{60} \beta_a 1\{AGE_{it} = a\} + \delta_t + e_{ijt}$$

Each model shows similar fits, though the cubic model with no linear term seems to penalize workers under 25 more. For my main results I exclude workers under 25.

To test if my functional form for experience biases the parameter estimates I modify the baseline simulation to incorporate alternative age effects. I generate the simulated outcome as

$$\log(\tilde{w}_{ijt}) = \tilde{\xi}_{jt} + \log(\tilde{h}_{ijt}) + \tilde{\zeta}_i + \tilde{\Gamma}_+^T X_{it,+} + \tilde{\epsilon}_{ijt}$$

I set all terms as in the baseline model except the age effect $\tilde{\Gamma}_+^T X_{it,+}$ equals either the above nonparametric age fit or cubic age fit with a linear term. I estimate the model using a cubic fit with no linear term.

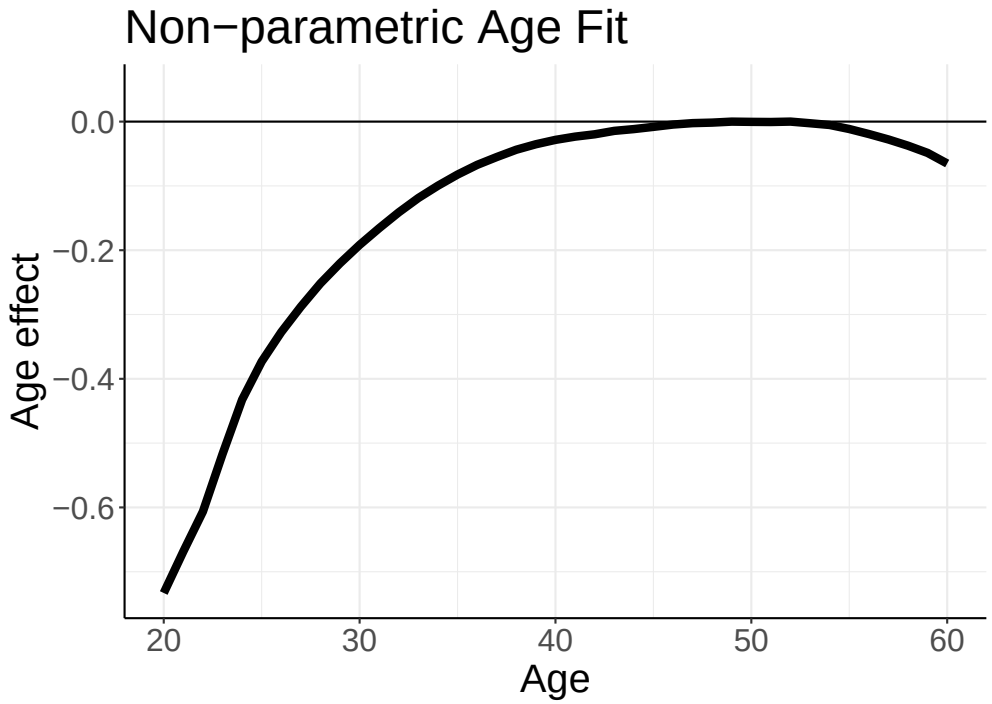


Figure B5: Nonparametric fit of earnings as a function of age, controlling for year fixed effects and normalizing to 0 at age 50.

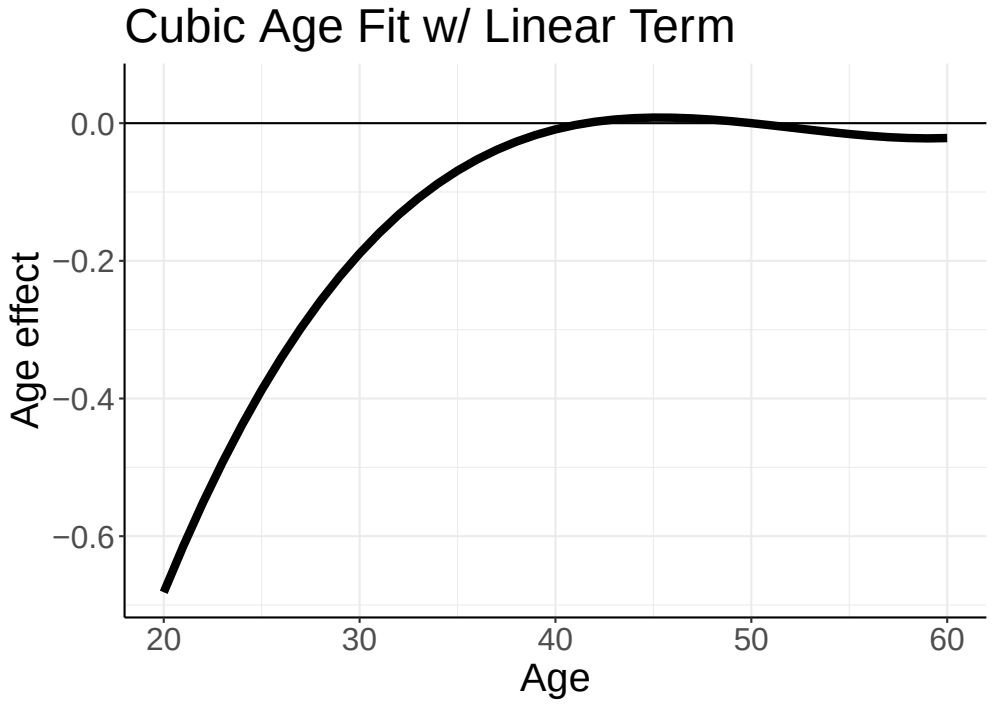


Figure B6: Cubic fit of earnings as a function of age, including a linear term. Controlling for year fixed effects and normalizing to 0 at age 50.

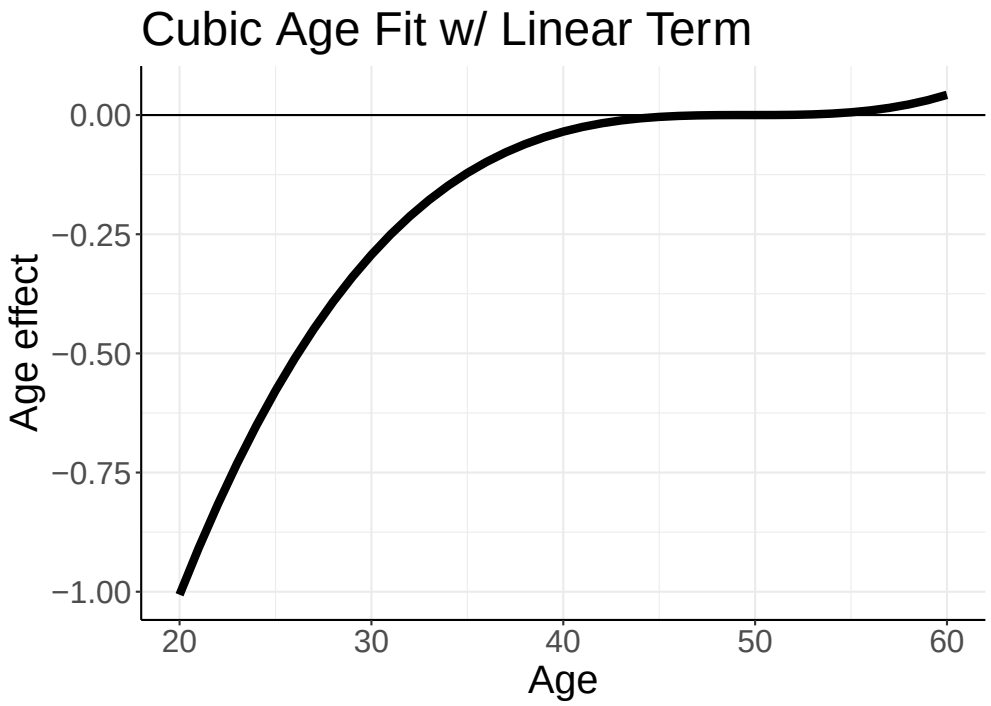


Figure B7: Cubic fit of earnings as a function of age, excluding a linear term. Controlling for year fixed effects and normalizing to 0 at age 50.

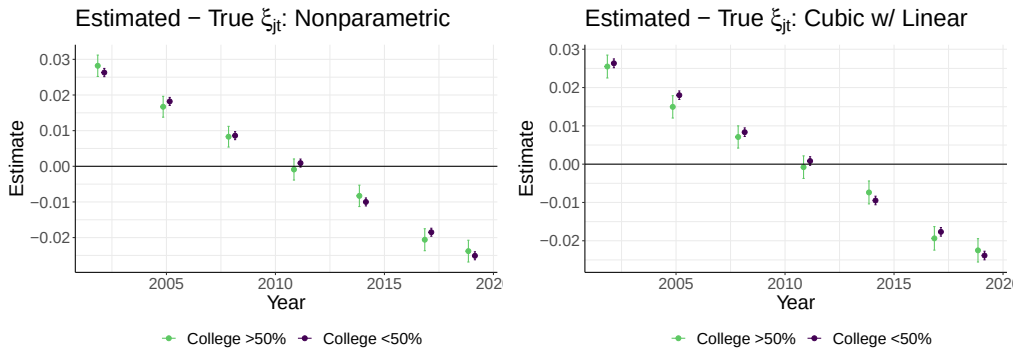


Figure B8: Average value of $\widehat{\xi}_{jt} - \tilde{\xi}_{jt}$ by year in the model. The model for the left panel simulates the age term with the nonparametric fit from Figure B5. The model for the right panel simulates the age term using cubic fit with a linear term from Figure B6. I estimate both models using a cubic fit with no linear term for age.

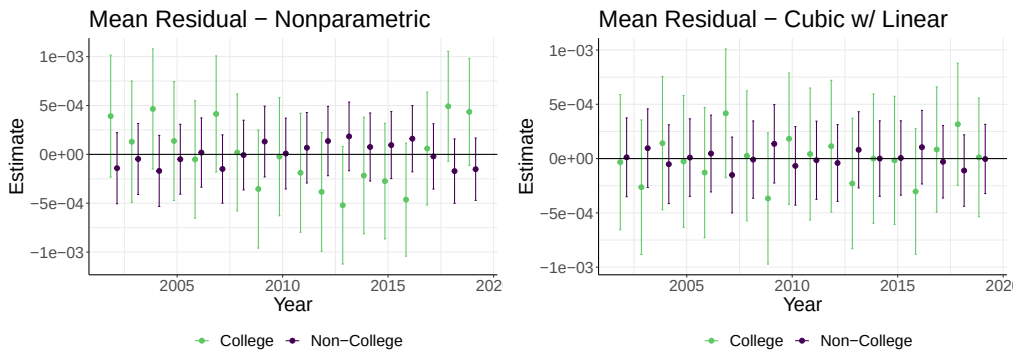


Figure B9: Average residual by year for college and non-college workers. The model for the left panel simulates the age term with the nonparametric fit from Figure B5. The model for the right panel simulates the age term using cubic fit with a linear term from Figure B6. I estimate both models using a cubic fit with no linear term for age.

I find that these changes cause downward trends in the estimated firm-time effects relative to their true values. These trends seem to be similar across firms. I find no trends in residuals.

The cubic functional form with no linear term may introduce biases in the parameter estimates. I test if using a more flexible function corrects these biases. In Figure B10 I use the same simulated outcome model but estimate it using a fifth order polynomial centered at age 50, excluding the linear term. With a more flexible polynomial the mean residual shrinks to almost zero with the nonparametric outcome model. The results suggests that potential errors may be corrected with a higher order polynomail model, though the error remains comparable to before for the cubic outcome model. Motivated by this exercise I replicate my main results using a fifth order polynomial with no linear term for age in Figure C18. I find very similar estimates quantitatively to my main results.

Finally, I consider allowing for different age-earnings profiles by education and sex. Figures B11 and B12 show nonparametric fits of earnings as a function of age by education and sex. The plots suggest greater gains from experience for college workers and for men.

To test if these features could lead to biases modify the simulated outcomes from the baseline model to have non-parametric age fits by education and sex. Figures B13 and B14 show these features may introduce biases when estimating the model with a single cubic fit with no linear term. Motivated by this exercise I replicate my main results when allowing

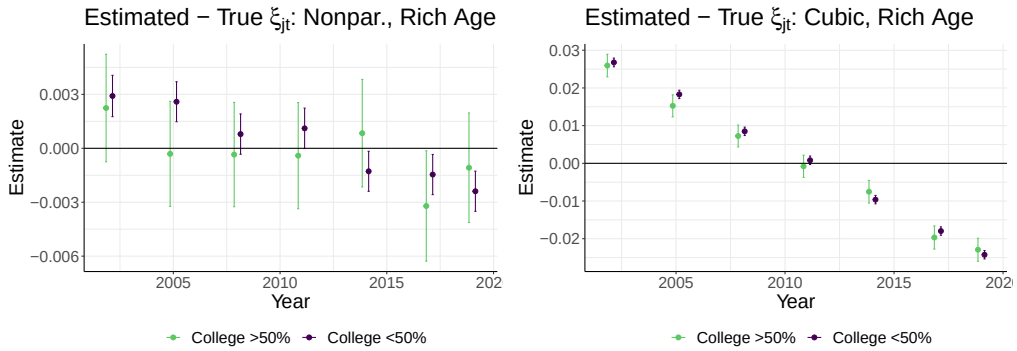


Figure B10: Average value of $\widehat{\xi}_{jt} - \tilde{\xi}_{jt}$ by year in the model. The model for the left panel simulates the age term with the nonparametric fit from Figure B5. The model for the right panel simulates the age term using cubic fit with a linear term from Figure B6. I estimate both models using a fifth order polynomial with no linear term for age.

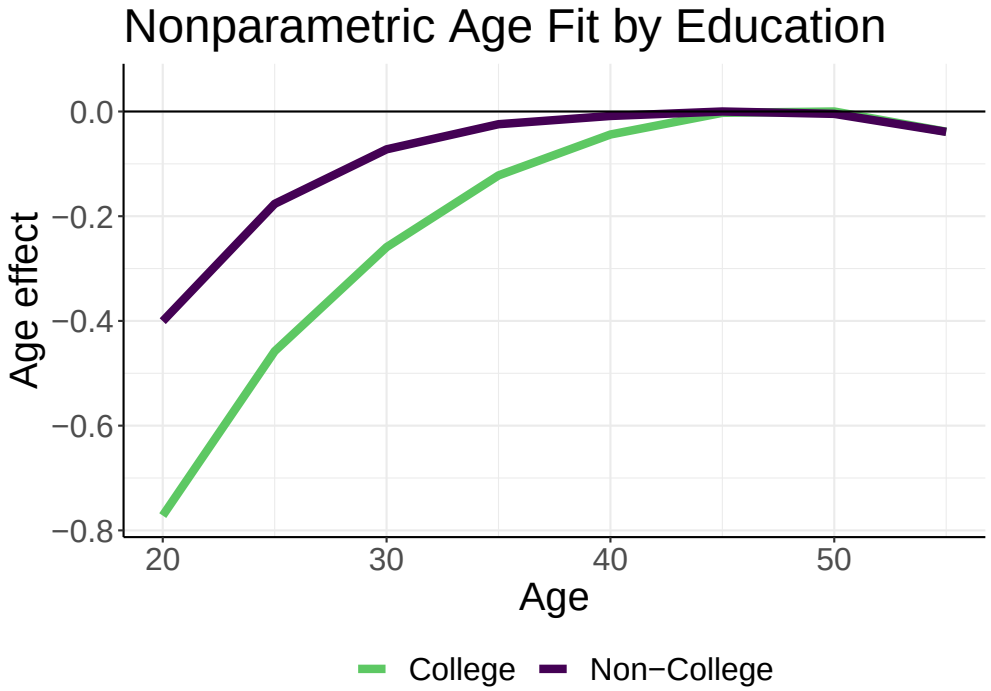


Figure B11: Nonparametric fit of earnings as a function of age, controlling for year fixed effects and normalizing to 0 at age 50. Shown separately for college and non-college workers.

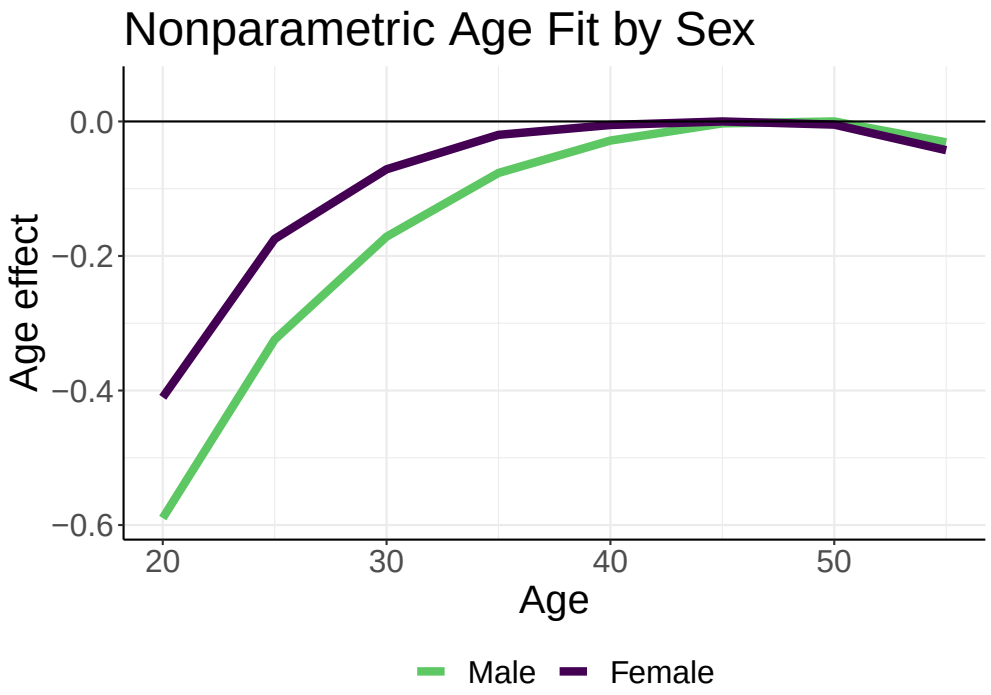


Figure B12: Nonparametric fit of earnings as a function of age, controlling for year fixed effects and normalizing to 0 at age 50. Shown separately for men and women.

for a separate fifth order polynomial by education and sex. Results in Figure C19 mimic my main results in Section 5.

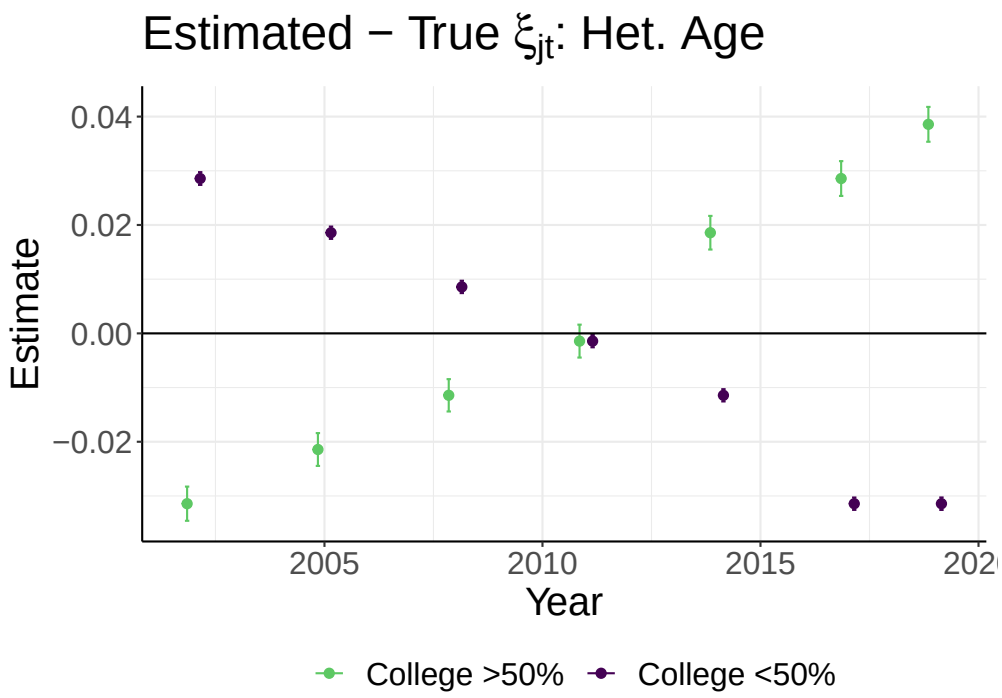


Figure B13: Average value of $\widehat{\xi}_{jt} - \tilde{\xi}_{jt}$ by year. Simulated outcomes have nonparametric age fits by education and sex.

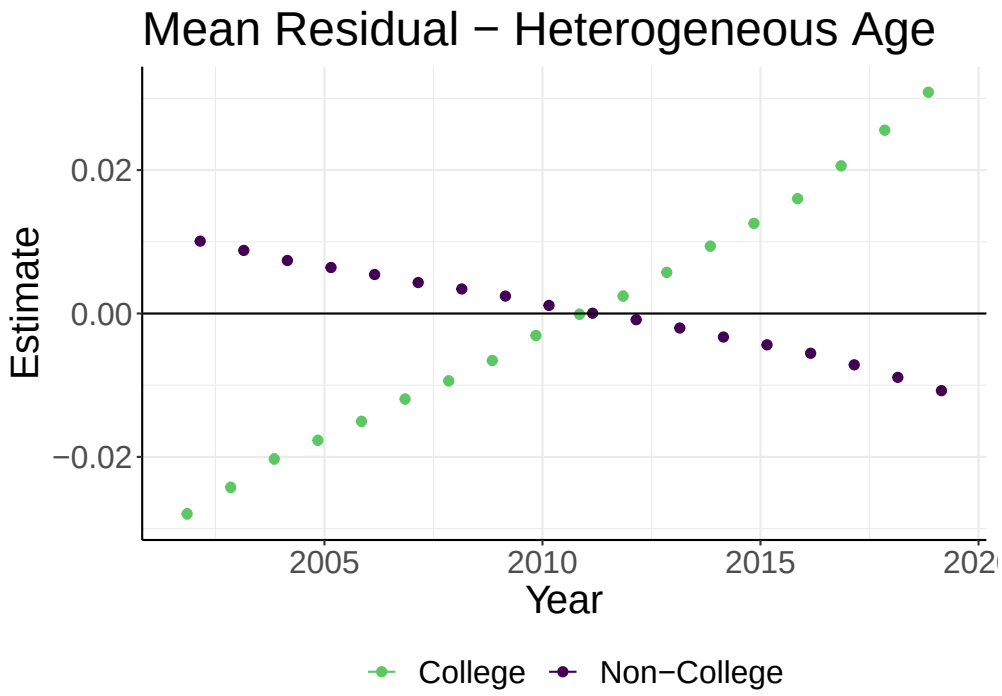


Figure B14: Average residual by year for college and non-college workers. Simulated outcomes have nonparametric age fits by education and sex.

C Additional Results

C.1 Event Study Graphs

In Figures C2 through C5 I show changes in average firm effects and average residuals when workers switch firms. I include all workers employed at a single firm from 2012 to 2013 who then move in 2014 and remain at a new firm from 2015 to 2017. In Figure C2 I show results for all workers. In Figure C3 I show results for college workers. In Figure C4 I show results for workers without a college degree. In Figure C5 I show results for workers employed in personal security, food preparation, and custodial work.

C.2 Comparing the CPS MORG to the LEHD-ACS

I use the CPS MORG data as a point of comparison for log hourly wages in my sample. I adjust wages using the PCE deflator and then limit to individuals with hourly wages between \$1 and \$100 in 1979 dollars. I drop observations with imputed values of earning or hours worked. In Figure C6 I show similar differences in average log wages for college and non-college workers in the CPS MORG data and in my main LEHD-ACS sample. In Table C1 I show the age composition of workers in the CPS MORG data over time.

C.3 Residual Trends

Figure C7 shows average residuals over time for college and non-college workers.

C.4 Including all workers in the ACS

In this section I redo my main analysis including workers in the ACS who work part time or for multiple firms in a year. For workers employed at multiple firms, I compute the pay premium as the earnings-weighted average premium across all jobs. Figure C12 shows results when including people who work at multiple firms but excluding people who fall below the

minimum earnings threshold, showing nearly identical results to my main estimation.

C.4.1 Age-Earnings Profiles

In Figure C13 I plot earnings as a function of age using the matched LEHD-ACS data, including workers who are part time or work for multiple firms in a year. Figure C14 shows hourly wages as a function of age.

C.5 Distributional Changes for Person Skill

Figure C15 shows population growth by decile of the person fixed effects, ζ_i. These results warrant particular caution because of potential changes in estimation error for the person effect. An alternative sampling scheme to test for changes in errors for the person effects would randomly sample matches instead of workers (Kline et al., 2020; Kline, 2024). I do not pursue that approach here.

C.6 Alternative Estimations

In this section I consider alternative ways to estimate average skill changes. Figure C16 averages firm effects across Samples A and B wheh computing average total skill. Figure C17 bins by sex, age, state, and year when controlling for hours in TV-AKM estimation, excluding imputed education. Figure C18 controls for a fifth order polynomial in age in the TV-AKM estimation. Figure C19 controls for a fifth order polynomial in age separately by education and sex. All of these alternative estimations show similar quantitative patterns to my main results.

C.7 Estimates in the Midwest Sample

I show my main results estimated only on the Midwest sample. Figure C20 shows results for average total skill for all workers, not taking a random sample before the TV-AKM esti-

mation. Figure C21 takes a 1/3 random sample of workers before the TV-AKM estimation. Figure C22 estimates the model on a sandwich sample. This drops workers who do not work at the firm in Q4 of the previous year or Q1 of the following year. I find similar results to my main analysis. Figures C23 through C25 show results for average person skill.

C.8 Separate Estimation for College and Non-College

In Figure C26 I show results for average total skill when estimating separate TV-AKM models for college and non-college workers in the Midwest Sample. For comparison I plot it side-by-side with the results from estimating a single TV-AKM model across education groups, equivalent to Figure C20. Results are comparable. Figures C27 show results for average person skill.

Estimating separate models for college and non-college workers changes the size of the largest connected set. The set of worker-year observations when estimating separate models is a subset of the worker-year observations with one joint model. In each year the total number of observations with the separate models is 95% or greater of the total number of observations with the joint model, indicating only a small change in the sample. This shows a large amount of overlap in the worker-firm graphs for college and non-college workers.

C.9 Post-College

In Figures C28 through C31 I show results that split college workers between people with college degrees only and people with post-college schooling.

C.10 Gender and Race

In Figures C32 through C39 I show results separately by for men and women and for White and Black workers.

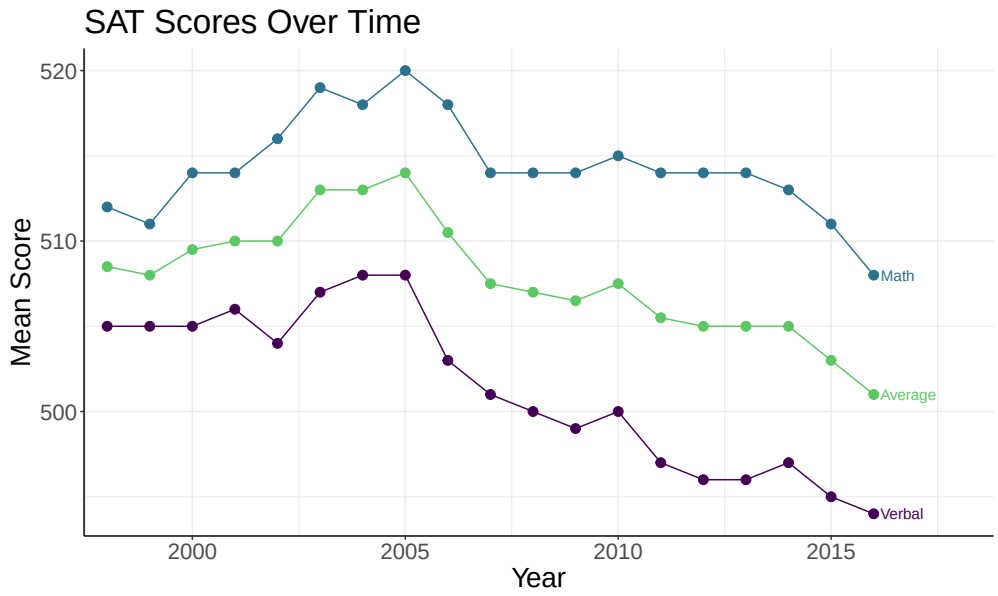


Figure C1: Average SAT scores over time. Data comes from College Board (2016).

C.11 Older and Younger Workers

In Figures C40 through C45 I show results separately for workers aged 25-40 and workers aged 41-60.

C.12 Estimation Error Over Time

I compute different estimates for each ξ_{jt} using the two random samples of the data described in Section 4. To test whether changes in estimation error over time bias my estimated skill deciles in Section 5.2 I compute quantiles of the absolute differences between my two estimates in 2002 and 2019. Figure C46 shows that the distribution of estimation error is stable over time.

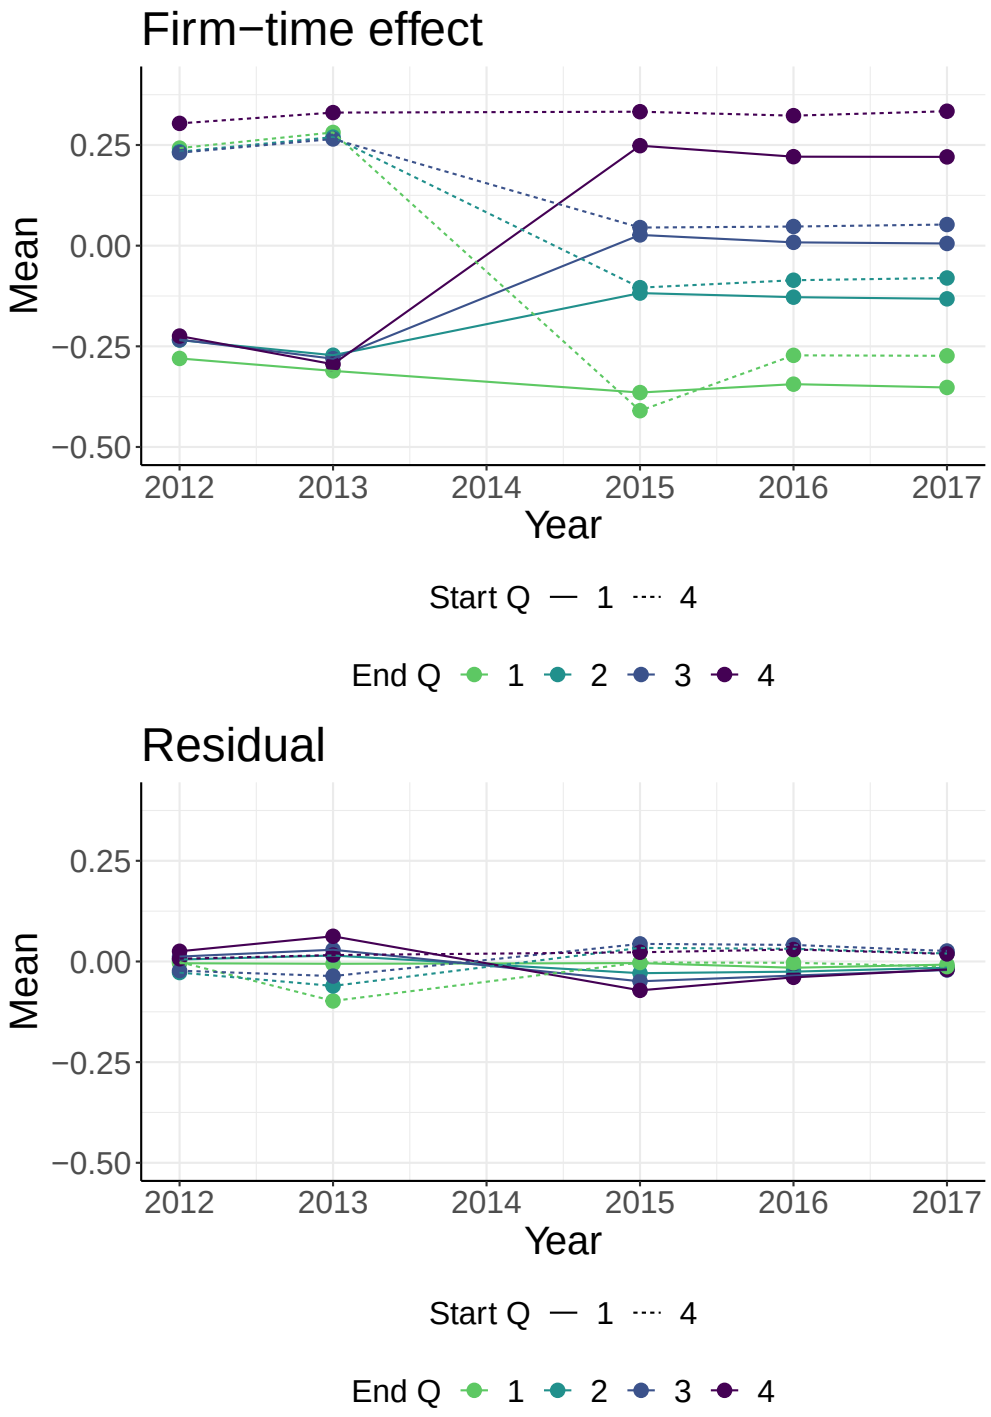


Figure C2: Changes in average firm effects and average residuals when workers switch firms. Start Q denotes the origin firm's pay premium quartile in 2013. End Q denotes the ending firm's pay premium quartile in 2015. Including all workers employed at a single firm from 2012 to 2013 who then move in 2014 and remain at a new firm from 2015 to 2017.

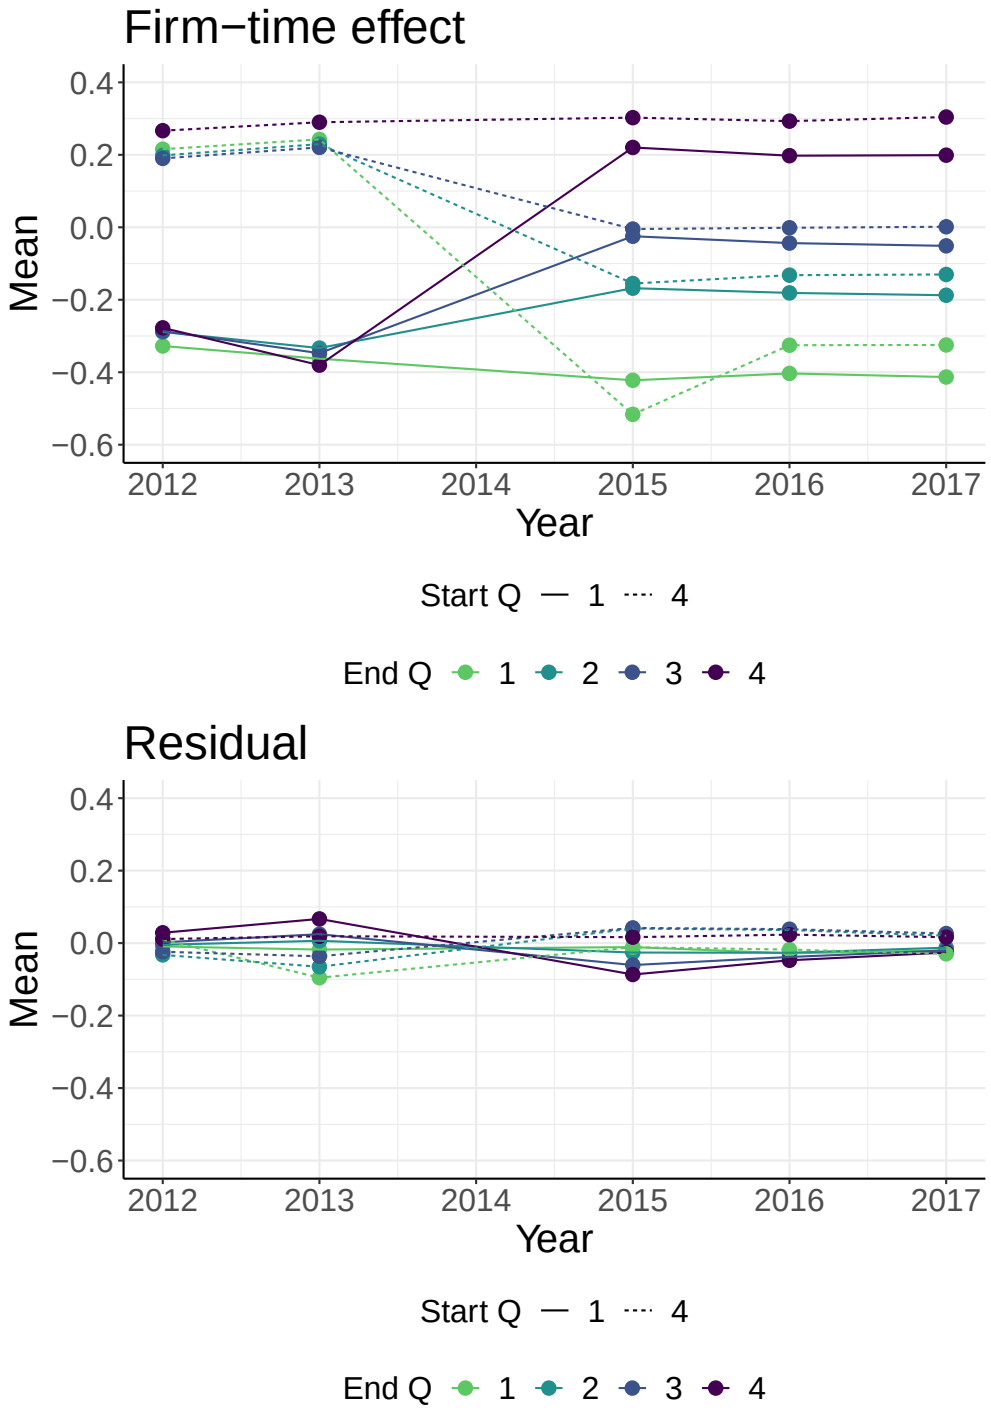


Figure C3: Changes in average firm effects and average residuals when workers switch firms.Start Q denotes the origin firm's pay premium quartile in 2013. End Q denotes the ending firm's pay premium quartile in 2015. Including all college workers employed at a single firm from 2012 to 2013 who then move in 2014 and remain at a new firm from 2015 to 2017. Using imputed college status from the LEHD.

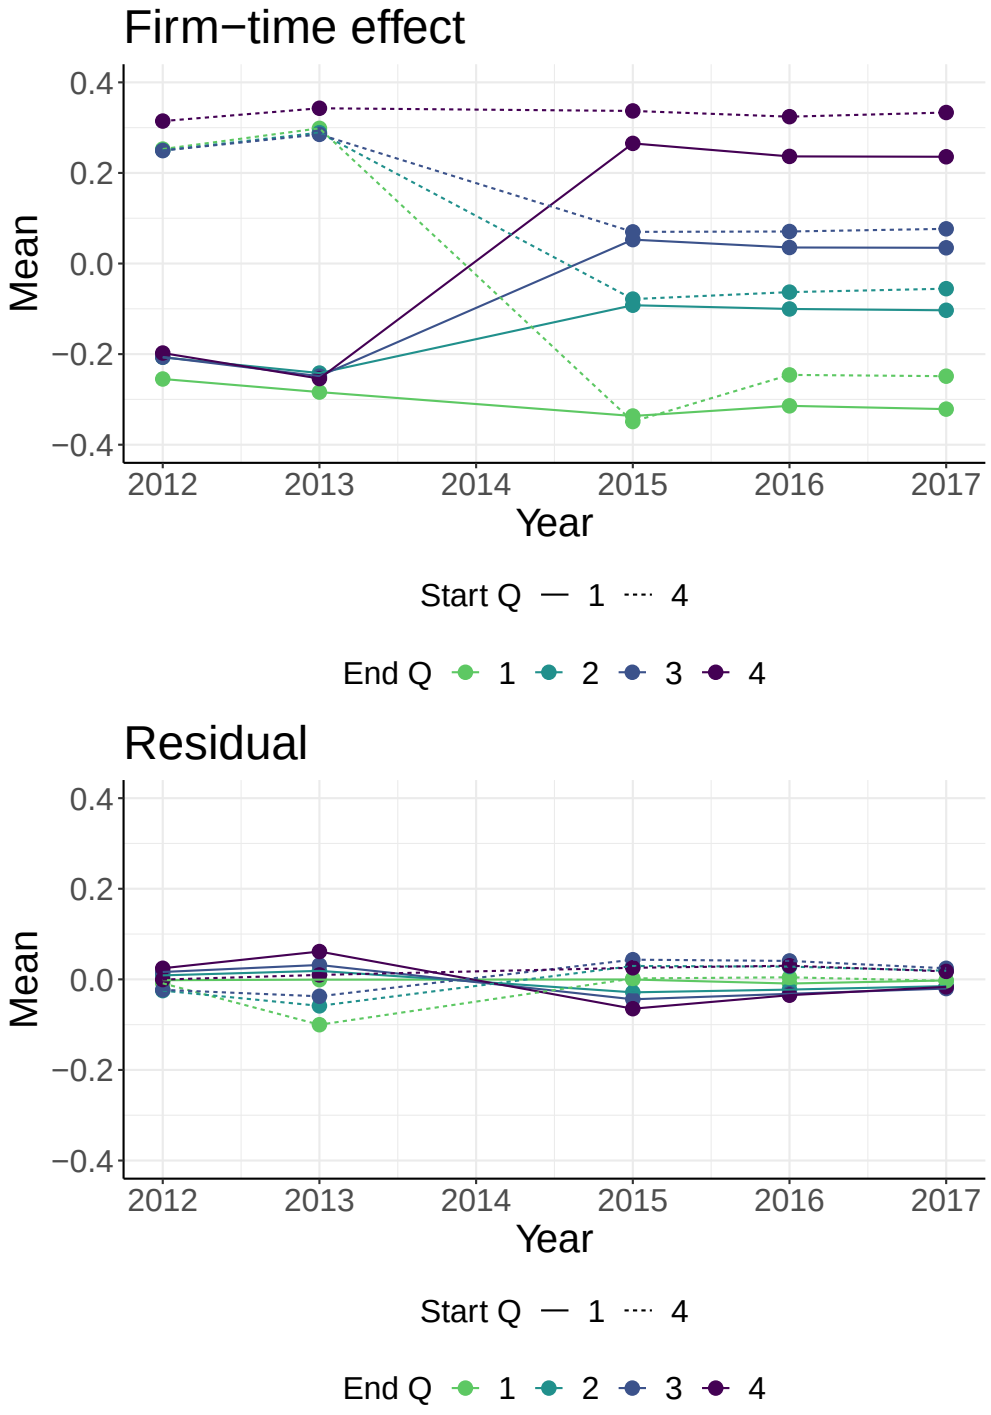


Figure C4: Changes in average firm effects and average residuals when workers switch firms. Start Q denotes the origin firm's pay premium quartile in 2013. End Q denotes the ending firm's pay premium quartile in 2015. Including all workers without a college degree employed at a single firm from 2012 to 2013 who then move in 2014 and remain at a new firm from 2015 to 2017. Using imputed college status from the LEHD.

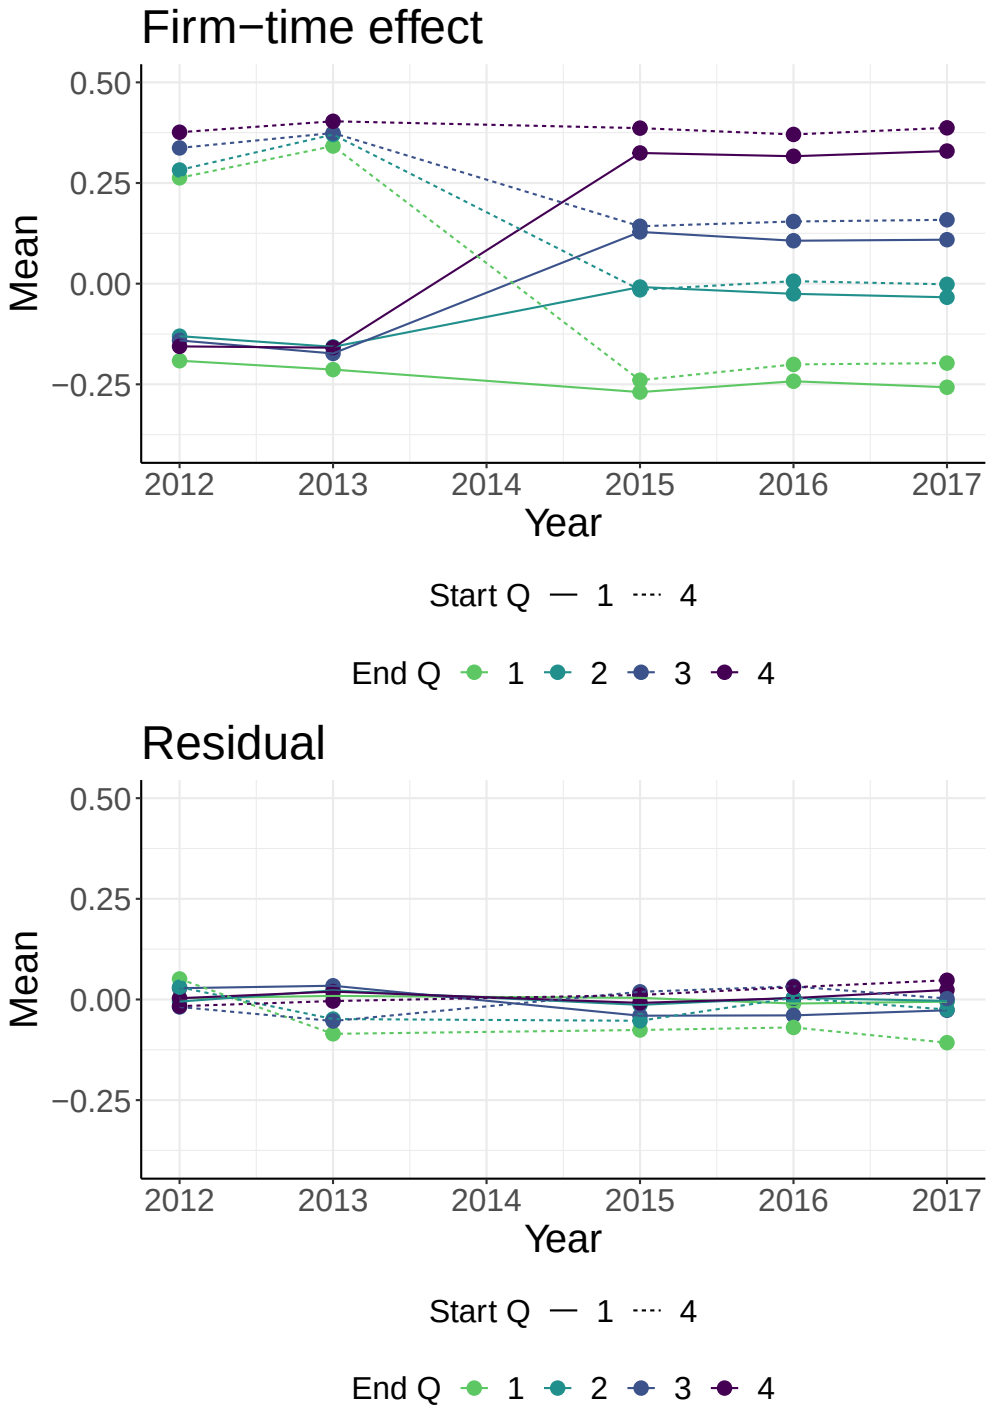


Figure C5: Changes in average firm effects and average residuals when workers switch firms. Start Q denotes the origin firm's pay premium quartile in 2013. End Q denotes the ending firm's pay premium quartile in 2015. Including all workers employed in personal security, food preparation, and custodial work employed at a single firm from 2012 to 2013 who then move in 2014 and remain at a new firm from 2015 to 2017.

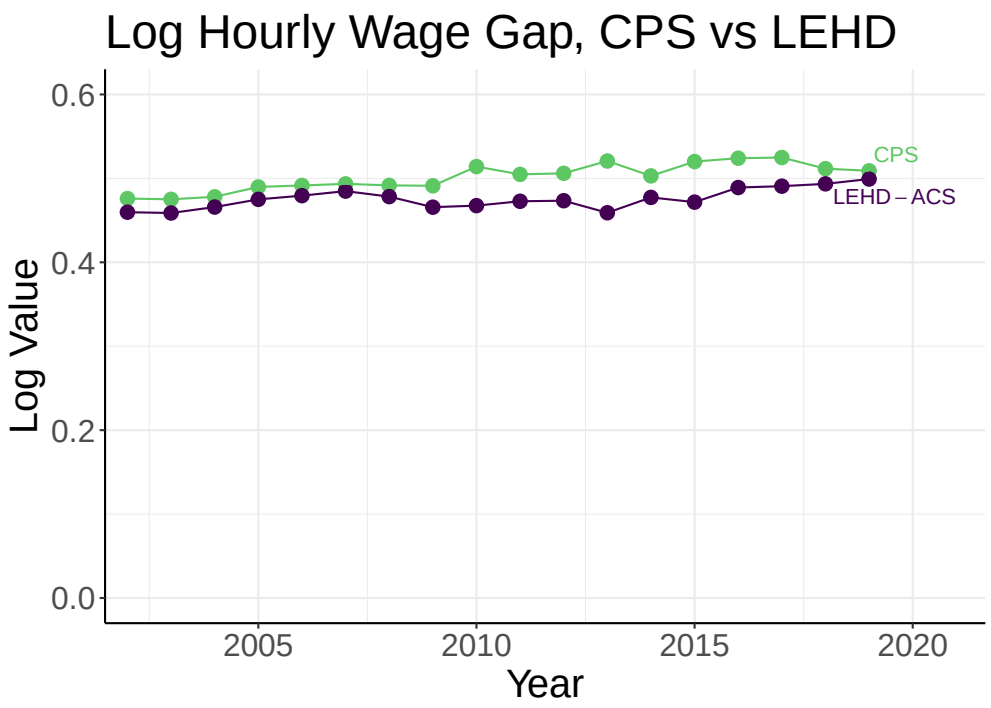


Figure C6: The difference in average log wages for college and non-college workers in the CPS MORG data and in my main LEHD-ACS sample.

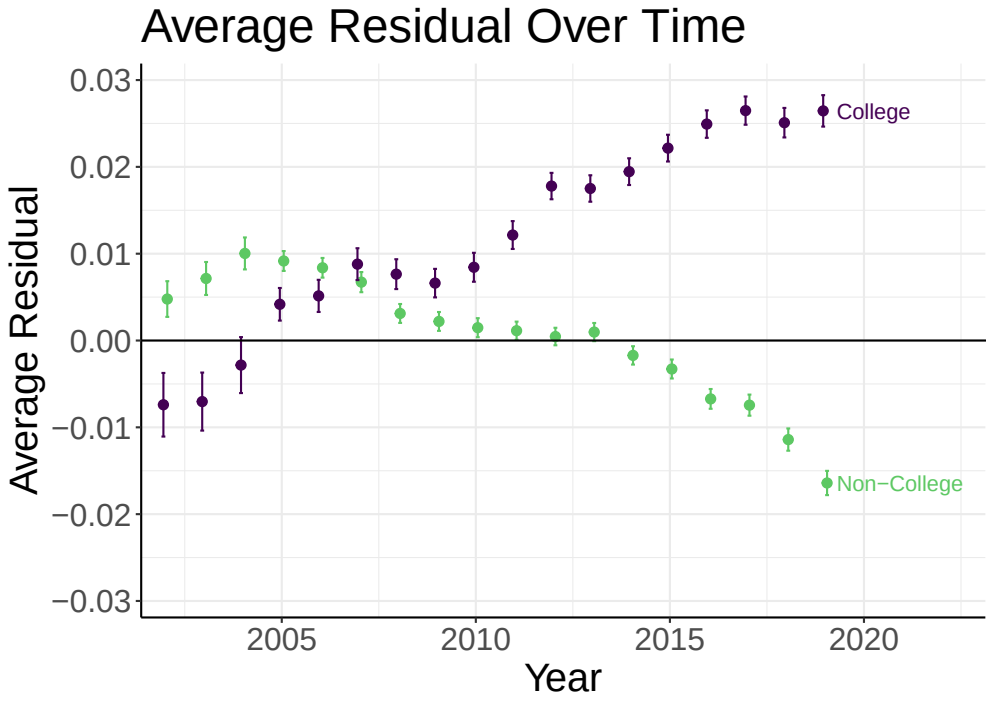


Figure C7: Mean residuals from the TV-AKM estimation for college and non-college workers in the LEHD-ACS data.

Table C1: Age composition over time in the CPS MORG data.

Age Composition Over Time in CPS MORG

Year	Mean Age	% 25-30	% 31-40	% 41-50	% 51-60
2002	41.47	15.98	30.80	32.60	20.63
2003	41.72	15.32	30.29	32.82	21.56
2004	41.87	15.52	29.51	32.50	22.46
2005	41.87	16.08	28.98	32.23	22.72
2006	41.99	16.56	28.07	31.66	23.71
2007	42.12	16.91	27.38	31.31	24.40
2008	42.26	16.56	27.25	30.85	25.35
2009	42.34	16.37	27.40	30.50	25.74
2010	42.20	17.00	27.64	29.64	25.72
2011	42.22	17.10	27.55	29.31	26.05
2012	42.27	16.90	27.63	28.82	26.65
2013	42.38	16.32	27.99	28.59	27.10
2014	42.31	16.75	28.30	27.58	27.36
2015	42.19	16.86	28.88	27.49	26.77
2016	42.12	17.30	28.94	26.84	26.92
2017	42.09	17.29	29.36	26.38	26.97
2018	42.04	17.22	29.71	26.90	26.16
2019	41.99	17.51	29.77	26.67	26.05

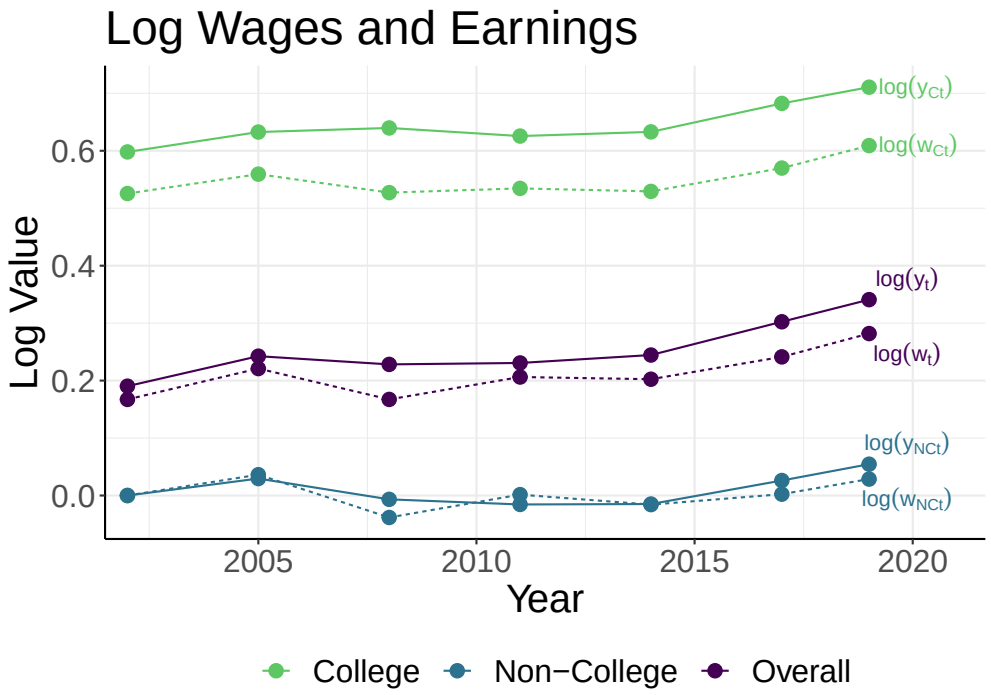


Figure C8: Mean log earnings and hourly wages for college and non-college workers in the LEHD-ACS data. Including workers in the ACS who work for multiple firms and fall under the minimum earnings threshold.

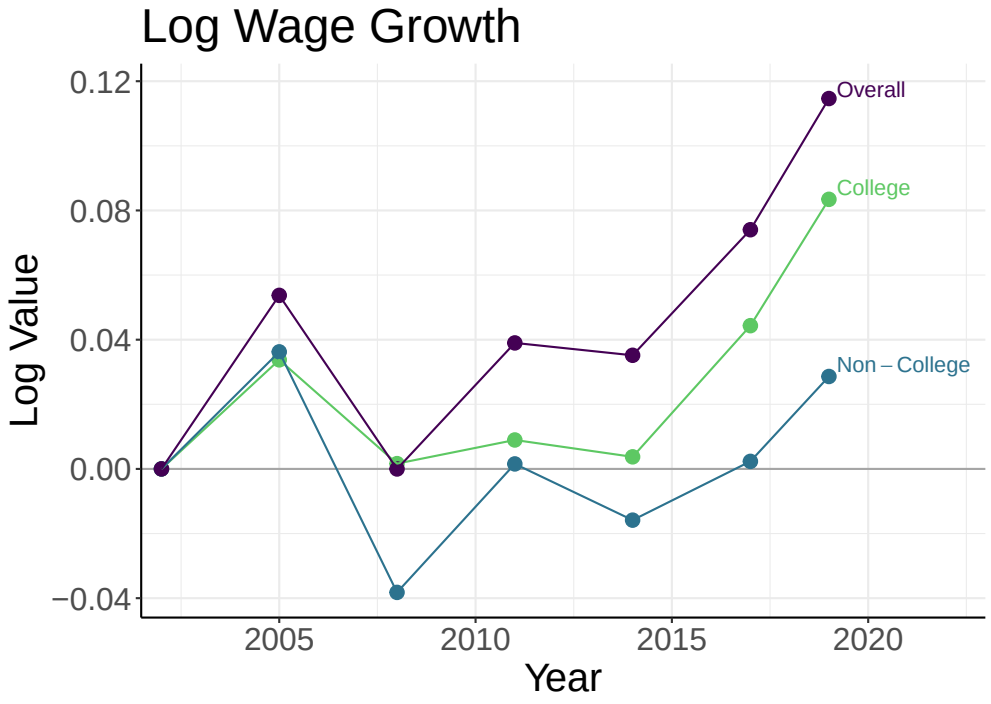


Figure C9: Growth in mean log hourly wages by education in the LEHD-ACS data. Including workers in the ACS who work for multiple firms and fall under the minimum earnings threshold.

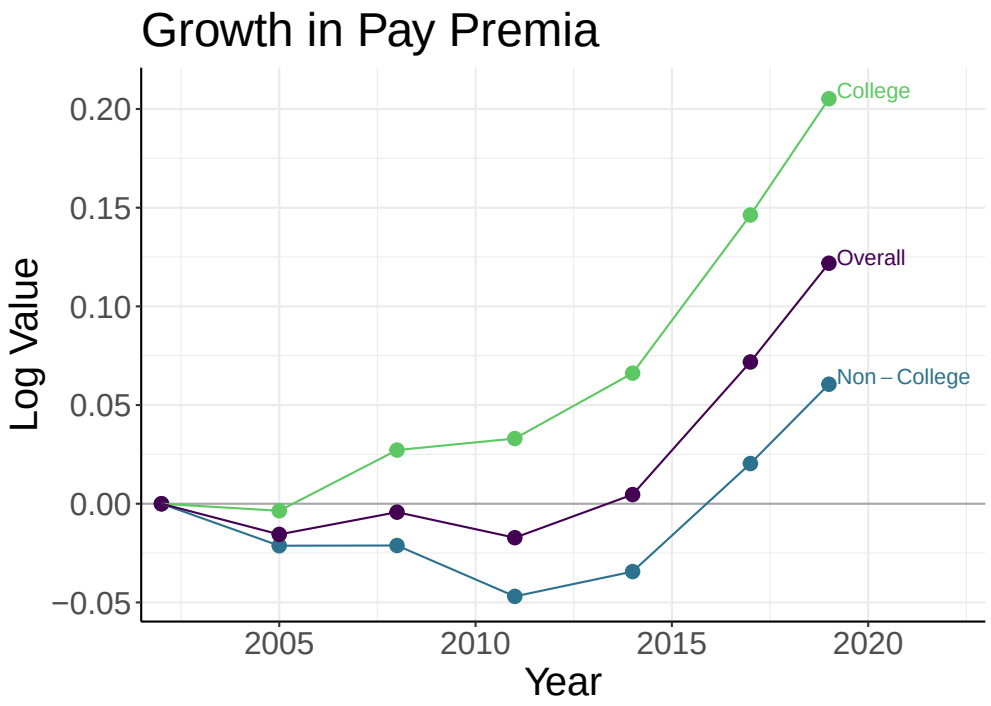


Figure C10: Growth in mean pay premium for the LEHD-ACS data by education. Including workers in the ACS who work for multiple firms and fall under the minimum earnings threshold.

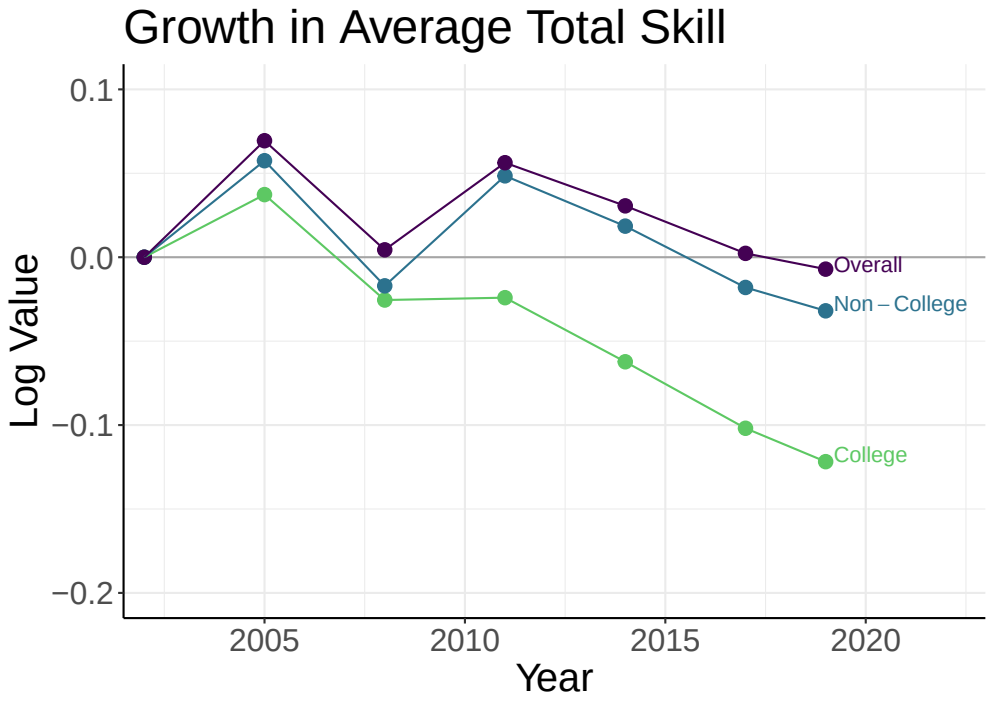


Figure C11: Growth in average total skill by education. Including workers in the ACS who work for multiple firms and fall under the minimum earnings threshold.

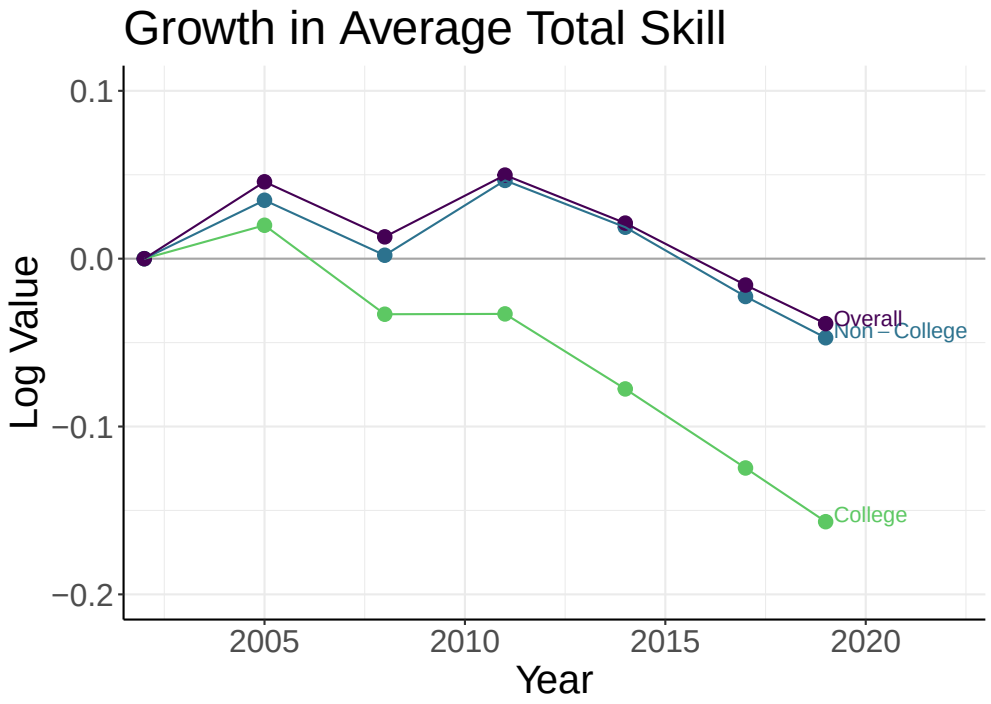


Figure C12: Growth in average total skill by education. Including workers in the ACS who fall under the minimum earnings threshold.

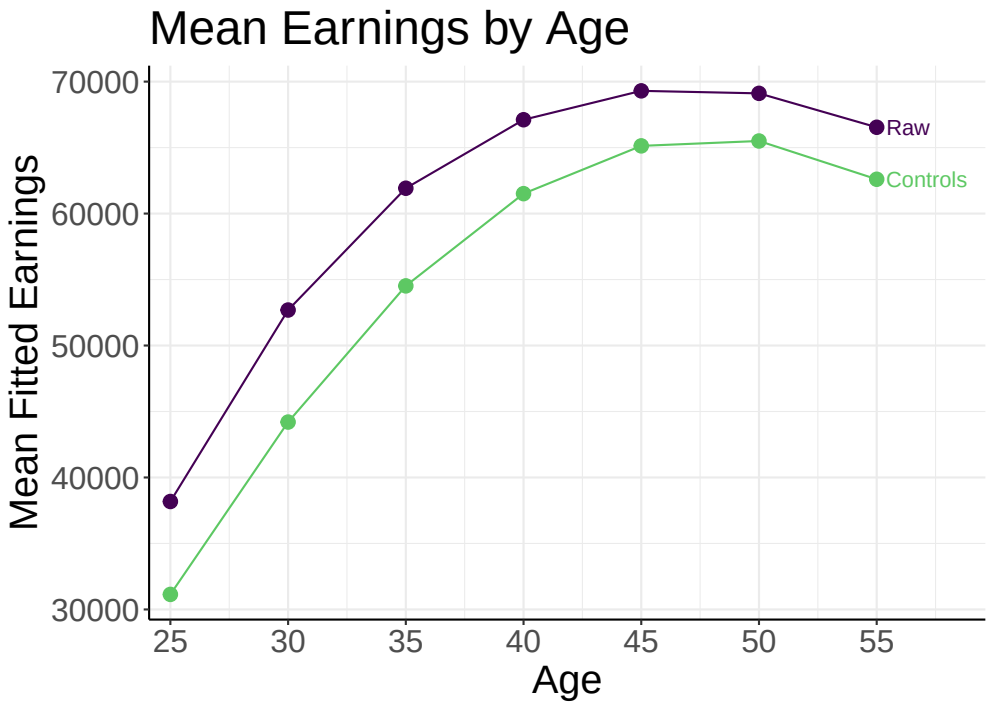


Figure C13: Earnings as a function of age. Controls include sex, education, and year from the ACS. I estimate the relationship nonparametrically with 5-year age bins (ages 25-30, 30-35, 35-40, etc.).

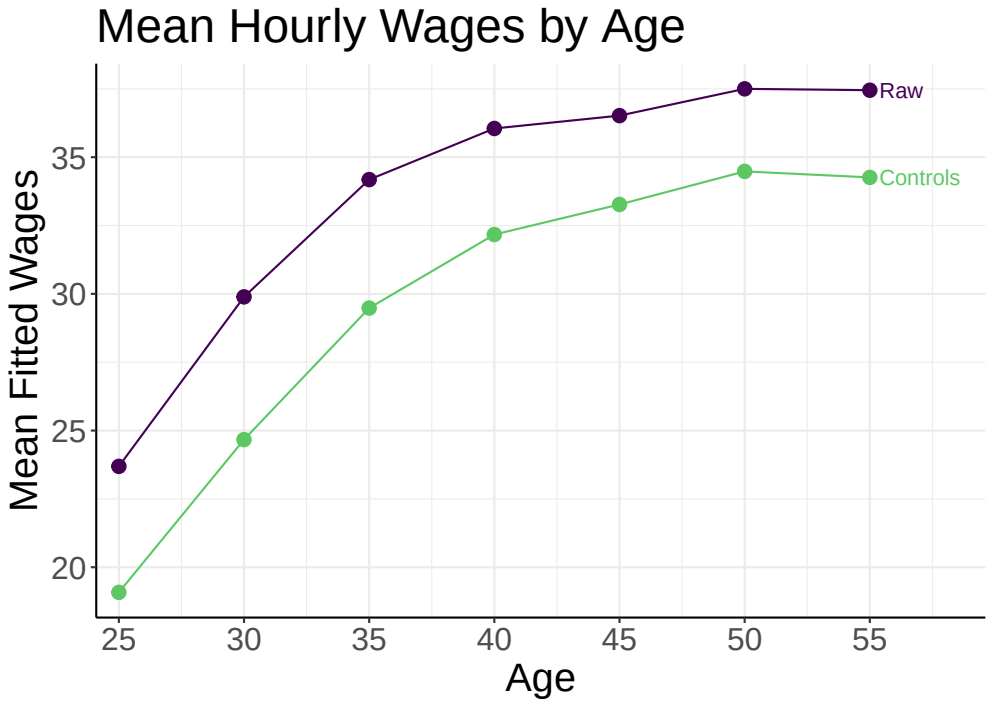


Figure C14: Hourly wages as a function of age. Controls include sex, education, and year from the ACS. I estimate the relationship nonparametrically with 5-year age bins (ages 25-30, 30-35, 35-40, etc.)

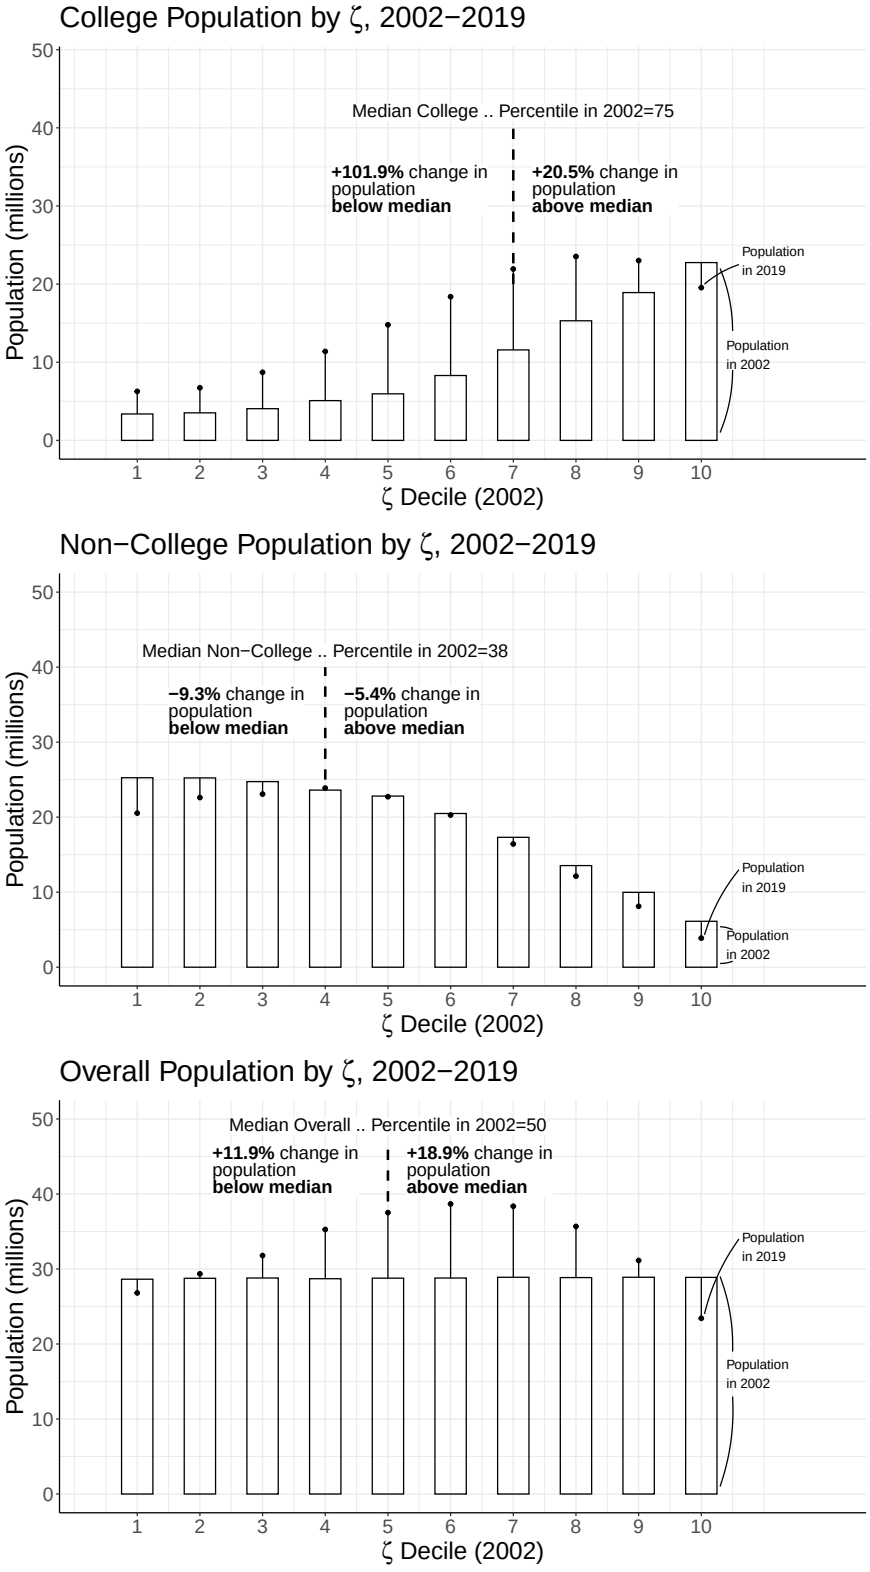


Figure C15: Changes in population by decile of ζ in 2002 for college workers, non-college workers, and all workers.

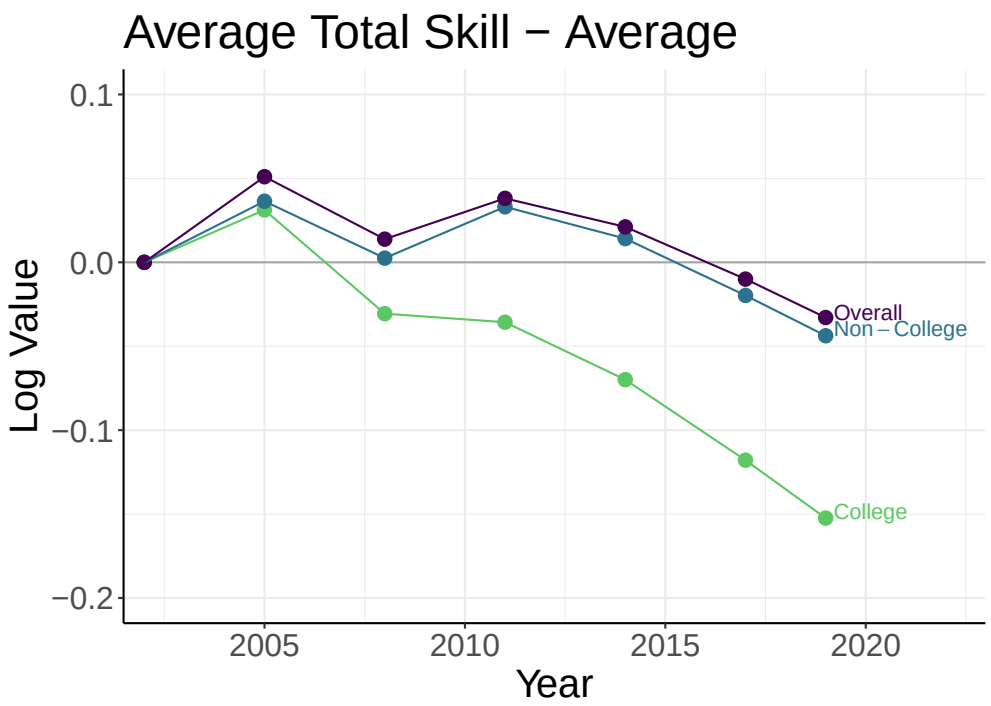


Figure C16: Growth in average total skill by education. Averaging across pay premia from Sample A and Sample B.

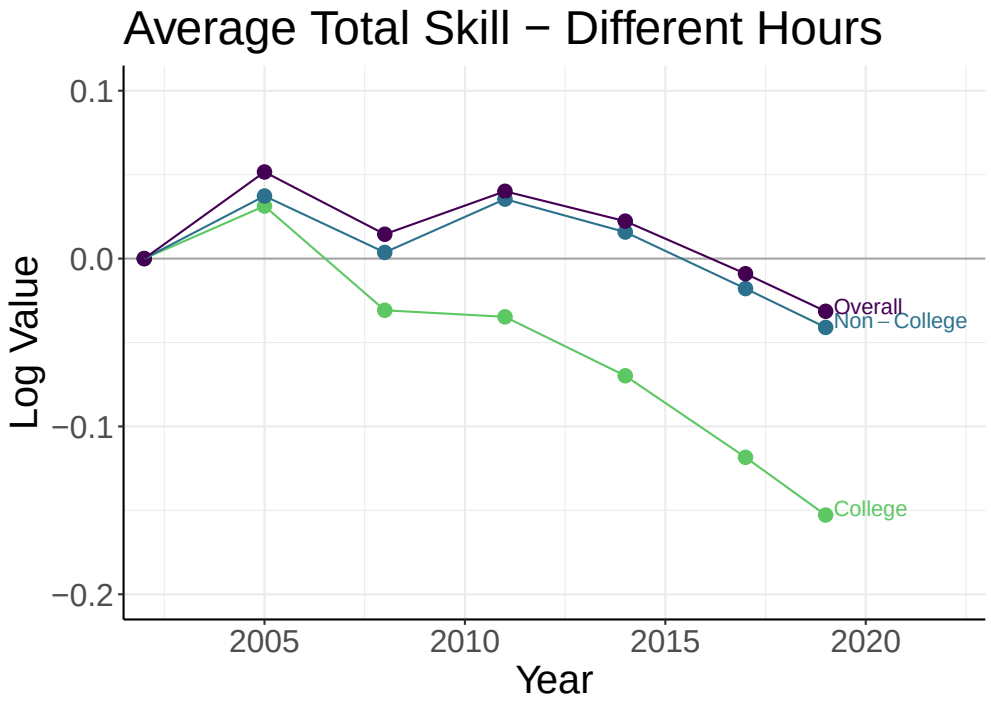


Figure C17: Growth in average total skill by education. Binning by sex, age, state, and year when controlling for hours in TV-AKM estimation.

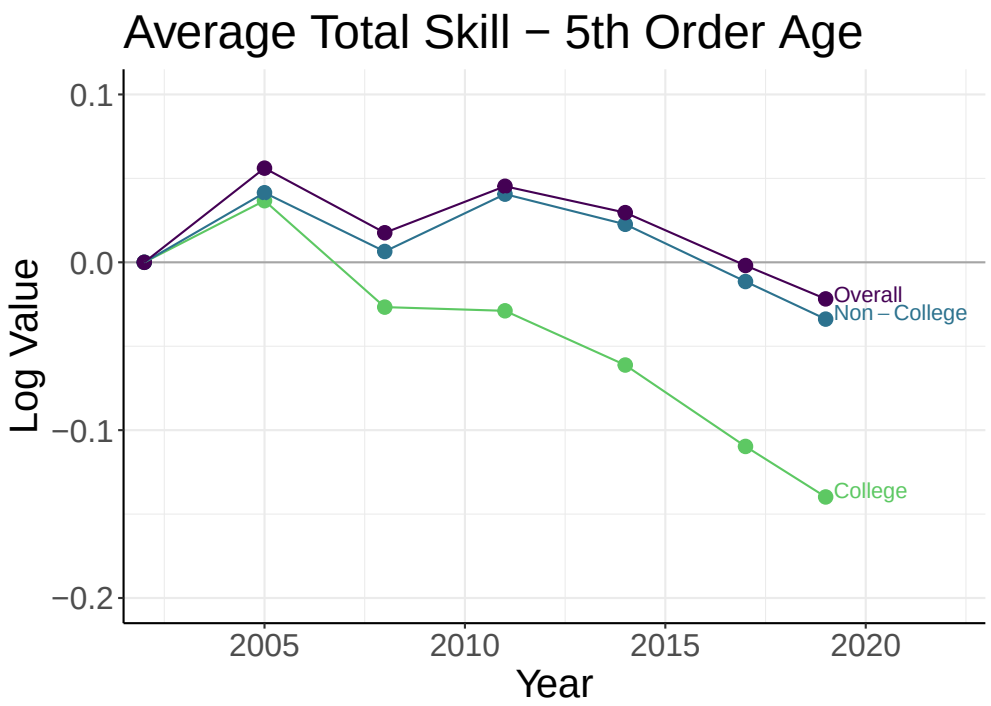


Figure C18: Growth in average total skill by education. Controlling for a fifth order polynomial in age in the TV-AKM estimation.

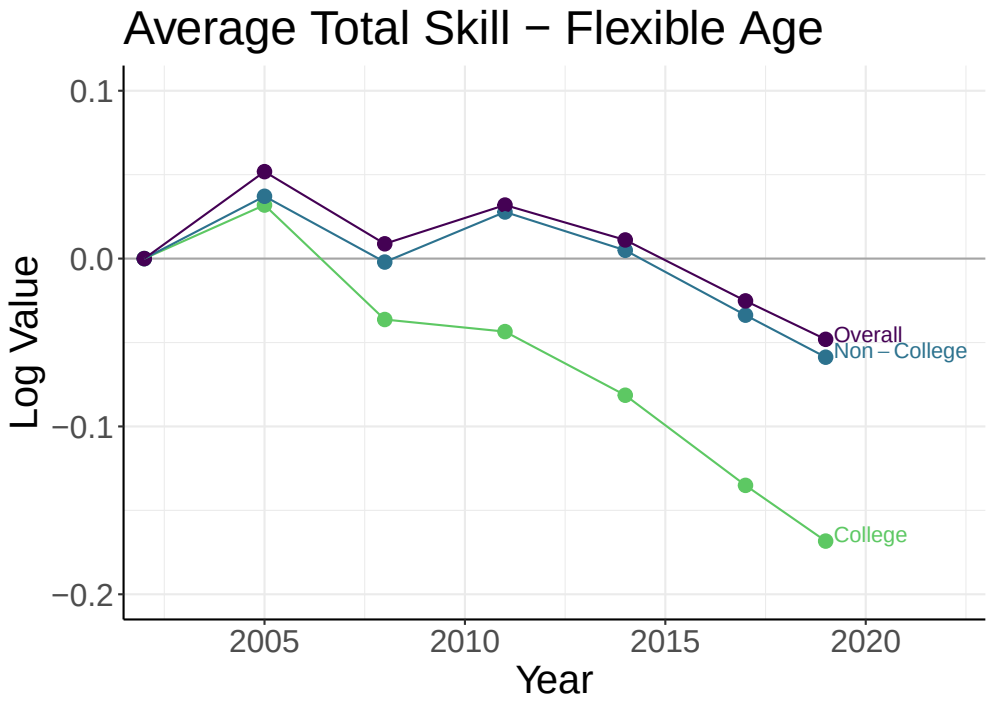


Figure C19: Growth in average total skill by education. Controlling for a fifth order polynomial in age separely by education and sex in the TV-AKM estimation.

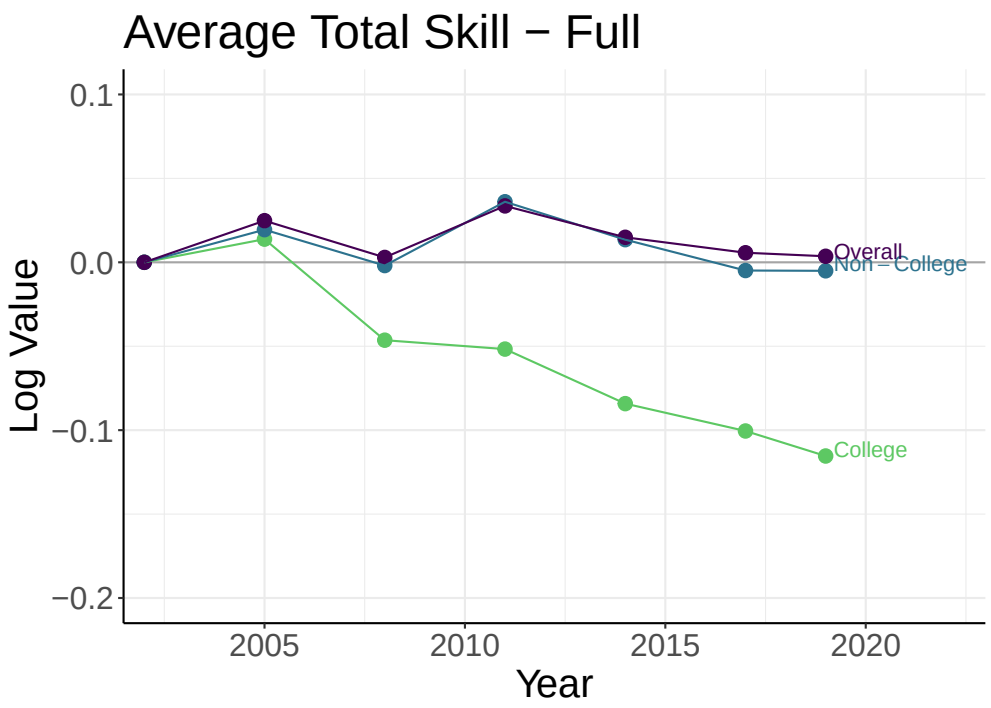


Figure C20: Growth in average total skill by education in the Midwest Sample. Including all workers.

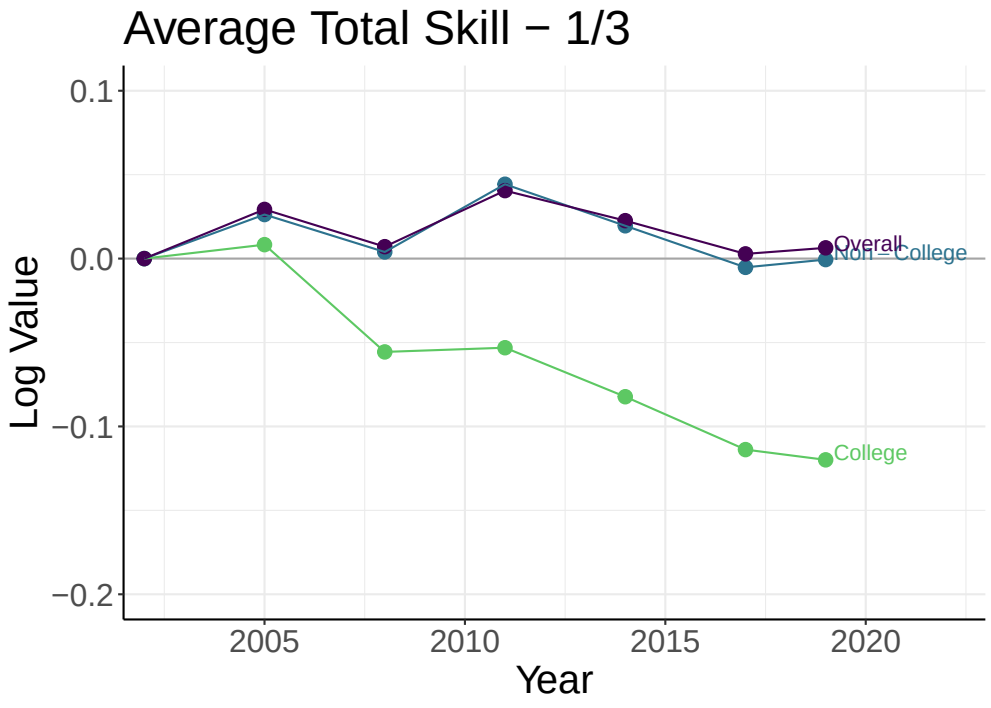


Figure C21: Growth in average total skill by education in the Midwest Sample. Including a 1/3 random sample of workers.

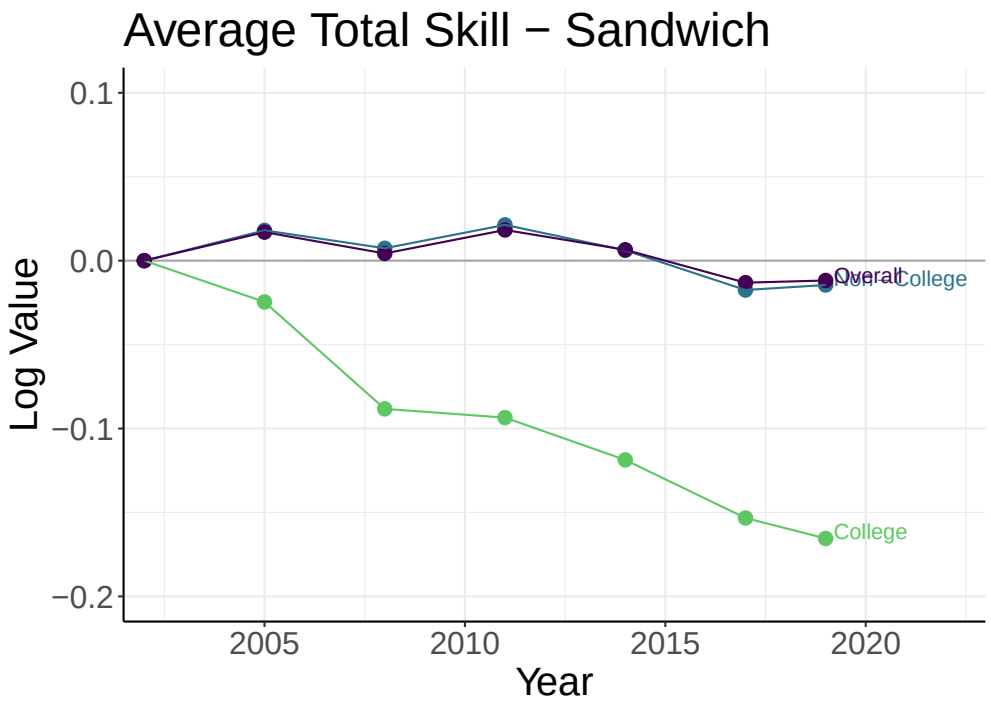


Figure C22: Growth in average total skill by education in the Midwest Sample. Including a sandwich sample of 1/3 of workers. The sandwich sample drops workers who do not work at the firm in Q4 of the previous year or Q1 of the following year.

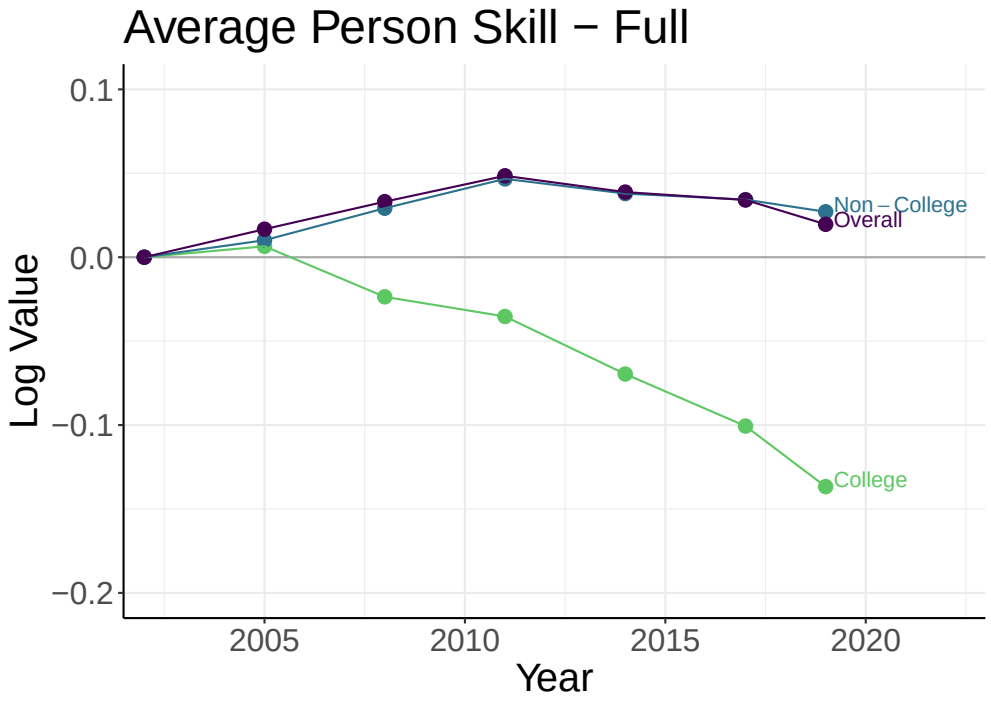


Figure C23: Growth in average person skill by education in the Midwest Sample. Including all workers.

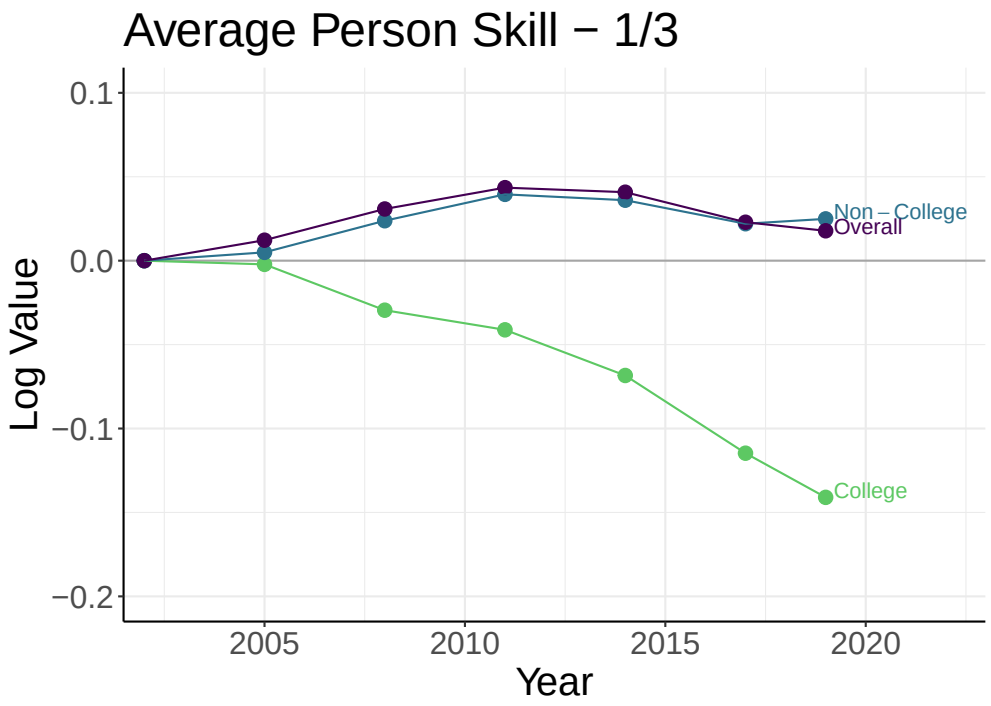


Figure C24: Growth in average person skill by education in the Midwest Sample. Including a 1/3 random sample of workers.

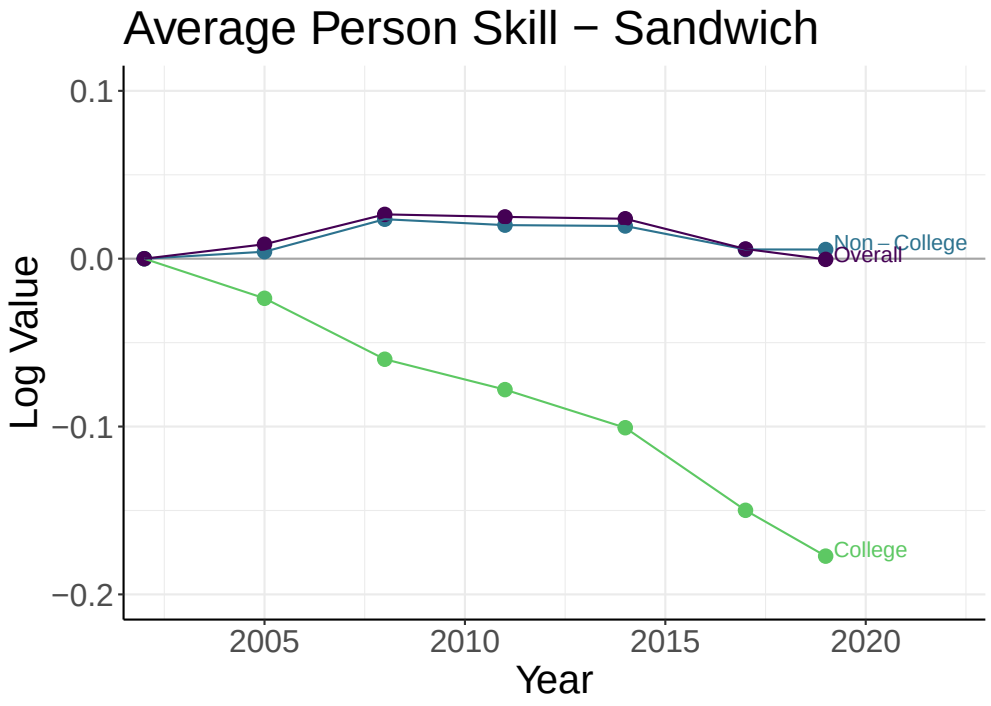


Figure C25: Growth in average person skill by education in the Midwest Sample. Including a sandwich sample of 1/3 of workers. The sandwich sample drops workers who do not work at the firm in Q4 of the previous year or Q1 of the following year.

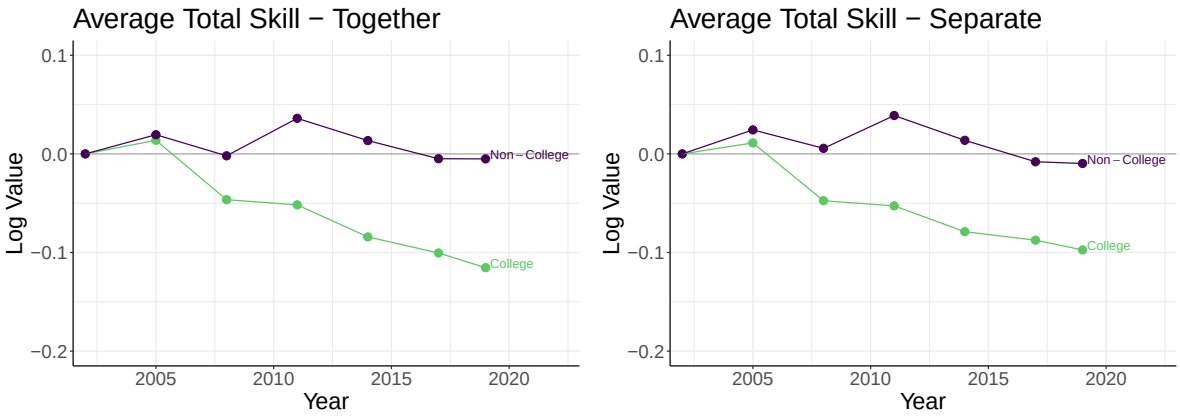


Figure C26: Left: Growth in average total skill by education in the Midwest Sample when estimating a single TV-AKM model across education groups. Including all workers. Equivalent to Figure C20. Right: Growth in average total skill by education in the Midwest Sample when estimating separate TV-AKM models for college and non-college workers.

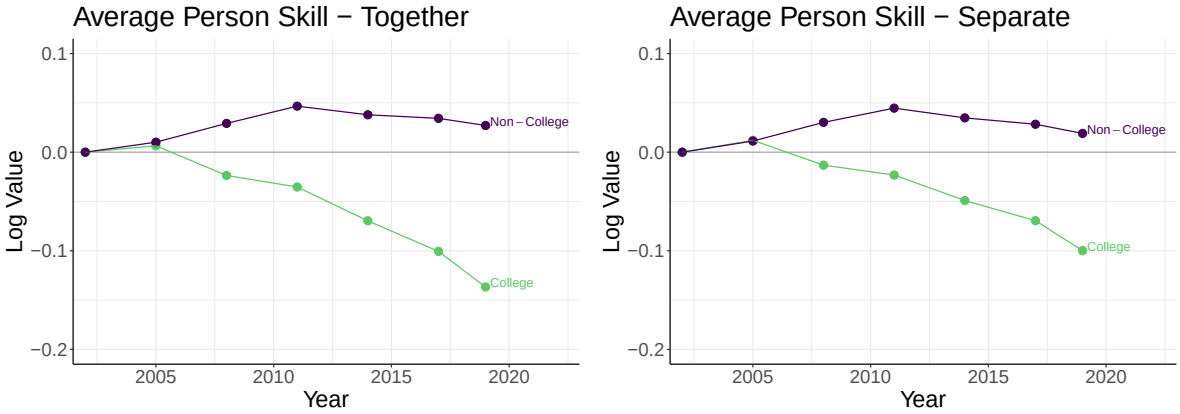


Figure C27: Left: Growth in average person skill by education in the Midwest Sample when estimating a single TV-AKM model across education groups. Including all workers. Equivalent to Figure C23. Right: Growth in average person skill by education in the Midwest Sample when estimating separate TV-AKM models for college and non-college workers.

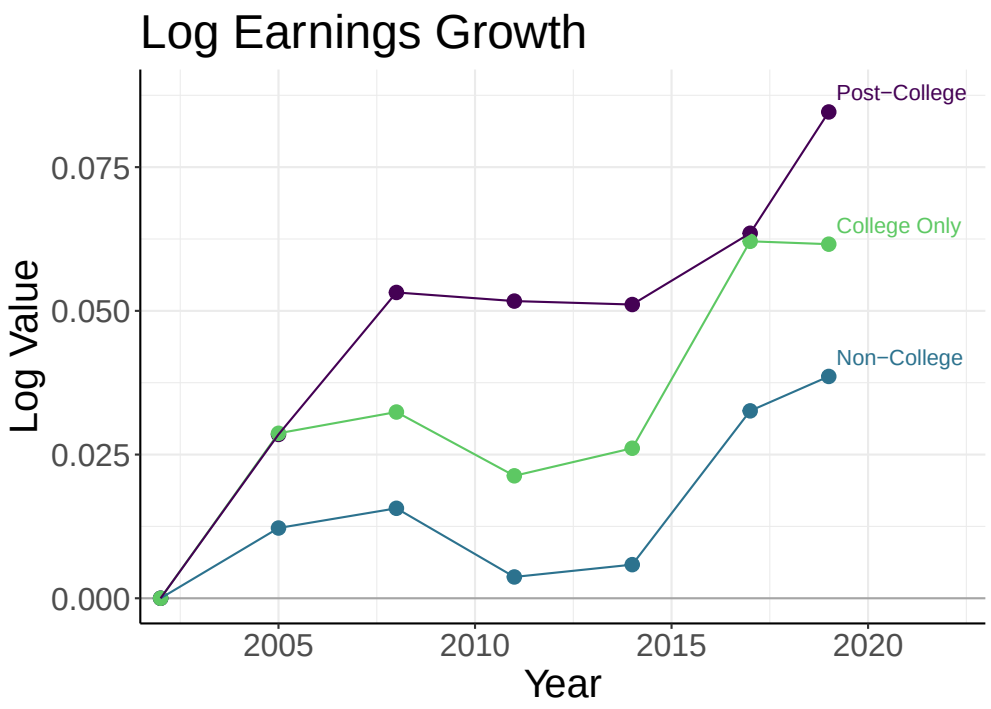


Figure C28: Growth in mean earnings by education.

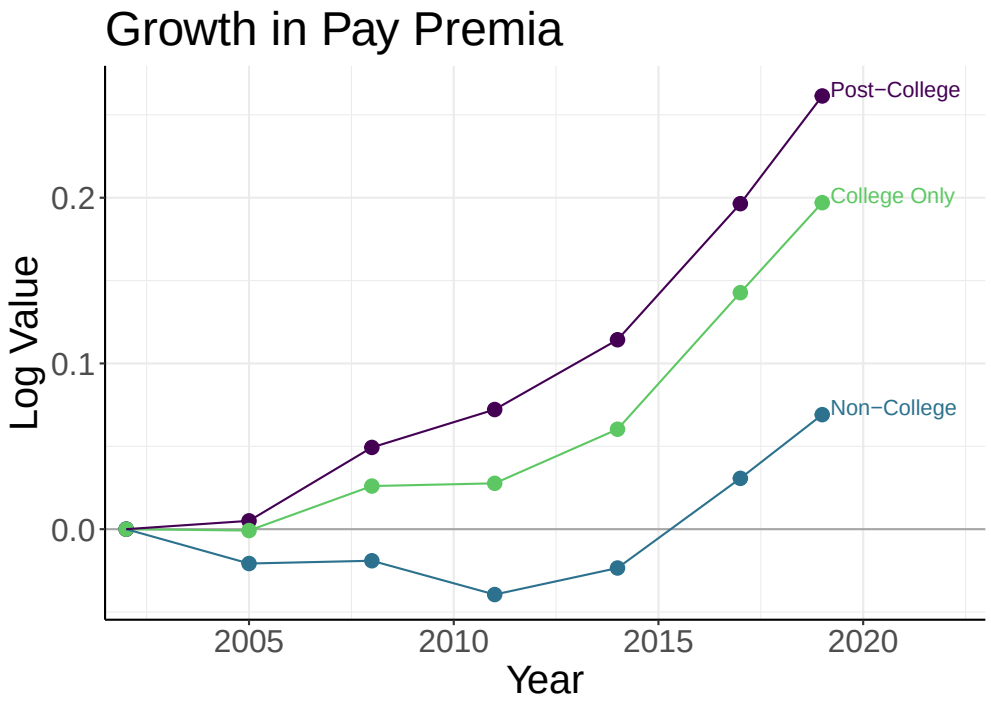


Figure C29: Growth in mean pay premium by education.

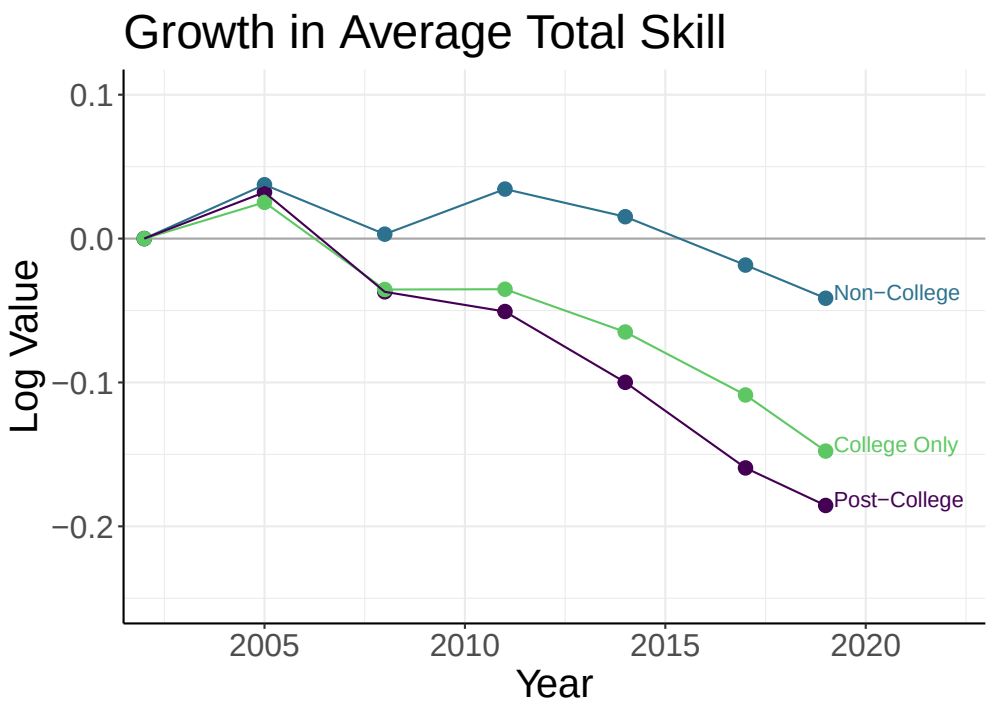


Figure C30: Growth in average total skill by education.

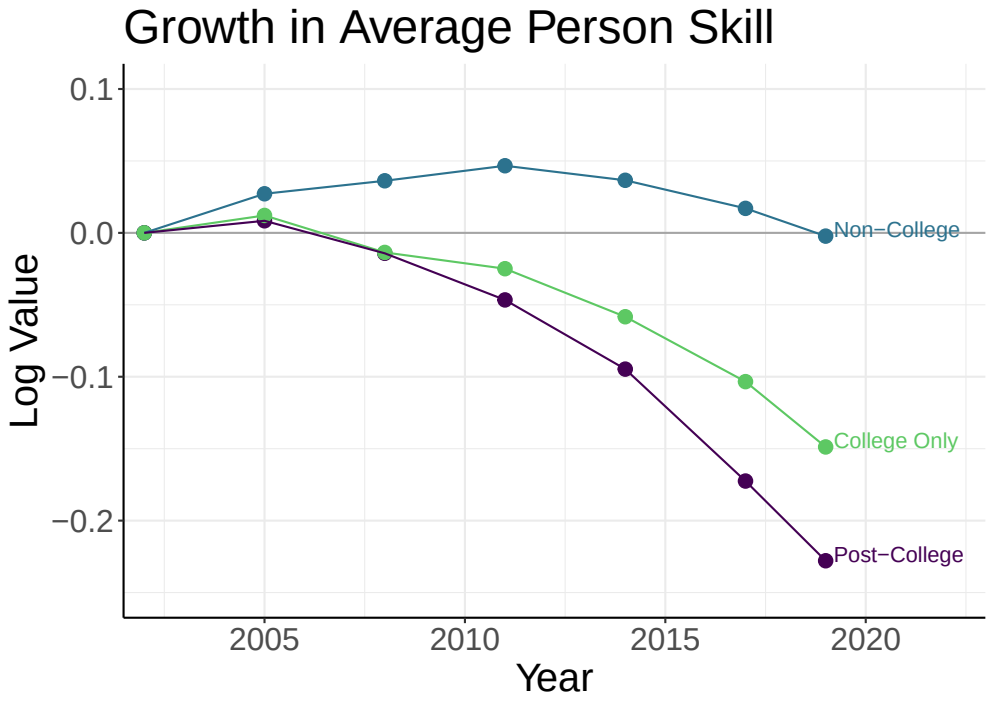


Figure C31: Growth in average person skill by education.

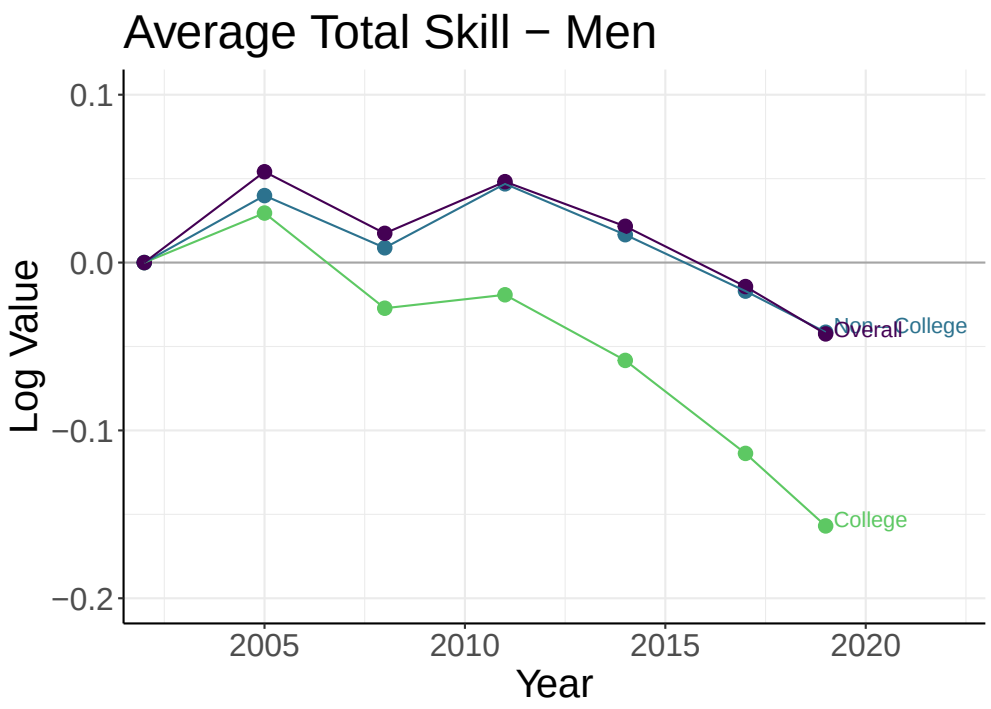


Figure C32: Growth in average total skill by education, for men.

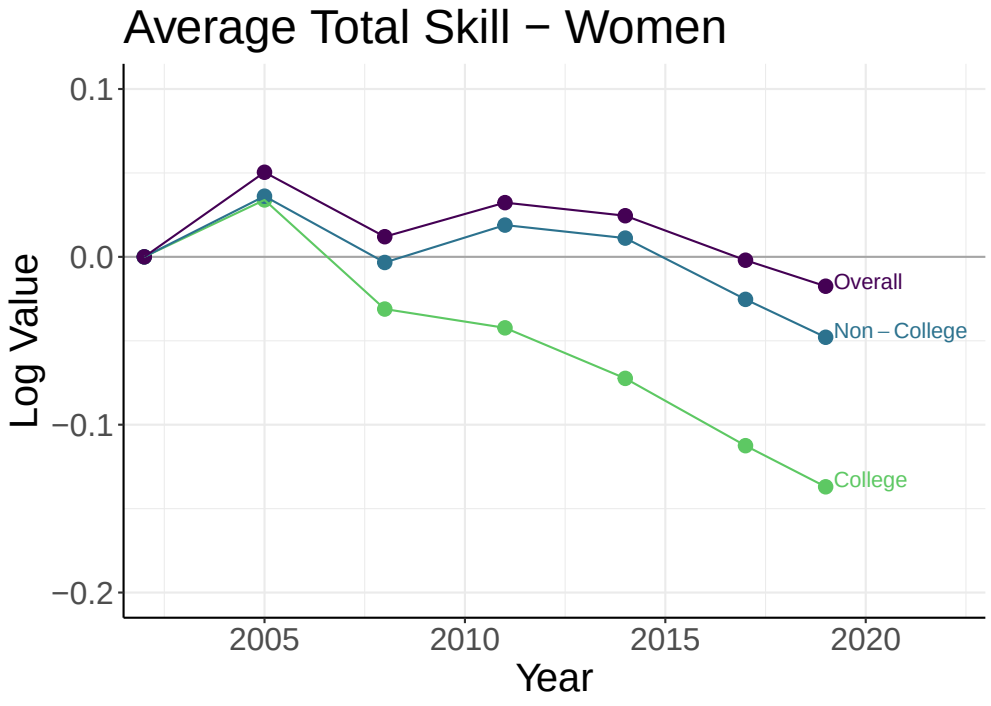


Figure C33: Growth in average total skill by education, for women.

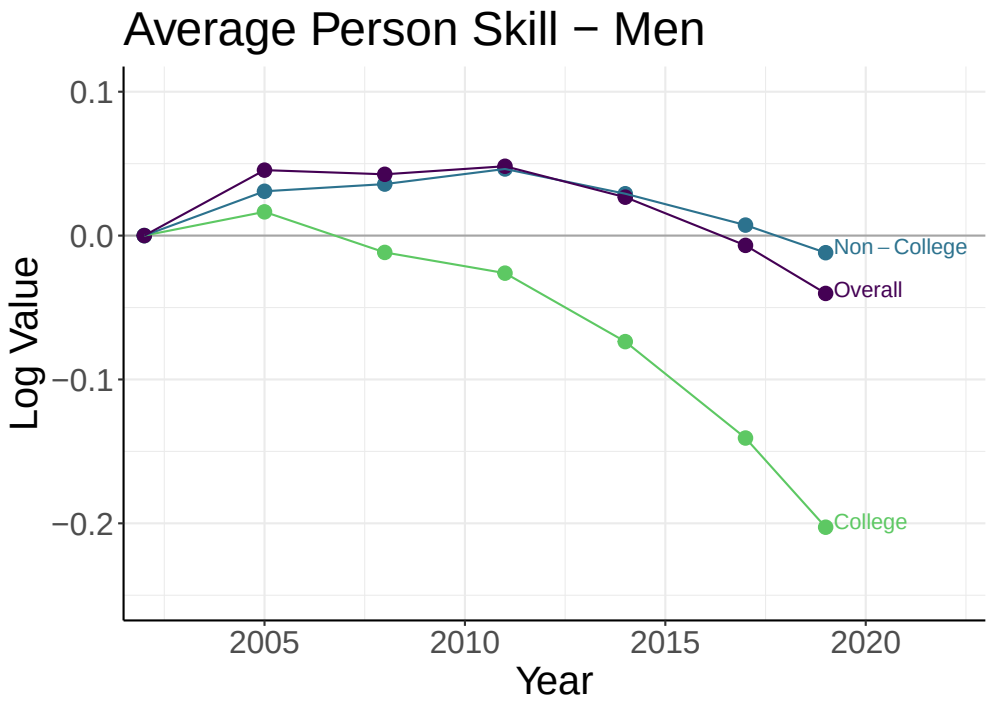


Figure C34: Growth in average person skill by education, for men.

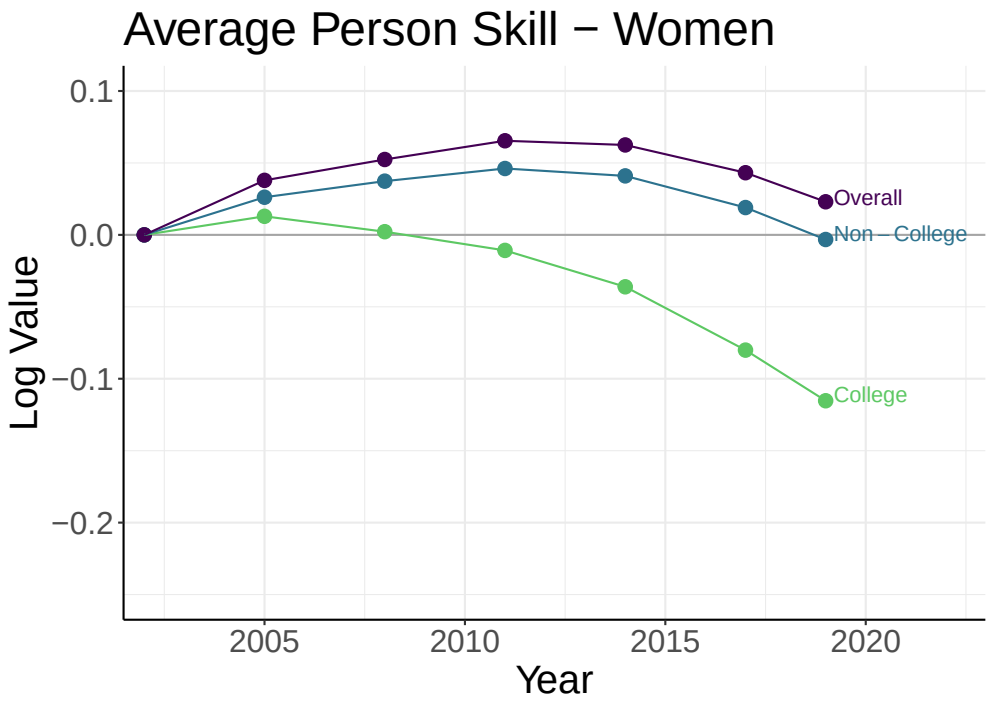


Figure C35: Growth in average person skill by education, for women.

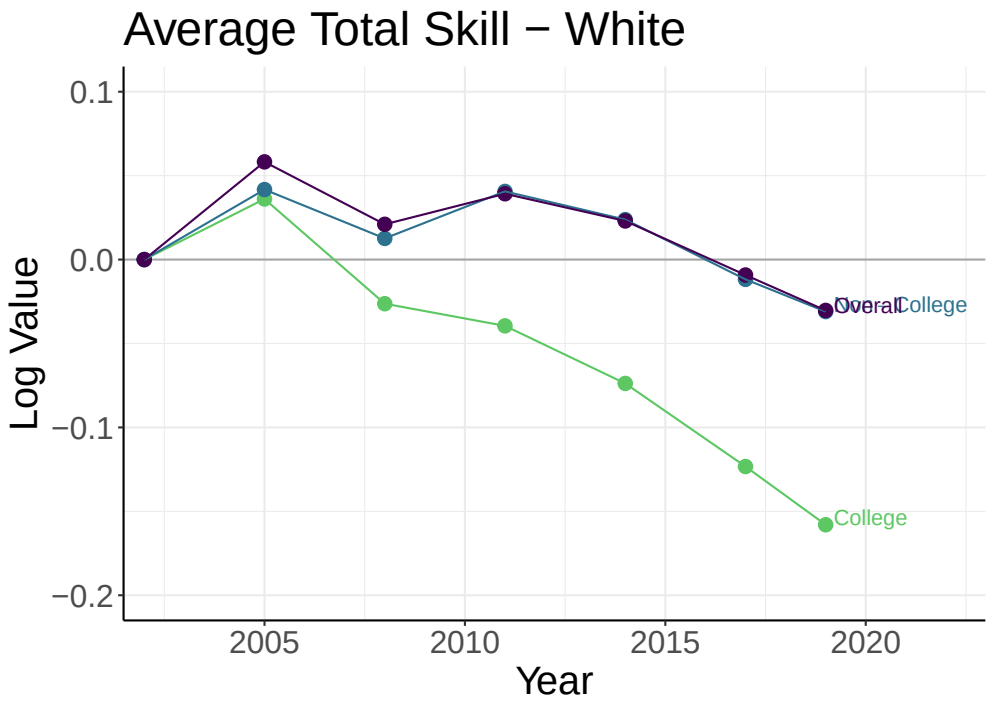


Figure C36: Growth in average total skill by education, for White workers.

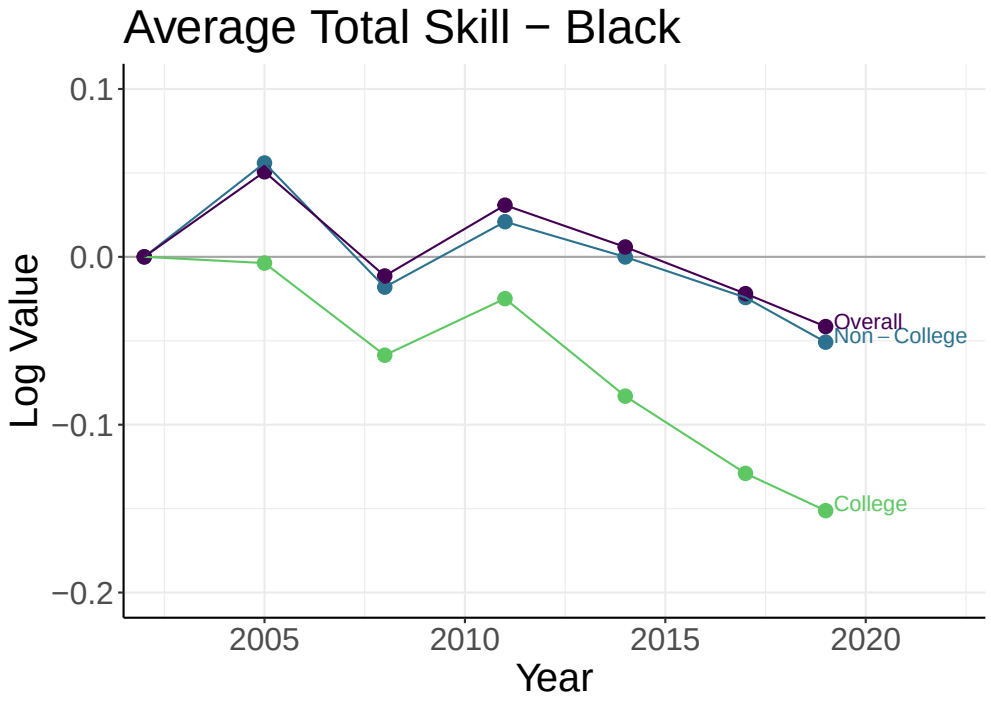


Figure C37: Growth in average total skill by education, for Black workers.

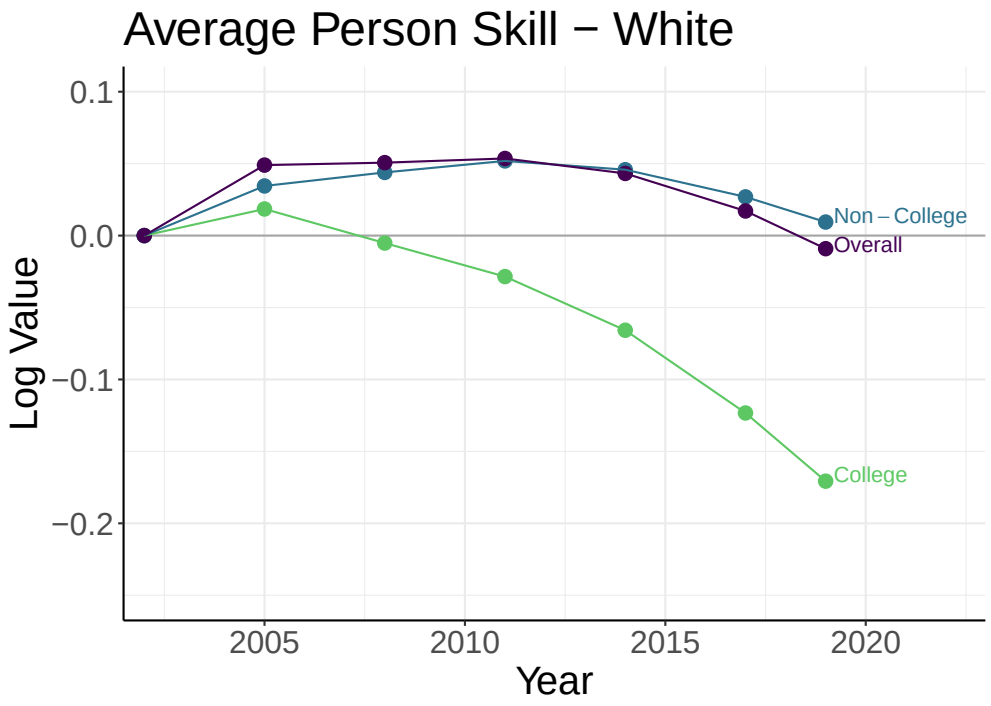


Figure C38: Growth in average person skill by education, for White workers.

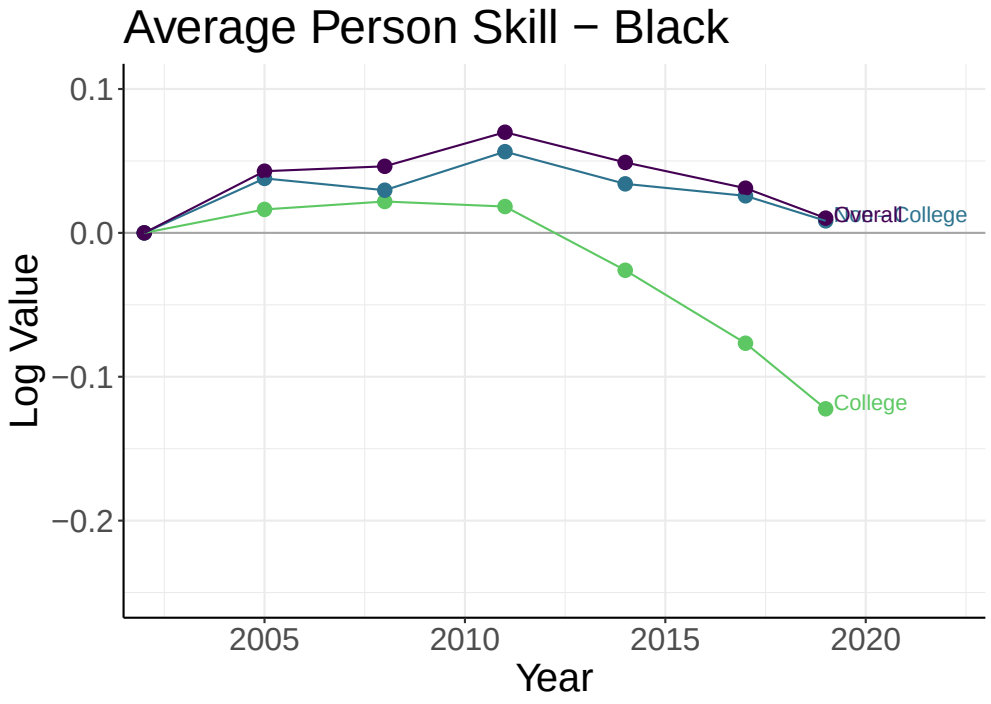


Figure C39: Growth in average person skill by education, for Black workers.

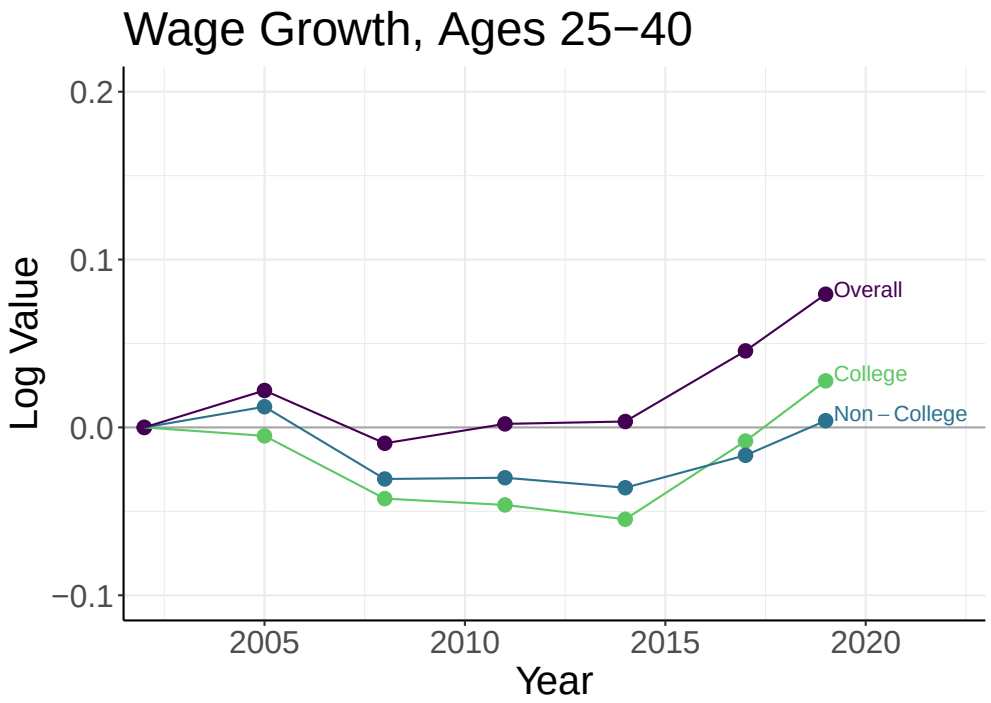


Figure C40: Growth in log hourly wages by education, for workers aged 25-40.

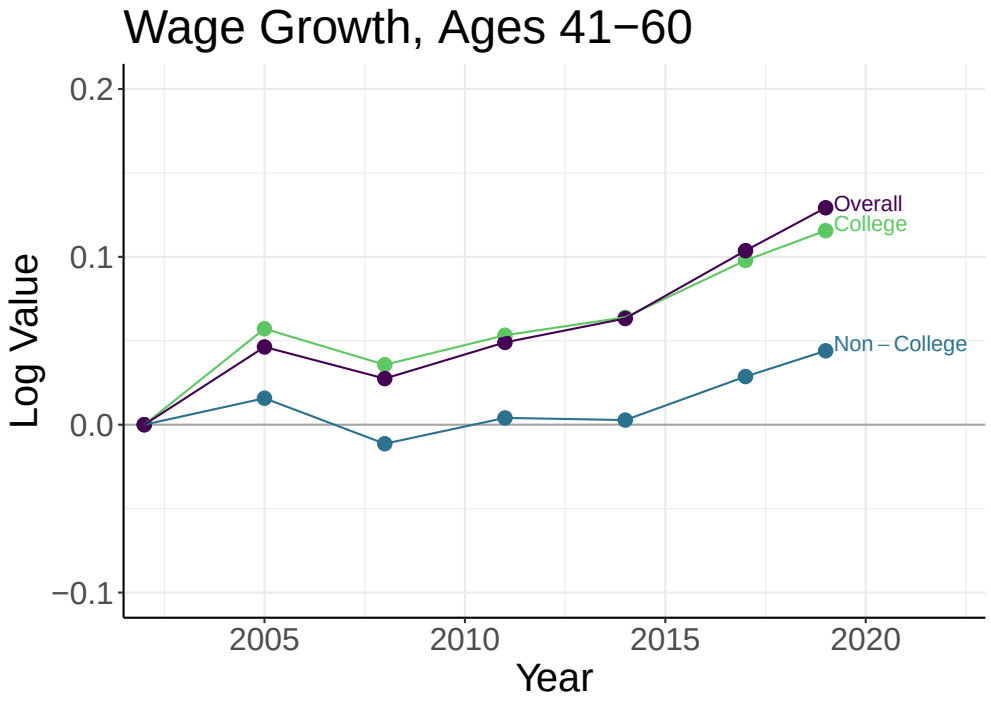


Figure C41: Growth in log hourly wages by education, for workers aged 41-60.

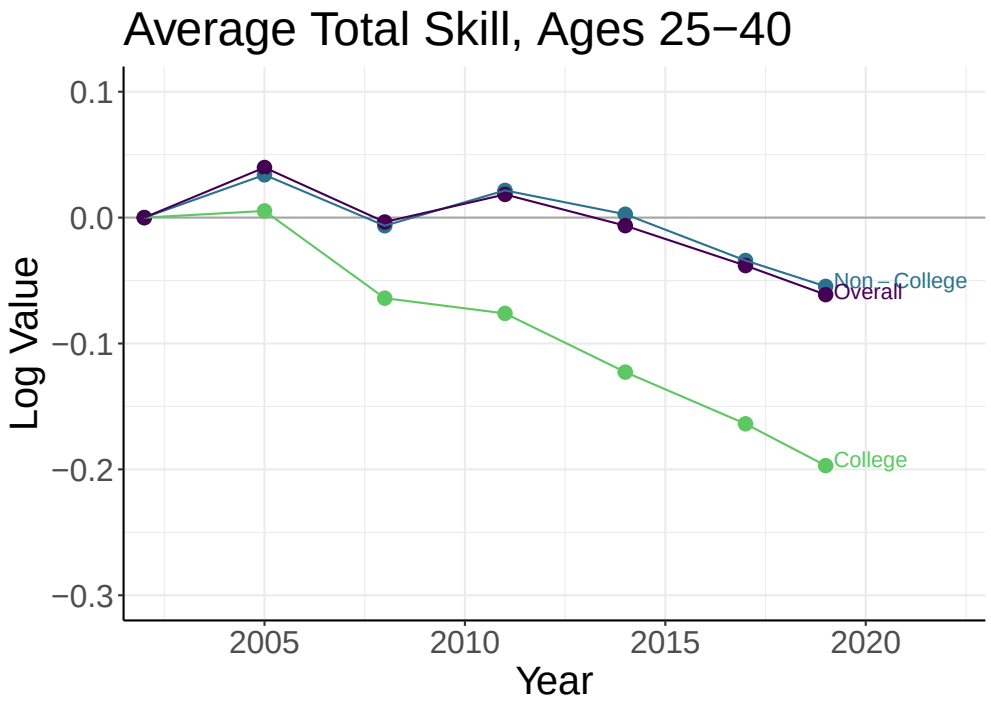


Figure C42: Growth in average total skill by education, for workers aged 25-40.

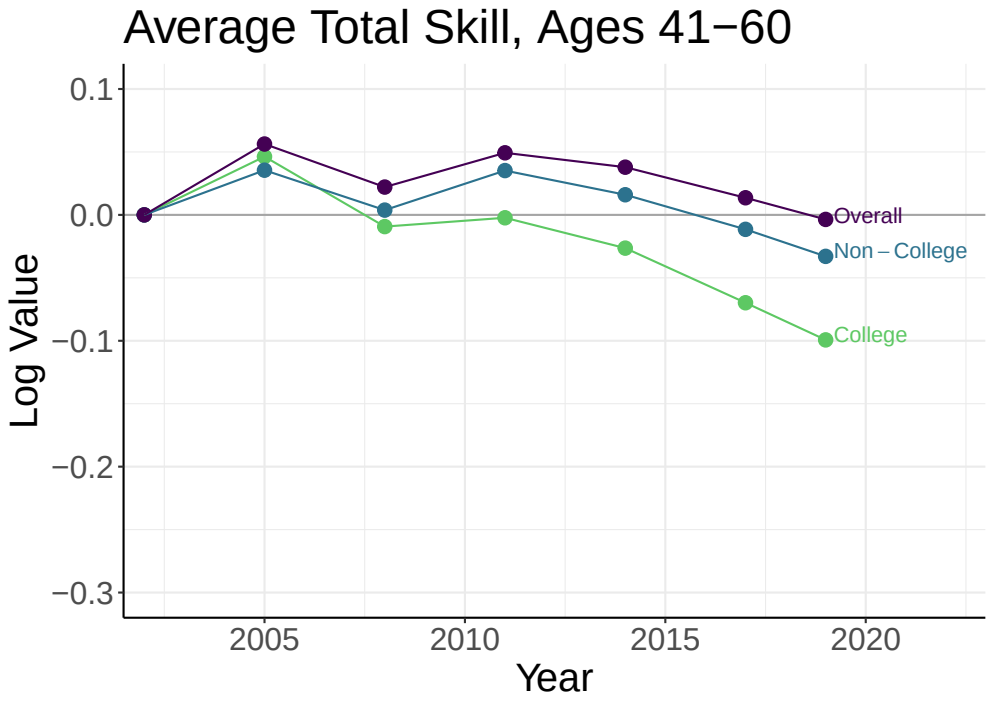


Figure C43: Growth in average total skill by education, for workers aged 41-60.

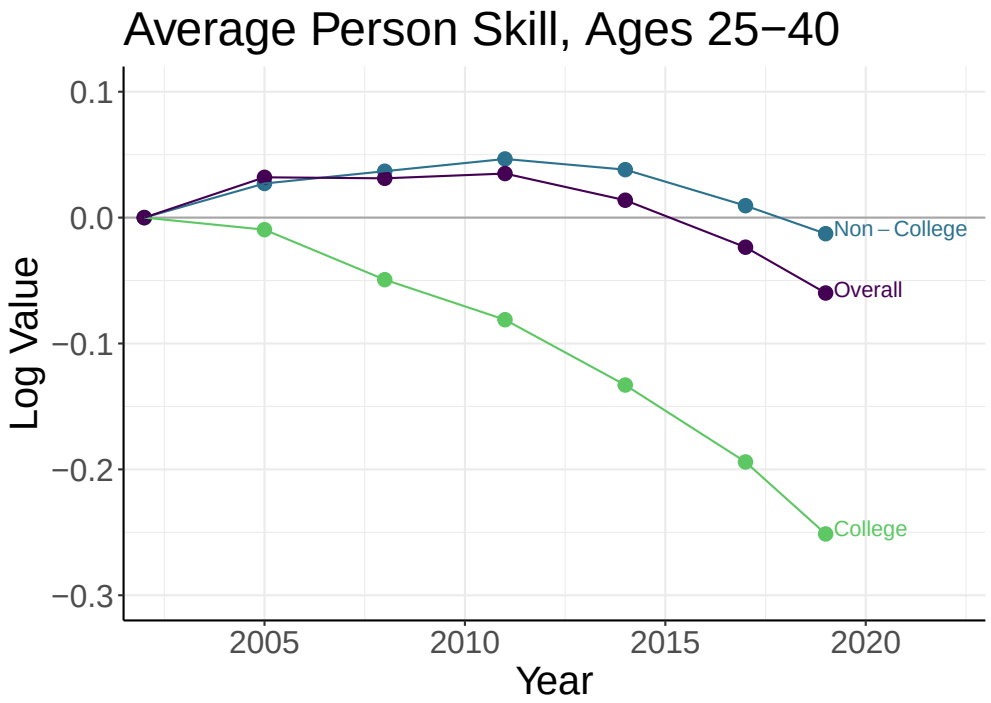


Figure C44: Growth in average person skill by education, for workers aged 25-40.

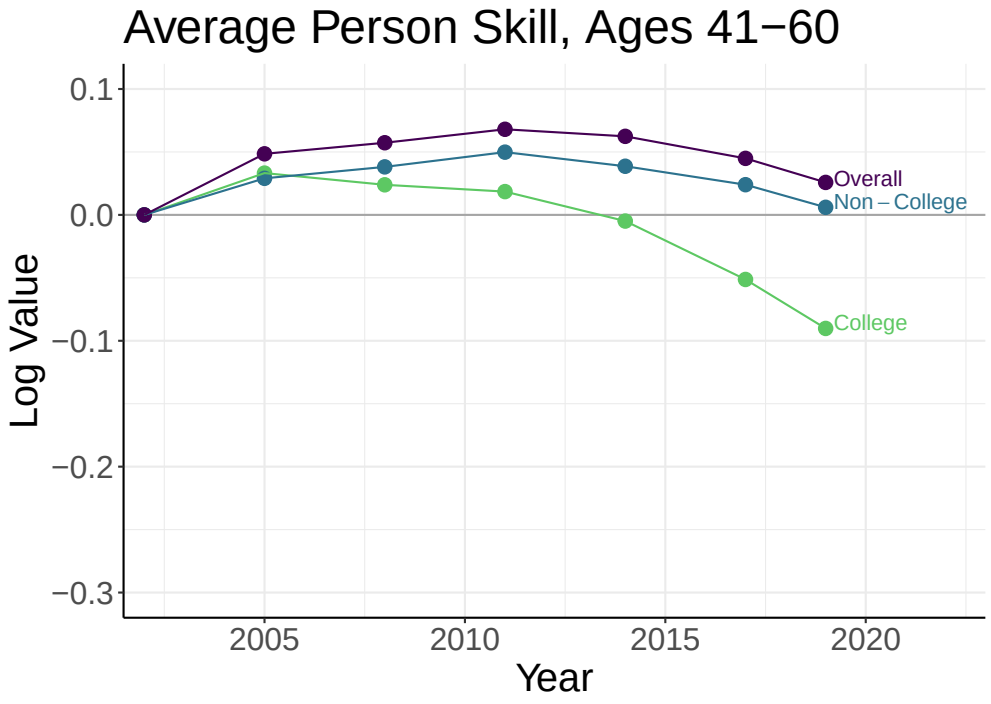


Figure C45: Growth in average person skill by education, for workers aged 41-60.

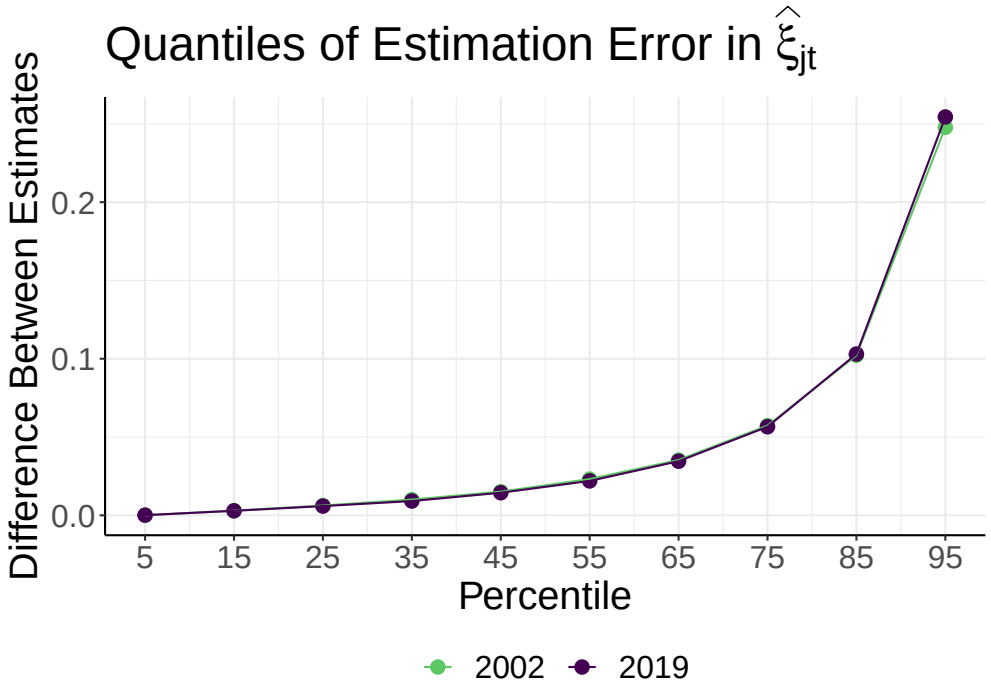


Figure C46: Quantiles of the distribution of the absolute difference between estimated ξ_{jt} in Samples A and B.

D Estimating Changes in Relative Demand

To compute changes in relative labor demand I follow the approach in Autor et al. (2008) following the replication files from Stephenson (2024). I use CPS ASEC data from the US Census. I filter to people with greater than 40 weeks worked and greater than 35 hours worked per week. I retrict to people between the ages of 16 and 64, exclude people in the military, exclude people who are self-employed, and drop observations with allocated earnings. I futher require that PCE adjusted weekly earnings exceed $136 in 2008 dollars and multiply top-coded earnings by 1.5. I define potential experience as age minus 6 minus the number of years of schooling. I subset to people with potential experience between 0 and 39.

I then sort workers into sex-education-experience groups. I use two sex bins, five education bins (dropped out of high school, high school degree, some college, college degree, post-college) and four potential experience bins (0-9, 10-19, 20-29, 30-39). I then regress log weekly wages separately in each year and sex on dummy variables for the education categories, a quartic in experience, the interaction of the experience quartic on the education dummies, race coded as dummy variables for White, Black, and other race, and regional dummies. The (composition-adjusted) mean log wage for each bin in a given year equals the predicted log wage from these regressions evaluated for Whites at the relevant experience level (5, 15, 25, or 35 years depending on the experience group). Mean log wages for broader groups in each year equal weighted averages of the relevant (composition-adjusted) cell means using a fixed set of weights, equal to the mean share of total hours worked by each group over the 2002 to 2019 sample period.

For the adjusted supply series, I again sort workers in each year into sex-education-experience bins. I define the sex and education groups as I did for the wage data. I use single-year categories for experience from 0 to 39. I measure labor supply in a year for broad education categories as weighted averages of labor supply in these cells. I used fixed weights equal to the average wage in each cell over the time period. In Table D1 I show the supply

and wage series over time.

$\sigma = 4.7$ equals the value I get from estimating σ using data fom 2002 to 2019. Specifically, rearranging Equation 7.7 and removing differences gives

$$\left[\overline{\log(w_C)} - \overline{\log(w_{NC})}\right]_t = \left(\frac{\sigma - 1}{\sigma}\right) \log\left(\frac{A_C}{A_{NC}}\right)_t - \left(\frac{1}{\sigma}\right) \log\left(\frac{N_C}{N_{NC}}\right)_t + \left(\frac{\sigma - 1}{\sigma}\right) \left[\overline{\log(z_C)} - \overline{\log(z_{NC})}\right]_t \tag{D.1}$$

As in Katz and Murphy (1992) I assume a log-linear increase in the relative demand for educated workers over time, $\log\left(\frac{A_C}{A_{NC}}\right)_t = \gamma_0 + \gamma_1 t$. Plugging this into Equation D.1 gives

$$\left[\overline{\log(w_C)} - \overline{\log(w_{NC})}\right]_t = \left(\frac{\sigma - 1}{\sigma}\right) (\gamma_0 + \gamma_1 t) - \left(\frac{1}{\sigma}\right) \log\left(\frac{N_C}{N_{NC}}\right)_t + \left(\frac{\sigma - 1}{\sigma}\right) \left[\overline{\log(z_C)} - \overline{\log(z_{NC})}\right]_t \tag{D.2}$$

I solve for the parameters that minimize the sum of squared differences between the left and right side of Equation D.2, getting that $\sigma = 4.7$.

Table D1: Changes in the log wage gap, log relative supply, and skill gap over time in the March CPS data described in Appendix D. Log Wage Gap equals growth in the wage gap between college and non-college workers, $\Delta \left[\overline{\log(w_C)} - \overline{\log(w_{NC})}\right]_t$, relative to 2002. Adj. Wage Gap adjusts the wage gap for demographic compositional changes similar to the procedure in Autor et al. (2008) and Acemoglu and Autor (2011). Log Relative Supply equals growth in the relative supply of college workers, $\Delta \log \left(\frac{N_C}{N_{NC}}\right)_t$. Adj. Relative Supply weights supply changes in different demographic groups by their relative wages. Skill gap equals my estimate for change in the average skill difference, $\Delta \left[\overline{\log(z_C)} - \overline{\log(z_{NC})}\right]_t$, from the LEHD data.

Changes in relative log wages, supply, and skill.

Year	Log Wage Gap	Adj. Wage Gap	Log Relative Supply	Adj. Relative Supply	Skill Gap
2002	0.00	0.00	0.00	0.00	0.00
2003	-0.01	-0.01	0.03	0.01	-0.01
2004	-0.01	-0.01	0.02	0.01	-0.01
2005	0.00	0.01	0.02	0.01	-0.01
2006	0.01	0.02	0.08	0.05	-0.01
2007	-0.01	0.01	0.12	0.08	-0.01
2008	0.01	0.04	0.14	0.10	-0.03
2009	0.01	0.03	0.19	0.14	-0.06
2010	0.01	0.03	0.23	0.18	-0.07
2011	0.01	0.04	0.25	0.19	-0.07
2012	0.02	0.03	0.26	0.21	-0.08
2013	-0.01	0.02	0.31	0.22	-0.09
2014	-0.02	0.04	0.31	0.23	-0.09
2015	-0.00	0.03	0.35	0.26	-0.10
2016	-0.00	0.03	0.37	0.27	-0.09
2017	-0.00	0.03	0.41	0.31	-0.10
2018	-0.01	0.03	0.44	0.34	-0.11
2019	-0.03	0.03	0.50	0.40	-0.11