Skill-Replacing Technology and Bottom-Half Inequality

Oren Danieli*
Tel Aviv University

September 5, 2024

Abstract

I propose a model of a skill-replacing routine-biased-technological-change (SR-RBTC). In this model, technology substitutes the usage of skill in routine tasks, in contrast to standard RBTC models which assume technology replaces the workers themselves. The SR-RBTC model explains three key trends that are inconsistent with standard RBTC models: why specifically middle-wages declined even though workers in routine occupations are dispersed across the entire bottom half of the wage distribution, why middle-wages stopped declining while the technological change continued, and why there is no substantial decline in the average wage of workers in routine occupations. I derive two new testable predictions from the model: a decrease in return to skill, and a decrease in skill level in routine occupations. I use an interactive-fixed-effect model to confirm both predictions. Since SR-RBTC violates the ignorability assumption required by standard decomposition methods, I introduce "skewness decomposition" to show that SR-RBTC is the main driver of bottom-half inequality trends.

*Email: orendanieli@tauex.tau.ac.il. I gratefully acknowledge financial support from the W.E. Upjohn Institute's Early Career Research Award program, Pinhas Sapir Center for Development, Foerder Institution, and Archimedes Center at Tel Aviv University. I thank Lawrence Katz, David Autor, Roland Fryer, Thomas Lemieux, Ro'ee Levy, Jonathan Roth, Itay Saporta-Eksten, Eran Yashiv, Nir Jaimovic, and four anonymous referees for very helpful comments and suggestions. Seminar participants at NBER Labor Studies, NBER Inequality and Macroeconomics, CEPR BOI Workshop, 25th IZA Conference in Labor Economics, BSE Summer Forum - Macroeconomics of Labor Markets, SOLE, EEA, ASSA, LERA, Harvard, Tel-Aviv University, Hebrew University, Princeton University, the department of environmental economics at HUJI, Haifa University, Harris School of Public Policy, Rutgers, Penn State, University of Toronto, University of Michigan, Brown University, Boston University, Harvard Kennedy School, Harvard Business School - Enterpreneurial Management, Federal Reserve Board and NY Fed provided useful feedback. Elad Guttman, Edan Kasan, Maya Orenstein, Aviya Tsoref, and Ofir Zeira provided outstanding research assistance. This paper was previously titled "Revisiting U.S. Wage Inequality at the Bottom 50%".

Recent trends in U.S. wage inequality are frequently attributed to technological advancements. In particular, several studies have proposed the model of routine-biased technological change (RBTC), which posits that technology is substitutional to routine workers since their performed tasks are easier to automate (Acemoglu and Autor, 2011).

However, the RBTC model cannot explain recent changes in the U.S. wage distribution, especially at the bottom. Figure 1 plots the evolution of the 90/50 and 50/10 wage ratios over time. While inequality at the top has steadily increased, inequality at the bottom has fluctuated. Even under more nuanced RBTC models, it remains unclear why bottom-half inequality decreased in the 1990s, and even less clear why it increased again after 2000. Thus, the prominent theories regarding the impact of technology on the wage distribution do not align with the observed U.S. data.

In this paper, I propose a new theory for how new automation technology has affected the bottom of the U.S. wage distribution in recent decades. My theory diverges from previous RBTC models by arguing that automation technology is skill-replacing (SR-RBTC). Instead of assuming that new technology is replacing workers, I assume that it replaces their skill. For example, calculators replace the need for arithmetic skills, enabling workers with limited arithmetic abilities to perform calculations just as effectively. This model can explain the recent trends of bottom-half inequality. I supplement this model with new empirical evidence showing a decline in the return to skill in routine occupations, as predicted. Furthermore, I find that routine occupations are increasingly filled by low-skilled workers, as skill becomes less necessary in these roles. Finally, I use a "skewness decomposition" to quantify that SR-RBTC can explain 93% of the discussed wage trends.

I start by outlining the theoretical framework of the paper. I construct a model in which workers are characterized by a one-dimensional continuous skill. Workers are employed in one of three occupations that vary in their return to skill (similar to Jung and Mercenier, 2014). In equilibrium, workers are allocated to occupations based on comparative advantage. The lowest-skilled workers sort into the manual occupation, middle-skilled to the routine occupation, and the highest-skilled to the abstract occupation.

The model diverges from most previous literature by assuming that the new technology in the routine occupation substitutes the skill of workers in this occupation. This skill-replacing technology reduces the return to skill in the routine occupation, leading to larger wage decreases for higher-skilled routine workers. This differs from previous models (i.e., Acemoglu and Autor, 2011; Cortes, 2016) which assume a skill-neutral technological change, where the wage effects are identical to all routine workers.It also differs from skill-enhancing models (Jung and Mercenier, 2014) that make the opposite assumption that technology increases the return to skill for routine workers. It is conceptually similar to Downey (2021) who argues new technology benefits low-skill workers.

There are numerous examples of skill-replacing technology in routine occupations. Cashiers today do not need any arithmetic skills, as all calculations are automated; administrators typically do not need to memorize any procedures or customer details because most are computerized; production workers rarely use physical strength anymore as machines can perform many physical tasks. Yet, counterexamples exist as well. Therefore, the empirical part of this paper tests how well the SR-RBTC model fits the data.

The SR-RBTC model explains three key stylized facts that previous (skill-neutral) RBTC models could not. These unresolved puzzles leave room for consideration of other explanations for the fluctuations in bottom-half inequality (Hunt and Nunn, 2022). First, as shown in Figure 1, starting from the late 1980s, the median wage declined relative to both high and low wages, a trend often referred to as "wage polarization". Earlier RBTC models argued that routine tasks, which are more easily automated, require middle-skill workers. Hence, wages relatively decline predominantly in the middle of the wage distribution. However, empirical evidence suggests that routine workers are dispersed almost equally across the entire bottom half of the wage distribution (Autor and Dorn, 2013). Therefore, a skill-neutral RBTC model would predict wage decreases across the entire bottom half of the income distribution, not just the middle, in contrast to the observed trend in the data.

By contrast, SR-RBTC predicts a decline in middle wages. A decrease in return to skill in the routine occupation would generate the largest wage decreases for the highest earning routine workers. Empirically, these highest earning routine workers used to be concentrated in the middle of the wage distribution. Consequentially, a decrease in their wages would generate wage polarization.

The second fact is that the decline in middle wages stopped around the year 2000. If RBTC was generating the decline in middle wages, it is unclear why this decline stopped. This is especially puzzling since the decline in routine employment suggests that RBTC continues long afterward (Autor, 2014).

This fact can be explained by SR-RBTC, which predicts a non-monotonic relationship between technological advancement and inequality at the bottom half of the distribution. At some point, the return to skill in the routine occupation would fall below its level in the manual occupation. When such a reversal of comparative advantage occurs, workers would reallocate and only the lowest-skilled workers would choose to work in the routine occupation. After this reversal, SR-RBTC would reduce wages for the remaining routine workers who would be at the bottom of the wage distribution. SR-RBTC would not affect middle-wage workers directly, as they would no longer work in the routine occupation.

Therefore, inequality at the bottom of the distribution would rise.

The third fact is that the average wage declines in routine occupations are relatively modest compared to the substantial decline in employment. It is unclear why the market adjusts to the decline in demand for routine workers almost mainly through quantities (employment), rather than prices (wages). Moreover, different decomposition methods found that the wage decrease in routine occupations is too small to account for the aforementioned wage trends (Autor et al., 2005; Firpo et al., 2013). In SR-RBTC, even though employment declines, the average wage in the routine occupation does not necessarily decrease. Wages fall for the highest-skilled workers in the routine occupation. However, lower-skilled routine workers may benefit from the change. As a result, the average routine wage may not decline.

To directly test the model, I derive two new predictions that can distinguish a skill-replacing RBTC from a skill-neutral or skill-enhancing RBTC. First, the model predicts a decrease in return to skill in the routine occupation. Second, it predicts a gradual decline in the skill level of workers in the routine occupation. These trends should continue throughout the entire period of RBTC, starting in the late 1980s. At some point, the return to skill in the routine occupation should fall below the return in the manual occupation, leading to a reversal of comparative advantage. Following such reversal, the average skill level should be lower in the routine occupation, compared to the manual occupation.

To test these predictions, I estimate an interactive fixed-effects model (IFEM). IFEM is a more general version of the standard fixed-effects model. It regresses log wages on a set of independent variables, including worker fixed effects that capture unobserved skill. The only difference from a standard fixed-effects model is that the worker fixed effects are interacted with the year and occupational category. This interaction allows the return to the unobserved skill to vary over time and across occupations, as the SR-RBTC model predicts. Since the unobserved skill is estimated with noise, I instrument for it with years of schooling to prevent an attenuation bias. For this exercise, I use data from the Panel Study of Income Dynamics (PSID) between 1980–2017.

The results from the IFEM generate new empirical facts that are consistent with both model predictions. I find a sharp decrease in the return to skill in routine occupations starting in the late 1980s, exactly when inequality at the bottom half of the distribution started to decline. The return to skill in routine occupations continued to decrease for more than two decades. I also find that the average skill level in routine occupations, as measured using the IFEM, steadily fell during this period. As a result, workers in routine occupations became more concentrated at the bottom income quintile, instead of working in middle-wage jobs. Previous work investigating the compositional change in routine occupations has focused mainly on the decline in employment in routine occupations (Goos and Manning, 2007; Goos et al., 2009; Goos et al., 2014), and the flow of workers in and out of routine occupations (Cortes, 2016; Cortes et al., 2020). Consistent with my findings, Dicandia (2023) finds that the share of white workers in routine occupations has also decreased. However, there has been little discussion on the impact of these employment trends on the average skill level in routine occupations.

Estimates from the IFEM are consistent with a reversal of comparative advantage at the bottom of the distribution. I find that around 1987, before inequality started to decline at the bottom of the distribution, the return to skill was slightly higher in routine occupations compared to manual occupations. During this time, the average skill level of workers in routine occupations was substantially higher than the skill level of manual workers. Over time, the return to skill in routine occupations fell far below its value in manual occupations. The average skill level of workers in routine occupations also declined and by 2015, it fell below the average skill level in manual occupations. As a result, workers in routine occupations and especially those in administrative or operator occupations have the lowest level of skill across all occupational categories.[1]

In the final part of the paper, I use a skewness decomposition to show that SR-RBTC is not only consistent with recent inequality trends, it is also substantial enough to account for almost the entire wage trend. I introduce a novel decomposition that is based on the skewness of the log wage distribution. In analogy to inequality that can be measured with the second moment of the log wage distribution, wage polarization can be measured with the third moment of that distribution, namely, skewness. When inequality increases at the top and decreases at the bottom, the log wage distribution becomes more positively skewed and this moment increases. For this analysis, I use the Current Population Survey Outgoing Rotation Groups (CPS-ORG). As expected, the skewness increases precisely when wage polarization occurs.

The main advantage of using skewness to measure wage polarization is that it can be decomposed into three independent components. In my main analysis, I decompose the rise in skewness by occupations. Similar to variance decomposition, skewness decomposition has a between-occupations and a within-occupations components. A skill-neutral RBTC is expected to reduce wages equally for all workers in routine occupations. Therefore, this model predicts that most of the rise in skewness would be driven by the between-occupations component. By contrast, in an SR-RBTC model, the third component is expected to rise. The third component captures the correlation between occupa-

[1]Manual workers still earned less than workers in routine occupations on average, despite having a higher skill level. One reason for this is that routine workers are more experienced (Autor and Dorn, 2009).

tion wage level and occupation inequality. Hence, if wage gaps decrease in lower-paying routine occupations, while they increase in higher-paying abstract occupations, as the SR-RBTC model predicts, this correlation will increase and generate a rise in skewness.

The decomposition results indicate that RBTC is skill-replacing. I find that 78% of the overall increase in skewness is driven by the correlation component. The correlation component increases mainly because of a decrease in inequality in routine occupations, exactly as predicted by the SR-RBTC model. For comparison, trying to decompose the increase in skewness by industries or education generates a much larger increase in the within component. This implies that the rise in skewness is driven primarily by an occupational trend, which also occurs within industries and education categories.

The results of skewness decomposition are distinct from previous decomposition attempts of wage polarization because it does not rely on the "ignorability assumption". Previous attempts to decompose wage polarization have found that technological changes and occupational trends in general, cannot generate wage polarization (Autor et al., 2005; Firpo et al., 2013).[2] The most common decomposition methods (e.g., Juhn et al., 1993; DiNardo et al., 1996; Firpo et al., 2009) rely on an assumption called ignorability (Fortin et al., 2011). As a result, these previous decomposition papers only quantified the decrease in average routine wages. However, the increase in skewness is driven by the decline in inequality in routine occupations. This trend was previously documented by Lemieux (2007), and causally identified by Gaggl and Wright (2017). Using skewness decomposition I find that the decrease in inequality in routine occupations is the main driver of wage polarization. Skewness decomposition was previously discussed in labor economics (Mincer, 1974) but was never applied to economics data.

I conclude this paper by discussing alternative explanations for these wage trends and why they are less consistent with my findings. Additionally, I briefly discuss bottom-half inequality trends in other developed countries and their potential explanation.

1 Model

1.1 Occupational Sorting by Skill

I outline a model that highlights the difference in return to skill in each occupation, building on earlier work by Jung and Mercenier (2014) and Cortes (2016). Assume that work-

[2]Acemoglu and Restrepo (2022), show that technological change can explain the majority of the rise of inequality between skill groups. I show that SR-RBTC can explain both the increase and the decrease in inequality, for the entire bottom-half of the distribution, and not just across skill groups.

ers have a one-dimensional skill, θ_i with some density function $f(\theta_i)$. This assumption is more general than the assumption of a discrete number of skill levels (Katz and Murphy, 1992; Autor et al., 2006; Acemoglu and Autor, 2011), but less general than assuming multidimensional skills (Roy, 1951), which I discuss in an extension in Appendix B.3.

Occupations differ in their return to skill. To simplify, I will assume three occupations: manual, routine, and abstract. In each occupation $j \in \{M, R, A\}$, workers produce an intermediate good with a production function $\varphi_j(\theta_i)$. Assume that in baseline

$$\forall\theta_i : \frac{\partial \log \varphi_M(\theta_i)}{\partial\theta_i} < \frac{\partial \log \varphi_R(\theta_i)}{\partial\theta_i} < \frac{\partial \log \varphi_A(\theta_i)}{\partial\theta_i} \tag{1}$$

so that the manual occupation has the lowest return to skill, and the abstract occupation has the highest. Assume also that for every θ_i and every occupation $\frac{\partial \log \varphi_j(\theta_i)}{\partial\theta_i} > c_j$ for some constant $c_j > 0$, implying that the return to skill is strictly positive at any level.

Under the assumption of perfect competition, wages are set at the marginal productivity. Let p_j be the price of the intermediate good in occupation j. Therefore, if worker i is working in occupation j, she will earn

$$w_j(\theta_i) = p_j\varphi_j(\theta_i)$$

Workers sort into occupations based on comparative advantage. Condition 1 guarantees the existence of two thresholds θ_0, θ_1 such that any worker with $\theta_i < \theta_0$ choose to work in the manual occupation, any worker with $\theta_0 < \theta_i < \theta_1$ choose the routine occupation, and any worker with $\theta_i > \theta_1$ choose the abstract occupation (Jung and Mercenier, 2014). Workers with a skill level that exactly equals the threshold will be indifferent; hence the following two equations hold in equilibrium:

$$\begin{aligned} p_M\varphi_M(\theta_0) &= p_R\varphi_R(\theta_0) \\ p_R\varphi_R(\theta_1) &= p_A\varphi_A(\theta_1) \end{aligned} \tag{2}$$

Figure A1 shows this graphically, by plotting the equilibrium log wages by skill level θ_i.

1.2 Routine-Biased Technological Change

I focus on technological change that improves productivity in the routine occupation. For simplicity, I assume that the technological change affects only φ_R directly, as this change is sufficient for explaining the inequality trends at the bottom of the wage distribution. Hence, φ_M, φ_A are left unchanged. However, wages in the manual and abstract occupations are affected as well in a general equilibrium.

Specifically, I assume that the production of a routine worker is a function of θ_i their skill, and τ the level of technology, $\varphi_R(\theta_i, \tau)$. I assume that production is monotonically increasing in both inputs, $\frac{\partial \varphi_R}{\partial \theta_i}, \frac{\partial \varphi_R}{\partial \tau} > 0$. RBTC would then be modeled as an increase in the technology level τ over time. This technology growth increases the productivity of every routine worker. Therefore, it enables the production of the same quantity with fewer workers.

While RBTC makes all workers in the routine occupation more productive, some workers may experience larger productivity gains than others. Formally, use ϵ to mark the effect of technology on the return to skill

$$\epsilon = \frac{\partial^2 \log \varphi_R}{\partial \theta_i \partial \tau}. \tag{3}$$

I assume that the sign of ϵ is the same for all workers for a given technology level.

I distinguish between three types of RBTC. If $\epsilon = 0$, RBTC is skill-neutral as in Cortes (2016). The effect of technology on log productivity ($\frac{\partial \log \varphi_R}{\partial \tau}$) would be the same for all workers in the routine occupation.[3] If $\epsilon > 0$, as hypothesized by Jung and Mercenier (2014), technology is skill-enhancing. That is, technology increases the productivity gaps by skill. If $\epsilon < 0$, technology is skill replacing, and the return to skill declines.

In Appendix B I discuss two alternative micro-foundations for RBTC that provide insight into when RBTC would be skill-neural, -enhancing, or -replacing. In Appendix B.1, an increase in τ represents the full automation of some of the tasks previously performed by workers in the routine occupation. Such automation allows these workers to allocate more time to other tasks. In this model, the RBTC type (skill-neutral/enhancing/replacing) is determined by the importance of skill in the automated task. If the automated task is more skill-intensive than the average task, RBTC is skill replacing. For example, for cashier workers, technology replaced the task of arithmetic calculations, which is relatively skill-intensive, and hence technology is skill-replacing. Alternatively, an increase in τ can also represent an improvement in the quality or quantity of computers or robots. Appendix B.2 discusses such a model when φ_R has constant returns to scale (CRS). In this case, the RBTC type depends on the elasticity of substitution between skill and technology. If skill and technology are substitutes, RBTC is skill replacing.

The type of RBTC determines the effect of technology on income gaps in the routine occupation. A skill-neutral RBTC would not affect inequality among workers in the routine occupation as stated in the following theorem.

[3] While the effect on productivity is skill-neutral, the effect on wages varies by skill (Cortes, 2016).

Theorem 1. *Let $\theta_a, \theta_b \in (\theta_0, \theta_1)$ be the skill levels of two workers in the routine occupation where $\theta_a < \theta_b$. Let w_a, w_b denote their corresponding equilibrium wages. The effect of an improvement in technology τ on the wage ratio $\frac{w_b}{w_a}$ depends on the sign of ϵ such that*

$$sign\left(\frac{\partial \frac{w_b}{w_a}}{\partial \tau}\right) = sign\left(\epsilon\right)$$

All proofs are given in Appendix C. In the next parts of the paper I show evidence that this sign is negative, and therefore RBTC is skill-replacing.

1.3 General Equilibrium

I assume that the three intermediate goods are used jointly to produce a final good. I also assume that workers with different skill levels are perfect substitutes in the production of the intermediate good. I use M, R, A to denote the total amount produced from each intermediate good, which equals

$$\begin{array}{rcll} M & = & \int_{\theta_{min}}^{\theta_0} & \varphi_M\left(\theta_i\right) \quad \mathrm{d}\theta_i \\ R & = & \int_{\theta_0}^{\theta_1} & \varphi_R\left(\theta_i\right) \quad \mathrm{d}\theta_i \\ A & = & \int_{\theta_1}^{\theta_{max}} & \varphi_A\left(\theta_i\right) \quad \mathrm{d}\theta_i \end{array} \tag{4}$$

The final good is the output of a CES function with $\rho < 0$,

$$Y = \left(M^\rho + R^\rho + A^\rho\right)^{\frac{1}{\rho}} \tag{5}$$

The three intermediate goods are complementary, as found by Jaimovich et al. (2021).

While RBTC increases the production of routine goods R, routine workers do not necessarily benefit. This depends on whether there is a sufficient demand increase for additional routine goods. The price of one unit of the routine good p_R would decrease due to the rise in quantity. Because of the complementarities ($\rho < 0$), the increased productivity in the routine occupation, increases demand for manual and abstract workers and raises the prices of the goods they produce. Overall the share of the total output that is spent on workers in the routine occupation $\frac{p_R R}{Y}$ declines (as found by Eden and Gaggl, 2018). This is summarized in the following theorem.

Theorem 2. *RBTC (i.e., an increase in τ) generates:*

1. *An increase in the production of the routine good ($\frac{dR}{d\tau} > 0$).*

2. *A decrease in the absolute price of the routine good ($\frac{dp_R}{d\tau} < 0$) and the relative price compared to the abstract/manual good ($\frac{dp_R/p_j}{d\tau} < 0$ for $j \in \{M, A\}$).*

3. *A decrease in the share of the total income that is spent on routine goods ($\frac{d\frac{p_R R}{Y}}{d\tau} < 0$).*

These predictions coincide with predictions of various other skill-neutral models for RBTC. In the next two sections, I will derive unique predictions for the case of a skill-replacing technology.

1.4 Skill-Replacing RBTC: First Phase

I now examine in more detail the case of SR-RBTC, where technology and skill are substitutes ($\epsilon < 0$). In contrast to other models of RBTC, in this model, the impact of technology on bottom-half inequality is non-monotonic. I start with the first phase, where the increase in τ is still relatively small, such that the comparative advantage at Condition 1 still holds. A small increase in τ generates wage polarization and additional predictions that can be tested against the data.

Theorem 3. *Assume a skill-replacing technology ($\epsilon < 0$). RBTC (i.e., an increase in τ) would generate the following:*

1. *A decrease in wage gaps between workers in the routine occupation who do not switch occupations.*

2. *The highest skill routine workers would leave the routine occupation ($\frac{\partial \theta_1}{\partial \tau} < 0$).*

3. *The wage for the highest-skilled routine worker (θ_1) would decrease relative to all other workers.*

Figure 2a illustrates the results of Theorem 3. Since technology is skill-replacing, the return to skill in the routine occupation becomes flatter. This generates lower gaps between workers who stay in the routine occupation. A relative drop in middle-wages occurs since the most significant wage drop is for the highest-earning routine workers, which (empirically) are concentrated in the middle of the overall distribution of skill. As the return to skill declines, some of the highest-skilled routine workers will have their comparative advantage in the abstract occupation, and so θ_1 will drop.

The effect on θ_0 could go either way. If ρ approaches $-\infty$ (Leonteif) θ_0 will increase, while if ρ is closer to 0 (Cobb–Douglas) θ_0 will decrease. Empirically, it seems that during the 1990s employment in manual jobs did increase, but not as fast as in abstract occupations (Acemoglu and Autor, 2011). In case θ_i is distributed uniformly, this could only occur if θ_0 increased, but by a smaller level compared to the decline in θ_1. The following theorem derives additional empirical implications for this particular case. Appendix C proves a more general version of this theorem for any continuous distribution of θ_i.

Theorem 4. *Assume a skill replacing technology ($\epsilon < 0$), $\theta_i \sim U\left[\underline{\theta}, \bar{\theta}\right]$, and $0 < \frac{d\theta_0}{d\tau} < \left|\frac{d\theta_1}{d\tau}\right|$. In the routine occupation, RBTC would generate a decrease in: (i) employment, (ii) within-occupation inequality, and (iii) mean skill level. Inequality within the abstract and manual occupation will rise. The overall inequality trend is asymmetric. Below θ_1, wage gaps would (weakly) decrease between every two workers. At the top, the wage gap between abstract workers and high-skill routine workers will increase. See Appendix C for formal definitions.*

The asymmetric trend in inequality can be seen in the difference between the red line and the black line in Figure 2a. The productivity increase for workers in the routine occupation is offset by the drop in prices. Therefore wages in the routine occupation fall relative to the other two occupations. Moreover, among workers in the routine occupation, the relative drop in wages is most significant for the highest-skilled workers. The abstract occupation expands and now includes some additional less-skilled workers, which increases its within-occupation inequality. Taken together these trends generate a U-shaped pattern where wages increase the most at the tails, and decrease the most around the middle of the skill distribution at the new value of θ_1.

In addition to the impact on wages, SR-RBTC also has an effect on employment in each occupation. Since there is not enough demand for all the new routine goods workers could potentially produce, some of them leave and employment in the routine occupation falls.[4] This decline in employment is driven primarily by the higher-skilled routine workers. As a result, workers in the routine occupation become less skilled on average.

Along with employment, inequality within the routine occupation also declines for two separate reasons. First, it declines directly due to the decrease in the productivity gap. Second, it declines indirectly due to the compositional changes that make the remaining workers in the routine occupation more homogenuous in their skill level.

1.5 Skill-Replacing RBTC: Second Phase

Wage polarization stops when middle-skilled workers' comparative advantage is no longer in the routine occupation. Assume that at some point the return to skill in the routine oc-

[4] This employment decline is sometimes referred to as "job polarization". However, since routine workers are dispersed across the entire bottom half of the income distribution, routine occupations are often not middle-wage occupations. Hence, a decrease in their employment might not generate "job polarization", consistent with the empirical evidence by Hunt and Nunn (2022).

cupation drops to a level that is below the return to skill in the manual occupation. At that point, the comparative advantage is reversed. The lowest-skilled workers sort into the routine occupation. Any further increase in τ will still reduce wage gaps among workers in the routine occupation. However, since the routine occupation employs the lowest-skilled workers, wages relatively decline for low-paid workers, as shown in Figure 2b. Hence, inequality at the bottom of the wage distribution could in fact increase.[5]

However, while wage trends change, the decline in employment continues as workers continue to leave the routine occupation. Since wages decline in the routine occupation, more workers would prefer to leave and join the manual occupation. These predictions are summarized in the following theorem.

Theorem 5. *Assume a skill-replacing technology ($\epsilon < 0$) and that there exists a $\widetilde{\tau}$ such that for any $\tau \geq \widetilde{\tau}$ and for any θ_i*

$$\frac{\partial \log \varphi_R(\theta_i; \tau)}{\partial \theta_i} < \frac{\partial \log \varphi_M(\theta_i)}{\partial \theta_i} \tag{6}$$

When $\tau \geq \widetilde{\tau}$ workers in the routine occupation earn the lowest wages. Any additional SR-RBTC ($\tau \uparrow$) decreases employment in the routine occupation ($\frac{d\theta_0}{d\tau}$<0), as well as wage gaps among workers in the routine occupation who do not switch occupations.

The key reason why the impact of SR-RBTC on the wage distribution changes over time is the change in the composition of workers in the routine occupation. At first, when workers in the routine occupation are middle-skilled, the main negative effect is concentrated around the median of the distribution. Later, when the routine occupation becomes a low-skilled job, the negative impact of SR-RBTC is concentrated at the bottom of the distribution.

This model of SR-RBTC is consistent with recent trends in bottom-half inequality, which were previously documented, but could not be explained in terms of a skill-neutral technological change. Specifically, a skill-neutral technological change cannot explain why inequality at the bottom of the distribution rose again after its initial decline. It also cannot explain why the relative wage decrease was concentrated in the middle of the distribution when most workers in routine occupations are concentrated below the median. The first phase of SR-RBTC corresponds to trends in the late 1980s and 1990s and the second phase to trends in the 2000s and onwards. Appendix B.3 presents a more general model in which workers use different skills in different occupations. The predictions of the more general model are also consistent with recent bottom-half inequality trends.

[5]Given that inequality declines among workers in the routine occupation while it increases between manual and routine workers the overall impact of SR-RBTC on inequality at this phase is ambiguous.

The model generates two new predictions that can be tested against the data. First, it predicts a decline in the return to skill in the routine occupation. This generates a decrease in wage gaps for routine workers who do not switch occupations. Second, it predicts a decline in the average skill level of workers in routine occupations. Both trends should be sufficiently large so that at some point a reversal of comparative advantage occurs whereby the return to skill and the average skill level become higher in the manual than in the routine occupation. In the following sections of the paper, I test these empirical predictions and show that they fit well with the data.

2 Methodology

This paper uses two separate empirical techniques. To test the predictions of the SR-RBTC model, I use an interactive fixed-effects model (IFEM). Then, to quantify the share of the overall wage trend that can be attributed to SR-RBTC, I use skewness decomposition.

2.1 Interactive Fixed-Effects Model

To test whether RBTC is skill-replacing, skill-enhancing, or skill-neutral, I estimate the return to skill directly, using an interactive fixed-effects model (IFEM). Specifically, I estimate the following equation for a worker i in occupation j in year t:

$$\log w_{ijt} = \beta_{jt} X_{it} + \lambda_{jt} + \alpha_{jt}\theta_i + \varepsilon_{ijt}, \tag{7}$$

where λ_{jt} are occupation-year fixed effects and X_{it} is an additional control for experience squared.[6] The individual fixed effects θ_i represent permanent wage differences across workers, which correspond to the notion of skill in the theoretical model in Section 1. The key parameters of interest are the coefficients α_{jt}, the return to skill in any combination of occupational category and year. I use either three occupational categories (abstract, routine, manual) or nine (defined by the first digit of the occupational code).

When the return to skill is constant across occupational categories and over time ($\alpha_{jt} = 1$), the model is identical to a standard fixed effects model. Worker fixed effects are frequently applied to account for permanent unobserved skill differences. It allows for overcoming differences in the composition of workers between occupations and over time. The concern of compositional change is particularly relevant given the substantial

[6] I do not control for education level and experience as they are collinear with θ_i and λ_{jt}.

decrease in employment in routine occupations, driven prominently by the highest- and lowest-earning workers in those occupations (Cortes, 2016; Böhm et al., 2019).

The IFEM model is a more general version of the fixed-effects model, which accommodates variation in the return to skill, as posited by the SR-RBTC model. The parameters α_{jt} capture changes in the return to skill across occupations and over time. Specifically, the model predicts that $\alpha_{R,t}$, the return to skill in routine occupations, will decrease as technology advances if and only if the technology is skill-replacing (Theorem 1).

I search for a combination of parameters that minimize the expected mean square errors, $E\left[\varepsilon_{ijt}^2\right]$. The first-order condition moments of this minimization problem imply that for every occupational category j and every year t

$$\mathbb{E}\left[\varepsilon_{ijt}|i \in E_{jt}\right] = \mathbb{E}\left[X_{it}\varepsilon_{ijt}|i \in E_{jt}\right] = \mathbb{E}\left[\theta_i\varepsilon_{ijt}|i \in E_{jt}\right] = 0, \tag{8}$$

where E_{jt} is the set of workers in occupational category j in year t. Moreover, deriving the first order conditions by θ_i imply that for every worker i, $\mathbb{E}\left[\alpha_{j(i,t)t}\varepsilon_{ij(i,t)t}|i\right] = 0$. From this moment, one can derive an estimator $\widehat{\theta}_i$, given the other parameters

$$\widehat{\theta}_i\left(\log w_i, X_i, \widehat{\alpha}, \widehat{\beta}, \widehat{\lambda}\right) = \frac{\sum_t \widehat{\alpha}_{j(i,t),t}\left(\log w_{ij(i,t)t} - \widehat{\beta}_{j(i,t)t}X_{it} - \widehat{\lambda}_{j(i,t)t}\right)}{\sum_t \widehat{\alpha}^2_{j(i,t),t}}. \tag{9}$$

Like a standard fixed-effects model, the interactive fixed-effects model also suffers from the incidental parameter problem (Lancaster, 2000). Each $\widehat{\theta}_i$ is estimated only from the finite-sample observations of a specific worker. Hence, the estimator $\widehat{\theta}_i$ would be noisy and would not converge to θ_i. Although the values of the θ_i parameters are not the focus of the analysis, this would still bias the estimates for α_{jt}. While in a standard fixed-effects model, the fixed-effects can be absorbed by demeaning the data, this approach would not work in an IFEM.

The estimates of α_{jt} will suffer from a measurement error. Formally, appendix D.1 shows that the while for the true parameters $E\left[\theta_i\varepsilon_{ijt}|i \in E_{jt}\right] = 0$, the empirical moment does not converge to zero ($E\left[\widehat{\theta}_i\widehat{\varepsilon}_{ijt}|i \in E_{jt}\right] \neq 0$). The measurement error in $\widehat{\theta}_i$ implies that the least square estimator for α_{jt} is inconsistent (Bound et al., 1994). Appendix D.1 provides an analytical expression for the bias.

A common solution to handle measurement error problems is to use an instrumental variable (Wald, 1940; Durbin, 1954).[7] In particular, let Z be an IV satisfying the following

[7]For example, Holtz-Eakin et al. (1988) estimate an IFEM with lagged outcomes as IVs. This assumes that ε_{ijt} are not serially correlated. Alternatively, Ahn et al. (2001) assume that ε_{ijt} has a constant variance

condition for all j, t,

$$E\left[Z_i \varepsilon_{ijt} | i \in E_{jt}\right] = 0. \tag{10}$$

Appendix D.1 shows that under strict exogeneity (similar to a standard fixed effects model as in Chamberlain, 1984) this condition implies that the IV is uncorrelated with the measurement error in $\widehat{\theta}_i$.

With the IV, the model parameters can be estimated using the method of moments. Specifically, I find the vector of parameters that solve the following equations for every combination of occupational category and year,

$$\begin{array}{rcll} m^1_{j,t}(\alpha,\beta,\lambda) & = & \frac{1}{|E_{jt}|}\sum_{i\in E_{jt}}\left(\log w_{ijt} - \beta_{jt}X_{it} - \lambda_{jt} - \alpha_{jt}\widehat{\theta}_i\left(\log w_i, X_i, \alpha,\beta,\lambda\right)\right) & = & 0 \\ m^X_{j,t}(\alpha,\beta,\lambda) & = & \frac{1}{|E_{jt}|}\sum_{i\in E_{jt}} X_{it}\left(\log w_{ijt} - \beta_{jt}X_{it} - \lambda_{jt} - \alpha_{jt}\widehat{\theta}_i\left(\log w_i, X_i, \alpha,\beta,\lambda\right)\right) & = & 0 \\ m^Z_{j,t}(\alpha,\beta,\lambda) & = & \frac{1}{|E_{jt}|}\sum_{i\in E_{jt}} Z_i\left(\log w_{ijt} - \beta_{jt}X_{it} - \lambda_{jt} - \alpha_{jt}\widehat{\theta}_i\left(\log w_i, X_i, \alpha,\beta,\lambda\right)\right) & = & 0 \end{array}. \tag{11}$$

The model includes $3 \times J \times T - 2$ independent parameters where J is the number of occupational categories, and T is the number of time periods.[8] Because of linear dependence between the equations, there are also exactly $3 \times J \times T - 2$ independent equations, and the model is exactly identified.[9]

While the estimation procedure is not identical to two-stage least squares (TSLS), the intuition is very similar. In Appendix D.1 I show that the α estimates satisfy

$$\widehat{\alpha_{jt}} = \frac{COV\left(Z_i, \widetilde{\log w_i} | i \in E_{jt}\right)}{COV\left(Z_i, \widetilde{\widehat{\theta}_i} | i \in E_{jt}\right)}, \tag{12}$$

where $\widetilde{\log w_i}, \widetilde{\widehat{\theta}_i}$ are the residuals of $\log w_i, \widehat{\theta}_i$ from a regression on all interactions of X_{it}. occupational categories and year dummies. This estimator is the coefficient on the IV in the reduced form (where the outcome is log wages), divided by the coefficient on the IV in the first stage (where the outcome is $\widehat{\theta}_i$). Since the IV is uncorrelated with the residuals or the measurement error, it yields a consistent estimator for α_{jt}.[10]

I use years of schooling as the IV. Based on Equation 12, the estimator $\widehat{\alpha_{jt}}$ is the pre-

instead of using an IV.

[8]There are two degrees of freedom since θ_i can be identified only up to a linear transformation. Therefore, I pin $\alpha_{Abstract,1980} = 1$ and $\lambda_{Abstract,1980} = 0$.

[9]From the construction of $\widehat{\theta}_i$, there exists a linear combination of the moments $m^1_{j,t}(\alpha,\beta,\lambda)$ as well as $m^Z_{j,t}(\alpha,\beta,\lambda)$ that equals zero for each choice of parameters.

[10]A TSLS estimator cannot be applied since the value $\widehat{\theta}_i$ depends on the value of α_{jt}.

mium for a year of schooling in occupational category j in year t, divided by the link between years of schooling and the estimated skill $\widehat{\theta}_i$ in the same category and year. Therefore, the link between schooling and skill is allowed to vary across occupations and change over time. This is a critical feature of this design, as the model emphasizes that the composition of skills is different across occupations and to change over time.[11] Intuitively, this estimator solves the measurement error bias by aggregating across many observations. While $\widehat{\theta}_i$ is noisy, the noise converges asymptotically to zero when averaging over all workers in the same schooling category.

I use a sensitivity analysis (Andrews et al., 2017) to show that the changes in the return to skill within an occupational category are estimated based on workers staying in this category. Intuitively, within an occupational category, α_{jt} would decrease (increase) over time if on average the wage gap between workers with different levels of θ_i would decrease (increase). The results are presented in Appendix D.2. I find that the changes in α_{jt} within an occupational category are almost entirely driven by workers in that category.

By contrast, differences in the return to skill across occupational categories are estimated based on movers between the categories. The sensitivity analysis shows that the difference in α_{jt} across occupations depends on workers in both occupational categories. Intuitively, the differences in α_{jt} across two categories depend on whether the wage gap between high- and low-skilled workers increases or decreases when they both switch between these two categories. These results clarify that if moves are not exogenous (a violation of strict exogeneity), this will bias the differences in α_{jt} between occupations. However, it will not affect the trend within the routine occupation as it is identified by workers who stay in the routine occupation.

The model is biased if the IV does not satisfy the condition in Equation 10. This will occur when the IV is correlated with skills that are not captured in θ_i. For example, the condition would be violated if θ_i captures mainly cognitive analytical skills, yet years of schooling is also correlated with other psychological skills that affect wages such as persistence or self-control.[12]

Even if the model assumption does not hold, the results would still be informative on whether technology is skill-replacing. Appendix D.3 derives an analytical expression for the α_{jt} parameters in a case of a simple multi-skill setting. In this case, θ_i is a weighted av-

[11] Carneiro and Lee (2011) show empirically that the composition of skill by years of schooling changes over time.

[12] Equation 10 does not imply that years of schooling must have a causal effect on wages. Since I am estimating the return to skill and not the return to education, concerns about selection to education are irrelevant. For example, in a pure signaling model in which workers choose years of schooling only based on their skill θ_i, and education has no causal effect on wages, this condition will hold.

erage of skills that have large impact on wage and whose returns experience similar time trends. However, when using an IV, the estimated returns to skill $\widehat{\alpha}_{jt}$ reflect the weighted average of returns to skills that are correlated with the IV. While this does not capture the return to θ_i, this still quantifies the trend in the return to skills that are correlated with years of schooling. Therefore, the estimates are still informative about the aggregate trend in return to skills across occupations.

The sensitivity analysis in Appendix D.2 can also be used to bound the potential bias. To generate a bias in the trend of $\alpha_{R,t}$, there must be a trend in the correlation of years of schooling with the error term. Given reasonable values for the overall correlation of years of schooling with log wages I calculate conservative bounds for the bias. I find that even a large bias is unlikely to change the estimated trend in $\alpha_{R,t}$ substantially.

2.2 Decomposing Wage Polarization

Even if all the predictions of the SR-RBTC model were corroborated by the data, it would not disqualify other mechanisms that are potentially occurring simultaneously. For example, institutional changes such as an increase in the real minimum wage (Piketty, 2014), or a decrease in the unionization rate (Firpo et al., 2013) can also coincide with SR-RBTC and potentially explain a significant portion of the trends as well.

Quantifying the importance of various potential explanations is often done using decomposition methods. These methods were proven especially useful in the study of the rise in income inequality in the 1980s. By showing that a large portion of the rise in inequality is driven by the rise in the return to education, they provided some of the most important evidence for skill-biased technological change.

Previous decomposition attempts have found that RBTC can account only for a small portion of the wage trends at the bottom of the distribution (Autor et al., 2005; Firpo et al., 2013). Earlier models of (skill-neutral) RBTC hypothesized that the recent wage trends were driven by changes in occupation premiums. Such changes are expected to be captured by the price component of various decomposition methods (e.g., Juhn et al., 1993; DiNardo et al., 1996; Firpo et al., 2009). Yet, the price component was not large enough to explain the main wage trends during this period, leaving room for other potential drivers.

Moreover, the decomposition methods discussed above are unable to quantify the impact of SR-RBTC, as SR-RBTC violates the ignorability assumption that underlies them. Commonly used decomposition methods assume that the distribution of wages conditional on observables does not change when the distribution of observables changes.[13] In

[13]Formally, ignorability assumes that the conditional distributions of wages $F_{w|X}(w|X=x)$ does not vary

this context, the critical observables are occupations or occupation characteristics. This assumption is innocuous if wages decline uniformly for all workers in routine occupations, as the skill-neutral RBTC model predicts. However, in an SR-RBTC the distribution of wages within occupations changes considerably, violating this assumption. This is because both the distribution of skill and the return to skill are changing within occupations. This generates a change in the wage distribution that is driven by the interaction of an observed characteristic (occupations) and an unobserved characteristic (skill). Most decomposition methods cannot accommodate such interactions without making strong assumptions, such as ignorability, which are violated in SR-RBTC (Fortin et al., 2011).

Igorability is also assumed in a recentered influence function (RIF) regression. This method was used by Firpo et al. (2013) to show that at least some share of wage polarization can be attributed to RBTC. Firpo et al. (2013) were the first to document that inequality trends within occupations are asymmetric, and inequality drops in routine occupations, as predicted by the SR-RBTC model. They also suggested a model where the return to skill varies by occupation. However, the RIF regression they used cannot fully account for the impact of these trends on the overall wage polarization, due to the ignorability assumption. Specifically, RIF regression is valid when either skill and occupations are independent, or when the distribution of unobserved skill is held fixed within occupations (Firpo et al., 2009). These assumptions are violated in the model, as well as in the data, as I will show.

To address this problem, I use a different decomposition, based on the skewness of the log wage distribution. Wage polarization can be measured with skewness, the third standardized moment.[14] For a random variable Y skewness is defined as

$$S(Y) = E\left[\left(\frac{Y - E\left[Y\right]}{\sigma}\right)^3\right] \tag{13}$$

It provides a measure of the asymmetry of the distribution relative to the mean. Appendix Figure A2 demonstrates the link between skewness and wage polarization by plotting the derivative of the empirical influence function at each quantile for a standard normal distribution. Intuitively, the figure shows the effect of a small increase in log wages on the skewness, for each quantile of the distribution, when log wages are normally distributed.

over time. This assumption implies invariance to conditional distributions, where $F_{w|X}\left(w|X = x\right)$ does not change when the marginal distribution of X (F_X) changes.

[14]In other contexts, polarization is typically measured with the fourth moment of the distribution (kurtosis). However, the term wage polarization refers to the polarization of the change in wages, where wages increase mostly at the top and at the bottom. The log wage distribution itself is not becoming more polarized or bipolar, and therefore the kurtosis will not necessarily change.

In particular, it shows that skewness increases exactly when wages at the edges increase relative to the middle. This pattern aligns quite well with the observed trends in wages by quantile that were shown by Autor et al. (2006) and I replicate in Section 6.1.

The main advantage of using skewness is that it has a simple decomposition. Letting Y be the standardized logarithmic wages, X be the category by which we wish to decompose, and μ_3 be the third centralized moment ($\mu_3(Z) = E\left[Z^3 - E[Z]^3\right]$), we get

$$S(Y) = \mu_3(Y) = \underbrace{E\left[\mu_3(Y|X)\right]}_{Within} + \underbrace{\mu_3(E[Y|X])}_{Between} + \underbrace{3COV(E[Y|X], V[Y|X])}_{Correlation}. \tag{14}$$

This decomposition was discussed by Mincer (1974), but was never used on economics data. It is the third-moment equivalent of the variance decomposition formula.[15]

The first and second components are quite standard. The first component $E[\mu_3(Y|X)]$ can be thought of as a "within" component. It captures the remaining skewness within each category. This component increases when the division into categories is orthogonal to the increase in skewness, and therefore can be thought of as a residual component. The second component, $\mu_3(E[Y|X])$ captures skewness between groups, which is the skewness due to differences between group averages. This component increases if wage polarization is due to a similar change in wages for all workers in a group, compared to other groups (e.g., a relative decline in routine wages).

The third component captures the correlation between wage levels and inequality in each categorical group. Formally, this component measures the covariance between the conditional mean and variance for each value of X. When highly paid groups also have larger inequality, inequality will be higher at the top than at the bottom of the overall distribution, making the distribution more positively skewed.

With this covariance component, we can capture trends that violate ignorability. The covariance component allows us to have interactions between unobserved characteristics (e.g., skill), and observed characteristics (e.g., occupation). Hence it can quantify changes to the wage structure that cannot be detected by other methods.

Thus, the covariance component allows us to measure the wage impact of SR-RBTC. According to SR-RBTC, inequality increases in the abstract occupation, since skill gaps increase when lower-skilled workers join this occupation. By contrast, inequality decreases in the routine occupation (Theorem 4). The effect on inequality in the manual occupation could go either way, yet since manual occupations are only a small portion of all occupa-

[15] Variance can be decomposed into $V(\log w) = \underbrace{E[V(\log w|X)]}_{Within} + \underbrace{V(E[\log w|X])}_{Between}$.

tions, their overall impact would be small. Taken together, we expect a wage trend that is exactly captured by this component: inequality is rising in the higher-paying abstract occupations and declining in the lower-paying routine occupations. Indeed, the covariance component will turn out to be responsible for most of the increase in skewness during the period of wage polarization.

Skewness decomposition also provides another suggestive test of whether RBTC is skill-replacing or skill-neutral. In SR-RBTC, most of the increase in skewness is due to the increase in the covariance component. By contrast, in a skill-neutral RBTC, most of the increase in skewness should be due to the between component. This is because according to a skill-neutral RBTC, recent wage trends are driven mostly by the decrease in price of routine goods p_R. Such a decrease in price has an identical effect on all workers in a given occupation. This is exactly the case when we expect a large effect on the between component. However, this test on its own is insufficient, since the covariance component can also increase in a skill-neutral RBTC due to some particular compositional changes. This is why I also estimate the IFEM, which shows that inequality within routine occupations declines regardless of any compositional changes.

Skewness decomposition has several important properties that make its results more robust. It allows quantifying wage polarization using a single index. Similar to variance decomposition, skewness decomposition breaks skewness into independent components. This means that there is no problem of path dependence nor any need to arbitrarily define a baseline year, as in other popular decomposition methods (Fortin et al., 2011).

While in this paper I use skewness decomposition to study wage polarization, it could also be applied to any distribution where the third moment is of interest. There are various cases in economics where skewness has important implications. Some examples are the distribution of the return to patents, firm productivity, the distribution of capital ownership, and raw wages (without logs). Any variation in these distributions over time or across places can be analyzed with skewness decomposition. To simplify and encourage the usage of skewness decomposition by more researchers, I provide an R package that implements it.[16]

3 Data

This paper combines three data sources. To estimate the interactive fixed-effects model panel data is required. I use the Panel Study of Income Dynamics (PSID) between 1980–

[16]The package implements both skewness and variance decomposition and provides an analytical calculation of the standard errors.

2017. This data was chosen due to its long panel. I measure income using hourly wage (annual income divided by hours worked) as this best captures the real price of labor, which is the focus of the model.I use the full core sample (SRC) without weights for every individual whose wage is available.

Whenever a panel structure is not needed, including the skewness decomposition exercise, I use a larger dataset from the Current Population Survey Outgoing Rotation Group (CPS-ORG). The CPS-ORG provides the most accurate representative sample of hourly wages (Lemieux, 2006a). I use the same sample definition as given in Acemoglu and Autor (2011). Observations with missing wages are dropped. The main results hold when using imputations instead. Sampling weights are used in all analyses.

One important limitation of hourly wage data is its relatively high level of measurement errors. This problem is particularly severe at both tails of the distribution. Misreporting of working hours could lead to extremely high or extremely low values of hourly wages. Therefore, I drop the top and bottom 5% of the positive wages throughout the paper. The level of 5% minimizes the loss of data, without generating substantial fluctuations between consecutive years in the skewness estimator. It is also similar to the data cuts made in earlier papers in this literature (Katz and Murphy, 1992; Autor et al., 2008). Smaller cuts also yield similar but noisier results, particularly for skewness estimates.[17]

Most of the skewness decomposition analysis is focused on the years between 1992–2002. This is due to a significant revision of the occupational classification system that took place before and after this period, which makes comparisons to other years less precise. As I will show, most of the increase in polarization occurred during this time period. For robustness, I also implement an analysis over a longer period using the occupational crosswalk constructed by Autor and Dorn (2013), and show the main results hold.

I maintain a consistent definition for routine occupations, similar to earlier papers in the literature. I first translate all versions of occupational coding into a uniform coding, using the Autor and Dorn (2013) crosswalk. I then define all administrative, operator, and production occupations as routine, based on their 1-digit category. All managerial, professional, and technician occupations are classified as abstract. Sales, services, and agricultural occupations are classified as manual.[18] This is a similar classification to that used in previous studies (e.g., Acemoglu and Autor, 2011) with one exception: I do not classify sales occupations as routine occupations.[19] The analysis by 1-digit occupational

[17]Cornfeld and Danieli (2015) analyze skewness in Israeli data where measurement errors are less severe, using the entire distribution. Their results are similar to the U.S. results I document in this paper.

[18]Unlike some of the other papers in this literature, I do not exclude agricultural workers.

[19]Classifying sales as a routine occupation does not strongly affect the results as it is a small share of workers compared to the other routine occupations.

category is fully consistent with previous literature.

Finally, I use data from the Occupational Information Network (O*NET) to measure the routine intensity of each occupation more accurately, in cases where a continuous index can be accommodated. This dataset contains 400 scales to describe various aspects of each occupation, based on a worker survey. I use the same index for routine intensity as Acemoglu and Autor (2011). This index summarizes six questions that proxy the routine level of the job.

More details on the data are provided in the data appendix (E).

4 The Decline in Return to Skill in Routine Occupations

The main prediction of the SR-RBTC model is that the return to skill declines in the routine occupation. Based on Theorem 1, such a decrease in the skill gap is only consistent with a skill-replacing RBTC ($\epsilon < 0$). In this section, I provide empirical evidence for this prediction. I first show reduced-form evidence that the education premium has declined in routine occupations. I then use an interactive fixed-effects model to get a direct estimate of the return to skill and its trends in all occupational categories.

4.1 The Education Premium in Routine Occupations

In cross-sectional data, the education premium reflects not only the return to education but also potential differences in unobserved ability between education groups (Card, 1999). Therefore, any changes in the education premium over time could reflect both changes in the return to skill and changes in the skill composition of workers across education levels. In order to focus on the changes in the return to skill, I measure the changes in the education premium in a panel setting, controlling for differences in ability.

I estimate the education premium using a (standard) fixed-effects model. I define education level by years of schooling. Specifically, I use the following model to estimate the education premium for the subsample of workers in routine occupations in the PSID

$$\log w_{it} = \gamma_t S_i + \psi_t + \theta_i + \rho_t X_{it} + \varepsilon_{ijt} \tag{15}$$

where S_i measures years of schooling, ψ_t are year fixed effects, θ_i are worker fixed effects, and X_{it} is an additional control for experience squared. I focus on the coefficient γ_t, which captures how much the wage gap between more- and less-educated routine workers changes over time.

This specification focuses on changes in the return to education, holding skill composition fixed. This is done by focusing on wage changes over time. Assuming that the skill differences between workers is constant over time, the fixed effects (θ_i) guarantee that any trend in γ_t is not driven by compositional changes, which are controlled for. For instance, if $\gamma_t > \gamma_{t+1}$, it implies that for a given set of workers, the gap between more and less educated workers is decreasing over time.

I find that the wage gap between more and less educated workers has decreased since the late 1980s. Appendix Figure A3 plots the estimated coefficient γ_t relative to its highest value in 1987.[20] The results indicate that since the late 1980s the education premium has declined in routine occupations. This is consistent with the timing when bottom-half inequality starts to decrease. In order to account for skill differences within education levels, compare the trends in routine occupations to abstract and manual occupations, and improve precision, I estimate the interactive fixed-effects model.

4.2 IFEM Results

The IFEM estimation results are consistent with the predictions of the SR-RBTC model. I estimate Equation 7 as described in Section 2.1. I start by analyzing the estimation results in detail for a specific year, before analyzing the full sample period.

In 1987, approximately the last year before bottom-half inequality starts to decline, the estimation results are consistent with the initial pre-SR-RBTC model predictions. Panel A of Figure A4 plots the expected log wage of workers in 1987 as a function of their skill θ_i in the three different occupational categories. This figure highly resembles the theoretical prediction in Figure 2a. Return to skill, α_{jt}, which corresponds to the slopes in the graph is highest in the abstract category, lower in the routine category, and lowest in the manual category. As a result, the lowest-skilled workers can earn their highest wage in manual occupations, the highest-skilled workers can earn the highest wage in abstract occupations and middle-skilled workers earn the highest wage in routine occupations. Moreover, the indifference point between routine and abstract occupations is very close to zero, which is the average skill level in the economy. Hence, if workers sort optimally into occupations, the highest-skilled routine workers have an average skill level.

Figure 3 extends the analysis over more years and shows that the return to skill has steadily declined in routine occupations since the late 1980s. This is approximately when wage polarization and the decline in routine employment start. The figure plots $\alpha_{j,t}$ in log units for the three broad occupational categories. Since there is a degree of freedom

[20] Since this specification is only focused on changes, it cannot estimate the absolute return to education.

in this estimation, I pin $\log \alpha_{A,1980}$ to 0. The figure shows that the return to skill in routine occupations has dropped substantially. The value of $\log \alpha_{R,t}$ decreased by more than 0.7, which correspond to a 50% cut between its peak value in 1987 and 2017. This means that conditional on experience, the average return to skill was reduced by more than half. Hence, skill gaps were substantially compressed in routine occupations among stayers.

The other two occupational categories did not see a similar sharp decline. For manual occupations, $\log \alpha_{M,t}$ remains very stable at around -0.3. In the late 1980s, the return to skill in manual occupations was below that in routine occupations, as assumed in the model (Condition 1). Yet because the return to skill declined in routine occupations while remaining relatively stable in manual occupations, their ranking reversed in the 1990s. This matches the prediction of Theorem 5. Abstract occupations also see some decline in return to skill, mostly after 1994, supporting recent evidence on a reversal in demand for cognitive skills (Beaudry et al., 2016). While interesting in itself, this decline is significantly smaller compared to the decline in routine occupations and is not large enough to change the ranking of occupational categories based on their return to skill.[21]

The same pattern of results emerges when using nine occupational categories based on 1-digit occupational coding. I estimate Equation 7 allowing the return to skill (α_{jt}) to vary by 1-digit occupation category and year. Figure 4 plots the coefficient for α_{jt} in log units for each 1-digit occupation category in three years: 1985, 1997, 2011.[22] In 1985, before wage polarization starts, the return to skill is in accordance with the assumption of the model: the return to skill is largest in the abstract occupations (managers, professionals, and technicians), lowest in the service occupations, and in between for routine occupations (administrative, operator and production). One noticeable exception is sales occupations, which seem to have a return to skill in the range of the abstract occupations despite often being classified as routine.

The return to skill then drops only in the routine occupations. All four routine occupations, including sales, experienced a decline in return to skill between 1985–1997. At the same time, the other four occupations (managers, professionals, technicians, and services) experience an increase in their return to skill. Later, between 1997–2011, there is a decline in the return to skill in all occupations, but it is sharper in the routine ones, and especially in the administrative category. By 2011, the categories with the lowest return

[21]Figure 3 documents a decline only in return to skill in abstract occupations ($\frac{\partial \log \varphi_A(\theta_i)}{\partial \theta_i}$), and not a general decline in the occupation premium (p_A). Wages in abstract occupations are still higher relative to other occupations in the end of the period. Inequality within abstract occupations is also still rising, possibly due to lower-skilled workers joining these occupations. See Section 6.2 for further discussion.

[22]The trend for agricultural workers, who comprise only a small share of the labor force, is similar to other manual workers in the service sector, and reported in Appendix Figure A5.

to skill are the four routine occupations. Appendix Figure A5 shows the trends in more detail and plots the value of $\alpha_{j.t}$ for each 1-digit occupational category by year.

Overall, these results fit well with the predictions of the skill-replacing RBTC model. Theorem 1 shows that a decrease in the skill premium is consistent only with a skill-replacing technology (and not with a skill-neutral or skill-enhancing technology). The IFEM shows since the late 1980s, wages of workers who stay in routine occupations has become more similar over time, consistent with the prediction of Theorems 3 and 5.

The IFEM assumes a one-dimensional skill as in the model in Section 1. This implies that the wages of workers in one occupational category would be correlated with their wages if they move to another occupational category. In other models, this prediction would not necessarily hold. In a multi-dimensional skill model (Appendix B.3), the wages of workers would be correlated across occupations only if skills are correlated. Similarly, Acemoglu and Restrepo (2024) propose a model in which automation decreases inequality among workers in routine occupations by decreasing rents. If higher wages in routine occupations reflect mostly rents, then the high-earning workers in routine occupations would not be high earners in another occupational category.

Appendix F shows two empirical exercises that support the unidimensional skill assumption. First, I measure the rank correlation of movers across the three occupational categories. Second, I estimate a more general IFEM with multidimensional skills in which θ_{ij} can vary by occupational category. I find that both wages and skills are highly correlated across occupational categories. This is consistent with a decrease in returns to a skill that is at least correlated with the relevant skills in other occupations.

5 The Reversal of Comparative Advantage

This section presents evidence that the employment decline in routine occupations was predominantly driven by higher-skilled routine workers. As a result, routine occupations' average skill level fell below that of manual workers. This explains why inequality at the bottom of the distribution stopped declining and started rising back, even though SR-RBTC continued.

I estimate each worker's skill using the interactive fixed-effects model. For each worker i, I estimate $\widehat{\theta}_i$ using Equation 9. Since $\widehat{\theta}_i$ is estimated separately for each worker, it is based only on a small number of observations, making it a very noisy estimate. To solve this problem, I analyze the average value of $\widehat{\theta}_i$ for large groups of workers.

Specifically, I examine the average value of $\widehat{\theta}_i$ for each occupational category in a given year. I do this either by dividing occupations into the three main categories (abstract, routine, manual) or by the 1-digit classification. I normalize $\widehat{\theta_i}$ to have a mean of zero for each cohort based on the year in which the worker entered the labor market. Therefore this is an analysis of relative skill within cohorts.

I find a substantial and steady decline in the average skill level of workers in routine occupations. Figure 5 plots the average skill level by occupational category and year. At the beginning of the sample period, in the early 1980s, workers in routine occupations were middle-skilled. Their average $\widehat{\theta_i}$ was very close to zero, which is the population average. Over the next three decades, the skill composition of workers in routine occupations steadily declined, reaching -0.2 at the end of the period. Since the return to skill in routine occupations in 2017 ($\alpha_{R,2017}$) was about 0.3 (-1.2 in log units), it follows that if workers in routine occupations in 2017 had been as skilled as they were in 1980 their wages would have been 6% higher.

At the end of the sample period, routine occupations employed the lowest-skilled workers. While the average skill level of workers in routine occupations declined, the average skill level of manual and abstract workers remained fairly stable. As a result, in 2015, the average skill level of routine workers fell for the first time below that of manual workers.[23] In appendix G I show evidence that the decline in the average skill in routine is driven primarily by a decrease in the share of middle-skill workers who join routine occupations when they enter the labor force.

I find very similar trends using the 1-digit classification of occupations. Figure 6 plots the average skill level in 1985, 1997, and 2011. At the beginning of the period, routine occupations were middle-skilled, where all four routine occupational categories (including sales) had a skill level between -0.1 to 0.1. In the following periods, administrators, operators, and production workers became significantly less skilled.

Other occupations, such as services, saw an increase in the average skill level of their workers. In 2011, service workers had higher skills than administrative workers and operators. This fits well with the prediction of the model that since the lowest-skilled workers were now employed in routine occupations, manual occupations such as services would see an increase in the skill level of their workers. Appendix Figure A6 plots the results for all years.

I also find that while in 1990 many middle-wage workers worked in routine occupations, by 2010 this is no longer the case. Figure 7 plots the average routine intensity index

[23] Workers in the routine occupations still earn higher wages than those in manual occupations. One reason for this is their higher level of experience (Autor and Dorn, 2009), which is not reflected in this within-cohort analysis.

for 20 quantiles of the wage distribution. This index is based on the routine intensity of the occupations of workers in this quantile (see Appendix E for more details on the index). Between 1990 and 2000, routine intensity fell mostly for wages above the 40th percentile. In the following decade between 2000–2010, routine employment fell mostly between the 20th and 40th wage percentiles, perhaps because not many workers in routine occupations were left in higher percentiles.

There are at least two potential explanations for this decline in the number of middle-wage routine workers. First, the decline could be driven by middle-skilled workers leaving or never joining routine occupations. Workers in the middle and upper half of the wage distribution may have switched to occupations with a lower routine intensity. This corresponds to a decline in θ_1 in the SR-RBTC model, as predicted by Theorems 3 and 5. Second, the decline could be driven by lower wages in routine occupations. This corresponds to a decline in p_R, as predicted by Theorem 2 and empirically shown by Cortes (2016). In either case, workers in routine occupations are now concentrated in much lower percentiles of the wage distribution than they were in the past. Therefore, any further RBTC is not expected to generate a decline in middle wages.

These findings fit very well with the model's predictions. Since wages decline mostly for the highest-skilled routine workers, they are the first to leave these occupations, as predicted by Theorem 3. As a result, the overall skill level decreases, as predicted by Theorem 4. At some point, the average skill level of workers in routine occupations falls below that of manual workers, as predicted by Theorem 5.[24]

This compositional change explains why inequality at the bottom of the distribution stopped decreasing and started increasing. Once middle-wage workers were no longer employed in routine occupations, they were no longer affected by SR-RBTC as before. Since workers in routine occupations were now the lowest-skilled workers, any further SR-RBTC was working mostly against the lowest earning workers. This generated an increase in inequality at the bottom half of the wage distribution.

6 Quantifying the Overall Impact of SR-RBTC

So far, I have shown that the main predictions of the model are consistent with the data. This section shows that SR-RBTC is also substantial enough to account for almost the entire trend of wage polarization. Using skewness decomposition, I show that wage polarization is driven almost entirely by occupational trends. Moreover, the effect is not

[24] Under particular parameters, the average skill could also decline in a skill-neutral RBTC. However, a skill-neutral RBTC is inconsistent with the decline in return to skill (Figure 3).

driven by the drop in the premium in routine occupations as predicted by a skill-neutral RBTC. Instead, I find that the decline in inequality in low-paying routine occupations is the main driver of wage polarization, consistent with the SR-RBTC model.

6.1 Evidence From Skewness Decomposition

I start by showing that skewness is indeed a good measure of wage polarization. Figure 8 shows the trend in skewness between 1979–2012. The rise in skewness aligns very well with the timing of wage polarization as depicted in Figure 1. Skewness increased between the late 1980s and the early 2000s, exactly when the 90/50 gap was rising and the 50/10 gap was falling.

The rise in skewness is driven by trends in all parts of the wage distribution. An increase in skewness occurs when the distribution becomes more tilted toward the left-hand side. This corresponds to an increase in the gap between middle and high wages and a decrease in the gap between middle and low wages. Appendix Figure A7 presents a bin scatter of the change in wages between 1992–2002, for 20 quantiles. The figure shows a U shape pattern, as previously shown by Autor et al. (2006) and Autor et al. (2008). The U shape received qualitatively resembles the EIF derivative plotted in Figure A2. This suggests that skewness rose in this period because of the rise in wages both at the top and at the bottom of the distribution, making it a good fit to measure wage polarization.

I decompose the rise in the skewness of the distribution into three components, namely, within, between, and covariance components, as described in Equation 14, for different choices of categorical groups (X). I first focus on the period between 1992-2002 since data on other years uses different occupational coding (see Section 3). As Figure 8 shows, this time period includes a big portion of the overall increase in skewness.

Decomposing by occupations can explain almost the entire rise in skewness. Table 1 presents the decomposition of the rise in skewness between 1992–2002 by 3-digit occupational coding. I report the values of each component and its overall contribution to the rise in skewness. Figure 9 depicts the annual change in each component, as well as in the sum of the three, which equals the total change in skewness.[25] The first conclusion from this exercise is the importance of occupational trends in explaining the rise in skewness. The within component, which captures the part that is unrelated to occupational trends, explains only 7% of the overall increase. That small share might also be the result of classification errors. Hence 93% of the rise in skewness is related to occupational trends.

[25] Appendix Figure A8 performs the same exercise using imputed wages for observations where wages are not reported and reaches very similar results.

Most of the increase in skewness is driven by the covariance component. This is indicated by the blue area in Figure 9, which captures 79% of the rise in skewness. These results imply that the rise in skewness is driven primarily by the increasing correlation between the mean and the variance of log wages in occupations. In other words, the rise in skewness is due to the growing correlation between wage levels and inequality levels in each occupation. As I discussed in Section 2, this type of correlation is not captured by other decomposition methods, which is why earlier work potentially underestimated the contribution of occupational trends.

These results align better with the hypothesis that RBTC is generating wage polarization than with institution-related hypotheses. The theory of RBTC argues that its effect is driven predominantly through occupations. Therefore, the fact that wage polarization, as measured with skewness, is driven by occupations greatly supports this hypothesis. By contrast, institutional changes do not operate directly through occupations.

Moreover, the results are most consistent with a skill-replacing RBTC. Earlier models of a skill-neutral RBTC (Autor et al., 2006; Acemoglu and Autor, 2011) argue that there is a drop in the price of routine tasks, which makes wages fall equally for all workers in routine occupations. Such a trend would have been captured by the between component as it generates the same effect for all workers in the same occupation. However, this component generates only 15% of the overall rise in skewness. Instead, the substantial rise in the covariance component suggests that the effect is mostly driven by the asymmetric trends within occupations. This is consistent with a decrease in the return to skill in low-paying routine occupations, as described in the SR-RBTC model.

The results are not driven by any other worker characteristic in the data. Since occupations are correlated with workers' skill levels or industries, I verify that occupations are not proxying for some other worker characteristics. In Appendix Figures A9 and A10 I show the same decomposition results by industry, education, and experience. Clearly, in those cases the within component is much larger, suggesting that a great portion of the trend in skewness is unrelated to these categories. Moreover, most of the increase in the between and covariance components in those decompositions is due to their correlation with occupations. Appendix H discusses how to decompose by more than one category using a linear model. I use this method to decompose jointly by occupation and industry or education. The results in Appendix H show that the increase in skewness is driven almost entirely by occupation and not other observables.

Looking at a longer time period yields similar results. Appendix Figure A11 plots a decomposition by occupations between 1988, when skewness starts to rise, and 2012. Within this time period the occupational coding changes and therefore I use the Autor and Dorn (2013) occupational crosswalk that generates a unified coding across periods. However, changes in the baseline coding might still generate a measurement error when the coding changes between 1991–1992 and between 2002–2003. With that caveat in mind, we still see a similar pattern in the earlier period between 1988–1992. Most of the increase in skewness is driven by the covariance component. In the period after 2002, when wage polarization stops, skewness is stable, as are the three different components.

6.2 The Decline in Inequality Within Routine Occupations

The increase in correlation between wage levels and inequality that drives wage polarization could be attributed to different explanations. The increase could be due to trends in wage levels, wage inequality, or perhaps the composition of workers in each occupation. The following section presents evidence that the main driver is the decrease in inequality in low-paying routine occupations, as predicted by Theorem 4.

During the 1990s, inequality trends within occupations were strongly correlated with the wage levels in those occupations. High-paying occupations saw an increase in inequality while low-paying occupations saw a decrease (Lemieux, 2007). Figure A12 replicates this by plotting the change in the variance of log wages from the beginning of the studied period (1992/3) to its end (2001/2) as a function of mean log wages. Changes in inequality, measured with the variance of log wages, are correlated with the occupation wage levels.

In fact, the trends in within-occupation inequality can explain the full rise in the covariance component. The covariance component equals

$$COV\left(E\left[Y|X\right],V\left[Y|X\right]\right)=\sum_{x}\Pr\left(X=x\right)E\left[Y|X=x\right]V\left(Y|X=x\right)$$

where in this case Y is log wages and X is 3-digit occupations. Most of the increase stems from changes in the variance of log wages in different occupations, $V\left(Y|X=x\right)$. To show this, I fix the share of workers and the expected log wage in each occupation to their averages throughout the period. Thus, I allow only the variance to vary between years. Formally, I calculate the following counterfactual partial-equilibrium covariance for t between 1992 and 2002:

$$\widetilde{COV}\left(\overline{E\left[Y_t|X_t\right]},V\left[Y_t|X_t\right]\right)=\sum_{x}\overline{\Pr\left(X=x\right)}\;\overline{E\left[Y|X=x\right]}V\left(Y_t|X_t=x\right)\tag{16}$$

where $\overline{\Pr\left(X=x\right)}$ and $\overline{E\left[Y|X=x\right]}$ are simple averages of the share of workers and mean log wages over all years between 1992–2002.

I find that the asymmetric trends in within-occupation inequality can explain the entire increase in covariance. Figure 10 compares the real value of the covariance to its counterfactual value from Equation 16. The counterfactual trend closely follows the real trend. Therefore, if the share of workers and the mean log wage in each occupation were held fixed, we would still get the same increase in the covariance, and hence the same increase in skewness and wage polarization. Letting the share of workers or the expected log wage vary while holding other factors fixed does not yield any similar results. This exercise demonstrates that the increase in covariance, and hence the increase in wage polarization, is mostly the result of the asymmetric changes in within-occupation inequality, as measured with the variance of log wages. I.e., inequality is increasing in high-paying occupations and decreasing in low paying occupations.

The drop in inequality in low-paying occupations is driven mostly by routine occupations. Figure A13 presents a bin scatter plot of the changes in within occupation variance for routine and non-routine occupations. I divide occupations into 10 bins separately for routine and non-routine occupations, based on their initial wage decile in 1992. I then plot the mean change in the variance of log wages between 1992–2002. While there is some drop in inequality in low-paying occupations that are non-routine, the trend is substantially stronger for routine occupations. This is consistent with the findings in Firpo et al. (2013) who show using O*NET data that routine occupations tend to have a stronger decrease in variance.

Overall, these findings fit well with the predictions of the model. Most of the wage polarization is related to occupational trends, which supports the explanation of RBTC. The trend is driven mostly by the asymmetric trends in within-occupational inequality. Inequality is decreasing in low-paying, mostly routine occupations, while it is increasing in high-paying occupations. This fits well with the predictions of the SR-RBTC model in Theorem 4, and explains why we see a U-shaped wage trend during the 1990s.

7 Discussion and Alternative Explanations

I conclude this paper by summarizing the empirical puzzles that the SR-RBTC model is able to address, as well as the new empirical facts I document, which also align with this model. I then discuss alternative explanations and highlight which empirical facts they are unable to explain.

The SR-RBTC model explains three empirical facts that could not be explained with a skill-neutral RBTC model (e.g., Acemoglu and Autor, 2011). First, it explains why middle wages declined in the 1990s, even though workers in routine occupations were dispersed across the entire bottom half of the income distribution. The SR-RBTC model predicts that wages would decrease for the highest-earning routine workers, who are empirically located exactly in the middle of the income distribution. It also explains why middle wages stopped declining around the year 2000. The SR-RBTC model predicts that, over time, middle-wage workers would prefer to work in the abstract occupation, as observed in the data. Therefore they are no longer negatively affected by RBTC. Finally, the model explains why we observe only mild wage decreases for workers in routine occupations on aggregate— while wages fall for routine workers with higher skill they increase for routine workers with lower skill. Hence, most of the wage adjustment occurs not for the average routine wages, but for the inequality level within the routine occupation.

I document new empirical facts that are consistent with the SR-RBTC model. I find that the return to skill declined in routine occupations significantly more than it did in other occupations. I also show evidence that the average skill level declined substantially in routine occupations. Together, this leads to a concentration of workers in routine occupations in lower parts of the income distribution. Finally, using skewness decomposition, I show that wage polarization is driven primarily by the inequality trends within occupations—in particular, the decrease in inequality in low-paying routine occupations.

Other explanations do not fit these empirical patterns as well as SR-RBTC. Several alternative explanations for the decline in bottom-half inequality in the 1990s focus on institutional changes, such as an increase in the real minimum wage (Piketty, 2014) or a decline in unionization (Lemieux, 2007). Other explanations focus on high growth and low unemployment rates as potential drivers for the increase in lower wages. Finally, it is possible that trade shocks have led to some of the changes in the wage distribution (Autor et al., 2013). However, none of these explanations is expected to work through occupations more than through education levels or industries. For example, while some occupations are more unionized than others, industries are likely better proxies for unionization status. Moreover, while these mechanisms could generate a decrease in inequality within lower-paying occupations, as part of the overall decrease in lower-half inequality, it is unclear why their impact would be mostly on low-paying routine occupations and not, for example, service occupations.

Among theories that focus on occupational-related trends, SR-RBTC best fits the empirical findings. Generally, theories related to technology seem to fit the data better, as most of the trends are related particularly to routine occupations, which can be automated more easily (Autor et al., 2006; Goos et al., 2014). Skill-neutral or skill-enhancing RBTC models are inconsistent with the clear decline in return to skill in routine occupations. They also do not predict the decrease in skill level in routine occupations.

Jaimovich et al. (2021) suggest a multi-skill model with a skill-neutral technological change. The multi-skill model is potentially more realistic and explains several empirical facts that I did not focus on in this paper. However, it does not explain some of the key facts that were the focus of this paper, including the decline in middle wages, the halt of this decline around 2000, the decrease in the return to skill, and the decline in the average skill level for workers in routine occupations. To explain these facts, the multi-skill model needs to accommodate a skill-replacing technology, as in the multi-skill model in Appendix B.3. The results in Appendix F suggest that the correlation between skills is likely high, at least for workers who switch occupational categories (which are the majority of workers). This correlation also supports the interpretation of a decline in return to skill, rather than a decline in rents (Acemoglu and Restrepo, 2024).

Another theory that could potentially explain the bottom-half inequality trends is a positive demand shock for service occupations (Autor and Dorn, 2013). SR-RBTC also predicts an increase in demand for manual occupations due to the complementarities between occupations. Therefore, several predictions of the SR-RBTC model overlap with the predictions in Autor and Dorn (2013). But one important distinction is the effect on the skill composition of workers in routine occupations. A demand shock in service occupations should attract more workers from the bottom of the skill distribution and so reduce the share of low-skilled routine workers, as they have a comparative advantage in service jobs. However, most of the decline in employment in routine occupations is driven by the highest-skilled routine workers, making it more consistent with SR-RBTC.

Institutional explanations are potentially more relevant for bottom-half inequality in other countries. Other developed countries have more dominant labor market institutions compared to the U.S. (Blau and Kahn, 2002). Such institutions tend to have a larger impact on bottom half compared to upper-half inequality. For example, Broecke et al. (2016) find that minimum wage levels are substantially more associated with bottom half inequality than with upper half inequality. As a result, technological changes could have less effect on bottom-half inequality in other developed countries. This could explain why similar inequality patterns are not seen in other countries, despite having similar patterns of employment decline in routine occupations (Naticchioni et al., 2014; Goos et al., 2014). Interestingly, in Israel, which is one of the few countries to experience similar fluctuations in bottom-half inequality, similar patterns are detected (Cornfeld and Danieli, 2015).

While this paper does not provide causal identification, my findings are consistent with those of previous papers that have studied the causal effect of RBTC on firm wage distribution. Gaggl and Wright (2017) exploit a natural experiment where exposure to technology varies by firm. They find that the new technology generates wage compression among workers in routine occupations in a given firm. In this paper, I show that this wage compression is the main driver of wage polarization and not a side effect.

It is possible that other technological advancements that are not only automating routine tasks have also a skill-replacing nature. This could create a decrease in inequality for workers in non-routine occupations. For example, Noy and Zhang (2023) use a lab experiment to show that generative AI yields a larger productivity boost for lower-skilled workers. Beaudry et al. (2016) argue that after a technology is adopted, the demand for high-skill abstract workers declines. Whether new technology is skill-replacing or not is a critical question that could determine the future of inequality in the labor market.

References

Acemoglu, D. and D. Autor (2011). "Skills, Tasks and Technologies: Implications for Employment and Earnings". In: ed. by O. Ashenfelter and D. Card. Vol. 4, Part B. Handbook of Labor Economics. Elsevier. Chap. 12, pp. 1043–1171.

Acemoglu, D. and P. Restrepo (2022). "Tasks, automation, and the rise in US wage inequality". In: *Econometrica* 90.5, pp. 1973–2016.

— (2024). "Automation and Rent Dissipation: Implications for Wages, Inequality, and Productivity". In.

Ahn, S. C., Y. H. Lee, and P. Schmidt (2001). "GMM Estimation of Linear Panel Data Models with Time-Varying Individual Effects". In: *Journal of Econometrics* 101.2, pp. 219–255.

Andrews, I., M. Gentzkow, and J. M. Shapiro (2017). "Measuring the sensitivity of parameter estimates to estimation moments". In: *The Quarterly Journal of Economics* 132.4, pp. 1553–1592.

Autor, D. (2014). *Polanyi's Paradox and the Shape of Employment Growth*. Working Paper 20485. National Bureau of Economic Research.

Autor, D. and D. Dorn (2009). "This job is "Getting Old": Measuring changes in job opportunities using occupational age structure". In: *American Economic Review* 99.2, pp. 45–51.

— (2013). "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market". In: *American Economic Review* 103.5, pp. 1553–1597.

Autor, D., D. Dorn, and G. H. Hanson (Oct. 2013). “The China Syndrome: Local Labor Market Effects of Import Competition in the United States”. In: *American Economic Review* 103.6, pp. 2121–2168.

Autor, D., L. Katz, and M. S. Kearney (Sept. 2005). *Rising Wage Inequality: The Role of Composition and Prices*. Working Paper 11628. National Bureau of Economic Research.

Autor, D., L. F. Katz, and M. S. Kearney (2006). “The Polarization of the US Labor Market”. In: *American Economic Review* 96.2, pp. 189–194.

— (Apr. 2008). “Trends in U.S. Wage Inequality: Revising the Revisionists”. In: *Review of Economics and Statistics* 90.2, pp. 300–323.

Autor, D., F. Levy, and R. Murnane (2003). “The skill content of recent technological change: An empirical exploration”. In: *The Quarterly Journal of Economics* 118.4, pp. 1279–1333.

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

Blau, F. D. and L. M. Kahn (2002). *At home and abroad: US labor market performance in international perspective*. Russell Sage Foundation.

Böhm, M. J., H.-M. v. Gaudecker, and F. Schran (2019). “Occupation, growth, skill prices, and wage inequality”. In.

Bound, J., C. Brown, G. J. Duncan, and W. L. Rodgers (1994). “Evidence on the validity of cross-sectional and longitudinal labor market data”. In: *Journal of Labor Economics* 12.3, pp. 345–368.

Broecke, S., G. Quintini, and M. Vandeweyer (Feb. 2016). *Wage Inequality and Cognitive Skills: Re-Opening the Debate*. Tech. rep.

Card, D. (1999). “The Causal Effect of Education on Earnings”. In: *Handbook of labor economics*.

Carneiro, P. and S. Lee (2011). “Trends in quality-adjusted skill premia in the United States, 1960-2000”. In: *American Economic Review* 101.6, pp. 2309–49.

Chamberlain, G. (1984). “Panel data”. In: *Handbook of econometrics* 2, pp. 1247–1318.

Cornfeld, O. and O. Danieli (2015). “The Origins of Income Inequality in Israel-Trends and Policy”. In: *Israel Economic Review* 12.2, pp. 51–95.

Cortes, G. M. (2016). “Where Have the Middle-Wage Workers Gone? A Study of Polarization Using Panel Data”. In: *Journal of Labor Economics* 34.1, pp. 63–105.

Cortes, G. M., N. Jaimovich, C. J. Nekarda, and H. E. Siu (2020). “The Dynamics of Disappearing Routine Jobs: A Flows Approach”. In: *Labour Economics* 65, p. 101823.

Dicandia, V. (2023). “Technological change and racial disparities”. In.

DiNardo, J., N. M. Fortin, and T. Lemieux (1996). “Labor Market Institutions and the Distribution of Wages, 1973–1992: A Semiparametric Approach”. In: *Econometrica*, pp. 1001–1044.

Downey, M. (2021). “Partial automation and the technology-enabled deskilling of routine jobs”. In: *Labour Economics* 69, p. 101973.

Durbin, J. (1954). “Errors in variables”. In: *Revue de l'institut International de Statistique*, pp. 23–32.

Eden, M. and P. Gaggl (2018). “On the Welfare Implications of Automation”. In: *Review of Economic Dynamics* 29, pp. 15–43.

Firpo, S., N. M. Fortin, and T. Lemieux (2009). “Unconditional Quantile Regressions”. In: *Econometrica* 77.3, pp. 953–973.

— (2013). “Occupational Tasks and Changes in the Wage Structure”. In: *University of British Columbia Working Paper*.

Fortin, N. M., T. Lemieux, and S. Firpo (2011). “Decomposition Methods in Economics”. In: *Handbook of Labor Economics*. Ed. by O. Ashenfelter and D. Card. Vol. 4. Elsevier, pp. 1–102.

Gaggl, P. and G. C. Wright (2017). “A Short-Run View of What Computers Do: Evidence From a UK Tax Incentive”. In: *American Economic Journal: Applied Economics* 9.3, pp. 262–94.

Goos, M. and A. Manning (Feb. 2007). “Lousy and Lovely Jobs: The Rising Polarization of Work in Britain”. In: *Review of Economics and Statistics* 89.1, pp. 118–133.

Goos, M., A. Manning, and A. Salomons (2009). “Job Polarization in Europe”. In: *American Economic Review* 99.2, pp. 58–63.

— (Aug. 2014). “Explaining Job Polarization: Routine-Biased Technological Change and Offshoring”. In: *American Economic Review* 104.8, pp. 2509–2526.

Holtz-Eakin, D., W. Newey, and H. S. Rosen (1988). “Estimating Vector Autoregressions with Panel Data”. In: *Econometrica* 56.6, pp. 1371–1395.

Hunt, J. and R. Nunn (2022). “Has US employment really polarized? A critical reappraisal”. In: *Labour Economics* 75, p. 102117.

Jaeger, D. A. (1997). “Reconciling the old and new census bureau education questions: Recommendations for researchers”. In: *Journal of Business & Economic Statistics* 15.3, pp. 300–309.

Jaimovich, N., I. Saporta-Eksten, H. Siu, and Y. Yedid-Levi (2021). “The macroeconomics of automation: Data, theory, and policy analysis”. In: *Journal of Monetary Economics* 122, pp. 1–16.

Juhn, C., K. M. Murphy, and B. Pierce (1993). "Wage Inequality and the Rise in Returns to Skill". In: *Journal of Political Economy* 101.3, pp. 410–442.

Jung, J. and J. Mercenier (2014). "Routinization-Biased Technical Change and Globalization: Understanding Labor Market Polarization". In: *Economic Inquiry* 52.4, pp. 1446–1465.

Katz, L. F. and K. M. Murphy (1992). "Changes in Relative Wages, 1963–1987: Supply and Demand Factors". In: *Quarterly Journal of Economics* 107.1, pp. 35–78.

Lancaster, T. (2000). "The incidental parameter problem since 1948". In: *Journal of econometrics* 95.2, pp. 391–413.

Lemieux, T. (June 2006a). "Increasing Residual Wage Inequality: Composition Effects, Noisy Data, or Rising Demand for Skill?" In: *American Economic Review* 96.3, pp. 461–498.

— (2006b). "The "Mincer Equation" Thirty Years after Schooling, Experience, and Earnings". In: *Jacob Mincer, A Pioneer of Modern Labor Economics*, pp. 127–145.

— (Sept. 2007). "The Changing Nature of Wage Inequality". In: *Journal of Population Economics* 21.1, pp. 21–48.

— (2010). "What Do We Really Know about Changes in Wage Inequality?" In: *Labor in the New Economy*. Ed. by K. G. Abraham, J. R. Spletzer, and M. Harper. National Bureau of Economic Research, pp. 17–59.

Mincer, J. (Jan. 1974). *Schooling, Experience, and Earnings*. Vol. 4. 1.

Naticchioni, P., G. Ragusa, and R. Massari (2014). "Unconditional and conditional wage polarization in Europe". In.

Noy, S. and W. Zhang (2023). "Experimental evidence on the productivity effects of generative artificial intelligence". In: *Science* 381.6654, pp. 187–192.

Piketty, T. (2014). *Capital in the 21st Century*. Harvard University Press.

Roy, A. D. (1951). "Some Thoughts on the Distribution of Earnings". In: *Oxford Economic Papers* 3.2, pp. 135–146.

Wald, A. (1940). "The fitting of straight lines if both variables are subject to error". In: *The annals of mathematical statistics* 11.3, pp. 284–300.

Figures and Tables

Tables

Year	Total Skewness	Within	Between	3COV
1992	.0946	.0334	-.04171	0.1030
$N = 130,502$		(.0037)	(.0032)	(.0053)
2002	.1898	.0397	-.0276	0.1777
$N = 108,145$		(.0045)	(.0036)	(.0062)
Δ 1992-2002	.0951	.0063	.0141	0.0747
		(.0058)	(.0048)	(.0081)
	100%	6.6%	14.9%	78.5%

Table 1: Skewness Decomposition by 3-Digit Occupation

Skewness decomposition based on Equation 14. The three components sum to the overall skewness (Equation 14). Wages at the top and bottom 5% were dropped (see Section 3). Standard errors are calculated analytically (using the delta method).
Source: CPS Outgoing Rotation Groups

Figures

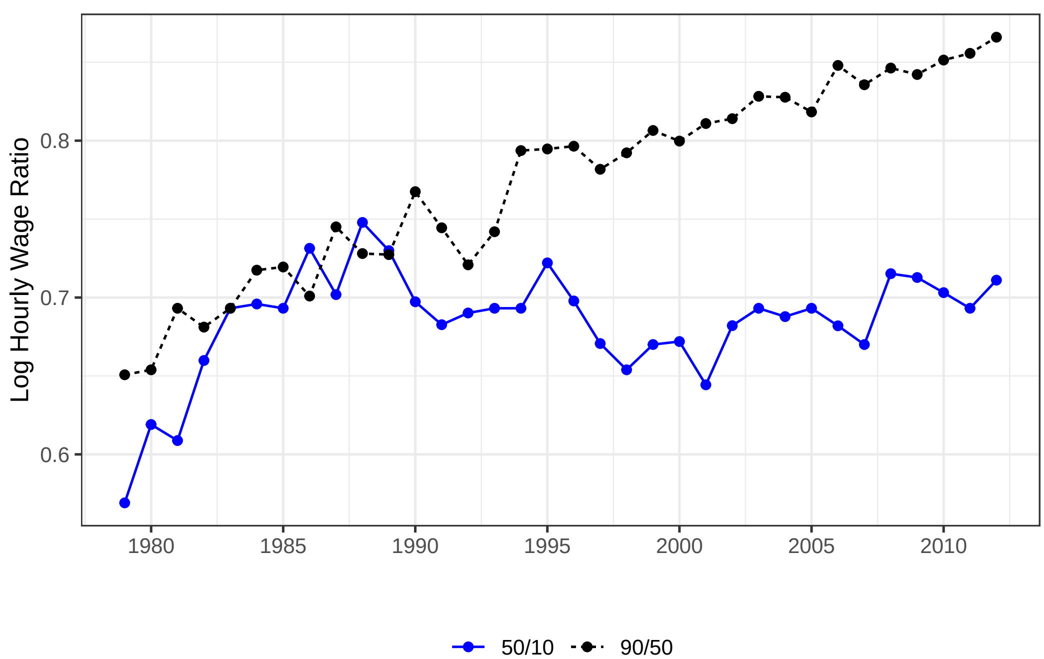

Figure 1: 90/50 and 50/10 Log Hourly Wage Ratio

Quantiles are calculated for all workers with positive earnings at the hours level, using sample weights multiplied by hours worked.
Source: CPS Outgoing Rotation Groups ($N = 4,401,711$)

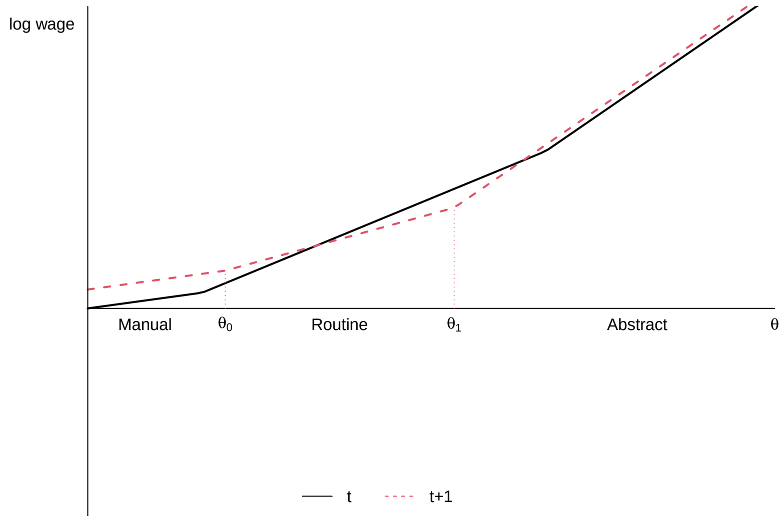

(a) Small SR-RBTC

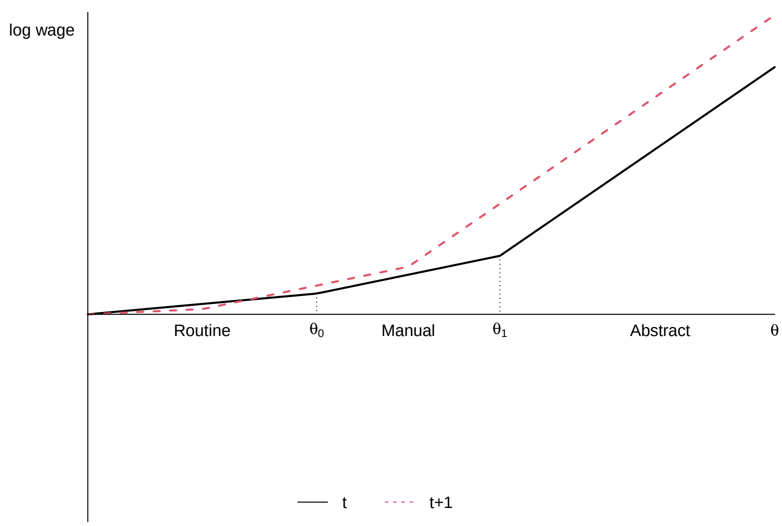

(b) Large SR-RBTC

Figure 2: Illustrated Changes in Log Wages by Skill

These figures illustrate the equilibrium sorting of workers into occupations and their log wages as a function of their skill θ_i. The dashed red line represents equilibrium log wages in a later time period when technology has further advanced (increase in τ). Panel A represents a small technological change, that reduces the slope of log wages as a function of θ_i only in the routine occupation. Panel B describes the equilibrium after a large technological advancement and a reversal of comparative advantage such that the slope in the routine occupation is lower than the slope in the manual occupation (Condition 6 replaces Condition 1).

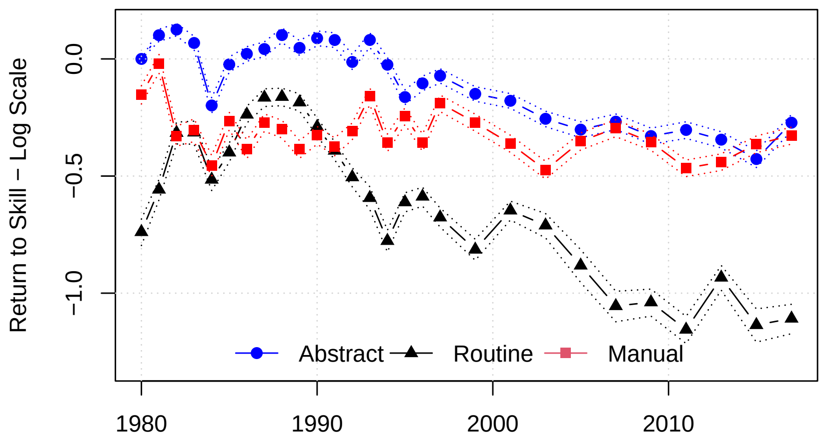

Figure 3: Return to Skill (a_{jt}) by Occupational Category

This figure presents the return to skill (α_{jt}) in log units for the three occupational categories. Return to skill is calculated using an interactive fixed-effects model (Equation 7). The log return to skill in the abstract occupation in 1980 is fixed to zero, hence all other values are relative to that year and occupational category. Routine workers are defined as workers in administrative, production, or operator occupations, classified by the first occupational coding digit. Abstract workers include managers, technicians, and professionals. Manual includes service, sales and agricultural occupations. Autor and Dorn (2013) occupational crosswalk is used for a consistent definition of occupations over time. Dashed lines represent 95% confidence intervals.
Source: PSID ($N = 122,162$)

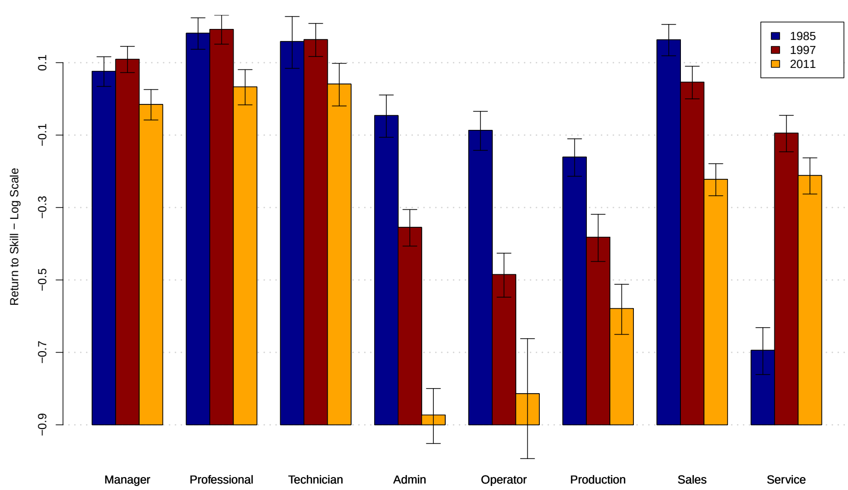

Figure 4: Return to Skill (α_{jt}) by 1-Digit Occupational Category

This figure presents the return to skill (α_{jt}) in log units for eight 1-digit occupational categories. The log return to skill in administrator occupations in 1980 is fixed to zero, hence all other values are relative to that year and occupational category. Results for all years are available in Appendix Figure A5. Return to skill are calculated in an interactive fixed-effects model (α_{jt}, using Equation 7). α_{jt} varies by 1-digit occupation and year. Autor and Dorn (2013) occupational crosswalk is used for a consistent definition of occupations over time.
Source: PSID ($N = 105,248$)

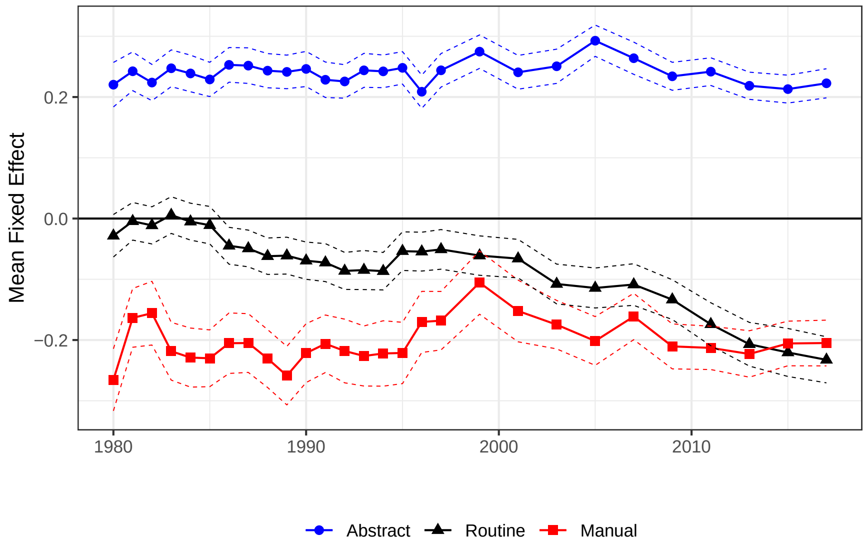

Figure 5: Mean Skill Level ($\widehat{\theta}_i$) by Occupational Category

Mean level of $\widehat{\theta}_i$ by occupational category and year. $\widehat{\theta}_i$ is calculated using Equation 9, and demeaned at the cohort level, where cohorts are defined based on year of entry into the labor market. Routine workers are defined as workers in administrative, production, or operator occupations, classified by the first occupational coding digit. Abstract workers include managers, technicians, and professionals. Manual workers include service, sales, and agricultural occupations. Autor and Dorn (2013) occupational crosswalk is used for a consistent definition of occupations over time.
Source: PSID ($N = 124,407$)

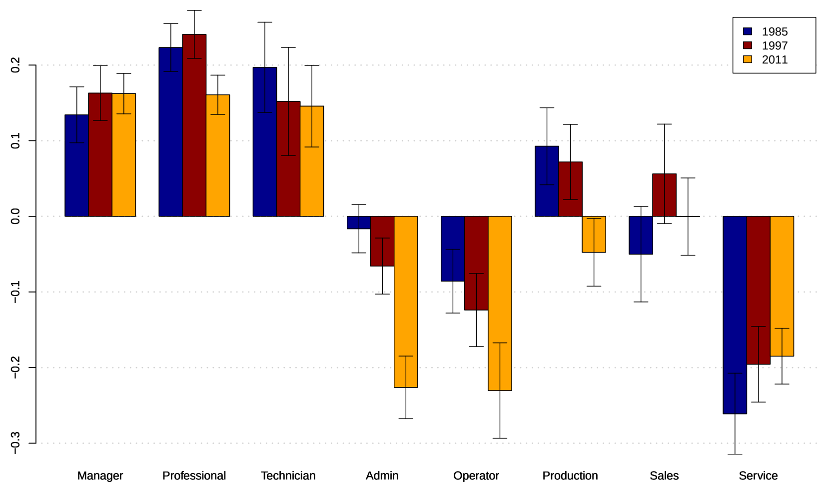

Figure 6: Mean Skill level ($\widehat{\theta}_i$) by1-Digit Occupational Category

Mean level of $\widehat{\theta}_i$ by occupational category and year. $\widehat{\theta}_i$ is calculated using Equation 9 and demeaned at the cohort level, where cohorts are defined based on year of entering the labor market. The Autor and Dorn (2013) occupational crosswalk is used for a consistent definition of occupation over time.
Source: PSID ($N = 108,413$)

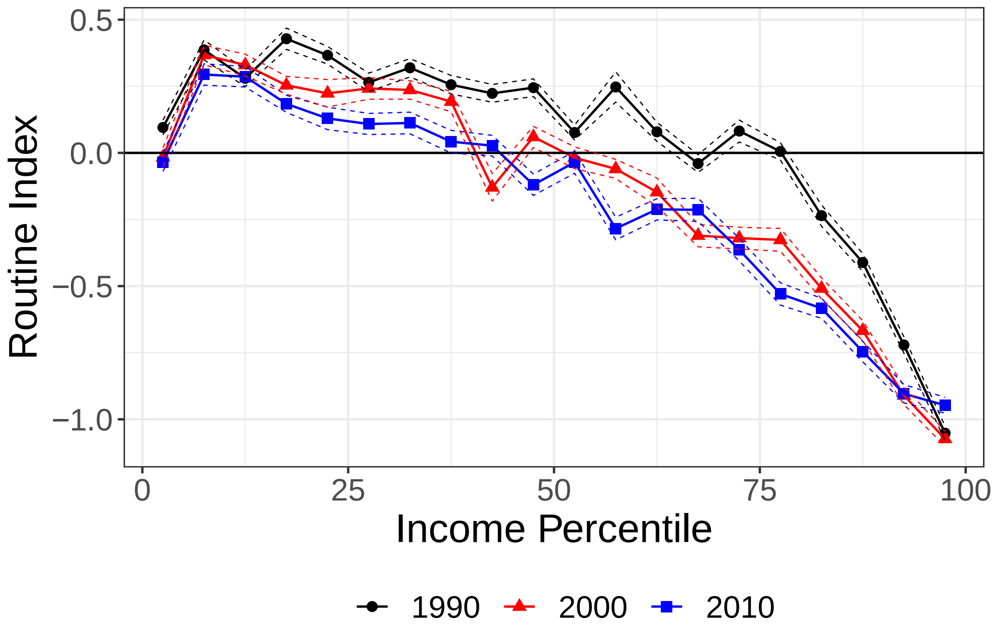

Figure 7: Routine Intensity of Occupation by Wage Percentile
This figure plots the average routine intensity by wage bins for 20 equal-sized bins. Bins are based on workers' hourly wages. The routine intensity is calculated at the occupational level as in Acemoglu and Autor (2011). It is the average of routine manual and routine cognitive indices, both standardized, such that the population average is 0. More details are provided in Appendix E. I use the occupation classification in Autor and Dorn (2013) for consistency across decades. Sample weights are used.
Source: CPS Outgoing Rotation Groups and O*NET ($N_{1990} = 147,851; N_{2000} = 105,461; N_{2010} = 101,915$)

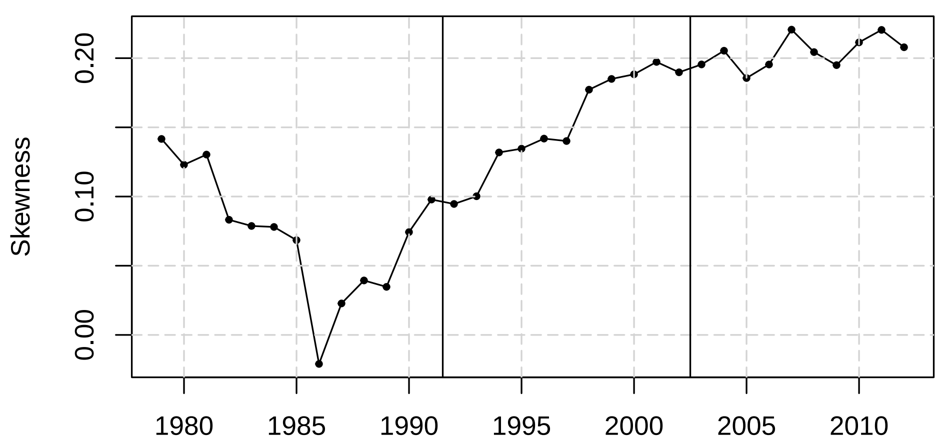

Figure 8: Skewness of Log Hourly Wage
Skewness (Equation 13) of the log wage distribution by year. Sample weights are used. Vertical lines represent changes in occupational coding. Wages at the top and bottom 5% were dropped (see Section 3).
Source: CPS Outgoing Rotation Groups ($N = 4,401,711$)

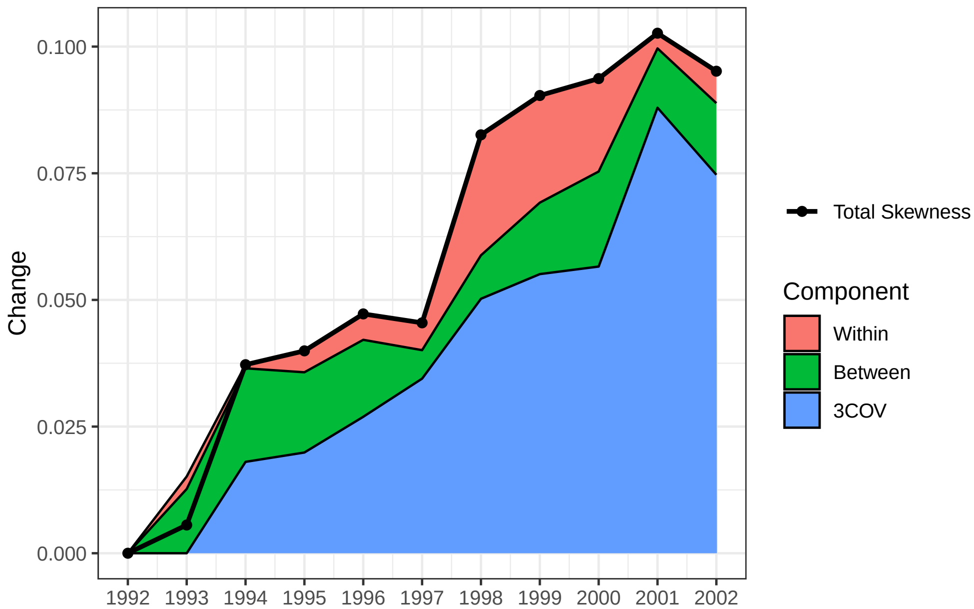

Figure 9: Skewness Decomposition by 3-Digit Occupation

Skewness decomposition based on Equation 14. Changes in each component (within, between, covariance) are plotted relative to the baseline year (1992). The three components sum to the overall skewness (Equation 14). Wages at the top and bottom 5% were dropped (see Section 3).
Source: CPS Outgoing Rotation Groups ($N = 1,208,151$)

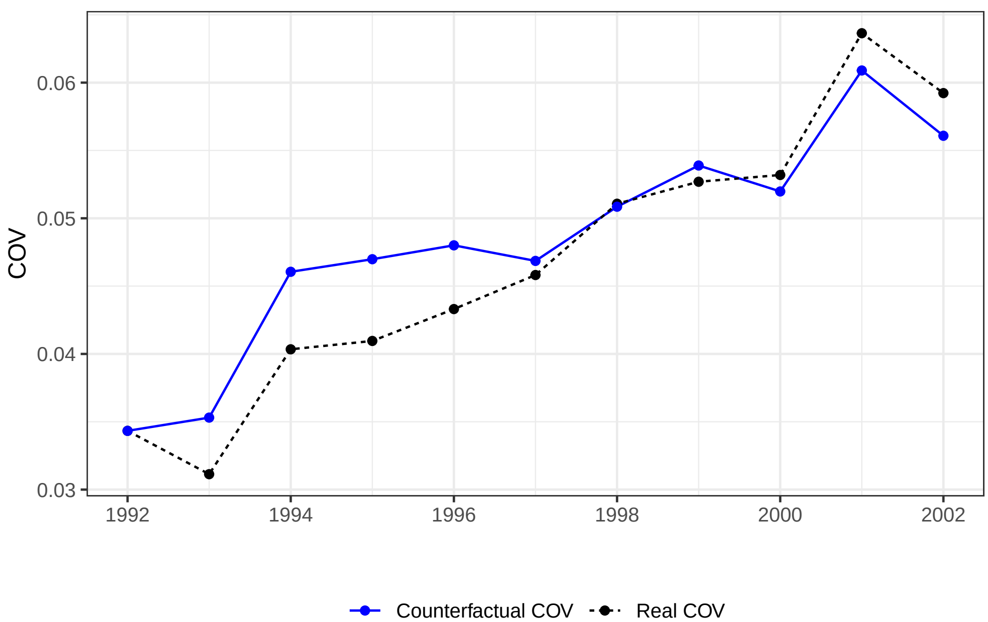

Figure 10: Covariance of Expectation and Variance of Log Wages by Occupation

This figure plots the covariance of mean log wage and variance of log wage by occupation, $COV\left(E\left[\log w|occ\right], V\left(\log w|occ\right)\right)$ (black line). The counterfactual covariance (in blue) is calculated by fixing $E\left[\log w|occ\right]$, and the share of workers in each occupation to their average throughout the period, allowing only the variance within each occupation to change (Equation 16.) Wages at the top and bottom 5% were dropped (see Section 3).
Source: CPS Outgoing Rotation Groups ($N = 1,208,151$)

A Appendix Figures and Tables

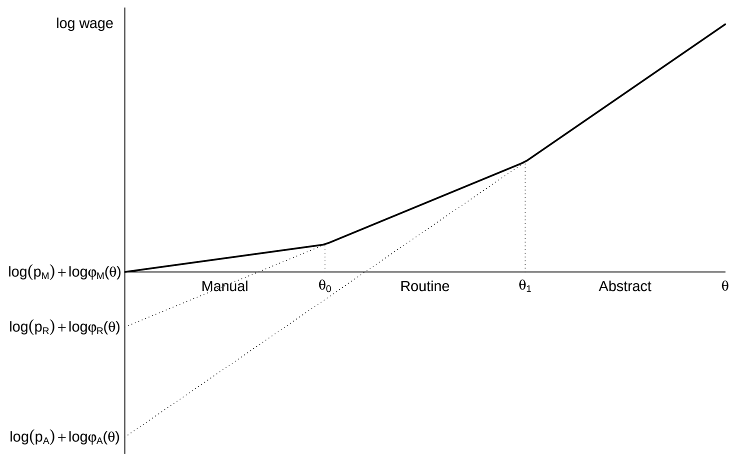

Figure A1: Equilibrium Log Wage by Skill

Sorting into occupations in a Jung and Mercenier (2014) model. The bold line represents equilibrium log wages as a function of θ_i. Dashed lines are off-equilibrium wages in other (suboptimal) occupations. θ_0, θ_1 are the threshold skill levels in which workers are indifferent between two occupations.

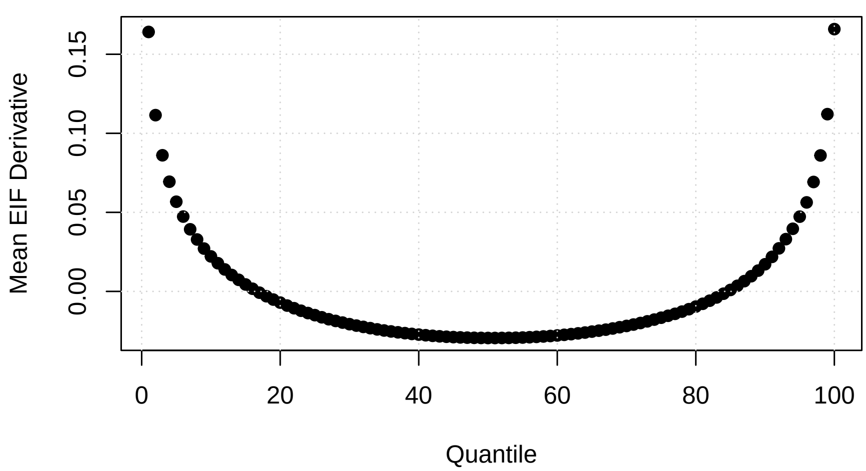

Figure A2: Derivative of EIF on Skewness for Standard Normal Distribution

The empirical influence function is a function from the value of a given observation x_i to some statistic $T_n(x)$ (in this case, the empirical skewness), taking the other observations x_{-i} as given. I calculate this for a sample of $n = 100$. I sample 1,000 samples of 100 observations from a standardized Normal distribution, and calculate numerically the derivative at the kth-order statistic at the sample point. The figure shows the mean over the 1,000 samples of this derivative.

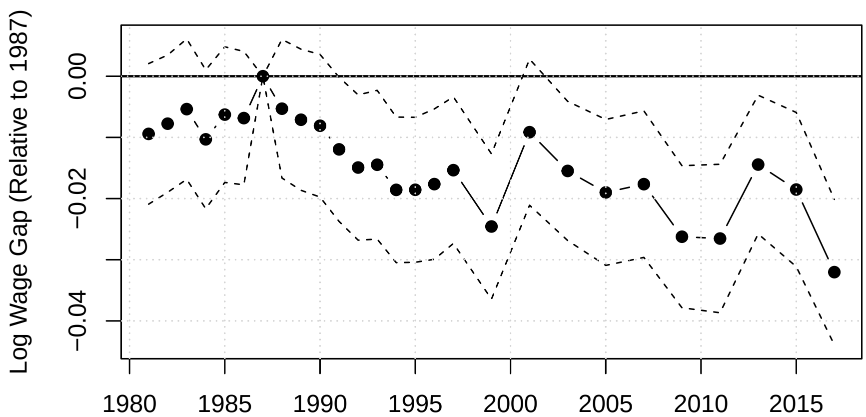

Figure A3: Education Premium in Routine Occupations

This figure plots the change in wage gaps between workers with different years of schooling relative to 1987. Education premium is estimated with coefficient γ_t in Equation 15. Data includes all routine workers in the PSID. See Section 3 for a definition of routine occupations.
Source: PSID

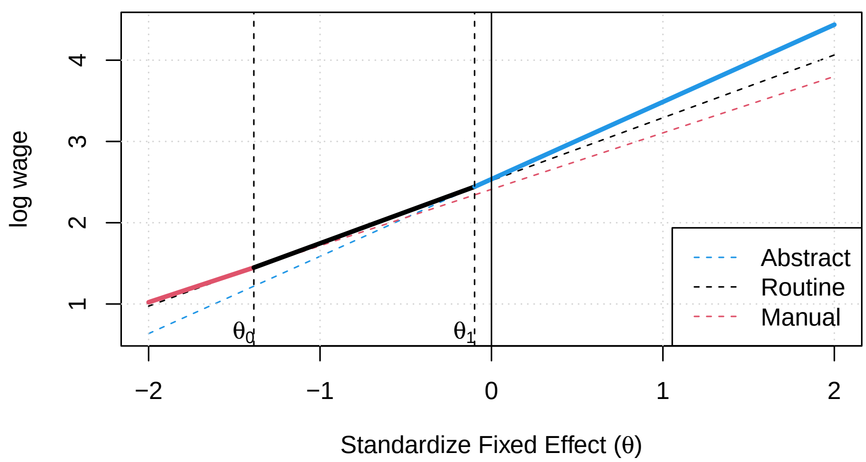

Figure A4: IFEM Expected Log Wage by Occupational Category - 1987

This figure plots the expected log wage as a function of the standardized worker fixed effect θ_i and occupational category for the years 1987 and 2011 using the estimation of the IFEM (Equation 7). The slope in each occupation is determined by the parameter $\alpha_{j,t}$, where a higher slope implies a larger return to skill. Worker fixed effects are standardized to have a mean of zero in every cohort and a standard deviation of 1 overall. The bold line represents the highest expected wage for each skill level. The dashed vertical lines mark the indifference points between two occupational categories that correspond to θ_0 and θ_1 in the model.
Source: PSID

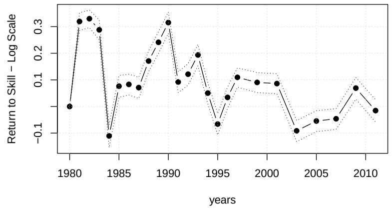

(a) Managers

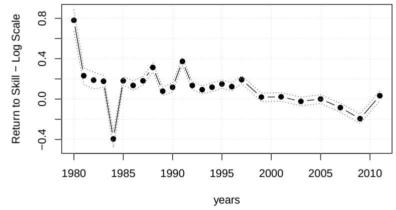

(b) Professional

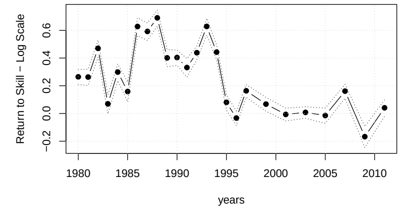

(c) Technicians

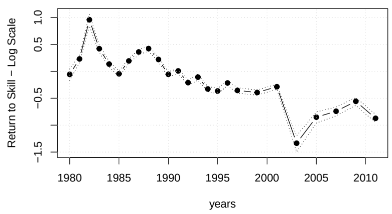

(d) Administrators

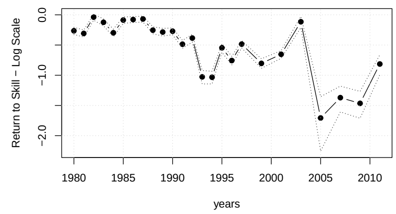

(e) Operators

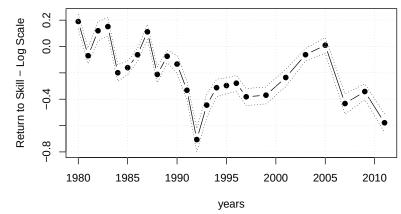

(f) Production

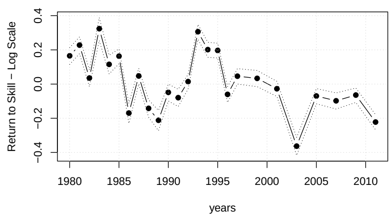

(g) Sales

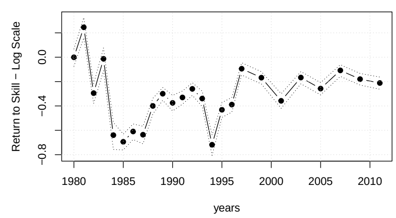

(h) Services

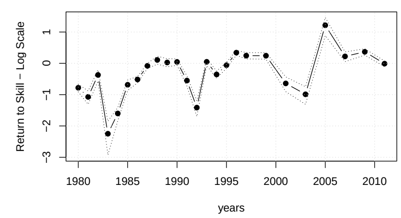

(i) Agriculture

Figure A5: Return to Skill (α_{jt}) by1-Digit Occupation

Returns to skill are calculated in an interactive fixed-effects model (α_{jt}, using Equation 7). α_{jt} varies by 1-digit occupation and year. The Autor and Dorn (2013) occupational crosswalk is used for a consistent definition of occupations over time.
Source: PSID

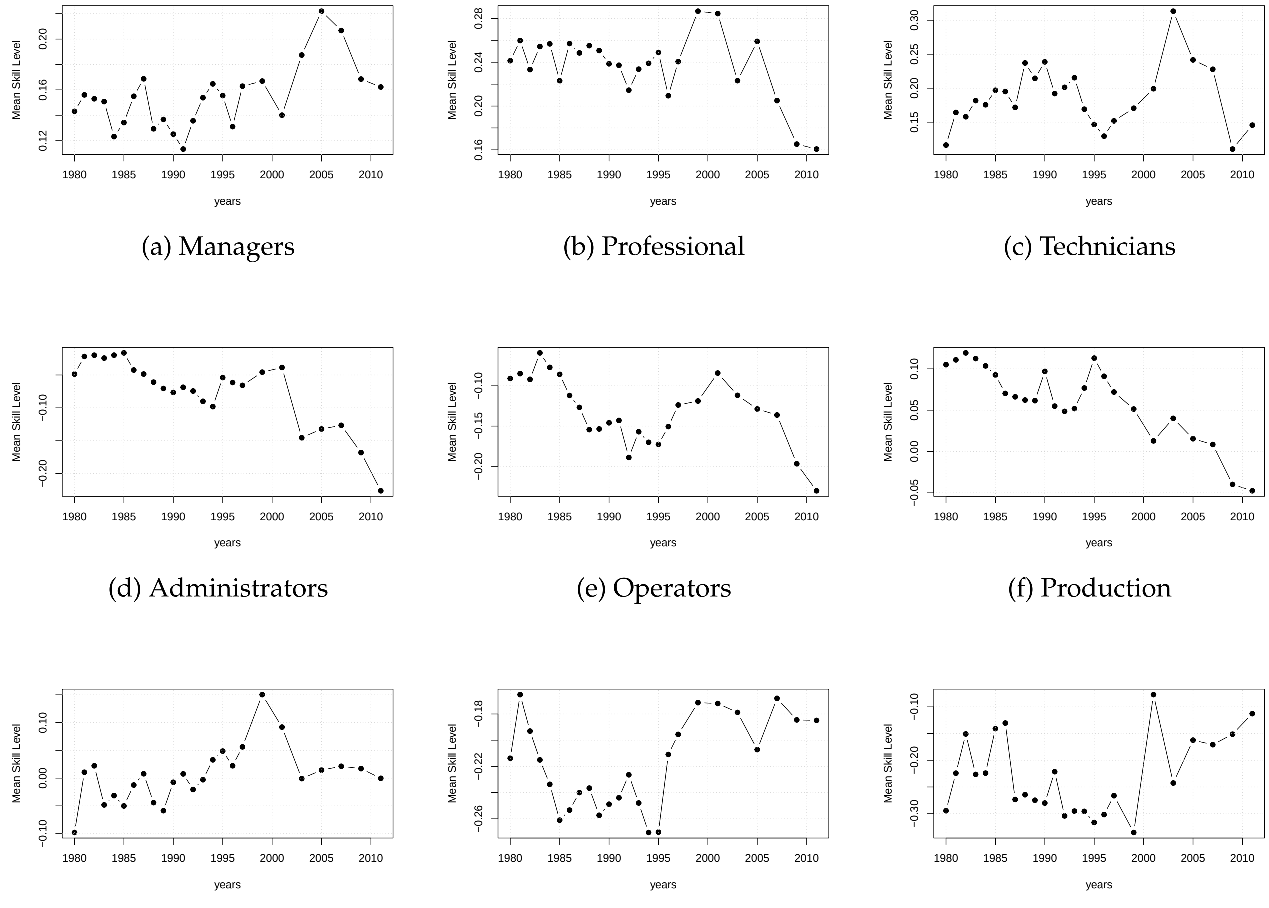

(a) Managers

(b) Professional

(c) Technicians

(d) Administrators

(e) Operators

(f) Production

(g) Sales

(h) Services

(i) Agriculture

Figure A6: Mean Skill Level ($\widehat{\theta}_i$) by1-Digit Occupation

Mean level of $\widehat{\theta}_i$ by occupational category and year. $\widehat{\theta}_i$ is calculated using Equation **??**, and demeaned at the cohort level, where cohorts are defined based on year of entering the labor market. The Autor and Dorn (2013) occupational crosswalk is used for a consistent definition of occupations over time.
Source: PSID

Figure A7: Bin Scatter Plot - Change in Log Wages 1992–2002

Change in log wages in each of 20 equal-sized quantiles. Quantiles are calculated separately for both 1992 and 2002. The x-axis shows the value of the mean log wage in each quantile. The y-axis plots the difference in mean log wages in each of the 20 quantiles between 1992–2002. Sample weights are used.
Source: CPS Outgoing Rotation Groups

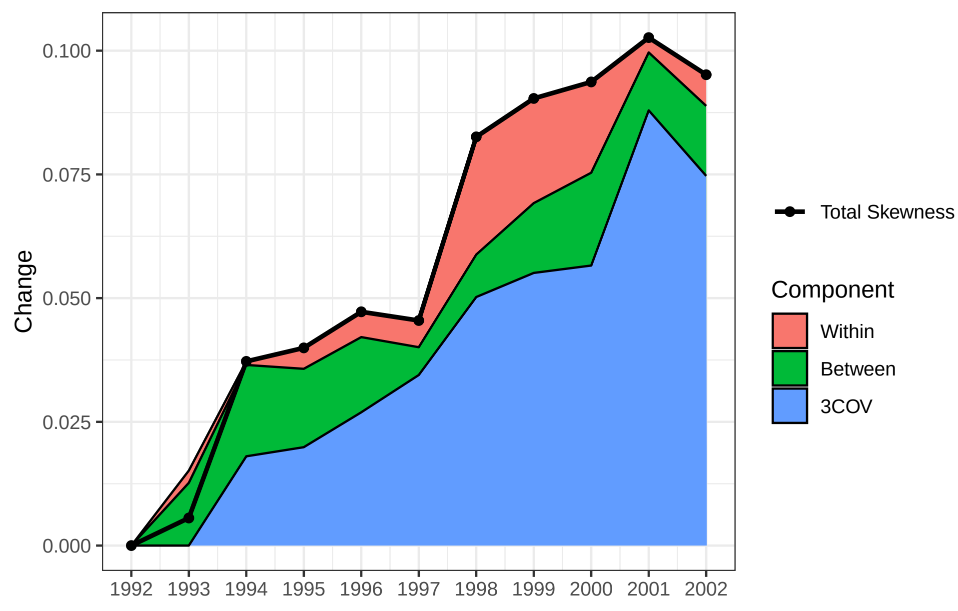

Figure A8: Skewness Decomposition by 3-Digit Occupation with Imputed Wages

Skewness decomposition based on Equation 14. This figure replicates the results in Figure 9 including imputed wages. Changes in each component are plotted relative to the baseline year (1992). Wages at the top and bottom 5% were dropped (see Section 3).
Source: CPS Outgoing Rotation Groups

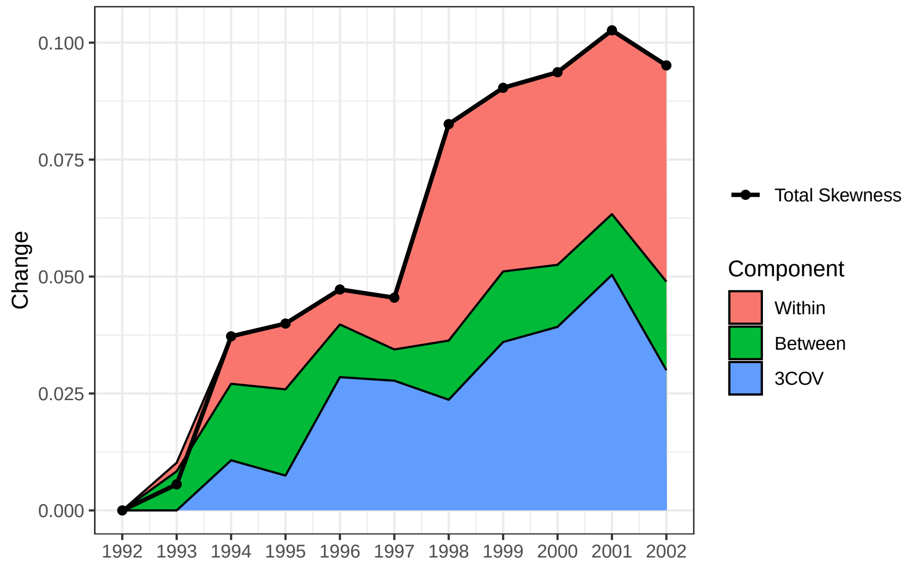

Figure A9: Skewness Decomposition by 3-Digit Industry

Skewness decomposition (Equation 9) by 3-digit industry categories. Changes in each component are plotted relative to the baseline year (1992). Wages at the top and bottom 5% were dropped (see Section 3).
Source: CPS Outgoing Rotation Groups

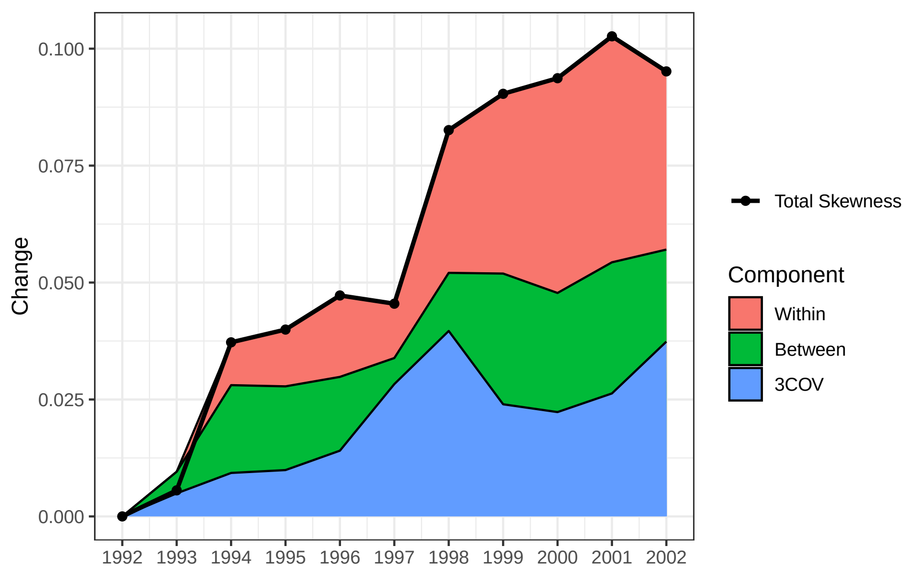

Figure A10: Skewness Decomposition by Education and Experience

Skewness decomposition (Equation 9) by the interaction of years of schooling and years of experience. Changes in each component are plotted relative to the baseline year (1992). Wages at the top and bottom 5% were dropped (see Section 3).
Source: CPS Outgoing Rotation Groups

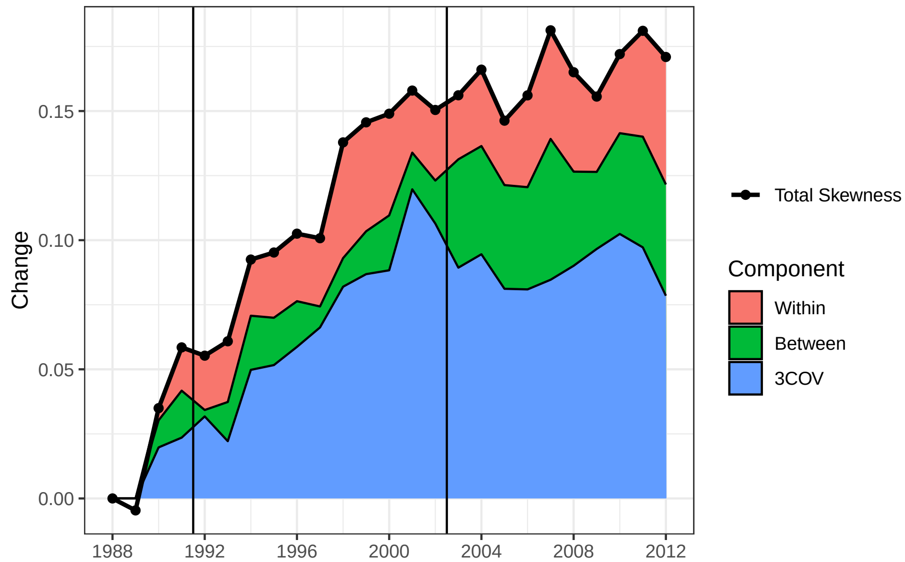

Figure A11: Skewness Decomposition by Occupation 1988–2012

This figure repeats the exercise of Figure 9 for a longer time period. Decomposition is based on Equation 36. Vertical lines represent changes in 3-digit occupational coding. Occupational coding is based on the Autor and Dorn (2013) crosswalk. Changes are reported relative to the baseline year (1988), which is approximately the beginning of the rise in skewness. Wages at the top and bottom 5% were dropped (see Section 3).
Source: CPS Outgoing Rotation Groups

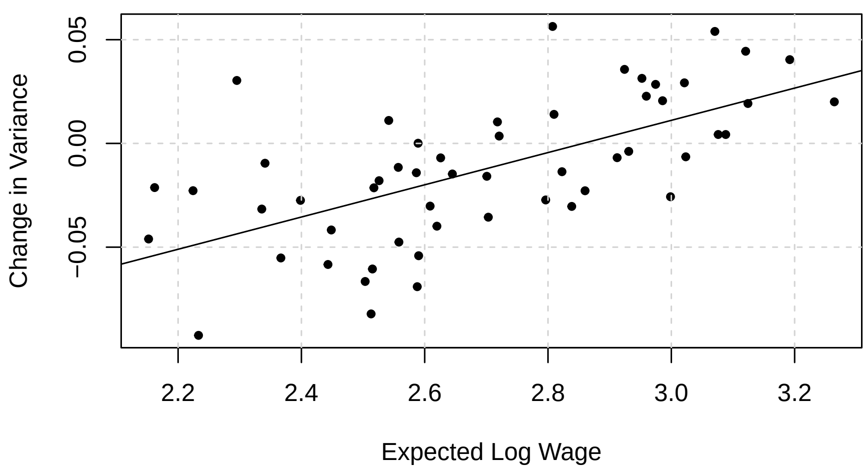

Figure A12: Changes in Occupational Inequality 1992/3–2001/2 by Mean Log Wages

This figure plots all occupations with at least 0.5% of the total working hours (top 47 out of 501 occupations that include 53% of the total working hours). The expected log wage (X-axis) is the average log wage in an occupation during the entire period (1992–2002). Change in Variance (Y-axis) is the difference between the average of the first and last two years (I pool two years together to reduce errors due to small sample size). The line is the best linear fit to the points. Wages at the top and bottom 5% were dropped (see Section 3).
Source: CPS Outgoing Rotation Groups

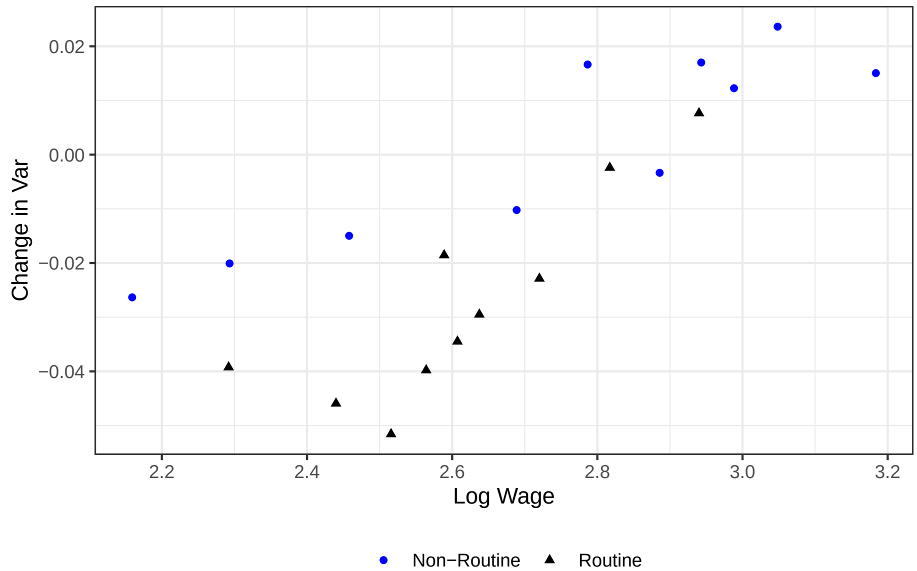

Figure A13: Binned Changes in Occupational Inequality by Mean Log Wages 1992-2002

This figure plots the change in occupational variance as a function of occupation mean separately for routine and non-routine 3-digit occupations. Occupations are binned based on occupation mean log wage separately for routine and non-routine occupations using 10 equal-sized bins (deciles) of occupations, weighted by occupation size. Each point displays the mean log wage in the baseline year 1992, and change in the variance between 1992–2002 when occupational coding is fixed. Routine occupations include all occupations classified as administrators, operators, and production workers based on 1-digit occupation coding.
Source: CPS Outgoing Rotation Groups

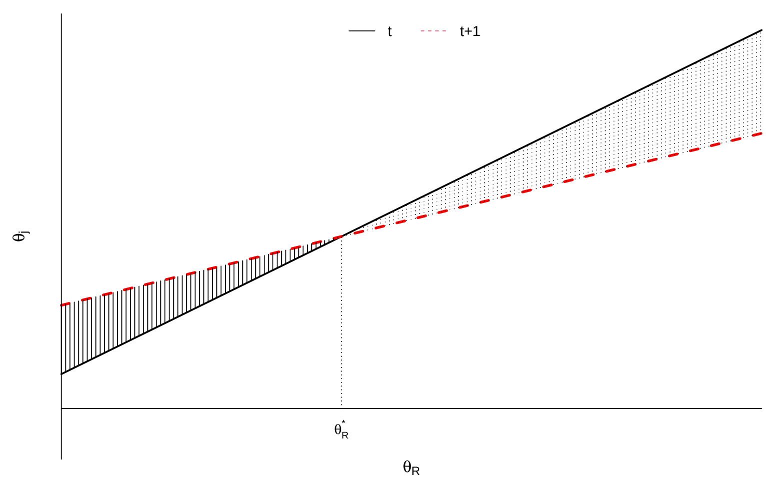

Figure A14: Indifference Curve in Multidimensional model

This figure illustrates the indifference curve between the routine occupation, and another non-routine occupation j (manual or abstract) in the multi-skill model in Section B.3. The figure shows the solutions to Equation 18, which are combinations of routine skill $\theta_{R,i}$ and j skill $\theta_{j,i}$ for which workers would be indifferent between the two occupations. Workers below the indifference curve would prefer the routine occupation, while workers above the indifference curve would prefer occupation j. The black solid line depicts the indifference curve in the initial point. The dashed red line represents the indifference curve in a later time period when technology has advanced further (increase in τ). The dotted (striped) area represents workers who sort into the routine (j) occupation at period t and sort into occupation j (routine) at period $t + 1$.

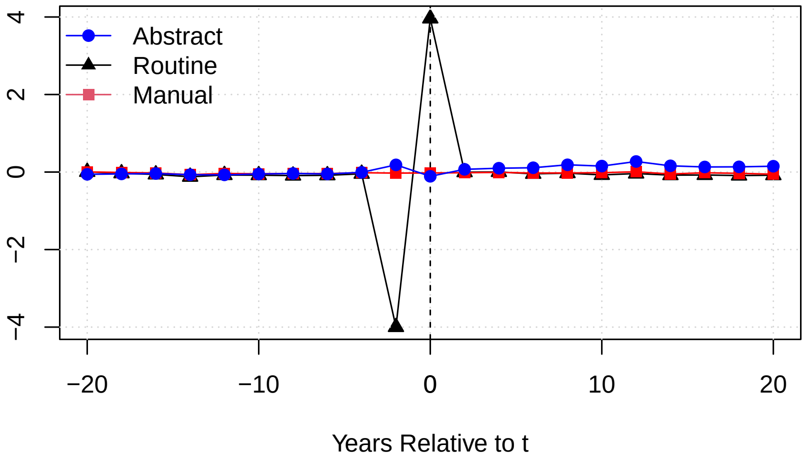

(a) $\log \alpha_{R,t} - \log \alpha_{R,t-2}$

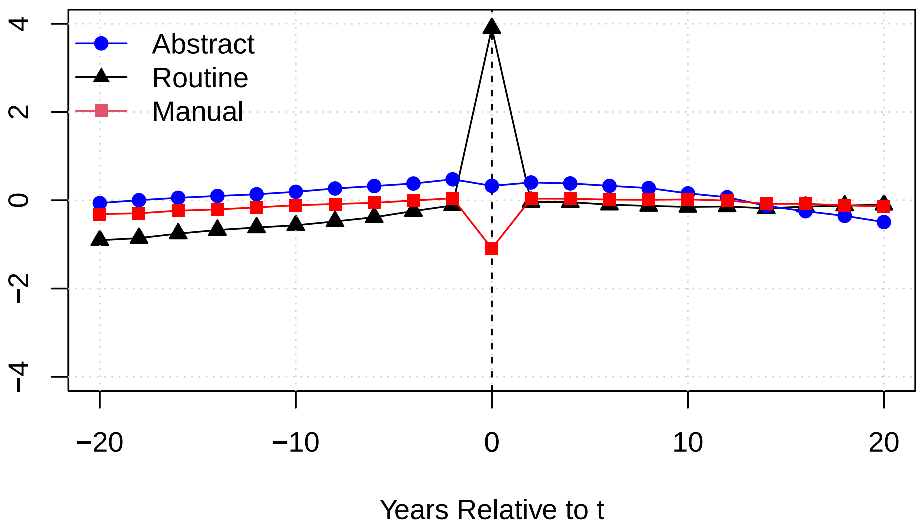

(b) $\log \alpha_{R,t} - \log \alpha_{M,t}$

Figure A15: Sensitivity Analysis

These figures report results of the Andrews et al. (2017) sensitivity analysis for changes in the return to skill. Panel A measures the sensitivity for changes in routine occupations ($\log \alpha_{R,t} - \log \alpha_{R,t-2}$) and Panel B measures the sensitivity for the differences between the routine and manual occupations ($\log \alpha_{R,t} - \log \alpha_{M,t}$). Each point represents the average sensitivity of the log ratios to a correlation of the IV (years of schooling) with the error term in that occupational category and year. The bias can be calculated using $\gamma^Z_{j't+k}$ (Equation 31), the coefficient of the IV on the error term in occupational category j', and year $t+k$. For a speculated value of $\gamma^Z_{j't+k}$, the average bias is $\gamma^Z_{j't+k}$ times the reported sensitivity in the figure. The figure shows the average sensitivity over all years t. For each year t, I calculate the sensitivity as $\Omega\Lambda v$ where Λ is the sensitivity matrix, Ω the covariance matrix of $\{1, X_{it}, Z_i\}$ interacted with an indicator for $i \in E_{jt}$, and v is the gradient of the log ratios in all model parameters. See Andrews et al. (2017) for details. I then average over all the sensitivity values calculated in each year, for any given combinations of occupation category j' and years $t+k$, where k is the distance from year t.

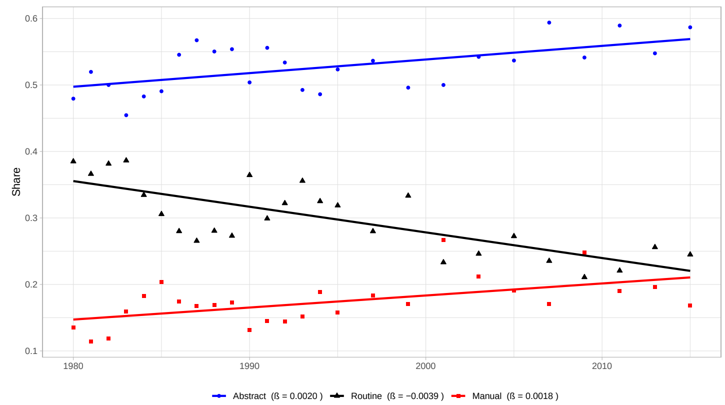

(a) Top Third

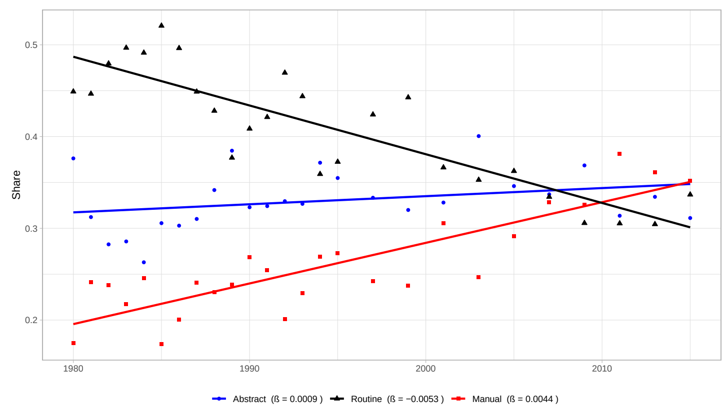

(b) Middle Third

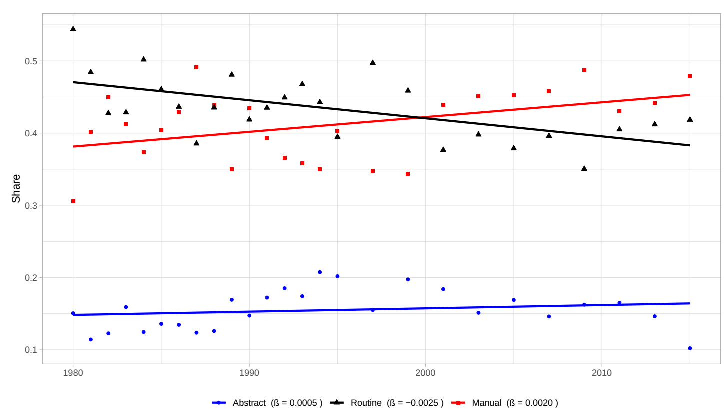

(c) Bottom Third

Figure A16: Share of Workers Who Join Each Occ. Category From Outside the LF

This figure plots the share of transitions from non-employment into each of the three occupational categories by year, separately for workers for each third of the skill distribution. Workers are assigned to each third based on their estimated skill ($\widehat{\theta}_i$), residualized by cohort. For each year t, I follow all workers who did not report income in that year, and report their occupational category two years after, conditional on reporting that they were employed. Each dot represents the share of workers in each occupational category, for all workers who joined the labor market. The lines report the best-fitted line for this category from a linear regression of the shares in this category on the year.
Source: PSID

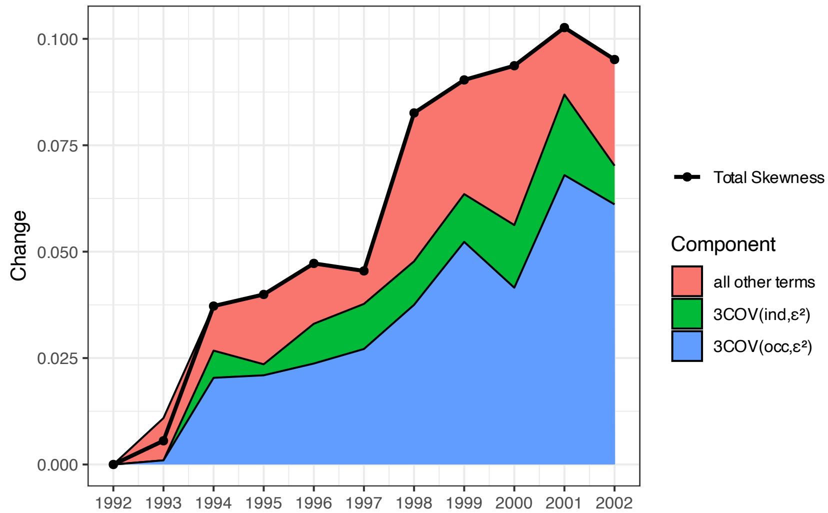

Figure A17: Skewness Decomposition by Occupation and Industry

Linear skewness decomposition (See Appendix H) by occupation and industry (Equation 36). $COV\left(occ, \varepsilon^2\right)$ (in blue) and $COV\left(ind, \varepsilon_i^2\right)$ (in green) are the covariance of occupation and industry premiums with the unexplained variance and are plotted separately. All other terms are aggregated (in red). Changes in each component are plotted relative to the baseline year (1992). Wages at the top and bottom 5% were dropped (see Section 3).
Source: CPS Outgoing Rotation Groups

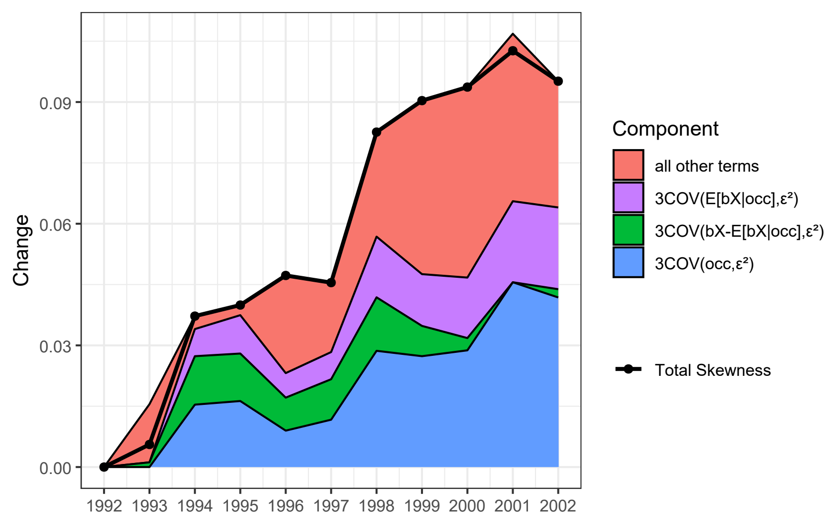

Figure A18: Skewness Decomposition by Occupation, School, Experience

Linear skewness decomposition (see Appendix H) by occupations, years of schooling and years of experience, where education and experience premiums are separated into the occupation mean and a residual (Equation 37). The covariances with the unexplained variance are plotted separately while all other terms are aggregated. occ is the occupation premium for each 3-digit occupation (conditional on education and experience), $E\left[\beta X|_i occ\right]$ is the mean of education and experience premiums in each 3-digit occupation, and $\beta X_i - E\left[\beta X_i | occ_i\right]$ are the demeaned premiums for education and experience. Changes in each component are plotted relative to the baseline year (1992). Wages at the top and bottom 5% were dropped (see Section 3).
Source: CPS Outgoing Rotation Groups

	Wage			Fixed Effects		
	Abstract (1)	Routine (2)	Manual (3)	Abstract (4)	Routine (5)	Manual (6)
Abstract	0.50	0.53	0.44	1		
Routine	0.48	0.52	0.40	.74	1	
Manual	0.45	0.35	0.51	.83	.69	1

Table A1: Correlation of Wages and Fixed Effects of Movers

This table shows the correlation in wages and fixed effects for workers who switch occupational categories in the PSID. Columns 1–3 present the rank correlation of wages before and after the move, for workers who switched their three-digit occupation classification within 6 years. The rows represent the occupational category before the move and the columns are the occupational category after the move. In each combination of origin and destination categories, and in every year between 1980-2011 in the PSID data, I calculate the rank correlation between wages before and after the move. I then present the average correlation across all years. Workers who remained in the same three-digit occupation are excluded from the analysis. Columns 4–6 present the Pearson correlation coefficient between $\widehat{\theta_{ij_0}}, \widehat{\theta_{ij_1}}$ for pairs of occupational categories. The values of $\widehat{\theta_{ij}}$ are estimated using Equation 7, allowing θ_i to vary by the three occupational categories. Routine workers are defined as workers in administrative, production or operator occupations, classified by the first occupational coding digit. Abstract workers are defined as workers in managerial, professional and technician occupations. Manual workers are defined as workers in service, sales, and agriculture. Each correlation is calculated using all workers who ever worked in both categories.

B Model Extensions

B.1 Task Model for RBTC

In this section, I present a model of RBTC as task automation, where the effect of RBTC on the return to skill depends on the type of tasks that are automated. In this model, workers allocate their time across different tasks, each with a different return to skill. The tasks are complementary and each worker is required to perform all of them. Automation replaces workers in a subset of these tasks, allowing them to allocate more time to the other tasks, thus increasing their productivity. The effect of RBTC on the return to skill depends on whether the automated tasks have a relatively high or low return to skill.

Assume that the production function in the routine occupation φ_R aggregates a continuum of tasks indexed by $k \in [0,1]$. Workers with skill θ_i allocate $t(k,\theta_i)$ of their time to task k. Each task has the following production function

$$y^k(\theta_i) = t(k,\theta_i)\exp[\alpha(k)\theta_i].$$

The parameter $\alpha(k) \geq 0$ represents the skill intensity in this task. For example, in the context of cashier workers, arithmetic calculations would have a higher α than bagging, as they have a higher return to skill.

Assume that some tasks can also be automated and performed using technology instead of workers. Let $\tau \in [0,1]$ denote the share of tasks that can be automated. Without loss of generality, assume that tasks are ordered by their automation order such that tasks with a lower k index are automated first. Therefore, for a given value of τ tasks $[0,\tau]$ are automated, while tasks $(\tau,1]$ are not. An increase in τ implies that more tasks can be automated. Also, assume that once a task is automated, employers would always prefer to use automation, as machines are more cost-effective than humans.

Assume that there are strong complementarities between the different tasks, which is represented in a Leontief production function,

$$\varphi_R(\theta_i, \tau) = \min_k y^k(\theta_i).$$

Workers allocate their time between tasks, under the constraint that their total working time is one unit, $\int_\tau^1 t(k, \theta_i)\, d_k = 1$. To maximize productivity, workers would set

$$t(k, \theta_i) = \frac{\exp(-\alpha(k)\,\theta_i)}{\int_\tau^1 \exp(-\alpha(k')\,\theta_i)\, dk'}.$$

Intuitively, lower-skilled will allocate a relatively larger portion of their time to tasks that depend more on skill. By contrast, high-skilled workers can complete tasks that depend on skill more quickly and allocate more time for other tasks. For example, high-skilled cashiers would spend relatively less time on arithmetic calculations, allowing them to devote more time to other tasks and serve more customers.

In this setting, the overall production for a given worker is

$$\varphi_R(\theta_i, \tau) = y^{k'}(\theta_i) = \left[\int_\tau^1 \exp(-\alpha(k)\,\theta_i)\, dk\right]^{-1}.$$

The return to skill in this model is a weighted average of $\alpha(k)$, where the weights depend on the time spent on each task. I denote this weighted average by $\overline{\alpha}(\theta_i, \tau)$.

$$\frac{d \log \varphi_R(\theta_i, \tau)}{d\theta_i} = \int_\tau^1 \alpha(k)\, t(k, \theta_i)\, dk = \overline{\alpha}(\theta_i, \tau). \tag{17}$$

Whether new technology is skill-replacing, skill-enhancing or skill-neutral depends on the return to skill in the marginal task, compared to the average return to skill. Formally, the type of RBTC depends on the sign of ϵ (Equation 3), the derivative of the return to skill in the routine occupation by technological progress, τ. If an increase in technology τ generates a decrease (increase) in the return to skill, then RBTC is skill-replacing (-enhancing), and if the return to skill does not change, RBTC is skill-neutral. Taking the derivative of the average return to skill (17) by the technology level, we get that

$$\epsilon = \frac{d^2 \log \varphi_R}{d\tau d\theta_i} = \frac{d\overline{\alpha}(\theta_i, \tau)}{d\tau} = (\overline{\alpha}(\theta_i, \tau) - \alpha(\tau))t(\tau, \theta_i).$$

Therefore, RBTC is skill-replacing ($\epsilon < 0$) if the average return to skill $\overline{\alpha}(\theta_i, \tau)$ is smaller than the return to skill in the marginal task $\alpha(\tau)$ (the task most recently automated). For simplicity, assume that α is monotonic in k, so the sign of $\alpha'(k)$ is constant. In this case,

$$sign(\epsilon) = sign(\alpha'(k)).$$

If $\alpha'(k) < 0$, the tasks with the highest return to skill are the first to be automated. This would decrease the return to skill and generate an SR-RBTC. By contrast, if $\alpha'(k) > 0$, the tasks with the lowest return to skill would be automated, which would increase the return to skill (skill-enhancing RBTC). If the return to skill is equal for all tasks, $\alpha'(k) = 0$, technology would be skill neutral. The empirical results in the paper are most consistent with skill-replacing RBTC.

Under a stronger assumption, this model would generate a reversal of comparative advantage, as described in Theorem 5. Assume a skill-replacing technology ($\alpha' < 0$). Assume also that for some task $k^* < 1$, the return to skill is lower than the lowest return to skill in the manual occupation (c_M). Marking $\tilde{\tau} = k^*$, yields the condition of Theorem 5.

B.2 CES Production

If workers' production has a constant elasticity of substitution (CES) between skill and technology, then the RBTC type depends only on this elasticity. Assume that production in the routine occupation follows the following CES function,

$$\varphi_R(\theta_i, \tau) = (\theta_i^{\frac{\eta-1}{\eta}} + \tau^{\frac{\eta-1}{\eta}})^{\frac{\eta}{\eta-1.}}$$

The parameter τ can represent the quantity or quality of technology such as computers or robots. The elasticity of substitution between skill and technology is $\eta > 0$.

In this setting, an increase in technology allows all workers to produce more ($\frac{\partial \log \varphi_R}{\partial \tau} > 0$). However, the increase in productivity can be different for workers with different skills. Therefore, ϵ can be positive, zero, or negative. Specifically, the value of ϵ satisfies

$$sign(\epsilon) = sign(1 - \eta).$$

If $\eta = 1$, the production function φ_R converges to a Cobb-Douglas function, and technology is skill-neutral. If $\eta < 1$, skill and technology are complements and the technology is skill-enhancing.

RBTC is skill-replacing in the case of $\eta > 1$. In this case, skill and technology are substitutes. An increase in technology would decrease the productivity (and therefore wage) gaps among workers in the routine occupation. Moreover, the return to skill declines to zero as technology progresses ($\lim_{\tau \to \infty} \frac{\partial \log \varphi_R(\theta_i, \tau)}{\partial \theta_i} = 0$). Therefore, at some technology level $\tilde{\tau}$, there must be a reversal of comparative advantage—the return to skill in the routine occupation would drop below the return to skill in the manual occupation, and Theorem 5 would hold.

B.3 Model with Multidimensional Skills

This section outlines a generalization of the model in Section 1, with multidimensional skills. In this version, workers use a different skill in each occupation. Hence, a worker can be very productive in one occupation but not in another. I find that the core empirical predictions of the model in Section 1 persist in this more general setting. Specifically, I replicate the results that an SR-RBTC reduces the skill premium in the routine occupation, and reduces wages for high-skilled routine workers relative to all other workers. Moreover, I show that high-skilled routine workers leave the routine occupation while low-skilled routine workers join it.

B.3.1 Setup

Assume that workers have a three dimensional skill vector $\theta_i = (\theta_{M,i}, \theta_{R,i}, \theta_{A,i})$. In each occupation $j \in \{M, R, A\}$, workers use only the relevant skill for production, and can produce an intermediate good $\varphi_j(\theta_i) = \varphi_j(\theta_{j,i})$, with $\varphi'_j > 0$.

Wages are set in perfect competition as before

$$w_j(\theta_i) = p_j \varphi_j(\theta_{j,i}).$$

Workers sort into the occupation that maximizes their wages. Since skill is multidimensional, there are indifference curves instead of a single threshold. The indifference curve between two occupations j, j' is the set of skill vectors θ_i that satisfy

$$p_j \varphi_j(\theta_{j,i}) = p_{j'} \varphi_{j'}\left(\theta_{j',i}\right).$$

RBTC is modeled similarly to Section 1. I assume the same production function in the routine occupation, which depends only on the routine skill and technology $\varphi_R(\theta_{R,i}, \tau)$. RBTC is an increase in τ as before. Once again, the effect on wage gaps in the routine occupation depends on ϵ, the derivative of the return to skill in the routine occupation by technological progress (Equation 3) in the same manner:

Theorem 6. *Let $\theta_{R,a}, \theta_{R,b}$ be the routine skill levels of two workers in the routine occupation where $\theta_{R,a} < \theta_{R,b}$. Let w_a, w_b denote their corresponding equilibrium wages. The effect of an improvement in technology τ on the wage ratio $\frac{w_b}{w_a}$ depends on ϵ such that*

$$sign\left(\frac{\partial \frac{w_b}{w_a}}{\partial \tau}\right) = sign(\epsilon)$$

All proofs are given in Section C.1. Therefore, in an SR-RBTC ($\epsilon < 0$), wage gaps in the routine occupation decline.

I define the total amount produced in each intermediate good by M, R, A as before (Equation 4), and the final good is produced using the same CES function as in Equation 5. The general equilibrium effects of an RBTC are exactly the same as before.

Theorem 7. *RBTC (i.e., an increase in τ) generates:*

1. *An increase in the production of the routine good ($\frac{dR}{d\tau} > 0$).*
2. *A decrease in the absolute price of the routine good ($\frac{dp_R}{d\tau} < 0$) and the relative price compared to abstract/manual good ($\frac{dp_R/p_j}{d\tau} < 0$ for $j \in \{M, A\}$).*
3. *A decrease in the share of the total income that is spent on routine goods ($\frac{d\frac{p_R R}{Y}}{d\tau} < 0$).*

B.3.2 Skill Replacing RBTC

I now examine the impact of a skill-replacing RBTC, which is an increase in τ when $\epsilon < 0$. This model yields similar trends as in the unidimensional model (Theorem 3).

Theorem 8. *Assume a skill-replacing technology ($\epsilon < 0$). RBTC (i.e., an increase in τ) would generate the following:*

1. *A decrease in wage gaps between workers in the routine occupation who do not switch occupations.*

2. *The highest-skilled routine workers are less likely to work in the routine occupation.*

3. *The wage for the highest-skilled routine workers would decrease relative to all other workers.*

Figure A14 shows the effect of an increase in τ on the indifference curve between the routine occupation and an alternative occupation j (either manual or abstract). The indifference curve can be written as the set of skill combinations $\theta_{j,i}, \theta_{R,i}$ for which

$$\log \varphi_j \left(\theta_{j,i}\right) = \log \frac{p_R}{p_j} + \log \varphi_R(\theta_{R,i}). \tag{18}$$

The initial indifference curve is plotted in black. Workers whose skill combination $(\theta_{j,i}, \theta_{R,i})$ falls below the indifference curve prefer the routine occupation over j, while those with skill above the curve prefer occupation j over routine (some workers on both sides of the curve may prefer the third option). The curve is upward sloping as higher routine skill generates higher routine wages, and therefore requires higher j skill to prefer occupation j.

The dashed red line plots the indifference curve in a later period. Since $\epsilon < 0$, an increase in τ decreases $\frac{d \log \varphi_R(\theta_R, \tau)}{d\theta_R}$ hence the slope of the indifference curve is lower. Let $\theta^*_{R,i}$ note the routine skill level where the two indifference curves cross (so $\theta_{j,i}\left(\theta^*_{R,i}\right)$ does not change when τ rises).

While there is no clear cutoff between occupations as in the unidimensional case, the multi-skill model preserves the result that high-skilled routine workers leave the routine occupation. The high-skilled routine workers (with $\theta_R > \theta^*_R$) experience a decrease of the indifference curve. Hence, some of these workers would move from R to j. These workers are marked in the dotted area in Figure A14, to the right of $\theta^*_{R,i}$.

Moreover the multi-skill model also preserves the result that low-skill workers join the routine occupation. The striped area in Figure A14, to the left of $\theta^*_{R,i}$ marks workers who would join the routine occupation from occupation j. In contrast to the single skill case, in the multi-skill model this is not a result of a reversal of comparative advantage. Therefore, there is an increase in the number of low-skill routine workers from the onset of the SR-RBTC process, and not only after a certain point in time. The following theorem summarizes these results:

Theorem 9. *Define $\Theta^{j_1}_{j_2}(\tau)$ as the set of skill vectors of workers who shift from occupation j_1 to j_2 for an infinitesimal rise in technology τ (see formal definition in the proof section, C.1). Then for every worker who shifts from routine to occupation j ($\theta_1 \in \Theta^R_j$) and every worker who shifts from occupation j to routine ($\theta_2 \in \Theta^j_R$) the routine skill is higher for those who leave:*

$$\theta_{R,1} > \theta_{R,2}.$$

To conclude, the multi-skill model yields similar predictions to the single-skill model. Therefore, it is also consistent with the empirical finding of the paper, including a decrease in the return to skill, and a decrease in the average skill level, both in the routine occupation.

Moreover, under certain parametric assumptions, this model can also generate fluctuations in bottom-half inequality, as documented in the data. This model predicts that high-skilled routine workers would suffer the largest relative decline in wages and consequentially would shift to other occupations. Therefore, wages are expected to initially decrease around the middle of the income distribution where the high-skilled routine workers are located. Over time, high-skilled routine workers are expected to be mainly in lower parts of the income distribution, which would generate a decrease in wages at lower percentiles.

C Proofs

Theorem 1 Let $\theta_a, \theta_b \in (\theta_0, \theta_1)$ be the skill levels of two workers in the routine occupation where $\theta_a < \theta_b$. Let w_a, w_b denote their corresponding equilibrium wages. The effect of an improvement in technology τ on the wage ratio $\frac{w_b}{w_a}$ depends on ϵ such that

$$sign\left(\frac{\partial \frac{w_b}{w_a}}{\partial \tau}\right) = sign\left(\epsilon\right).$$

Proof. The log wage ratio is

$$\log w_b - \log w_a = \log \varphi(\theta_b, \tau) - \log \varphi(\theta_a, \tau) = \int_a^b \frac{\partial \varphi(\theta_i, \tau)}{\partial \theta_i} d\theta_i.$$

Taking the derivative by τ yields

$$\frac{\partial \log w_b - \log w_a}{\partial \tau} = \int_a^b \frac{\partial^2 \varphi(\theta_i, \tau)}{\partial \theta_i \partial \tau} d\theta_i = \int_a^b \epsilon(\theta_i, \tau) d\theta_i.$$

Since the sign of ϵ is assumed to be the same for all workers,

$$sign\left(\frac{\partial \frac{w_b}{w_a}}{\partial \tau}\right) = sign\left(\frac{\partial \log w_b - \log w_a}{\partial \tau}\right) = sign\left(\epsilon\right).$$

□

Theorem 2 RBTC (i.e., an increase in τ) generates:

1. An increase in the production of the routine good ($\frac{dR}{d\tau} > 0$).
2. A decrease in the absolute price of the routine good ($\frac{dp_R}{d\tau} < 0$) and the relative price compared to the abstract/manual good ($\frac{dp_R/p_j}{d\tau} < 0$ for $j \in \{M, A\}$).
3. A decrease in the share of the total income that is spent on routine goods ($\frac{d\frac{p_R R}{Y}}{d\tau} < 0$).

Proof. The proof follows the order of the claims in the theorem:

1. Under the same allocation of workers, when τ increases R increases while M, A are not changed. Hence, Y must be larger in GE (otherwise, Y is not maximized). If R does not increase then M, A must increase in order for Y to increase. Assume without lost of generality that A increases. From the FOC of the CES function we have

$$\frac{p_R}{p_A} = \left(\frac{R}{A}\right)^{\rho-1}$$

Hence, if A increases and R decreases, p_R/p_A increases. Therefore, for the previous equilibrium level of θ_1

$$p_R \varphi_R(\theta_1) > p_A \varphi_A(\theta_1)$$

which implies that θ_1 increases until equality is reached. Therefore, A must decrease, in contradiction to the assumption.

2. From the FOC we have

$$p_R = \left(\frac{R}{Y}\right)^{\rho-1}$$

Hence, it is sufficient to show R increases more than both M, A such that $\frac{R}{Y}$ increases (and p_R declines). Assume without loss of generality that R/A decreases. By a similar argument as before, if R/A decreases, θ_1 increases and hence A decreases (contradiction). Hence, R/A must increase and hence so must $\frac{R}{Y}$.

Relative prices must also decrease since $\frac{p_R}{p_A} = \left(\frac{R}{A}\right)^{\rho-1}$ and $\rho - 1 < 0$.

3. The share is

$$\frac{p_R R}{Y} = \frac{R^\rho}{Y^\rho} = \frac{R^\rho}{M^\rho + R^\rho + A^\rho} = \frac{1}{\left(\frac{M}{R}\right)^\rho + \left(\frac{A}{R}\right)^\rho + 1}$$

since R increases relative to both M, A, $\left(\frac{M}{R}\right)^\rho$ and $\left(\frac{A}{R}\right)^\rho$ increase and therefore $\frac{p_R R}{Y}$ decreases.

□

Theorem 3 Assume a skill-replacing technology ($\epsilon < 0$). RBTC (i.e., an increase in τ) would generate the following:

1. A decrease in wage gaps between workers in the routine occupation who do not switch occupations.

2. The highest skill routine workers would leave the routine occupation ($\frac{\partial \theta_1}{\partial \tau} < 0$).

3. Wages for the highest skill routine worker (θ_1) would decrease relative to all other workers.

Proof. 1. By theorem 1.

2. By the following lemma (where $\varphi_\tau^R(\theta_0, \tau) = \frac{\partial \varphi^R(\theta_0,\tau)}{\partial \tau}$)

Lemma 10. $\frac{\varphi_\tau^R(\theta_0,\tau)}{\varphi^R(\theta_0,\tau)} > \frac{R_\tau}{R} = \frac{\int_{\theta_0}^{\theta_1}\varphi_\tau^R(\theta_i,\tau)}{\int_{\theta_0}^{\theta_1}\varphi^R(\theta_i,\tau)} > \frac{\varphi_\tau^R(\theta_1,\tau)}{\varphi^R(\theta_1,\tau)}$

Proof. We use $\epsilon = \frac{\partial^2 \log\varphi}{\partial\tau\partial\theta_i} < 0$. Hence, for any positive $b > a$, we have

$$\frac{\varphi_\tau(a,\tau)}{\varphi(a,\tau)} > \frac{\varphi_\tau(b,\tau)}{\varphi(b,\tau)}$$

$$\frac{\varphi_\tau(a,\tau)}{\varphi_\tau(b,\tau)} > \frac{\varphi(a,\tau)}{\varphi(b,\tau)}.$$

Defining $b = \theta_1$ and taking the integral over a between $[\theta_0, \theta_1]$ yields $\frac{R_\tau}{R} = \frac{\int_{\theta_0}^{\theta_1}\varphi_\tau^R(\theta_i,\tau)}{\int_{\theta_0}^{\theta_1}\varphi^R(\theta_i,\tau)} > \frac{\varphi_\tau^R(\theta_1,\tau)}{\varphi^R(\theta_1,\tau)}$. Similarly defining $a = \theta_0$, taking the inverse of the above inequality, and integrating over $b \in [\theta_0, \theta_1]$ yields $\frac{\varphi_\tau^R(\theta_0,\tau)}{\varphi^R(\theta_0,\tau)} > \frac{R_\tau}{R}$.

$\square$

Using this lemma we can show that θ_1 decreases. To do so, we use the equations

$$p_R\varphi_R(\theta_1) = p_A\varphi_A(\theta_1)$$

$$p_R\varphi_R(\theta_0) = p_M\varphi_M(\theta_0),$$

and the CES structure

$$\log\varphi_A(\theta_1) - \log\varphi_R(\theta_1,\tau) = \log p_R - \log p_A = (\rho-1)\log R - (\rho-1)\log A$$

and similarly for M. Define two functions of θ_0, θ_1, τ such that

$$f_M = \log\varphi_M(\theta_0) - \log\varphi_R(\theta_0,\tau) + (\rho-1)\,m(\theta_0) - (\rho-1)\,r(\theta_0,\theta_1,\tau)$$

$$f_A = \log\varphi_A(\theta_1) - \log\varphi_R(\theta_1,\tau) + (\rho-1)\,a(\theta_1) - (\rho-1)\,r(\theta_0,\theta_1,\tau)$$

where $m(\theta_0) = \log M(\theta_0)$, $r(\theta_0,\theta_1) = \log R(\theta_0,\theta_1)$, and $a(\theta_1) = \log A(\theta_1)$. In equilibrium $f_M, f_A = 0$ as the FOCs hold. Using the implicit theorem function we can derive $\frac{d\theta_1}{d\tau}$. Taking the derivative by τ we have

$$\frac{\partial f_M}{\partial\tau} = -\frac{\varphi_\tau^R(\theta_0,\tau)}{\varphi^R(\theta_0,\tau)} - (\rho-1)\frac{R_\tau}{R}$$

$$\frac{\partial f_A}{\partial\tau} = -\frac{\varphi_\tau^R(\theta_1,\tau)}{\varphi^R(\theta_1,\tau)} - (\rho-1)\frac{R_\tau}{R} > 0$$

where the last inequality is from the previous lemma and $\rho < 0$. Taking the derivative by θ_0, θ_1 and using Condition 1

$$\frac{\partial f_M}{\partial \theta_0} = \frac{\varphi_M'}{\varphi_M} - \frac{\varphi_R'}{\varphi_R} + (\rho - 1)\, m'(\theta_0) - (\rho - 1) \frac{R_{\theta_0}}{R} < 0$$

$$\frac{\partial f_A}{\partial \theta_1} = \frac{\varphi_A'}{\varphi_A} - \frac{\varphi_R'}{\varphi_R} + (\rho - 1)\, a'(\theta_1) - (\rho - 1) \frac{R_{\theta_1}}{R} > 0$$

$$\frac{\partial f_A}{\partial \theta_0} = -(\rho - 1) \frac{R_{\theta_0}}{R} < 0$$

$$\frac{\partial f_M}{\partial \theta_1} = -(\rho - 1) \frac{R_{\theta_1}}{R} > 0.$$

From the implicit function theorem we have

$$\nabla \theta_i(\tau) = -\begin{pmatrix} \frac{\partial f_M}{\partial \theta_0} & \frac{\partial f_M}{\partial \theta_1} \\ \frac{\partial f_A}{\partial \theta_0} & \frac{\partial f_A}{\partial \theta_1} \end{pmatrix}^{-1} \left(\frac{\partial f_M}{\partial \tau}, \frac{\partial f_A}{\partial \tau} \right)$$

or

$$\nabla \theta_i(\tau) = -\frac{1}{det} \begin{pmatrix} \frac{\partial f_A}{\partial \theta_1} & -\frac{\partial f_M}{\partial \theta_1} \\ -\frac{\partial f_A}{\partial \theta_0} & \frac{\partial f_M}{\partial \theta_0} \end{pmatrix} \left(\frac{\partial f_M}{\partial \tau}, \frac{\partial f_A}{\partial \tau} \right)$$

The determinant is negative. Hence, there is $c > 0$ such that

$$\frac{\partial \theta_1}{\partial \tau} = c * \left(-\frac{\partial f_A}{\partial \theta_0} \frac{\partial f_M}{\partial \tau} + \frac{\partial f_M}{\partial \theta_0} \frac{\partial f_A}{\partial \tau} \right)$$

$$\frac{\partial \theta_0}{\partial \tau} = c * \left(\frac{\partial f_A}{\partial \theta_1} \frac{\partial f_M}{\partial \tau} - \frac{\partial f_M}{\partial \theta_1} \frac{\partial f_A}{\partial \tau} \right).$$

Plugging in the values that were previously calculated and using Lemma 10 yields

$$\frac{\partial \theta_1}{\partial \tau} = \left(\frac{\varphi_M'}{\varphi_M} - \frac{\varphi_R'}{\varphi_R} + (\rho - 1)\, m'(\theta_0) \right) \frac{\partial f_A}{\partial \tau} - (\rho - 1) \frac{R_{\theta_0}}{R} \left(\frac{\partial f_A}{\partial \tau} - \frac{\partial f_M}{\partial \tau} \right) =$$

$$= \left(\frac{\varphi_M'}{\varphi_M} - \frac{\varphi_R'}{\varphi_R} + (\rho - 1)\, m'(\theta_0) \right) \frac{\partial f_A}{\partial \tau} - (\rho - 1) \frac{R_{\theta_0}}{R} \left(-\frac{\varphi_\tau^R(\theta_1, \tau)}{\varphi^R(\theta_1, \tau)} + \frac{\varphi_\tau^R(\theta_0, \tau)}{\varphi^R(\theta_0, \tau)} \right) < 0$$

Howevre, for θ_0 we have an expression with a sign that could go both ways,

$$\frac{\partial \theta_0}{\partial \tau} = \left(\frac{\varphi_A'}{\varphi_A} - \frac{\varphi_R'}{\varphi_R} + (\rho - 1)\, a'(\theta_1) \right) \frac{\partial f_M}{\partial \tau} - (\rho - 1) \frac{R_{\theta_1}}{R} \left(-\frac{\varphi_\tau^R(\theta_0, \tau)}{\varphi^R(\theta_0, \tau)} + \frac{\varphi_\tau^R(\theta_1, \tau)}{\varphi^R(\theta_1, \tau)} \right).$$

3. Among workers in the routine occupation, $\epsilon < 0$ implies that the largest decline in

wages is for the highest-skilled workers. Their wages also fall compared to those of abstract workers since abstract workers see a change only in p_A and at θ_1:

$$\frac{\partial \log p_R}{\partial \tau} + \frac{\partial \log \varphi_R}{\partial \tau}(\theta_1) - \frac{\partial \log p_A}{\partial \tau} + 0 < 0 \tag{19}$$

as otherwise θ_1 would not go down.

Finally, for manual workers, if $\frac{d\theta_0}{d\tau} \geq 0$ (employment in the manual occupation is weakly increasing) then at θ_0 we have

$$\frac{\partial \log p_R}{\partial \tau} + \frac{\partial \log \varphi_R}{\partial \tau}(\theta_0) - \frac{\partial \log p_M}{\partial \tau} + 0 < 0$$

and since $\frac{\partial \log \varphi_R}{\partial \tau}(\theta_1) < \frac{\partial \log \varphi_R}{\partial \tau}(\theta_0)$ we get

$$\frac{\partial \log p_R}{\partial \tau} + \frac{\partial \log \varphi_R}{\partial \tau}(\theta_1) < \frac{\partial \log p_M}{\partial \tau} \tag{20}$$

and so wages at θ_1 fall relative to all manual jobs.

If $\frac{d\theta_0}{d\tau} < 0$ then $\frac{dM}{d\tau} < 0$, and since $\frac{dA}{d\tau} > 0$ it must be that $\frac{d\left(\frac{M}{A}\right)^{\rho-1}}{d\tau} > 0$. Using $\frac{p_M}{p_A} = \left(\frac{M}{A}\right)^{\rho-1}$ we get that wages in the manual occupation increase faster than in abstract occupation and using 19 we get that 20 also holds.

□

Theorem 4 Assume a skill replacing technology ($\epsilon < 0$), and

$$\frac{\left|\frac{dF(\theta_1)}{d\tau}\right|}{\frac{dF(\theta_0)}{d\tau}} > \frac{E\left[\theta_i|\theta_0 < \theta_i < \theta_1\right] - \theta_0}{\theta_1 - E\left[\theta_i|\theta_0 < \theta_i < \theta_1\right]}. \tag{21}$$

In the **routine occupation**, RBTC would generate a decrease in :

1. Employment ($F(\theta_1) - F(\theta_0)$)
2. Within-occupation inequality ($V\left[w_i|\theta_i \in [\theta_0, \theta_1]\right]$)
3. Mean skill level ($E\left[\theta_i|\theta_i \in [\theta_0, \theta_1]\right]$).

In the **other occupations**, inequality within the abstract occupation ($V\left[w_i|\theta_i > \theta_1\right]$) and manual occupation ($V\left[w_i|\theta_i < \theta_0\right]$) will rise.

The **overall inequality trend** is asymmetric. Wage gaps are decreasing at the bottom such that for every two workers with skill level $\theta_a < \theta_b \leq \theta_1$, the wage gap between them decreases $\frac{d\frac{w_b}{w_a}}{d\tau} \leq 0$. At the top, the wage gap between abstract workers and high-skill routine workers increases. Formally, there exists a value $\theta_0 < \theta^* < \theta_1$ such that for every high skill routine worker $\theta^* < \theta_a < \theta_1$), and every abstract worker $\theta_b > \theta_1$, the wage gap between them increases $\frac{d\frac{w_b}{w_a}}{d\tau} \geq 0$.

Note: This is a more general version of the theorem in the main text. Specifically, when $\theta_i \sim U$, condition 21 is equivalent to $0 < \frac{d\theta_0}{d\tau} < |\frac{d\theta_1}{d\tau}|$.

Proof. In the **routine occupation** employment decreases since θ_1 decreases (Theorem 3) and θ_0 increases (by Equation 21).

Within the routine occupation, inequality would decrease since the skill distribution is now more equal ($V[\theta_i|\theta_i \in [\theta_0, \theta_1]]$ decreases), and conditional on skill wage gaps are smaller (Theorem 1). Same argument will also apply for other inequality measures.

Mean skill is $E[\theta_i|\theta_i \in [\theta_0, \theta_1]] = \frac{\int_{\theta_0}^{\theta_1} \theta_i f(\theta_i) d\theta_i}{\int_{\theta_0}^{\theta_1} f(\theta_i) d\theta_i}$. Taking the derivative yields

$$\frac{\left[\frac{d\theta_1}{d\tau}\theta_1 f(\theta_1) - \frac{d\theta_0}{d\tau}\theta_0 f(\theta_0)\right] \int_{\theta_0}^{\theta_1} f(\theta_i) d\theta_i - \int_{\theta_0}^{\theta_1} \theta_i f(\theta_i) d\theta_i \left(\frac{d\theta_1}{d\tau} f(\theta_1) - \frac{d\theta_0}{d\tau} f(\theta_0)\right)}{\left(\int_{\theta_0}^{\theta_1} f(\theta_i) d\theta_i\right)^2}$$

which is always positive when condition 21 holds.

Within the **other occupations**, inequality increases as there is a larger variation in skill.

The **overall inequality trend** is asymmetric. This is because for manual workers, the wage effect is $\frac{d\log p_M}{d\tau}$. For workers in the routine occupation, it is $\frac{d\log p_R}{d\tau} + \frac{d\log \varphi_R}{d\tau}$. And for abstract workers, it is $\frac{d\log p_A}{d\tau}$. The wage increase in the abstract occupation is larger than that in the routine occupation close to θ_1 because from the decrease in θ_1 if follows that

$$\frac{d\log p_R}{d\tau} + \frac{d\log \varphi_R(\theta_1)}{d\tau} < \frac{d\log p_A}{d\tau}.$$

Among workers in the routine occupation, wages increase relatively for the lower-skilled since $\epsilon < 0$. And manual workers see a larger increase from the wage change in θ_0 since

$$\frac{d\log p_R}{d\tau} + \frac{d\log \varphi_R(\theta_0)}{d\tau} < \frac{d\log p_M}{d\tau}$$

as otherwise θ_0 would not increase.

Within the manual and abstract occupations, wage impact is the same. □

Theorem 5 Assume a skill-replacing technology ($\epsilon < 0$) and that there exists a $\tilde{\tau}$ such that for any $\tau \geq \tilde{\tau}$ and for any θ_i

$$\frac{\partial \log \varphi_R(\theta_i; \tau)}{\partial \theta_i} < \frac{\partial \log \varphi_M(\theta_i)}{\partial \theta_i}.$$

When $\tau \geq \tilde{\tau}$ workers in the routine occupation earn the lowest wages. Any additional SR-RBTC ($\tau \uparrow$) decreases employment in the routine occupation ($\frac{d\theta_0}{d\tau}$<0), as well as wage gaps within workers in the routine occupation who do not switch occupations.

Proof. Wage gaps in the routine occupation fall based on Theorem 1.

I now prove that $\frac{d\theta_0}{d\tau} < 0$. The equilibrium condition that determines θ_0 is

$$\log \varphi_M(\theta_0) - \log \varphi_R(\theta_0, \tau) = \log p_R - \log p_M = (\rho - 1)\log R - (\rho - 1)\log M$$

Defining

$$f(\theta_0, \tau) = \log \varphi_M(\theta_0) - \log \varphi_R(\theta_0, \tau) - (\rho - 1)\log R + (\rho - 1)\log M$$

and using the implicit function theorem we get

$$\frac{\partial f}{\partial \tau} = -\frac{\varphi_\tau^R(\theta_0, \tau)}{\varphi^R(\theta_0, \tau)} - (\rho - 1)\frac{R_\tau}{R} > 0$$

where the last inequality follows from Lemma 10 (for which now the upper bound of the routine occupation is θ_0).

Taking the derivative by θ_0, we have

$$\frac{\partial f}{\partial \theta_0} = \frac{\varphi_M'}{\varphi_M} - \frac{\varphi_R'}{\varphi_R} + (\rho - 1)m'(\theta_0) - (\rho - 1)\frac{R_{\theta_0}}{R} > 0$$

where $m'(\theta_0) < 0$ as a result of the following lemma:

Lemma 11. *When $\tau \geq \tilde{\tau}$ $sign\left(\frac{dM}{d\tau}\right) = sign\left(\frac{dA}{d\tau}\right)$*

Proof. When $\tau \geq \tilde{\tau}$, θ_1 separates between manual and abstract workers. Using the equilibrium condition we have

$$\log \varphi_M(\theta_1) - \log \varphi_A(\theta_1) = \log p_A - \log p_M = (\rho - 1)\log A - (\rho - 1)\log M$$

If $\frac{dM}{d\tau} > 0 \Rightarrow \frac{dA}{d\tau} > 0$; else $\frac{d\frac{p_A}{p_M}}{d\tau} > 0 \Rightarrow \frac{d\log\varphi_M(\theta_1) - \log\varphi_A(\theta_1)}{d\tau} > 0 \Rightarrow \frac{d\theta_1}{d\tau} < 0$, where the last part is because

$$\frac{d\log \varphi_M(\theta_1) - \log \varphi_A(\theta_1)}{d\tau} = \frac{\partial \log \varphi_M(\theta_1) - \log \varphi_A(\theta_1)}{\partial \theta_1}\frac{\partial \theta_1}{\partial \tau}$$

□

By this lemma, if θ_0 increases and M increases, A increases as well, implying that θ_1 decreases, which is a contradiction (M cannot increase if θ_0 increases and θ_1 decreases).

Taken together

$$\frac{d\theta_0}{d\tau} = -\frac{\frac{\partial f}{\partial \tau}}{\frac{\partial f}{\partial \theta_0}} < 0$$

□

C.1 Multidimensional Model Proofs

Theorem 6

Proof. Similar to Theorem 1. □

Theorem 7

Proof. We first prove that $\frac{R}{A}$ and $\frac{R}{M}$ both increase with τ.

Assume without lost of generality that $\frac{M}{A}$ weakly decreases. Assume by contradiction that $\frac{R}{A}\downarrow$. From the FOC of the CES function we have

$$\frac{p_R}{p_A}=\left(\frac{R}{A}\right)^{\rho-1},\frac{p_M}{p_A}=\left(\frac{M}{A}\right)^{\rho-1}. \tag{22}$$

Since $\rho<1$, p_R/p_A increases and p_M/p_A weakly increases. The indifference curve between the two occupations R,A is

$$\varphi_A\left(\theta_{A,i}\right)=\frac{p_R}{p_A}\varphi_R\left(\theta_{R,i},\tau\right),$$

and between M,A

$$\varphi_A\left(\theta_{A,i}\right)=\frac{p_M}{p_A}\varphi_M\left(\theta_{M,i}\right).$$

Both $\frac{p_R}{p_A}$, $\varphi_R\left(\theta_{R,i}\right)$ increase. Hence the indifference abstract skill level $\theta_{A,i}\left(\theta_{R,i}\right)$ increases for every $\theta_{R,i}$. Similarly $\theta_{A,i}\left(\theta_{M,i}\right)$ is weakly increasing. Therefore, for every pair $\left(\theta_{M,i},\theta_{R,i}\right)$, the employment level at A are lower. Hence, the overall intermediate good A is decreasing. Since $\frac{M}{A}$, $\frac{R}{A}$ are assumed to decrease, this implies that both M and R are also decreasing. Since all intermediate goods decline Y must be lower. However, under the same allocation of workers into occupations, when τ increases R increases while M,A are unchanged. Hence, Y must be larger in GE (otherwise, Y is not maximized), a contradiction.

Therefore $\frac{R}{A}$ is increasing, and since $\frac{M}{A}$ is weakly decreasing, $\frac{R}{M}=\frac{R}{A}\frac{A}{M}$ is also increasing. If $\frac{M}{A}$ is weakly increasing, a similar argument applies.

We can now prove the theorem following the same order.

1. If R does not increase then either M or A must increase in order for Y to increase. Assume without lost of generality that A increases. Hence $\frac{R}{A}\downarrow$, contradicting our previous claim.

2. From the FOC we have

$$p_R=\left(\frac{R}{Y}\right)^{\rho-1}=\left(\frac{R}{\left(M^\rho+R^\rho+A^\rho\right)^{\frac{1}{\rho}}}\right)^{\rho-1}=\left(\left(\frac{M}{R}\right)^\rho+1+\left(\frac{A}{R}\right)^\rho\right)^{\frac{1-\rho}{\rho}}.$$

   Since $\frac{M}{R}$, $\frac{A}{R}$ decrease and $\rho<0$ p_R decreases.

   Relative prices must also decrease since $\frac{p_R}{p_A}=\left(\frac{R}{A}\right)^{\rho-1}$ and $\rho-1<0$.

3. The share is $$\frac{p_R R}{Y} = \frac{R^\rho}{Y^\rho} = \frac{R^\rho}{M^\rho + R^\rho + A^\rho} = \frac{1}{\left(\frac{M}{R}\right)^\rho + \left(\frac{A}{R}\right)^\rho + 1}$$
since R increases relative to both M, A, $\left(\frac{M}{R}\right)^\rho$ and $\left(\frac{A}{R}\right)^\rho$ increase and therefore $\frac{p_R R}{Y}$ decreases.

$\square$

Theorem 8 Assume a skill-replacing technology ($\epsilon < 0$). An increase in τ would generate the following:

1. A decrease in wage gaps between workers in the routine occupation who do not switch occupations.
2. The highest skill routine workers are less likely to work in the routine occupation. Formally, use $r(\theta_R)$ to note the probability of workers with routine skill level θ_R to choose to work in the routine occupation. $\exists \theta_R^*$ such that
$$\forall \theta_R > \theta_R^* : \frac{dr(\theta_R)}{d\tau} < 0.$$
3. The wage for the highest-skilled routine workers would decrease relative to all other workers. Formally, $\exists \theta_R^{**}$ such that the set of "high-skilled routine workers",
$$\overline{R} = \{\theta_i = (\theta_{M,i}, \theta_{R,i}, \theta_{A,i}) \,|\theta_{R,i} \geq \theta_R^{**}; j(\theta_i) = R\}$$
satisfies
$$\forall (\theta_1, \theta_2) \in \overline{R} \times \overline{R}^C : \frac{d \log w(\theta_1)}{d\tau} < \frac{d \log w(\theta_2)}{d\tau}$$
where $\overline{R}^C$ is the complement set of $\overline{R}$ and $w(\theta_i)$ is the wage at the chosen occupation.

Proof. 1. By theorem 6.

2. Define $\theta_R^{\max} = \sup \theta_{R,i}$. Assume WLOG that $\frac{d\frac{p_M}{p_A}}{d\tau} \geq 0$. Therefore we have
$$\frac{d \log\left(\frac{p_R}{p_A}\varphi_R\left(\theta_R^{\max}, \tau\right)\right)}{d\tau} \geq \frac{d \log\left(\frac{p_R}{p_M}\varphi_R\left(\theta_R^{\max}, \tau\right)\right)}{d\tau}.$$
Assume by contradiction that
$$\frac{d \log\left(\frac{p_R}{p_A}\varphi_R\left(\theta_R^{\max}, \tau\right)\right)}{d\tau} \geq 0. \tag{23}$$

Because $\epsilon < 0$ ($\frac{d^2 \log \varphi_R(\theta_R,\tau)}{d\tau d\theta_R} < 0$) we get,

$$\forall \theta_{R,i} < \theta_R^{\max} : \frac{d \log\left(\frac{p_R}{p_A}\varphi_R\left(\theta_{R,i},\tau\right)\right)}{d\tau} > \frac{d \log\left(\frac{p_R}{p_A}\varphi_R\left(\theta_R^{\max},\tau\right)\right)}{d\tau} \geq 0. \tag{24}$$

The indifference curve $\theta_{A,i}\left(\theta_{R,i}\right)$ can be written as the solution to

$$\log \varphi_A\left(\theta_{A,i}\right) = \log\left(\frac{p_R}{p_A}\varphi_R\left(\theta_{R,i},\tau\right)\right). \tag{25}$$

From Equation 24, the RHS of Equation 25 is increasing in τ for every value of $\theta_{R,i}$. Since $\log \varphi_A$ is monotone then $\theta_{A,i}(\theta_{R,i})$ is also increasing. By the assumption that $\frac{d\frac{p_M}{p_A}}{d\tau} \geq 0$ then the indifference curve $\theta_A\left(\theta_M\right)$ also increases for every θ_M. Together, this implies that the total production in A declines ($\frac{dA}{d\tau} < 0$).
However, rewriting Equation 23 we get

$$\frac{d \log \varphi_R\left(\theta_R^{\max},\tau\right)}{d\tau} \geq \frac{d \log \frac{p_A}{p_R}}{d\tau}.$$

From Lemma 10 and using $\rho < 0$ and $\frac{d \log \frac{p_A}{p_R}}{d\tau} > 0$ (Theorem 7) we can write

$$\frac{d \log R}{d\tau} > \frac{d \log \varphi_R\left(\theta_R^{\max},\tau\right)}{d\tau} \geq \frac{d \log \frac{p_A}{p_R}}{d\tau} > \frac{1}{1-\rho}\frac{d \log \frac{p_A}{p_R}}{d\tau}.$$

Therefore, using Equation 22

$$\frac{d \log A}{d\tau} = \frac{d \log R\left(\frac{p_R}{p_A}\right)^{\frac{1}{1-\rho}}}{d\tau} = \frac{d \log R}{d\tau} - \frac{1}{1-\rho}\frac{d \log \frac{p_A}{p_R}}{d\tau} > 0.$$

Therefore, $\frac{dA}{d\tau} > 0$, a contradiction. Hence,

$$0 > \frac{d \log\left(\frac{p_R}{p_A}\varphi_R\left(\theta_R^{\max},\tau\right)\right)}{d\tau} \geq \frac{d \log\left(\frac{p_R}{p_M}\varphi_R\left(\theta_R^{\max},\tau\right)\right)}{d\tau}$$

and therefore both the indifference curves decline at $\theta_R^{\max}$ when τ increases

$$\frac{d\theta_A\left(\theta_R^{\max}\right)}{d\tau}, \frac{d\theta_M\left(\theta_R^{\max}\right)}{d\tau} < 0.$$

Therefore $\frac{dr\left(\theta_R^{\max}\right)}{d\tau} < 0$. From continuity of r, $\exists\theta_R^*$ such that $\forall\theta_R > \theta_R^* : \frac{dr(\theta_R)}{d\tau} < 0$.

3. Assume again WLOG that $\frac{d\frac{p_M}{p_A}}{d\tau} \geq 0$. Define θ_R^{**} as the solution to

$$\frac{d\log\frac{p_R}{p_A}\varphi_R\left(\theta_R^{**},\tau\right)}{d\tau} = 0.$$

If such a solution does not exist then define $\theta_R^{**} = \inf\theta_R$. From the previous proof $\frac{d\log\frac{p_R}{p_A}\varphi_R\left(\theta_R^{\max},\tau\right)}{d\tau} < 0$. Therefore $\forall\theta_{R,i} \in \left(\theta_R^{**},\theta_R^{\max}\right)$

$$\frac{d\log\frac{p_R}{p_A}\varphi_R\left(\theta_{R,i},\tau\right)}{d\tau} < 0,$$

and hence

$$\frac{d\log p_R\varphi_R\left(\theta_{R,i},\tau\right)}{d\tau} < \frac{d\log p_A}{d\tau} < \frac{d\log p_M}{d\tau}.$$

Since $\frac{d\log w_R\left(\theta_{R,i}\right)}{d\tau} = \frac{d\log p_R\varphi_R\left(\theta_{R,i},\tau\right)}{d\tau}$ and $\frac{d\log w_j\left(\theta_{j,i}\right)}{d\tau} = \frac{d\log p_j}{d\tau}$ for $j \in \{M, A\}$ we can write

$$\frac{d\log w_R\left(\theta_{R,i}\right)}{d\tau} < \frac{d\log w_A}{d\tau} < \frac{d\log w_M}{d\tau}.$$

Hence wages for workers in the routine occupation with $\theta_R > \theta_R^{**}$ are falling relative to wages of abstract and manual workers. Wages also fall relative to all workers in the routine occupation with $\theta_R \leq \theta_R^{**}$ based on part 1. □

Theorem 9 Use $j\left(\theta_i,\tau\right)$ to mark the set of occupations that maximizes wages for a worker with a skill vector θ_i for for a given technology level τ.[26] Define the set of skill vectors of workers who shift from occupation j_1 to j_2 for an infinitesimal rise in technology τ as

$$\Theta_{j_2}^{j_1}\left(\tau\right) = \left\{\theta | j_1 \in j\left(\theta;\tau\right), j_2 \in j\left(\theta;\tau+\varepsilon\right)\right\}$$

for $\varepsilon > 0$ when $\varepsilon \to 0$. Then for every workers who shifts from routine to occupation j with skill vector $\theta_1 \in \Theta_j^R$, and every worker who shifts from occupation j to routine with skill vector $\theta_2 \in \Theta_R^j$, the routine skill is higher for those who leave the routine occupation

$$\theta_{R,1} > \theta_{R,2}.$$

[26]While $j\left(\theta_i,\tau\right)$ is typically a singleton, it could contain two or more occupations for workers whose skill levels are on the indifference curve.

Proof. Since $\theta_1 \in \Theta_j^R, \theta_2 \in \Theta_R^j$ the indifference curve declines at $\theta_{R,1}$ and increases at $\theta_{R,2}$

$$\frac{d \log p_R + \log \varphi_R(\theta_{R,1};\tau) - \log p_j}{d\tau} < 0$$
$$\frac{d \log p_R + \log \varphi_R(\theta_{R,2};\tau) - \log p_j}{d\tau} > 0.$$

Therefore

$$\frac{d \log \varphi_R(\theta_{R,2};\tau)}{d\tau} > \frac{d \log \varphi_R(\theta_{R,1};\tau)}{d\tau}$$

and since $\epsilon < 0$, $\theta_{R,1} > \theta_{R,2}$. □

D Interactive Fixed Effects Model Appendix

D.1 Estimation Details

Estimation Goal The goal of the IFEM estimation is to find the set of parameters that minimizes the expected squared error. That is, solve

$$\alpha, \beta, \lambda, \theta = \arg \min_{\alpha,\beta,\lambda,\theta} E\left[\left(\log w_{ijt} - \beta_{jt} X_{it} - \lambda_{jt} - \alpha_{jt}\theta_i\right)^2\right]. \tag{26}$$

Writing E_{jt} as the set of workers in occupation j in year t, these parameters provide the best approximation of the conditional expectation of $E\left[\log w_{ijt} | X_{it}, i \in E_{jt}\right]$ given this specification. The first-order conditions are

$$E\left[\varepsilon_{ijt} | i \in E_{jt}\right] = E\left[X_{it}\varepsilon_{ijt} | i \in E_{jt}\right] = E\left[\theta_i \varepsilon_{ijt} | i \in E_{jt}\right] = E\left[\alpha_{j(i,t)t}\varepsilon_{i,j(i,t)t,} | i\right] = 0,$$

where $\varepsilon_{ijt} = \log w_{ijt} - \beta_{jt}X_{it} - \lambda_{jt} - \alpha_{jt}\theta_i$ is the minimized error. From the last condition, we can find an expression for the least-squares estimator of $\widehat{\theta}_i$ given the other parameters (Equation 9).

Measurement Error While the true parameters would satisfy the first-order conditions in Equation 26, they would not satisfy the corresponding empirical moments due to measurement errors. Under the true parameters

$$\widehat{\theta}_i\left(\log w_i, X_i, \alpha, \beta, \lambda\right) = \theta_i + \frac{\sum_t \alpha_{j(i,t),t}\varepsilon_{ij(i,t)t}}{\sum_t \alpha^2_{j(i,t),t}} = \theta_i + \nu_i,$$

where $\nu_i = \frac{\sum_t \alpha_{j(i,t),t}\varepsilon_{ij(i,t)t}}{\sum_t \alpha^2_{j(i,t),t}}$ is the error in the estimation of θ_i, and is a linear function of all residuals ε_{ijt} for person i. Then, the empirical estimator of $\widehat{\varepsilon}_{ijt}$ is

$$\widehat{\varepsilon_{ijt}}\left(\log w_i, X_i, \alpha, \beta, \lambda\right) = \log w_{ijt} - \beta_{jt}X_{it} - \lambda_{jt} - \alpha_{jt}\widehat{\theta}_i\left(\log w_i, X_i, \alpha, \beta, \lambda\right) = \varepsilon_{ijt} - \alpha_{jt}\nu_i. \tag{27}$$

The moments $E\left[\widehat{\theta}_i\widehat{\varepsilon}_{ijt}|i\in E_{jt}\right]$ do not equal to zero in expectation since

$$E\left[\widehat{\theta}_i\widehat{\varepsilon}_{ijt}|i\in E_{jt}\right]=COV\left(\nu_i,\varepsilon_{ijt}-\alpha_{jt}\nu_i\right)=-\alpha_{jt}V\left(\nu_i\right)+COV\left(\nu_i,\varepsilon_{ijt}\right)\neq 0.$$

When the number of observations for every individual workers goes to infinity, $E\left[\widehat{\theta}_i\widehat{\varepsilon}_{ijt}|i\in E_{jt}\right]\to 0$, and the model is consistent. However, many workers are observed only in a small number of periods. Therefore, a least square estimator would be asymptotically biased. In particular the least square estimator for α_{jt} is

$$\widehat{\alpha}_{jt}^{LS}=\frac{\sum_{i\in E_{jt}}\left(\log w_{ij(i,t)t}-\widehat{\beta}_{jt}X_{it}-\widehat{\lambda}_{jt}\right)\widehat{\theta}_i}{\sum_{i\in E_{i,t}}\widehat{\theta}_i^2}.$$

This estimator is not consistent as it converges to

$$\begin{aligned}\text{plim}_{N\to\infty}\widehat{\alpha}_{jt}^{LS}&=\frac{E\left[\left(\alpha_{jt}\theta_i+\varepsilon_{ijt}\right)\widehat{\theta}_i\mid i\in E_{jt}\right]}{E\left[\widehat{\theta}_i^2\mid i\in E_{jt}\right]}\\&=\alpha_{jt}\underbrace{\frac{E\left[\theta_i^2\mid i\in E_{jt}\right]}{E\left[\theta_i^2\mid i\in E_{jt}\right]+V\left(\nu_i\mid i\in E_{jt}\right)}}_{\text{attenuation bias}}+\underbrace{\frac{\text{COV}\left(\varepsilon_{ijt},\nu_i\mid i\in E_{jt}\right)}{E\left[\theta_i^2\mid i\in E_{jt}\right]+V\left(\nu_i\mid i\in E_{jt}\right)}}_{\text{correlated errors}}.\end{aligned}$$

Instrumental Variable An IV Z_i can be used to solve the asymptotic bias. Equation 10 guarantees that the IV is uncorrelated with the error ε_{ijt}. Guaranteeing that Z_i is also uncorrelated with the measurement error ν_i requires strict exogeneity, similar to standard fixed-effects models. In standard fixed-effect models, the error is assumed to be mean independent of all control variables (Chamberlain, 1984). Together with the independence from the IV, this is

$$E\left[\varepsilon_{ijt}|C_{i1},...,C_{iT},Z_i\right]=0, \tag{28}$$

where C_{it} are the control variables in time t (X_{it} and indicators for employment in each occupational category). Under this assumption, the IV is uncorrelated with the measurement error,

$$E\left[Z_i\nu_i\right]=E\left[Z_i\frac{\sum_t\alpha_{j(i,t),t}\varepsilon_{ij(i,t)t}}{\sum_t\alpha_{j(i,t),t}^2}\right]=E_{C_{i1},...,C_{iT}}\left[\frac{\sum_t\alpha_{j(i,t),t}E\left[Z_i\varepsilon_{ij(i,t)t}|C_{i1},...,C_{iT}\right]}{\sum_t\alpha_{j(i,t),t}^2}\right]=0,$$

where the second equality is by the law of iterated expectations and the third equality is from Equation 28. Therefore, under the true parameters, an IV would satisfy $E[Z_i\widehat{\varepsilon}_{ijt}|i\in E_{jt}]=0$.

The parameters can be estimated using the method of moments.

Theorem 12. *Let Z_i be an IV that satisfies strict exogeneity (Equation 28). The parameters that solve Equation 26 are the solution of the system of equations specified in Equation 11.*

Proof. Based on the law of large numbers, $\text{plim}_{N\to\infty} m^W_{jt} = E[W_i \widehat{\varepsilon}_{ijt} | i \in E_{jt}]$ for $W \in 1, X, Z$ (where m^W_{jt} is as defined in Equations 11). From the first-order conditions, the conditional expectations $E[\widehat{\varepsilon}_{ijt}|i \in E_{jt}], E[X_i \widehat{\varepsilon}_{ijt}|i \in E_{jt}]$ equal zero for the parameters that solve Equation 26. The conditional expectation $E[Z_i \widehat{\varepsilon}_{ijt}|i \in E_{jt}]$ is zero from Equation 28. □

TSLS Intuition The estimator received from solving the system of equations in 11 is not a TSLS estimator, since $\widehat{\theta}_i$ is a function of the other parameters. However, it can be written in a similar manner.

Theorem 13. *Let $\widehat{\alpha}_{jt}, \widehat{\beta}_{jt}, \widehat{\lambda}_{jt}$ be the solution to Equations 11 with a an IV Z that satisfies strict exogeneity (Equation 28). Then*

$$\widehat{\alpha_{jt}} = \frac{COV\left(Z_i, \widetilde{\log w_i} | i \in E_{jt}\right)}{COV\left(Z_i, \widetilde{\widehat{\theta}_i} | i \in E_{jt}\right)},$$

where $\widetilde{\log w_i}, \widetilde{\widehat{\theta}_i}$ are the residuals of $\log w_i, \widehat{\theta}_i$ from a regression on $1, X_{it}$ interacted with occupational category and year dummies.

Proof. We can write $\widehat{\theta}_i$ (given the other parameters $\widehat{\alpha}_{jt}, \widehat{\beta}_{jt}, \widehat{\lambda}_{jt}$) as

$$\widehat{\theta}_i = \phi_{jt} X_{_{it}} + \mu_{jt} + \rho_{jr} Z_i + \vartheta_{ijt}, \tag{29}$$

where $\phi_{jt}, \mu_{jt}, \rho_{jt}$ are the OLS parameters from a regression of $\widehat{\theta}_i$ on $1, X_{it}, Z_i$ interacted with occupational category and year dummies. This is similar to the first stage equation a 2SLS. Define the first-stage predictor of $\widehat{\theta}_i$ as $\widehat{\theta}^{FS}_{ijt} = \phi_{jt} X_{_{it}} + \mu_{jt} + \rho_{jr} Z_i$. We can also write a second-stage equation,

$$\log w_{ijt} = \widehat{\beta}_{jt} X_{it} + \widehat{\lambda}_{jt} + \widehat{\alpha}_{jt} \widehat{\theta}^{FS}_{ijt} + \zeta_{ijt}, \tag{30}$$

where $\zeta_{ijt} = \widehat{\alpha_{jt}}(\widehat{\theta}_i - \widehat{\theta}^{FS}_{ijt}) + \varepsilon_{ijt}$. By the construction of $\widehat{\theta}^{FS}_{ijt}$ as an OLS predictor, $(\widehat{\theta}_i - \widehat{\theta}^{FS}_i)$ is orthogonal to all the regressors of Equation 29. Since these parameters solve the system of equations in 11, ε_{ijt} is also orthogonal to the same regressors. Therefore, ζ_{ijt} is orthogonal to the regressors of Equation 29 as well. Because $\widehat{\theta}^{FS}_{ijt}$ is a linear combination of this set of regressors (Equation 29), it is orthogonal to ζ_{ijt} as well. Therefore, ζ_{ijt} is orthongal to all regressors in Equation 30 and the parameters $\widehat{\alpha}_{jt}, \widehat{\beta}_{jt}, \widehat{\lambda}_{jt}$ can be derived from an OLS estimate of this equation.

Since both Equations 29 and 30 represent OLS estimates, Frisch–Waugh–Lovell theorem implies that we can residualize both equations by any subset of the regressors. Residualizing by the interaction of $1, X_{it}$ with all occupational categories and year dummies yields

$$\begin{array}{rcl} \widetilde{\widehat{\theta}^{FS}_{ijt}} & = & \rho_{jt} \widetilde{Z}_i + \widetilde{\vartheta_{ijt}} \\ \widetilde{\log w_{ijt}} & = & \widehat{\alpha_{jt}} \widetilde{\widehat{\theta}^{FS}_{ijt}} + \widetilde{\zeta_{ijt}} \end{array}.$$

Plugging in the residualized first-stage equation in the second-stage equation yields the residualized reduced-form equation

$$\widetilde{\log w_{ijt}} = \widehat{\alpha_{jt}} \rho_{jt} \widetilde{Z}_i + \xi_{ijt},$$

with $\xi_{ijt} = \widetilde{\zeta_{ijt}} + \widehat{\alpha_{jt}} \widetilde{\vartheta_{ijt}}$. Therefore, $\widehat{\alpha}_{jt}$ is the division of the first-stage and reduced-form equations

$$\widehat{\alpha}_{jt} = \frac{COV\left(\widetilde{Z}_i, \widetilde{\log w_{ijt}}\right)}{COV\left(\widetilde{Z}_i, \widetilde{\widehat{\theta}}_i\right)} = \frac{COV\left(Z_i, \widetilde{\log w_{ijt}}\right)}{COV\left(Z_i, \widetilde{\widehat{\theta}}_i\right)}$$

where the last equality is because for any two variables W_1, W_2 $COV\left(\widetilde{W_1}, \widetilde{W_2}\right) = COV\left(W_1, \widetilde{W_2}\right)$. □

D.2 Sensitivity Analysis

In this section, I discuss a sensitivity analysis (Andrews et al., 2017) to study how the various moments in Equation 11 affect the estimated parameters. This analysis shows that changes in α_{jt} over time in a given occupational category are driven almost entirely by workers in this category. By contrast, differences in α_{jt} between two occupations are driven by workers in these two occupations. Moreover, this analysis shows that even under relatively large violations of the model assumptions, the key results would still hold.

The estimated error under the target parameters (Equation 27) can be written as

$$\widehat{\varepsilon}_{ijt} = \gamma^1_{jt} + \gamma^X_{jt} X_{it} + \gamma^Z_{jt} Z_i + \zeta_{ijt} \tag{31}$$

where ζ satisfies

$$E\left[\zeta_{ijt} | i \in E_{jt}\right] = E\left[X_{it}\zeta_{ijt} | i \in E_{jt}\right] = E\left[Z_i \zeta_{ijt} | i \in E_{jt}\right] = 0.$$

The γ parameters are the population-level OLS of the errors $\widehat{\varepsilon}_{ijt}$ on $1, X_{it}, Z_i$ for workers employed in occupational category j at time t. Under strict exogeneity (Equation 28) $\gamma = 0$ and the target parameters are identified. However, if strict exogeneity does not hold $\gamma \neq 0$, the error is correlated with the regressors and the estimated parameters would be biased. The coefficients $\gamma^1_{jt}, \gamma^X_{jt}$ are by construction small since the true parameters satisfy Equation 8.[27] I therefore focus on γ^Z_{jt} which potentially could have larger deviations from zero. This parameter is large when the IV is strongly correlated with the errors.

The sensitivity analysis by Andrews et al. (2017) enables quantifying the link between the bias indicated by γ^Z_{jt} and any function of the model parameters. This process involves calculating the sensitivity matrix Λ, which delineates the first-order relationship between

[27] Equation 8 guarantees that $E\left[\varepsilon_{ijt'} | i \in E_{jt}\right] = E\left[\varepsilon_{ijt'} X_{it} | i \in E_{jt}\right] = 0$ for $t' = t$. Since $\widehat{\varepsilon}_{ijt}$ is close to ε_{ijt} (Equation 27), $\gamma^1_{jt}, \gamma^X_{jt}$ will be close to zero. However, if strict exogeneity (Equation 28) does not hold, then for $t' \neq t$, $E\left[\varepsilon_{ijt'} | i \in E_{jt}\right]$ and $E\left[\varepsilon_{ijt'} X_{it} | i \in E_{jt}\right]$ might not equal zero and therefore $E\left[\widehat{\varepsilon}_{ijt'} | i \in E_{jt}\right], E\left[\widehat{\varepsilon}_{ijt'} X_{it} | i \in E_{jt}\right] \neq 0$.

each moment and each parameter. Since the number of parameters equals the number of moments, the sensitivity matrix is calculated directly from the Jacobian of the moments on the parameters G using $\Lambda = -G^{-1}$. I focus on the relative difference between two α_{jt} parameters: $\log \frac{\alpha_{j_0 t_0}}{\alpha_{j_1 t_1}}$. Therefore, I multiply the matrix Λ by vector v which is the derivative of $\log \frac{\alpha_{j_0 t_0}}{\alpha_{j_1 t_1}}$ by all the model parameters.[28] I follow Andrews et al. (2017) and multiply Λ by Ω, the covariance of all combinations of $1, X_{it}, Z_i$ interacted with an indicator for $i \in E_{jt}$. Andrews et al. (2017) show that the bias in $\log \frac{\alpha_{j_0 t_0}}{\alpha_{j_1 t_1}}$ is $\gamma' \Omega \Lambda v$, where γ is the vector of $\gamma^1_{jt}, \gamma^X_{jt}, \gamma^Z_{jt}$ for all j, t. Therefore, given a speculated bias level γ, one can asses the magnitude of the bias.

I first examine the potential bias in the estimate of a two-year relative change in α in routine occupations. For every year t, I calculate the bias in $\log \frac{\alpha_{R,t}}{\alpha_{R,t-2}}$ for a deviation from zero in $\gamma^Z_{j' t+k}$, for every possible combination of occupation category j' and even year distance k. That is, I calculate how much a correlation between the IV and the error term in occupational category j' and year $t+k$ would bias the estimate for the trend in the return to skill in routine occupations, α_R, between years $t-2$ and t. Figure A15a shows the average results over all years in the sample. Each point represents an occupational category j' and the distance years k in $\gamma^Z_{j' t+k}$.

I find that the trend in α_R is driven almost entirely by workers in routine occupations in the relevant years. Figure A15a shows that for the abstract and manual categories, as well as for other years in the routine category, the value is close to zero. This implies that even a very strong correlation between years of schooling and the residual $\widehat{\varepsilon}_{i,j',t+k}$ in these occupations or years will not change the results. This also implies that the estimate for $\log \frac{\alpha_{R,t+2}}{\alpha_{R,t}}$ is driven almost entirely by workers in routine occupations in those years. If years of schooling had a large correlation with the errors in other occupations or years (high value of $\gamma^Z_{j' t+k}$), the results will barely change. The results for abstract and manual occupations are qualitatively similar and are therefore unreported.

Figure A15a also shows that any bias in the trend in α in the routine occupation is likely small. A constant bias ($\gamma^Z_{Rt} = c$ for all t) would have no impact on the estimated trend. This is because $\gamma^Z_{R,t-2}$ and $\gamma^Z_{R,t}$ cancel each other out almost exactly. Therefore, if schooling also affects wages not through θ_i, and the effect is stable over time, the bias in the trend would be very small.

To generate a bias in $\log \frac{\alpha_{R,t}}{\alpha_{R,t-2}}$, the bias coefficient γ^Z_{Rt} must change over time. Consider a simple case where this change is linear and

$$\gamma^Z_{Rt} = a + b * t.$$

In this case the bias would be approximately $4\gamma^Z_{Rt} - 4\gamma^Z_{Rt-2} = 8b$ based on the values of Figure A15a.

Reasonable values for b imply that the bias would change the trend in α_R by less than 10%. The premium for one year of schooling is approximately .07 (Lemieux, 2006b). Most of the return to schooling represents permanent income differences and is therefore captured in θ_i. Therefore, the coefficient of $\widehat{\varepsilon}_{ijt}$ on years of schooling is likely much lower, and so presumably below .02. An extreme scenario where this coefficient doubled itself between 1987-2007 would imply that $b = 0.001$. In this case the bias is $8b = 0.008$. The overall bias in the course of these

[28] The vector v equals $\alpha^{-1}_{j_0,t_0}$ for α_{j_0,t_0}, $\alpha^{-1}_{j_1,t_1}$ for α_{j_1,t_1} and zero for every other parameter.

twenty years would be .08. This is approximately 10% of the measured decrease in $\alpha_{R,t}$ in this period.

Figure A15b shows that the differences between occupational categories are identified mainly by workers in these categories in that year. For every year t, I calculate the bias in $\log \frac{\alpha_{R,t}}{\alpha_{M,t}}$ for a deviation from zero in $\gamma^{Z}_{j't'}$, for every combination of occupational category and year, j', t'. Figure A15b then shows the average results over all years. This shows that the differences in α between routine and manual occupations in year t are mainly driven by the return to schooling in this occupation in that year. Similar results are received for the other two pairs of occupations and are not reported.

D.3 Interactive Fixed Effects Model With Multiple Skills

In this section, I derive an analytic solution to the IFEM parameters in a simple multi-skill setting. I show that the target IFEM parameters are the first principal component of the matrix of the returns to the different skills over time. The parameters estimated using an IV are a weighted average of the returns to different skills, where the weights are proportional to the correlation of each skill with the IV. Hence, while using years of schooling as an IV would not yield the target parameters, it would still be informative about the aggregate trends in the returns to skills that are correlated with education.

I analyze a model in which workers use multiple skills simultaneously to produce. The skills have different returns that change over time. Let μ_i be the vector of K skills for each worker i. Without loss of generality, assume that skills are uncorrelated and standardized such that $V(\mu_i) = I$. For simplicity, assume only one occupation and full employment in each period.[29] Let a_{kt} be the return to skill k in period t. The matrix A denotes the return to each skill in each period ($[A]_{kt} = a_{kt}$).

Log wages are the product of the skills and their return, together with an idiosyncratic zero-mean shock

$$\log w_{it} = a_t' \mu_i + \varepsilon_{it}.$$

In matrix form, the vector of log wages for worker i in all periods is

$$\log w_i = A^T \mu_i + \varepsilon_i.$$

The IFEM is misspecified under this data-generating process (DGP). However, the IFEM can find the best approximation of this DGP with only one skill. In particular, it searches for a vector $\alpha = (\alpha_1, ..., \alpha_T)$ of return to skill in each period, and a function $\theta(\mu_i) : \mathbb{R}^K \to \mathbb{R}$ that aggregates the skills into one dimension. Formally, it solves

$$\min_{\alpha,\theta} \sum_{t=1}^{T} \mathbb{E}\left[\left(a_t'\mu_i - \alpha_t \theta_i(\mu_i)\right)^2\right]. \tag{32}$$

The solution to this minimization problem can be derived analytically.

Theorem 14. *Let the vector of parameters α and function θ be the solution to Equation 32. Then α is the first (non-centralized) principal component of A^T, and θ is a linear combination of the different skills, where the coefficients are the loadings of the first PCA component.*

[29]This model is different from the model in Appendix B.3 as workers use multiple skills simultaneously in the same job.

Intuitively, this theorem implies that the IFEM estimates $\widehat{\theta}_i$ is a weighted average of skills and α is the weighted average of the returns of these skills. The weights are set such that the IFEM predictions are closest to the actual log wages. Hence, skills that have higher returns and that are more internally correlated will have larger weights in θ, and so larger impact on α.

Proof. Taking first order conditions by $\theta_i(\mu_i)$, yields

$$\theta_i(\mu_i) = \|\alpha\|^{-2}\alpha' A^T \mu_i,$$

therefore, θ_i is a linear function, $\theta(\mu_i) = b'\mu_i$, for some K-dimensional b vector. Define $B = b\alpha'$, where b, α are the solution to Equation 32. This matrix satisfies $rank(B) = 1$ by construction, and any matrix that satisfies $rank(B) = 1$ can be written as $B = b\alpha'$. Therefore, Equation 32 can be rewritten in matrix form as

$$\min_{B|rank(B)=1} \mathbb{E}\left[\mu_i'(A-B)(A-B)^T \mu_i\right].$$

We assumed that $V(\mu_i) = \mathbb{E}\left[\mu_i'\mu_i\right] = I$. Therefore, the minimized expression equals

$$trac\mathbb{E}\left[(A-B)(A-B)^T\right],$$

which is known as the Frobenius norm of $A - B$. Hence, the true IFEM parameters solve

$$\min_{B|rank(B)=1} \|A-B\|_F. \tag{33}$$

Equation 33 is the well-known "low-rank approximation" problem whose solution is given by the Eckart–Young–Mirsky theorem. The solution is given by the first dimension of the singular value decomposition of the return matrix A. Therefore, α is the first eigenvector of $A^T A$, which is also the first (non-centralized) principal component of matrix A^T. Hence, α is a linear combination of the returns to various skills and $\theta(\mu_i)$ is a linear combination of the same skills (the PCA loadings).

□

This model also yields an informative analytic solution for estimating IFEM using an IV.

Theorem 15. *Let $\widehat{\alpha}$ be the vector of IFEM parameters, estimated with an IV Z. Then*

$$\widehat{\alpha} = A \cdot COV(\mu_i, Z_i).$$

This theorem implies that when using an IV, $\widehat{\alpha}$ is a weighted average of the return to different skills. The weights are set by $COV(\mu_i, Z_i)$, and therefore depend on the correlation of the IV with different skills. For example. if the IV is years of schooling, the vector $\widehat{\alpha}$ represents the trend in the average return to skill for skills correlated with schooling. The estimated parameter would not be equal to the target parameter, as it does not necessarily minimize the EMSE (Equation 32). However, it still captures an aggregate trend in the returns to the skills most correlated with schooling.

Proof. When using an IV Z, the IFEM finds $\widehat{\alpha}, b$ which solve

$$\mathbb{E}\left[\left(A - \widehat{\alpha} b'\right) \mu_i Z_i\right] = 0. \tag{34}$$

The expression $\mathbb{E}\left[\mu_i Z_i\right] = COV\left(\mu_i, Z_i\right)$ is the K-dimensional correlation of the IV with each one of the skills. The expression $b' COV\left(\mu_i, Z_i\right)$ is a scalar. Since there is a degree of freedom, this scalar can be set to 1. Therefore, Equation 34 can be written as

$$\widehat{\alpha} = A \cdot COV\left(\mu_i, Z_i\right).$$

□

E Data Appendix

E.1 CPS-ORG

The CPS-ORG provides the most accurate representative sample of hourly wages (Lemieux, 2006a). I use the same sample definitions as Acemoglu and Autor (2011), who kindly made their cleaned data files available online. See Acemoglu and Autor (2011) data appendix for exact definitions of the sample. Observations with missing wages are dropped. The main results hold when using imputations instead. Sampling weights are used in all CPS-ORG analyses. Education categories are equivalent to those employed by Autor et al. (2003) based on the consistent classification system proposed by Jaeger (1997).

One important limitation of the CPS data is its relatively high level of measurement errors. This problem is particularly severe at both tails of the distribution. Misreporting of working hours could lead to extremely high or extremely low values of hourly wages. Moreover, the CPS applies top coding to prevent the identification of individuals with extremely high income. Therefore, I drop the top and bottom 5% of the positive wages. Similar methods have been applied in previous work that used this data (Katz and Murphy, 1992; Autor et al., 2006; Autor et al., 2008; Acemoglu and Autor, 2011).

Since my analysis focuses on hourly wages, I multiply the CPS weights by the number of hours worked to obtain the real price of an hour of labor, as explained in Lemieux (2010). This procedure is also consistent with the literature.

E.2 PSID

I merge data from the individual survey and the family survey. The over-sampling of low-income households and immigrants samples are not used as it was added only in the 1990s. Therefore, I do not use sampling weights (similar to Cortes, 2016). I drop observations with hourly wages at the top or the bottom 5%, as with the CPS. Observations in which wage is imputed are also omitted.

The education variable is defined based on a survey question that inquires about the highest grade/years of schooling that the respondent has completed. A value of 16 indicates a college graduate. This variable is capped at 17. Occupations are coded using the Autor and Dorn (2013) occupational crosswalk for census coding.

E.3 O*NET

Routine Index: Acemoglu and Autor (2011) construct two indices: routine-manual and routine-cognitive. These indices are based on occupational averages of survey responses in the O*NET database (version 14). I take the average of both (standardized) indices. The routine manual index includes questions on:

- Pace determined by speed of equipment.
- Controlling machines and processes.
- Spend time making repetitive motions.

The routine cognitive index includes questions on:

- Importance of repeating the same tasks.
- Importance of being exact or accurate.
- Structured v. unstructured work (reverse).

I thank the authors for sharing their data with me.

F Cross-Occupation Wage Rank

In this section, I test to what extent are wage ranks preserved across the three different occupational categories. The results are presented in Appendix Table A1. In columns 1–3, I present the rank correlation of wages before and after the move, for workers who changed their three-digit occupation classification (51% of the sample). The table reports the average correlation across all survey years. Occupational switches are generating some change in ranking even for workers who remain in the same occupational category, with a rank correlation of around 0.5. When workers switch to a different occupational category, the rank correlations are only slightly lower, ranging between 0.35–0.53.

One reason why ranking changes across workers is short-term fluctuations in wages. To examine longer-term rankings, I estimate Equation 7 allowing θ_i to vary by the three broad occupational categories. This allows for a different skill to be used in each occupational category, as in the multi-skill model in Appendix B.3. I find that the (Pearson) correlations between θ_{ij} for a given value of i are between 0.69–0.83 as shown in columns 4–6 of Table A1. Since θ_{ij} is measured with a high level of noise, the results are downward biased.

These relatively high levels of correlations suggests that the long-term wage gaps within occupations seem to be a result of skill differences, that are relevant across occupations. Differences in rents (as in Acemoglu and Restrepo, 2024) are more likely to be occupation-specific. While workers are likely using different skills in different occupations, these skills are sufficiently correlated, at least for movers, who are the majority of the sample.

G Occupational Choice of New Hires

In this section, I show a substantial decline in the share of workers who join the labor market and choose to work in routine occupations, especially among middle-skilled workers. Previous work has found that the decline in routine employment is primarily driven by entry and exit from the labor force, and not by direct occupation transitions (Cortes et al., 2020). In an unreported analysis, I also find only a small number of switches between occupational categories. However, I do not find an increase in the share of exits from routine occupations outside of the labor market. Therefore, I focus primarily on the changes in entries.

Figure A16 looks at the share of new hires in each occupational category by year. I divide workers into three equal-sized bins based on their estimated $\widehat{\theta}$, redisualized by the cohort of entry to the labor force. For each bin, I plot the share of workers in each occupational category as a share of the total number of workers who were out of the labor force in that year and joined the labor force two years later.

Panel B of Figure A16 finds a substantial decline in the share of middle-skilled workers who join the labor force and are employed in routine occupations. In the early 1980s, almost half of the middle-skilled workers (based on their estimated $\widehat{\theta}_i$) joined routine occupations. This number decreased to around 33% after 2010. At the same time, there has been a substantial increase in the likelihood of middle-skilled workers joining manual occupations. After 2010, middle-skill workers are almost equally likely to join each of the occupational categories.

Panels A and C show smaller declines in the share of new routine workers from the top and bottom thirds of the skill distribution. In both panels, there is a decrease in the share of workers who join the labor market and work in a routine occupation. There is also an increase in the share of workers who join manual occupations. However, these trends are not as substantial as for middle-skilled workers. Estimating a linear trend for the share of middle-skilled workers who join a routine occupation when they enter the labor market, I find a decrease of 0.53 percentage points decline per year. This is compared to a 0.39 percentage point decline per year in the top third and a 0.25 percentage point decline per year in the bottom third.

These results imply that the decline in the average skill of workers in routine occupations (Figure 5) is primarily driven by the composition of new entries from outside the labor market. Over time, workers who join routine occupations from outside the labor market are more likely to arrive from the bottom of the skill distribution.

H Decomposing by More Than One Category

As variance decomposition, skewness decomposition can also be easily extended to accommodate linear models. Assume the following simple linear model when Y is standardized:

$$Y = \sum_i X_i$$

Using simple algebra we get

$$\mu_3 (Y) = \sum_i \mu_3 (X_i) + \sum_i \sum_{j \neq i} COV \left(X_i^2, X_j \right) + \sum_i \sum_{j \neq i} \sum_{k \neq i,j} E \left[X_i X_j X_k \right] \tag{35}$$

Therefore, we can decompose the skewness of Y into a linear combination of: (i) the skewness within the linear components, (ii) the covariance of the second and first moments of the linear components, and (iii) the triple multiplication of all three distinguished components. Though this decomposition includes a large number of different terms, many of them equal zero.

For example, writing Y as the sum of its conditional expectation in X and a residual ε

$$Y = E[Y|X] + \varepsilon$$

and using Equation 35 yields Equation 14, by the law of iterated expectations over X.

Linear skewness decomposition can be applied to decompose any linear model. This is useful for comparing occupations directly to other categories, such as industry and education. I show this for the equation

$$\ln w_i = occ_i + ind_i + \varepsilon_i, \tag{36}$$

where occ_i and ind_i are occupation and industry dummies. I then decompose the increase in skewness using Equation 35. Figure A17 presents the results. Most of the increase in skewness is generated by the increase in the correlation between the occupation premium occ_i and the residual variance ε_i^2. The equivalent component for industries (in green) is negligible. All other components, such as the skewness between occupations or industries, the correlation between the occupation premium and the industry premium variance and others are aggregated and plotted in red. Altogether, they comprise only a small share of the increase.

To do the same exercise for occupations with observable skills I estimate a Mincer equation with occupational dummies

$$\ln w_i = occ_i + \beta X_i + \varepsilon_i \tag{37}$$

where X_i includes years of schooling, experience and experience squared.

In this estimation, the occupation premiums are conditional on the workers' observed skills. Therefore, I decompose βX_i into mean occupation skill level and within-occupation skill difference

$$E[\beta X_i|occ_i] + (\beta X_i - E[\beta X_i|occ_i]),$$

such that the first component captures the average skill level in an occupation and the second component captures the skill part that is orthogonal to the occupation.

I implement a linear skewness decomposition into four components using the equation

$$\ln w_i = occ_i + E[\beta X_i|occ_i] + (\beta X_i - E[\beta X_i|occ_i]) + \varepsilon_i.$$

Figure A18 plots the results. I find that the two main components are the correlation of ε^2 with both occ_i and $E[\beta X_i|occ_i]$. This means that the correlation of the inequality of the unobservables (the variance of ε) with occupational wage levels is due to both occupation premium (occ_i) and the mean skill level at the occupation ($E[\beta X_i|occ_i]$). Hence, inequality is large in occupations that pay more and have higher-skilled workers, consistent with the SR-RBTC model. Categories that are unrelated to occupations are still negligible.