Skill-Replacing Technology and Bottom-Half Inequality

Oren Danieli*

June 14, 2026

Abstract

I propose a model of skill-replacing routine-biased technological change (SR-RBTC). In this model, technology substitutes for the use of skill in routine tasks, in contrast to standard RBTC models, which assume that technology replaces the workers themselves. The SR-RBTC model explains three key trends that are inconsistent with standard RBTC models: 1) why specifically middle wages declined even though workers in routine occupations are dispersed across the entire bottom half of the wage distribution, 2) why middle wages stopped declining while technological change continued, and 3) 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 the return to skill and a decrease in skill level in routine occupations. I use an interactive fixed-effects model to confirm both predictions. Since SR-RBTC violates the ignorability assumption required by standard decomposition methods, I introduce a "skewness decomposition" to show that SR-RBTC is the main driver of bottom-half inequality trends.

*Tel Aviv University, CEPR and IZA. 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, Thomas Lemieux, Chiara Farronato, Roland Fryer, 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 Wage Dynamics in the 21st Century, 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–Entrepreneurial Management, Federal Reserve Board, NY Fed, Yale, IAB, and Columbia 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%."

Over the past decades, the U.S. wage distribution has experienced profound shifts, which have substantially affected the population's welfare. These trends in inequality are frequently attributed to technological advancements. In particular, several studies have proposed a model of routine-biased technological change (RBTC), which posits that technology is a substitute for routine workers because the tasks they perform 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 argue that RBTC is skill-replacing, and that this change to the RBTC model can explain the recent trends in bottom-half inequality. Instead of assuming that new technology directly replaces workers, the skill-replacing RBTC (SR-RBTC) model assumes that new technology replaces the workers' usage of skill. I test this model with new empirical evidence. I find a decline in the return to skill in routine occupations, consistent with the theory. Moreover, I show that routine jobs are increasingly employing low-skilled workers, as skill becomes less necessary. Finally, I use a "skewness decomposition" to quantify that SR-RBTC can explain 93 percent of the wage trends discussed.

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 of the previous literature by assuming that 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 (e.g., Acemoglu and Autor, 2011; Cortes, 2016) that assume a skill-neutral technological change, where the wage effects are identical for all routine workers.[1] It also differs from skill-enhancing models (Jung and Mercenier, 2014) that make the opposite assump-

[1] Acemoglu and Restrepo (2022) propose a skill-neutral model with heterogeneous wage effects as some workers are more exposed to automation for exogenous reasons.

tion that technology increases the return to skill for routine workers. It is conceptually similar to Downey (2021), who argues that new technology benefits low-skill workers.

There are numerous examples of skill-replacing technologies 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.

The SR-RBTC model explains three key stylized facts that previous (skill-neutral) RBTC models could not. These unresolved puzzles leave room for other explanations for the fluctuations in bottom-half inequality (Hunt and Nunn, 2022). First, as shown in Figure 1, starting in 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-skilled 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.

By contrast, SR-RBTC predicts a decline in middle wages. A decrease in the 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. As a result, 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 continued long afterward (Autor, 2014).

This fact can be explained by SR-RBTC, which predicts a non-monotonic relationship between technological advancement and inequality in the bottom half of the distribution. The decline in return to skill in the routine occupation incentivizes high-skilled workers to leave the routine occupation and low-skilled workers to join. SR-RBTC would still reduce the wages of the highest-earning routine workers. However, due to the compositional changes, these workers would be concentrated at the bottom of the wage distribution. SR-RBTC would stop affecting middle-wage workers, as they would no longer work in the routine occupation. Therefore, inequality at the bottom of the distribution could increase.

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 entirely 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 the 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.

To test these predictions, I estimate an interactive fixed-effects model (IFEM). The 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 the 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) for 1980–2017.

The IFEM results 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 in 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 decreased during this period. By the end of the studied period, the average skill of workers in routine occupations fell below the average skill in manual occupations. Workers in routine occupations, especially those in administrative or operator occupations, have the lowest level of skill across all occupational categories.[2] As a result, workers in routine occupations became more concentrated at the bottom income quintile,

[2] 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).

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). However, there has been little discussion of the impact of these employment trends on the average level of skill in routine occupations.[3]

In the final part of the paper, I use a skewness decomposition to quantify the role of SR-RBTC. 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, the 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 occupation. Similar to variance decomposition, skewness decomposition has a between-occupation component and a within-occupation component. In addition, skewness decomposition has a third component, which captures the correlation between occupation wage level and occupation inequality. 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-occupation component. By contrast, in an SR-RBTC model, the third component is expected to rise. In SR-RBTC, wage gaps decrease in low-paying routine occupations while they increase in high-paying abstract occupations. This increases the correlation between occupation wage level and occupation inequality.

The decomposition results are unequivocal and have two important implications. I find that 78 percent of the overall increase in skewness is driven by the correlation component. The correlation component increases mainly due to a decrease in inequality in routine occupations. This implies that wage polarization is not driven by a decline in average routine wages as previously hypothesized in skill-neutral models. Instead, these results imply that to understand wage polarization, we must understand the inequality trends within occupations. Theoretically, the decline in inequality could be the result of a more homogeneous composition of routine workers (consistent with the findings of

[3]Dicandia (2023) finds that the share of White workers in routine occupations has also decreased.

Cortes, 2016). I rule this out with the IFEM exercise, which shows that inequality trends are driven by a decline in return to skill, as predicted only by the SR-RBTC model.

The second implication is that SR-RBTC is not only consistent with recent inequality trends, it is also substantial enough to account for almost the entire wage trend. Hence, other hypotheses that are not expected to generate inequality trends particularly within occupations (e.g., changes in minimum wage, Piketty, 2014) could only account for a small part of the overall trend.

The results of the skewness decomposition are distinct from previous attempts to decompose wage polarization because skewness decomposition does not rely on the "ignorability assumption." Previous attempts to decompose wage polarization, based on more natural indices of wage polarization such as the 90–50 or 50–10 wage ratios, have found that technological changes and occupational trends in general cannot generate wage polarization (Autor et al., 2005; Firpo et al., 2013). 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 economic data. My results are similar to those of Acemoglu and Restrepo (2022) who 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.

I conclude this paper by discussing alternative explanations for these wage trends and why they are less consistent with my findings.

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 workers 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 the 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 the return to skill is strictly positive at any skill 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 $\theta_i < \theta_0$ chooses to work in the manual occupation, any worker with $\theta_0 < \theta_i < \theta_1$ chooses the routine occupation, and any worker with $\theta_i > \theta_1$ chooses 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 is then modeled as an increase in 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, I 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.[4] 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 potential micro-foundations for RBTC that provide insight into when RBTC would be skill-neutral, -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, in Appendix B.2, an increase in τ represents an improvement in the quality or quantity of automation technology such as computers or robots. 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 as stated in the following theorem.

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*

[4] While the effect on productivity in the routine occupation is skill-neutral, the effect on wages varies by occupation and therefore by skill (Cortes, 2016).

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 following sections 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 the 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

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

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 in its 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-skilled routine workers would leave the routine occupation* ($\frac{\partial \theta_1}{\partial \tau} < 0$).
3. *The wage of the highest-skilled routine worker (θ_1) would decrease relative to all other workers.*

Figure 2a illustrates the results of Theorem 3. Since the return to skill declines, the slope in the routine occupation becomes flatter. This generates smaller gaps between workers who stay in the routine occupation. A relative decline in middle wages occurs because the most significant wage drop is for the highest-earning routine workers, who (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 $-\infty$ (Leontief), 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 workers in the abstract and manual occupations will increase. The overall inequality trend is asymmetric. Below θ_1, wage gaps would (weakly) decrease between any two workers. At the top, the wage gap between abstract workers and high-skilled 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.[5] 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 homogeneous 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 occupation 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

[5] 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).

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.[6]

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 corresponds to trends in the 2000s and onward.

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. In the following sections of the paper, I test these empirical predictions and show that they fit the data.

[6]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.

Appendix B.3 presents a more general model in which workers use different skills in different occupations. This model generates similar predictions, despite being more general, and therefore is also consistent with recent bottom-half inequality trends. However, in the multi-skill model, a full reversal of the return to skill and average skill level between the routine and manual occupations is possible, but not necessary.

2 Empirical 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 a 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.[7] 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).

If 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. They allow for overcoming differences in the composition of workers between occupations and over time. The concern of compositional change is particularly relevant given the substantial decrease in employment in routine occupations, driven prominently by the highest- and lowest-earning workers in those occupations (Cortes, 2016; Böhm et al., 2024).

The IFEM model is a more general version of the fixed-effects model, which accommodates variations in the return to skill, as posited by the SR-RBTC model. The parameters

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

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 minimizes the expected mean squared errors, $\mathbb{E}\left[\varepsilon_{ijt}^2\right]$. 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, taking the first-order conditions with respect to θ_i implies 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

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

Similar to 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 the fixed effects can be absorbed by demeaning the data in a standard fixed-effects model, this approach would not work in an IFEM.

Least-squares estimates of α_{jt} will suffer from a measurement error. Formally, Appendix D.1 shows that even though for the true parameters $\mathbb{E}\left[\theta_i\varepsilon_{ijt}|i \in E_{jt}\right] = 0$, the corresponding empirical moment does not converge to zero ($\mathbb{E}\left[\widehat{\theta}_i\widehat{\varepsilon}_{ijt}|i \in E_{jt}\right] \neq 0$). The measurement error in $\widehat{\theta}_i$ implies that the least squares 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). In particular, let Z be an IV satisfying the following condition for all j, t,

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

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,

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

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

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

where $\widetilde{\log w_{ijt}},\widetilde{\widehat{\theta}_i}$ 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 premium for a year of schooling in occupational category j in year t, divided by the link between years of schooling and estimated skill $\widehat{\theta}_i$ in the same category and year. Therefore, the link between schooling and skill is allowed to vary between occupations and change over time. This is a critical feature of this design, as the model emphasizes that the composition of skills is different between occupations and changes over time.[11] Intuitively, this estimator solves the measurement error bias by aggregating across many

[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}.

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

observations. While $\widehat{\theta}_i$ is noisy, the noise converges asymptotically to zero when averaging over all workers in the same schooling category.[12]

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} between 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 necessarily affect the trend within the routine occupations as it is identified by workers who stay in these occupations.

The model is biased if the IV does not satisfy the condition in Equation (10). This will occur when the IV is correlated with factors that are not captured in θ_i. For example, the condition would be violated if θ_i captures primarily cognitive analytical skills, yet years of schooling is also correlated with social skills that also affect wages.[13]

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 parameters in case of more than one skill, in a simple setting with one occupational category. In this case, θ_i is a weighted average of skills that have a large impact on wages and whose returns experience similar time trends. However, when using an IV, the estimated $\widehat{\alpha}_{jt}$ reflects the weighted average of returns to skills that are correlated with the IV. While this estimator is biased because it does not capture the returns to skills that are not correlated with the IV, it is still informative. In this case, it quantifies the aggregate trend in the return to skills that are correlated with years of schooling.

With multiple occupational categories, the model could be biased if workers use different skills in different occupations. To accommodate this, I also estimate a more general

[12] Measurement error in the IV will not bias the results as long as the errors are not correlated with ε_{ijt}.

[13] 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.

IFEM, with a different fixed effect θ_{ij} for every worker in every occupational category. This allows workers to have different skill levels in different occupational categories, as in the multi-skill model in Appendix B.3. The results are discussed in Appendix G, and are similar to the main results.

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.

Previous papers have suggested alternative ways to estimate an IFEM. Bai (2009) analyzes an IFEM when the number of time periods is large, which is not the case in the PSID data. Other methods derive additional moments by making further assumptions about the error structure. Ahn et al. (2001) assume that ε_{ijt} has a constant variance. Holtz-Eakin et al. (1988) assume that the errors are not serially correlated and propose to estimate the IFEM with lagged outcomes as IVs.

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 for only a small portion of the wage trends at the bottom of the distribution. 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 (Autor et al., 2005). Firpo et al. (2013) documented 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 Recentered influence function (RIF) regression they used also cannot fully quantify the impact of these trends on the overall wage polarization.

The reason why the aforementioned decomposition methods are unable to quantify the impact of SR-RBTC is that SR-RBTC violates the ignorability assumption that underlies them. These decomposition methods assume that the distribution of wages conditional on observables does not change when the distribution of observables changes.[14] In this context, the critical observables are occupations or occupation characteristics. This assumption is satisfied if wages decline uniformly for all workers in routine occupations, as the skill-neutral RBTC model predicts. However, under 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).

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.[15] For a random variable Y skewness is defined as

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

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. 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

[14] Formally, ignorability assumes that the conditional distribution of wages $F_{w|X}(w|X = x)$ does not vary over time. This assumption implies invariance to conditional distributions, where $F_{w|X}(w|X = x)$ does not change when the marginal distribution of X (F_X) changes.

[15] 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.

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 ($\mu_3(Z) = \mathbb{E}\left[(Z - \mathbb{E}[Z])^3\right]$), we get

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

This decomposition was discussed by Mincer (1974) but was not applied to economic data.

The first and second components are equivalent to the variance decomposition formula.[16] The first component $\mathbb{E}\left[\mu_3(Y|X)\right]$ 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, $\mu_3(\mathbb{E}[Y|X])$, 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 category. 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.

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 occupations, 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 oc-

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

cupations 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.

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 in which we expect a large effect on the between component. However, this test alone is only suggestive since, under some specific compositional changes, the covariance component can possibly increase in a skill-neutral RBTC as well. This is why I also estimate the IFEM, which shows that inequality within routine occupations declines regardless of any compositional changes.

The skewness decomposition has several properties that make its results more robust. It allows quantifying wage polarization using a single index. Similar to the variance decomposition, the skewness decomposition breaks the 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, capital ownership, and raw wages (without logs). Any variation in these distributions (e.g., over time) can be analyzed with a skewness decomposition. To simplify and encourage the usage of skewness decomposition by more researchers, I provide an R package that implements it.[17]

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) for 1980–2017. This data was chosen because of 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.

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

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, 2006). 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 higher 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 percent of the positive wages throughout the paper. The level of 5 percent 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.[18]

Most of the skewness decomposition analysis is focused on the years 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 that 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.[19]

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.[20] The analysis by 1-digit occupational 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

[18] 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.

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

[20] Classifying sales as a routine occupation does not strongly affect the results.

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 changed 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 are 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 routine workers has decreased since the late 1980s. Figure 3 plots the estimated coefficient γ_t relative to its highest value in 1987.[21] 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. Appendix Figure A3 plots the same analysis for all three occupational categories, showing that the decline in the education premium is considerably larger in routine occupations. Appendix Figure A4 plots the analysis with additional controls for the interaction of education and age. 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. Figure A5 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.

Figure 4 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 α_{jt} in log units for the three broad occupational categories. Since there is a degree of freedom 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 corresponds to a 50 percent reduction 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.

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

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.[22] Appendix F shows that the estimation results are qualitatively similar when estimating the IFEM with a (biased) naive least-squares estimator.

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 5 plots the coefficient for α_{jt} in log units for each 1-digit occupation category in three years: 1985, 1997, 2011.[23] 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 and 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 and 2011, there is a decline in the return to skill in all occupations, but it is sharper in the routine ones, especially in the administrative category. By 2011, the categories with the lowest return to skill are the four routine occupations. Appendix Figure A6 shows the trends in more detail and plots the value of α_{jt} for each 1-digit occupational category by year.

Overall, these results fit well with the predictions of the skill-replacing RBTC model.

[22] Figure 4 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 at 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.

[23] 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 is reported in Appendix Figure A6.

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 that since the late 1980s, wages of workers who stay in routine occupations have 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 moved to another occupational category. In other models, this prediction would not necessarily hold. In a multidimensional 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 G 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. This is similar to the multi-dimensional skill model described in Appendix B.3. I find that both wages and skills are highly correlated across occupational categories. Moreover, the results of the more general model are very similar to the unidimensional skill results discussed here.

5 The Decline in Average Skill in Routine Occupations

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 again, 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 6 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 percent 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 below that of manual workers.[24]

I find very similar trends using the 1-digit classification of occupations. Figure 7 plots the average skill level in 1985, 1997, and 2011. At the beginning of the period, routine occupations were middle-skilled, and all four routine occupational categories (including sales) had a skill level between -0.1 and 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 A7 plots the results for all years.

The decline in the average skill in routine occupations is primarily driven by a decrease in the flow of middle-skilled workers from outside the labor force. Figure 8 divides newly-hired workers into three equal-sized bins based on their estimated $\widehat{\theta}$, residualized by the cohort of entry to the labor force. For each bin, I plot new hires 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 8 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,

[24]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.

almost half of the middle-skilled workers (based on their estimated $\widehat{\theta_i}$) joined routine occupations. This number decreased to around 33 percent after 2010. At the same time, there has been a substantial increase in the likelihood that middle-skilled workers to join manual occupations. After 2010, middle-skilled workers are almost equally likely to join each of the occupational categories.

Panels A and C of Figure 8 show smaller declines in the share of new routine workers from the top and bottom thirds of the skill distribution. 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 decline of 0.53 percentage points 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 is 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. In an unreported analysis, I find that transitions between occupational categories, or outside the labor force, do not have trends of similar magnitudes.

I also find that while in 1990 many middle-wage workers worked in routine occupations, by 2010 this is no longer the case. Figure 9 plots the average routine intensity index 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, 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 (or stop joining) 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.[25]

This compositional change explains why inequality in 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 in 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 significant enough to account for almost the entire trend of wage polarization. Using a skewness decomposition, I show that wage polarization is driven almost entirely by occupational trends. Moreover, the effect is not 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 10 shows the trend in skewness between 1979 and 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. Appendix Figure A8 presents a bin scatter of the change in wages between 1992 and 2002, for 20 quantiles. The figure shows a U-shaped pattern, as previously shown by Autor et al. (2006; 2008). The observed U-shape qualitatively resembles the EIF derivative plotted in Appendix Figure A2. This suggests that skewness rose in this period because of the rise

[25]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 4).

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—as described in Equation (14), for different choices of categorical groups (X). I first focus on the period 1992–2002 since data on other years uses different occupational coding (see Section 3). As Figure 10 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 and 2002 by 3-digit occupational coding. I report the values of each component and its overall contribution to the rise in skewness. Figure 11 depicts the annual change in each component, as well as in the sum of the three, which equals the total change in skewness.[26] 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 percent of the overall increase. That small share might also be the result of classification errors. Hence, 93 percent of the rise in skewness is related to occupational trends.

Most of the increase in skewness is driven by the covariance component. This is indicated by the blue area in Figure 11, which captures 79 percent 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 2.2, 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 would generate the same effect for all workers in the same occupation. However,

[26] Appendix Figure A9 shows similar results using imputation for missing wages.

this component generates only 15 percent of the overall rise in skewness. Instead, the substantial rise in the covariance component suggests that the effect is mainly driven by the asymmetric trends within occupations. This is more 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 A10 and A11, I show the same decomposition results by industry, as well as 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 A12 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 and 1992 and between 2002 and 2003. With that caveat in mind, we still see a similar pattern in the earlier period between 1988 and 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 A13 replicates this finding 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

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

where in this case Y denotes log wages and X denotes 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 initial values in 1992. 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:

$$\widetilde{COV}\left(\mathbb{E}\left[Y_{92}|X_{92}\right], V\left[Y_t|X_t\right]\right) = \sum_{x} Pr\left(X_{92} = x\right) \mathbb{E}\left[Y_{92}|X_{92} = x\right] V\left(Y_t|X_t = x\right), \quad (16)$$

where $\Pr\left(X_{92} = x\right)$ and $E\left[Y_{92}|X_{92} = x\right]$ are the share of workers and mean log wages in the initial year, 1992.

I find that the asymmetric trends in within-occupation inequality can explain the entire increase in covariance. Figure 12 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. That is, inequality is increasing in high-paying occupations and decreasing in low-paying occupations.[27]

The rise in covariance is driven by both parts of the wage distribution. Appendix Figure A14 examines the contribution of only low- and only high-paying occupations. The red line shows the counterfactual covariance increase when the variance in high-paying occupations is held constant at its level in 1992 as well (together with the share of

[27] Cashiers contribute the most to the rise in the counterfactual covariance. This is a large occupation (1.9% of the workforce) with very low average wages and a 0.08 decline in the variance of log wages.

workers and the expected log wages). The green line shows the results holding variance in low-paying occupations fixed. The increase in the counterfactual covariance is similar in both scenarios, suggesting that both parts of the distribution contribute about equally to the rise in skewness and wage polarization.

The drop in inequality in low-paying occupations is driven mostly by routine occupations. Appendix Figure A15 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 and 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 Firpo et al. (2013), who show that routine occupations tend to have a greater 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 empirically are 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 other occupations, 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 in the 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 in the average routine wages, but in 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 a 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 occupation-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 G suggest that the correlation between skills is likely high, at least for workers who switch occupational categories (the majority of workers). This correlation also supports the interpretation of a decrease 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 than the United States (Blau and Kahn, 2002). Such institutions tend to have a larger impact on bottom-half inequality than on 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 recent technological advances are able to automate even non-routine tasks. This would extend the relevance of the skill-replacing model to additional occupations that could also experience a decline in inequality. Recent research argues that generative AI yields a larger productivity boost for lower-skilled workers (Noy and Zhang, 2023; Althoff and Reichardt, 2025). Beaudry et al. (2016) argue that after a technology is adopted, the demand for high-skill abstract workers declines. Extending the skill-biased technological change framework to non-routine occupations could also potentially explain the decline in employment of recent college graduates (Coskun, 2024; Hosseini and Lichtinger, 2025), or the recent decline in inequality in several developed countries (Autor et al., 2023; Dustmann et al., 2024). Whether or not new technology is skill-replacing is a critical question that could determine the future of inequality in the labor market.

Data Availability Statement

The data and code required to replicate all results in this paper are publicly available at https://doi.org/10.5281/zenodo.20021422. The replication package contains all programs and instructions needed to reproduce the tables and figures in the paper.

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*. Working Paper 32190. National Bureau of Economic Research.

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.

Althoff, L. and H. Reichardt (2025). "Task-Specific Technical Change and Comparative Advantage". In.

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., A. Dube, and A. McGrew (2023). *The unexpected compression: Competition at work in the low wage labor market*. Tech. rep. National Bureau of Economic Research.

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.

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

Bai, J. (2009). “Panel data models with interactive fixed effects”. In: *Econometrica* 77.4, pp. 1229–1279.

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 (2024). “Occupation, Growth, Skill Prices, and Wage Inequality”. In: *Journal of Labor Economics* 42.1, pp. 201–243.

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.

Coskun, S. (2024). *Young, Educated, Unemployed*. Discussion Paper DP19351. CEPR Discussion Paper No. 19351. C.E.P.R.

Dicandia, V. (2023). *Technological Change and Racial Disparities*. Working Paper.

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.

Dustmann, C., C. Gergs, and U. Schönberg (2024). *The Evolution of the German Wage Distribution Before and After the Great Recession*. Tech. rep. mimeo, University College London.

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. 262–294.

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.

Hosseini, S. M. and G. Lichtinger (2025). “Generative AI as Seniority-Biased Technological Change: Evidence from US Resume and Job Posting Data”. In.

Hunt, J. and R. Nunn (2022). “Has US Employment Really Polarized? A Critical Reappraisal”. In: *Labour Economics* 75, p. 102117.

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

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

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*. Discussion Paper. IZA.

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	0.0946	0.0334	-.04171	0.1030
$N = 126,313$		(.0037)	(.0032)	(.0053)
2002	0.1898	0.0397	-.0276	0.1777
$N = 105,620$		(.0045)	(.0036)	(.0062)
Δ 1992-2002	0.0951	0.0063	0.0141	0.0747
		(.0058)	(.0048)	(.0081)
	100%	6.6%	14.9%	78.5%

Table 1: Skewness Decomposition by 3-Digit Occupation

Note: 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

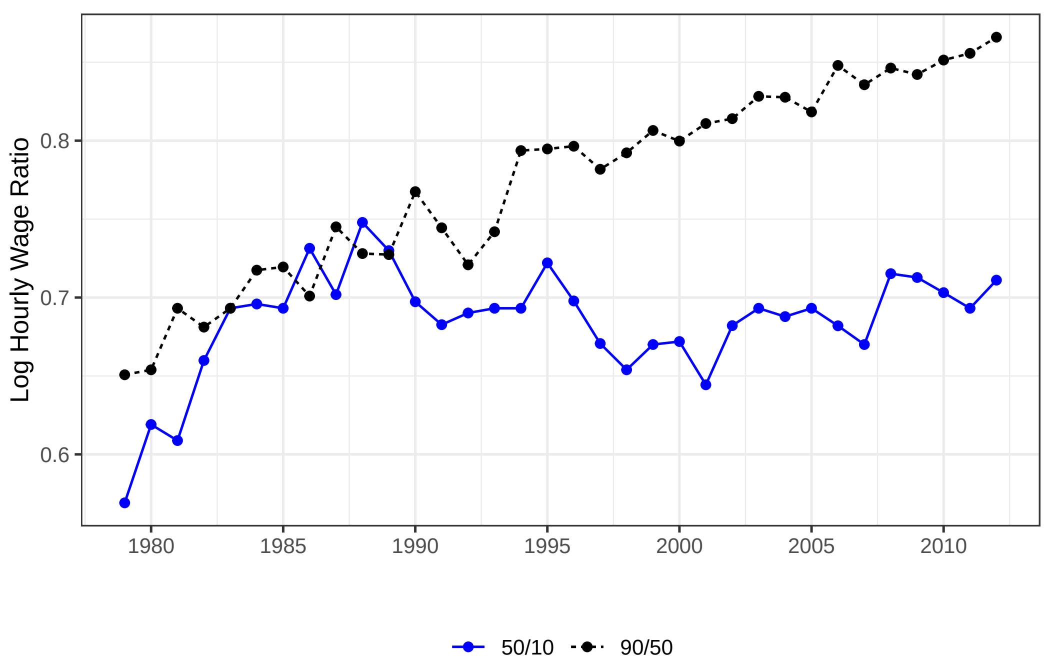

Figure 1: 90/50 and 50/10 Log Hourly Wage Ratio

Note: 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 ($N = 4,401,711$)

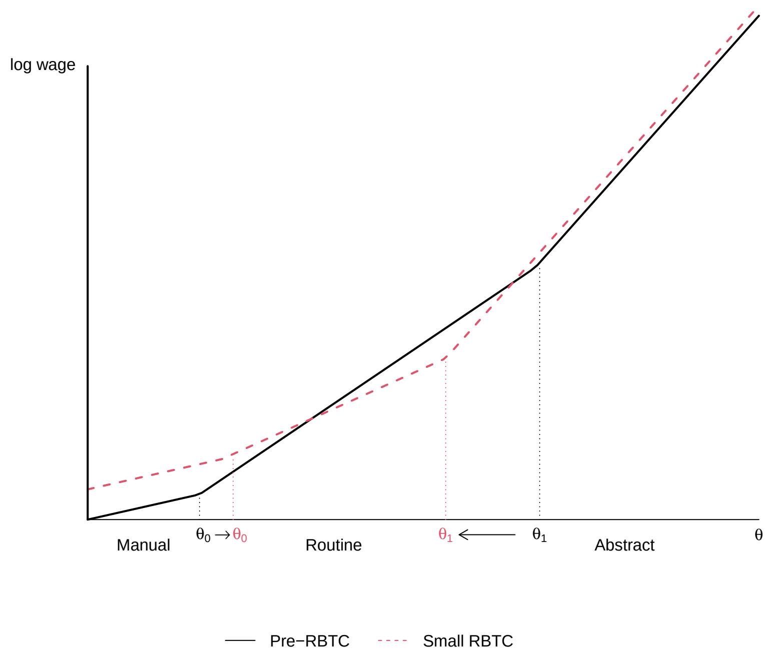

(a) Small SR-RBTC

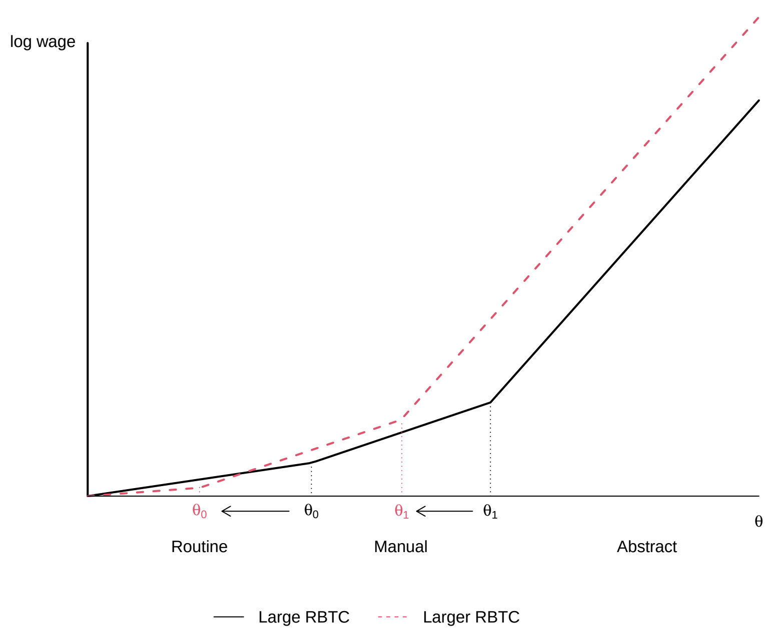

(b) Large SR-RBTC

Figure 2: Illustrated Changes in Log Wages by Skill

Note: 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 (an 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)).

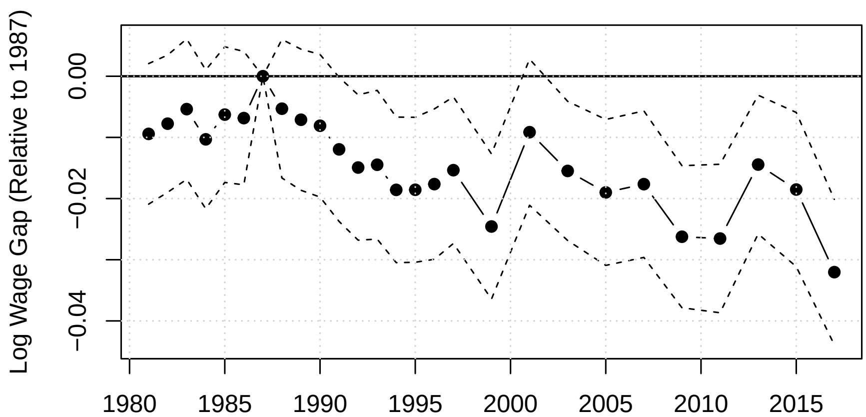

Figure 3: Education Premium in Routine Occupations

Note: 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.

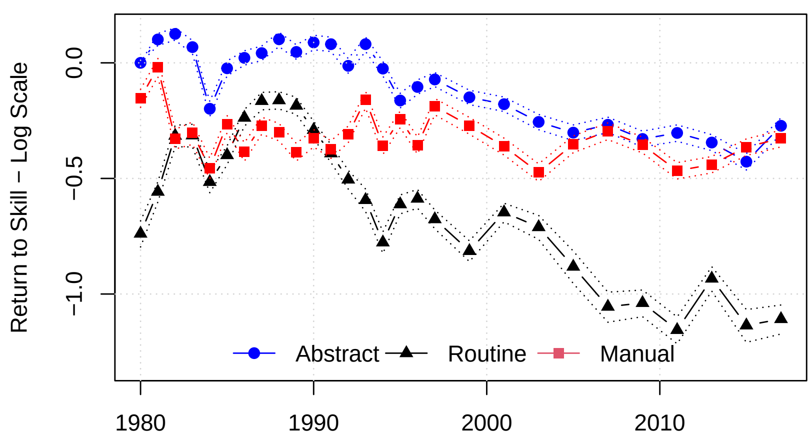

Figure 4: Return to Skill (a_{jt}) by Occupational Category

Note: 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. The Autor and Dorn (2013) occupational crosswalk is used for a consistent definition of occupations over time. Dashed lines represent 95% confidence intervals.
Source: PSID ($N = 122,162$)

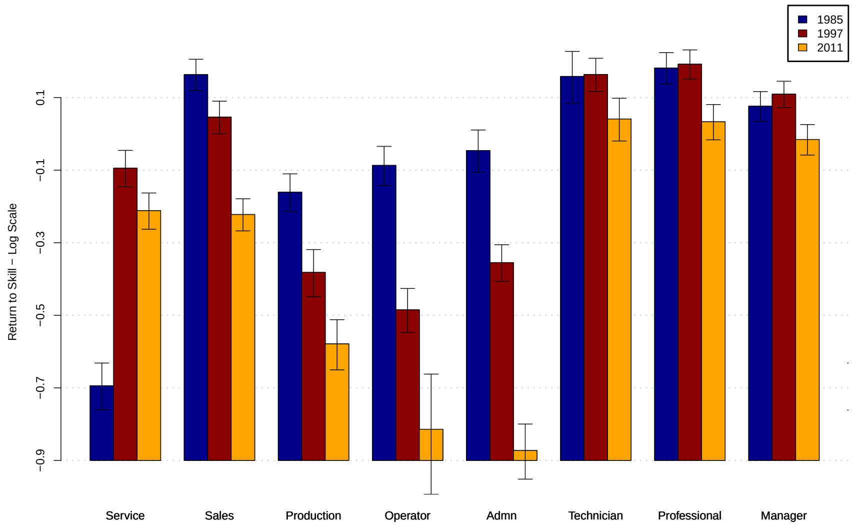

Figure 5: Return to Skill (α_{jt}) by 1-Digit Occupational Category

Note: 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 A6. Return to skill is 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. ($N = 105,248$)

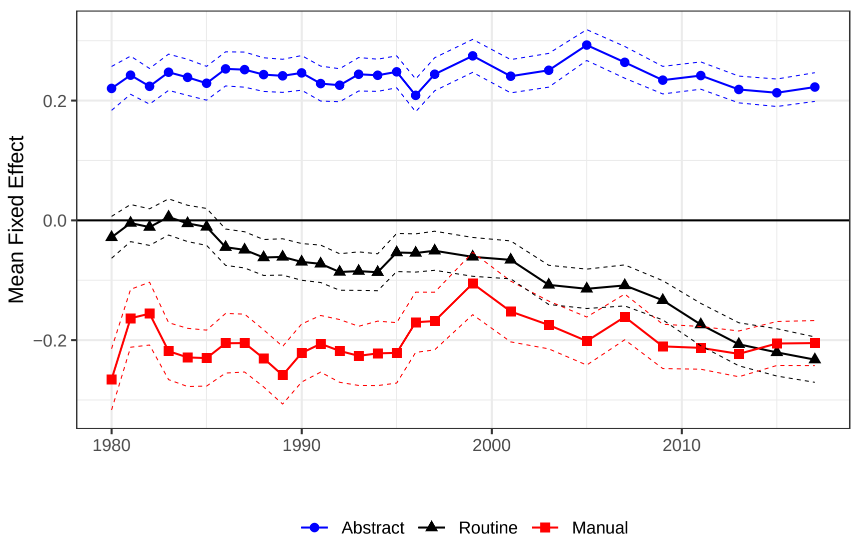

Figure 6: Mean Skill Level ($\widehat{\theta}_i$) by Occupational Category

Note: 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. The Autor and Dorn (2013) occupational crosswalk is used for a consistent definition of occupations over time.
Source: PSID ($N = 124,407$)

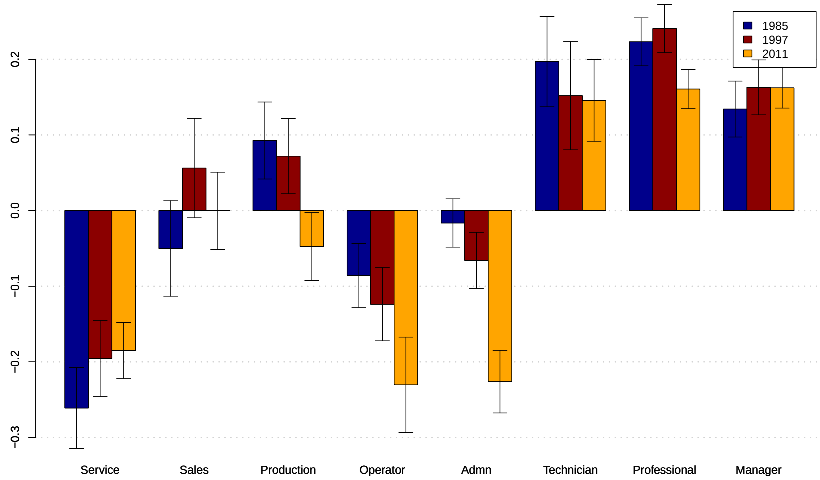

Figure 7: Mean Skill Level ($\widehat{\theta_i}$) by 1-Digit Occupational Category

Note: 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. ($N = 108,413$)

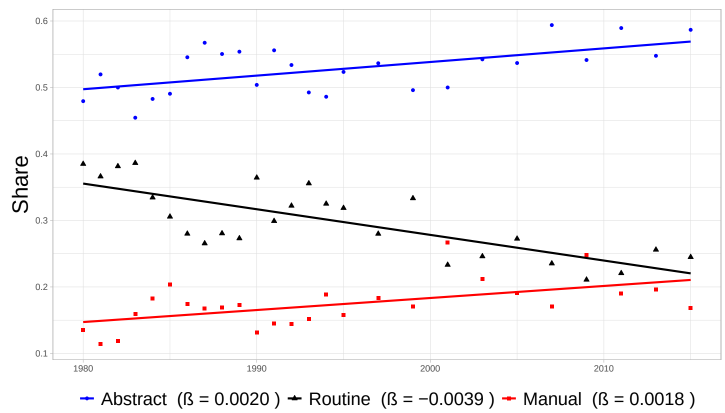

(a) Top Third

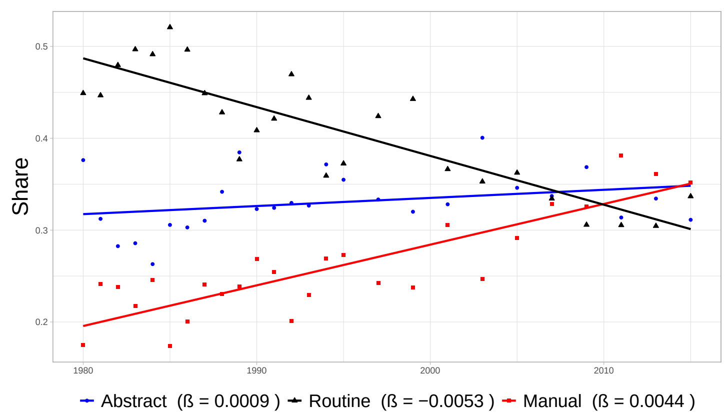

(b) Middle Third

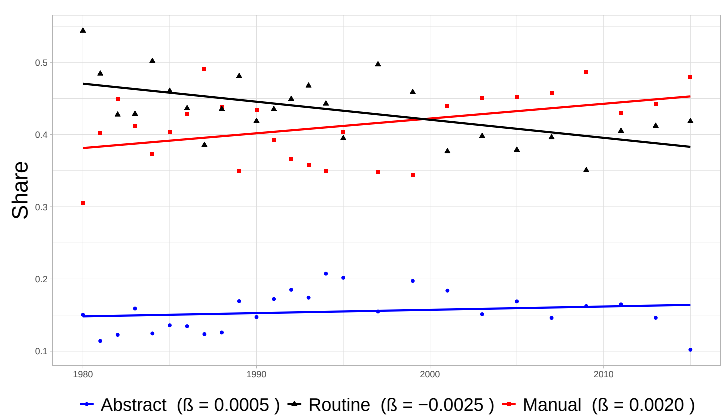

(c) Bottom Third

Figure 8: Share of Workers Who Join Each Occ. Category From Outside the LF

Note: This figure plots the share of transitions from non-employment into each of the three occupational categories by year, separately for workers in 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-fitting line for each category from a linear regression of the shares in this category on the year.
Source: PSID

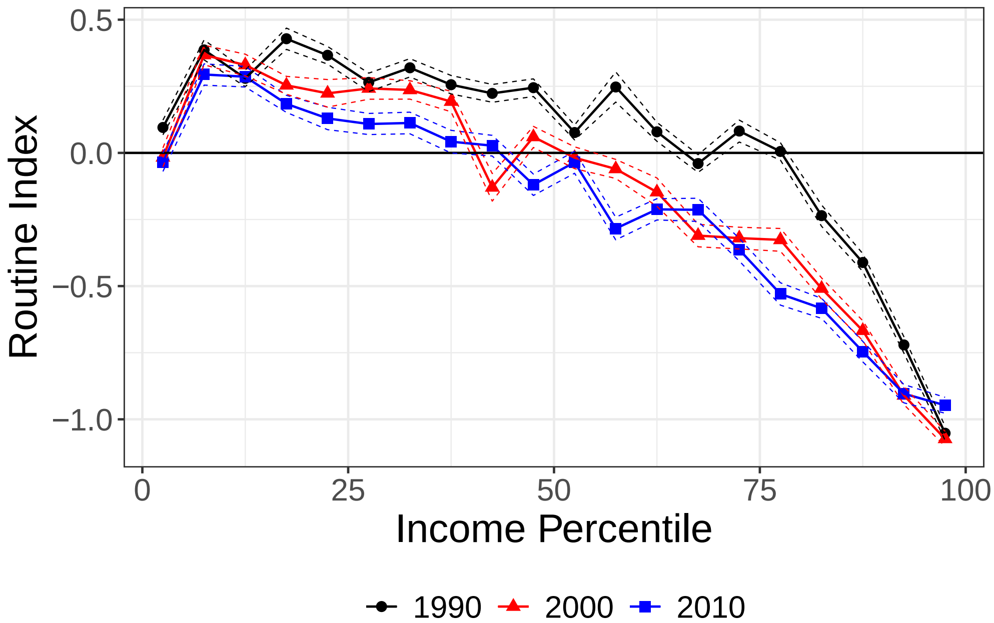

Figure 9: Routine Intensity of Occupation by Wage Percentile

Note: 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 ($N_{1990} = 147,851; N_{2000} = 105,461; N_{2010} = 101,915$).

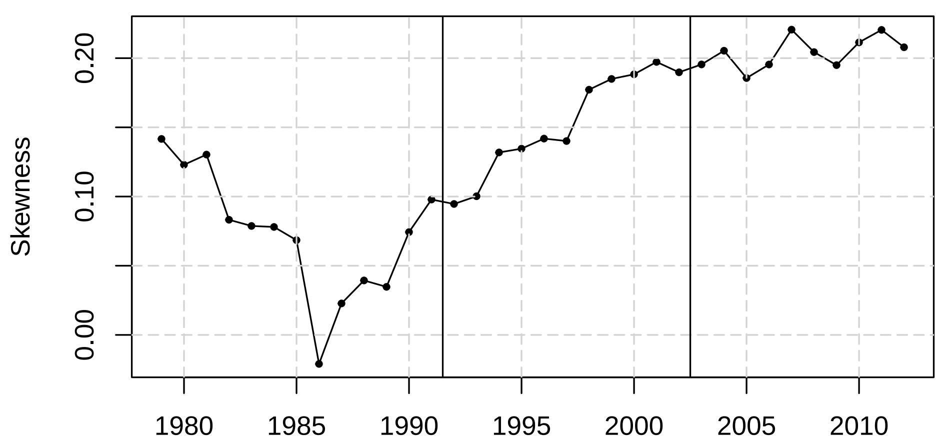

Figure 10: Skewness of Log Hourly Wage

Note: 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 ($N = 4,401,711$).

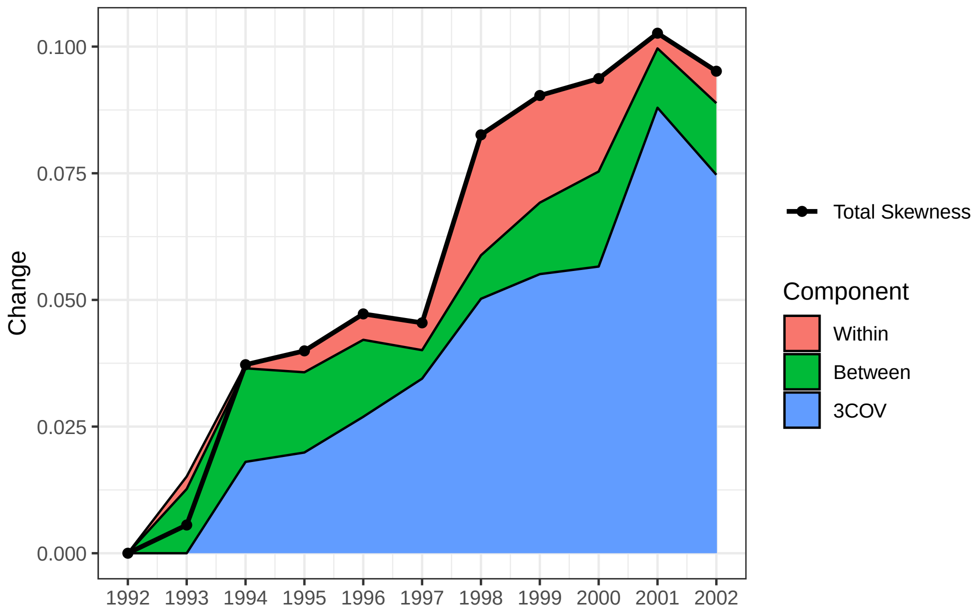

Figure 11: Skewness Decomposition by 3-Digit Occupation

Note: 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 ($N = 1,208,151$).

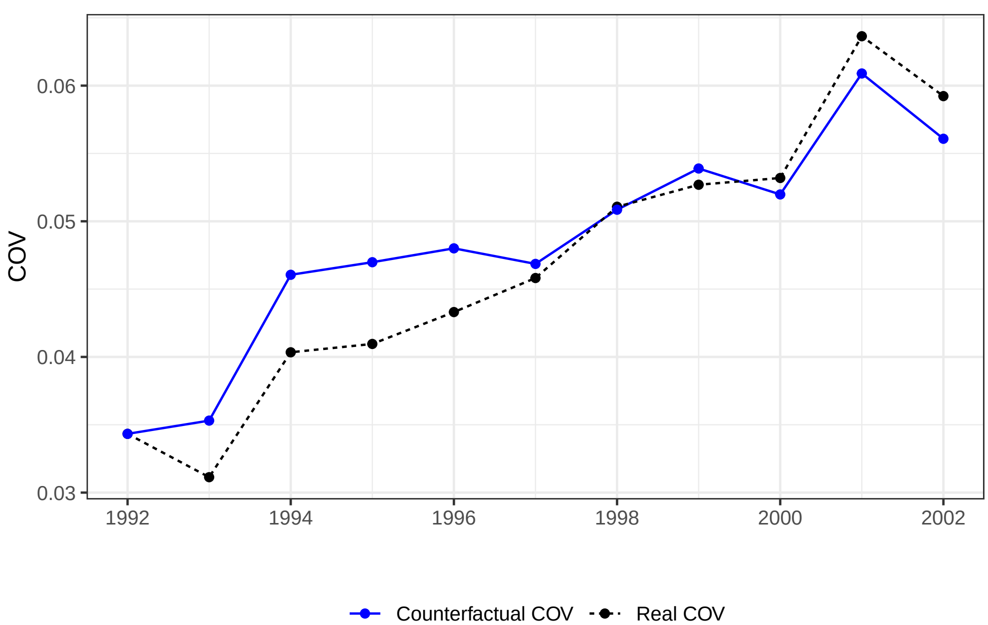

Figure 12: Covariance of Expectation and Variance of Log Wages by Occupation

Note: This figure plots the covariance of mean log wage and variance of log wage by occupation, $COV\left(\mathbb{E}\left[\log w|occ\right], V\left(\log w|occ\right)\right)$ (black line). The counterfactual covariance (in blue) is calculated by fixing $\mathbb{E}\left[\log w|occ\right]$ and the share of workers in each occupation at their value in 1992, 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 ($N = 1,208,151$).