Why Did Employment Polarize More Than Wages?

Davide Alonzo* Giovanni Gallipoli†

March 25, 2026

Preliminary and Incomplete

Abstract

The hollowing out of routine employment has been a defining feature of the U.S. labor market since 1980, yet the reallocation of workers across occupations has been far more pronounced than changes in the wage structure. We examine this disconnect in an equilibrium framework where jobs differ in both wages and unobservable latent values. Our estimates suggest that labor supply is significantly more elastic to latent values than to wages. Holding latent values fixed at their 1980 levels eliminates much of the observed employment polarization, while wage polarization remains largely intact. The model accounts for the resilience of routine wages without assuming downward rigidity. In equilibrium, the supply-driven outflow of workers raises the marginal productivity of those who remain, providing an endogenous floor that cushions wage declines. These results suggest that employment polarization and wage stability are joint outcomes of a single process driven by shifting non-wage job attributes.

*Department of Economics, Université de Montréal, 3150 Rue Jean-Brillant, Montréal, QC H3T 1N8, Canada. E-mail: davidealonzo03@gmail.com.

†Vancouver School of Economics, University of British Columbia, 6000 Iona Drive, Vancouver, BC V6T 1L4, Canada. E-mail: gallipol@mail.ubc.ca.

1 Introduction

The hollowing out of routine employment since 1980 is well documented (Autor et al., 2006; Acemoglu and Autor, 2011; Autor and Dorn, 2013; Cortes et al., 2017; Acemoglu and Restrepo, 2022). There is evidence that wages polarized too, but the shift in the occupational wage structure was considerably less striking.

The empirical evidence for this asymmetry is well-established, even if rarely foregrounded. Autor and Dorn (2013) present smoothed plots of changes in employment and hourly wages by skill percentile between 1980 and 2005 (their Figure 1): the U-shape in employment shares is considerably more pronounced than the corresponding U-shape in wages. Acemoglu and Autor (2011) document the same pattern across task categories and note that employment shares in routine occupations declined by 8–10 percentage points while wage differentials between routine and non-routine occupations moved by substantially less. More recently, Acemoglu and Restrepo (2022) show that between 50 and 70 percent of changes in the U.S. wage structure can be attributed to the displacement effects of automation, yet the employment reallocation driven by this displacement exceeds what the estimated wage effects would predict in standard models. Bárány and Siegel (2018) extend the evidence to 1950, showing that employment polarization is considerably more pronounced than wage polarization across all thirty-year windows since mid-century, with employment swings substantially larger than the corresponding wage changes in each decade they study. This evidence suggests that the bodies-vs-wages disconnect is a persistent feature of the U.S. labor market rather than the product of a specific episode.[1]

We provide direct evidence on the magnitude of this gap over our sample period (1980–2018). Figure 1 in Section 2 replicates the polarization exercise of Autor and Dorn (2013) by gender, using full-time workers aged 25 to 35 and extending the horizon to 2018. The focus on younger workers is intentional: workers who have spent decades in an occupation are less likely to respond to current incentives, so the employment patterns of younger cohorts offer a cleaner picture of active occupational sorting. The figure documents substantial employment reallocation across the skill distribution for both men and women. Wages tell a different story: they are essentially flat for most occupations, with a sharp increase confined to a narrow band at the very top of the skill distribution, more pronounced for women than for men. The employment reallocation is broad; the wage response is not. Figure 2 quantifies

[1]Hunt and Nunn (2022) argue that pre-2000 polarization reflects breaks in occupation coding, and that reallocation was concentrated in the 2000-2018 period; more details on post-2000 employment declines are in Charles et al. (2019). Our analysis indicates the employment-wage disconnect was more pronounced post-2000. Hunt and Nunn (2022) also argue that inequality grew primarily within rather than between occupations; we allow for within-occupation wage and employment variation across gender and education groups.

this asymmetry. We estimate a discrete-choice regression of changes in log employment shares on changes in log wages for every possible time interval in the data, separately by gender and education group. Wage changes account for a modest share of employment reallocation across all combinations. As Figure 2 shows, wages account for a modest share of employment variation across all groups and horizons, with the highest explanatory power among college-educated women. The evidence is consistent with the view that non-wage forces are a primary driver of occupational reallocation, and motivates the structural investigation that follows.

Can one resolve the employment-wage disconnect by looking beyond pecuniary returns? Recent work has shown that non-wage amenities and latent match values dictate a substantial share of employment choices (Mas and Pallais, 2017; Sorkin, 2018; Lamadon et al., 2022; Taber and Vejlin, 2020; Maestas et al., 2023; Bagga et al., 2023). To explore their role, we estimate an equilibrium model of occupational choice (building on Alonzo and Gallipoli, 2023)[2] and find that unobservable job values exhibit greater heterogeneity than observable wages and that labor supply is substantially more elastic to them. Our quantitative analysis suggests that the occupational polarization of the last four decades is largely a polarization of 'bodies' driven by worker responses to non-wage job attributes, rather than a straightforward wage-clearing reallocation. The estimated latent values are consistent with several interpretations: changes in the task content of occupations as technology evolved (Autor et al., 2003), shifts in search friction intensity (Martellini and Menzio, 2021), declining barriers to occupational entry (Hurst et al., 2024), or changes in the amenity content of jobs (Mas and Pallais, 2017; Maestas et al., 2023). The equilibrium also delivers a prediction on wage dynamics: muted declines in routine wages are a consequence of the employment outflow itself, which raises the marginal product of workers who remain. No assumption of downward wage rigidity is required.

The structural approach addresses a natural concern: that attributing reallocation to latent values amounts to naming the residual. Two features of the equilibrium discipline this. First, wages are determined by firm optimization on the production side; latent values enter only on the supply side, so any admissible estimate must simultaneously satisfy both. Fixing latent values at 1980 levels and re-solving the model yields joint predictions for employment and wages moving in opposite directions, with employment polarization largely disappearing while wage polarization survives. A residual extracted from the employment equation alone would be silent on wages by construction. Second, the equilibrium constrains

[2]The framework accommodates search frictions implicitly. More persistent frictions are embedded in the latent match values b_{ijt}: following McFadden (1974), the lower present value of jobs with higher access costs is reflected in workers' occupational choices. Short-run frictions are less binding at the decadal horizons we study, given the high rate of occupational mobility documented by Kambourov and Manovskii (2008). See Section 3.

latent values across demographic groups simultaneously: fitting women's reallocation while misfitting men's is inadmissible. The finding that women's supply is more responsive to non-wage attributes is what the equilibrium requires, not an assumption. This aligns with independent evidence from the amenity literature (Mas and Pallais, 2017; Maestas et al., 2023).

The paper proceeds as follows. Section 2 presents motivating reduced-form evidence on the employment-wage disconnect. Section 3 describes the proposed model modifications. Section 4 describes identification and estimation. Section 5 presents baseline results on the role of technology versus latent values. Section 6 formalizes why muted wage polarization is an equilibrium prediction rather than a puzzle. Section 7 examines gender heterogeneity. Section 8 concludes.

2 Motivating Facts

Figure 1 replicates the exercise of Autor and Dorn (2013): smoothed cumulative changes in employment shares and real hourly wages across the skill distribution, with occupations ranked by their 1980 mean log wage. Both figures use a locally weighted polynomial smoother applied to cumulative changes; the exercises are directly comparable. The original covers 1980–2005 and pools across gender. We extend the horizon to 2018 and disaggregate by gender, using full-time workers aged 25–35. The focus on younger workers is deliberate: workers who have spent decades in an occupation are less likely to respond to current incentives, so employment patterns among younger cohorts offer a cleaner picture of active occupational sorting.

The employment panel broadly replicates the U-shape of Autor and Dorn (2013), but the gender split reveals asymmetries the pooled figure obscures. Men show a sharp peak at the very bottom of the skill distribution that falls steeply before recovering; women show less growth at the bottom and a more monotone increase from the middle upward. The muted bottom-end growth relative to Autor and Dorn (2013) may partly reflect the sample restriction. Low-skill service occupations draw disproportionately from workers outside the 25–35 window: teenagers and workers in their early twenties account for a large share of employment in food service, cleaning, and personal care, while older displaced workers represent another substantial group. The full-time restriction compounds this, since part-time arrangements are especially prevalent at the bottom of the skill distribution. Restricting to full-time workers aged 25–35 could therefore thin out precisely the group where A&D document the most pronounced employment growth.

The wage panel departs more clearly from the original. Autor and Dorn (2013) find

wages roughly U-shaped or gently rising across the skill distribution. In Figure 1, wages are essentially flat from the 5th to the 85th percentile for both groups, with a sharp spike confined to a narrow band at the very top. That spike is considerably more pronounced for women than for men, a feature invisible in the pooled figure. Extending the horizon to 2018 amplifies the contrast: post-2005 top-end wage growth, particularly for women, accounts for much of the divergence from the original.

The central observation is the asymmetry between the two panels. Employment reallocation is broad, spanning most of the skill distribution for both men and women. The wage response is narrow, concentrated at the top and negligible elsewhere. The near-universal flatness of wages below the top percentiles is not simply a data regularity: Section 6 shows it is an equilibrium prediction. In occupations that shed workers, the exit of marginal employees raises the marginal product of those who remain, supporting wages. The model implies this force is strongest in the middle of the skill distribution, where routine employment declined most sharply; only at the top, where technology gains are large enough to dominate, does it fail to absorb the shock.

2.1 A Discrete Choice Approximation

To quantify the disconnect, we estimate how well wage changes can account for employment shifts across occupations and demographic groups. We consider a discrete choice model (McFadden, 1974; Berry, 1994) in which workers of demographic group i choose among N occupations indexed by j:

$$\max_{j=0,\ldots,N} u(w_j) + \epsilon_j$$

where ϵ_j is Type I extreme value and $j = 0$ is the outside option (non-employment). At equilibrium, the share of group i choosing occupation j is

$$s_j = \frac{\exp\left(\frac{u(w_j)}{\sigma}\right)}{\sum_{j'} \exp\left(\frac{u(w_{j'})}{\sigma}\right)}$$

Dividing by s_0 (occupation zero, the baseline) and taking logs:

$$\log\left(\hat{s}_j\right) = \frac{u(w_j) - u(w_0)}{\sigma}$$

where $\hat{s}_j = s_j/s_0$. Comparing two cohorts observed at times t_0 and t_1 and assuming the utility function is stable over time:

$$\Delta \log(\hat{s}_{jt_0}) = \frac{u(w_{jt_1}) - u(w_{0t_1}) - (u(w_{jt_0}) - u(w_{0t_0}))}{\sigma} \tag{1}$$

Adopting a flexible polynomial representation $u(w) = \sum_{n=0}^{N} \check{\alpha}_n w^n$, equation (1) becomes

$$\begin{aligned}
\Delta \log(\hat{s}_{jt_0}) &= \sum_{n=0}^{N} \alpha_n w_{jt_1}^n - \sum_{n=0}^{N} \alpha_n w_{0t_1}^n - \left(\sum_{n=0}^{N} \alpha_n w_{jt_0}^n - \sum_{n=0}^{N} \alpha_n w_{0t_0}^n \right) \\
&= \sum_{n=0}^{N} \alpha_n \tilde{w}_{jt_1}^n - \sum_{n=0}^{N} \alpha_n \tilde{w}_{jt_0}^n \\
&= \sum_{n=0}^{N} \alpha_n \Delta \tilde{w}_{jt_0}^n
\end{aligned}$$

where $\alpha_n = \frac{\check{\alpha}_n}{\sigma}$ and $\tilde{w}_{jt}^n = w_{jt}^n - w_{0t}^n$. We estimate this equation separately for each demographic group i. This is a conservative choice as it gives wage changes the best opportunity to explain employment reallocation within each group. The estimating equation is:

$$\Delta \log(\hat{s}_{jt_0}) = \sum_{n=0}^{N} \alpha_n \Delta \tilde{w}_{jt_0}^n + \epsilon_{jt} \tag{2}$$

Using this relationship, we compute the share of explained variation in equation (2) for different demographic groups and time horizons, using a fourth-degree polynomial ($N = 4$). The sample is restricted to full-time workers aged 25-35, so that each cross-section captures a distinct cohort actively making occupational choices; younger workers are better able to respond to wage signals, particularly over longer horizons. Figure 2 presents the results. A low share of explained variation is direct evidence that changes in relative wages cannot account for observed shifts in employment shares (the "polarization of bodies" puzzle).

3 Model

The model builds on Alonzo and Gallipoli (2023). We describe its main elements below.

Workers. A continuum of workers of measure S populates the labor market. Each worker ι belongs to one of I demographic groups indexed by i, with μ_i denoting the measure of workers in group i. Workers are endowed with unobserved innate ability a, distributed lognormally in the population with mean one and standard deviation σ_a, so $\ln a \sim \mathcal{N}(0, \sigma_a^2)$.

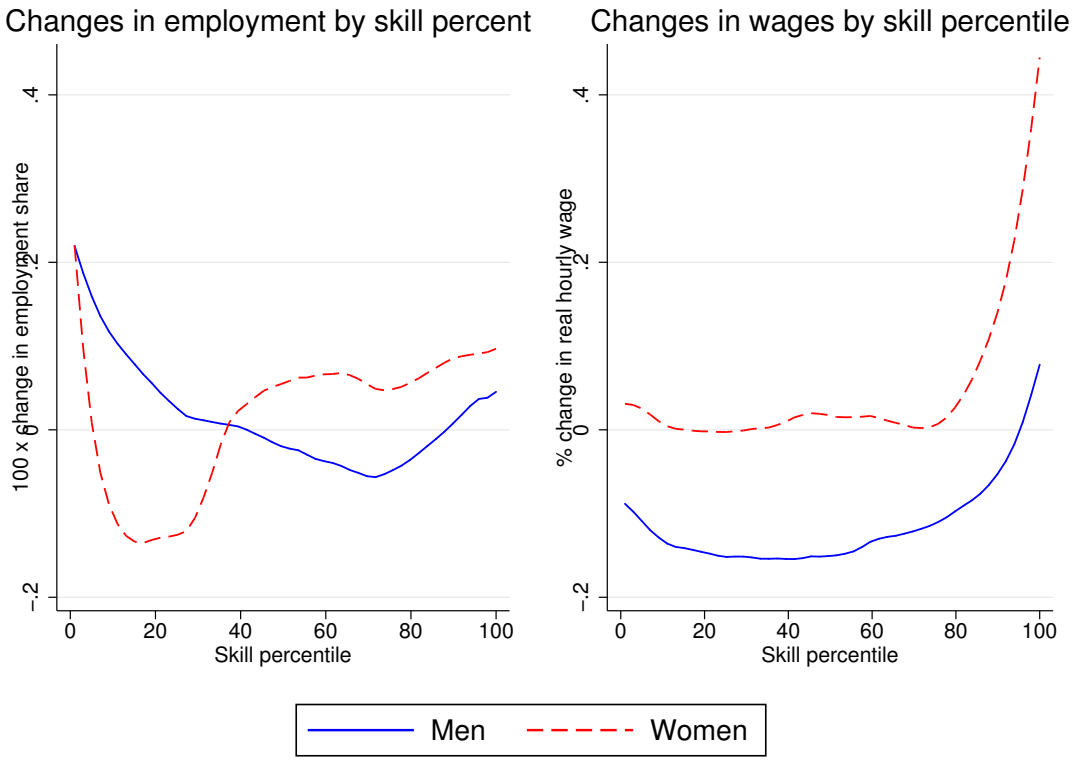


Figure 1: Smoothed Changes in Employment and Wages by Skill Percentile and Gender, 1980–2018. Panel A plots the change in employment share for men and women across 331 detailed occupations. Occupations are ranked on the x-axis by their 1980 skill level, which is approximated by the mean log hourly wage of workers in each occupation in 1980. Panel B plots the corresponding percentage changes in real hourly wages by 1980 occupational skill percentile. Both panels apply locally weighted polynomial regressions to smooth cumulative changes over the sample period, 1980–2018.

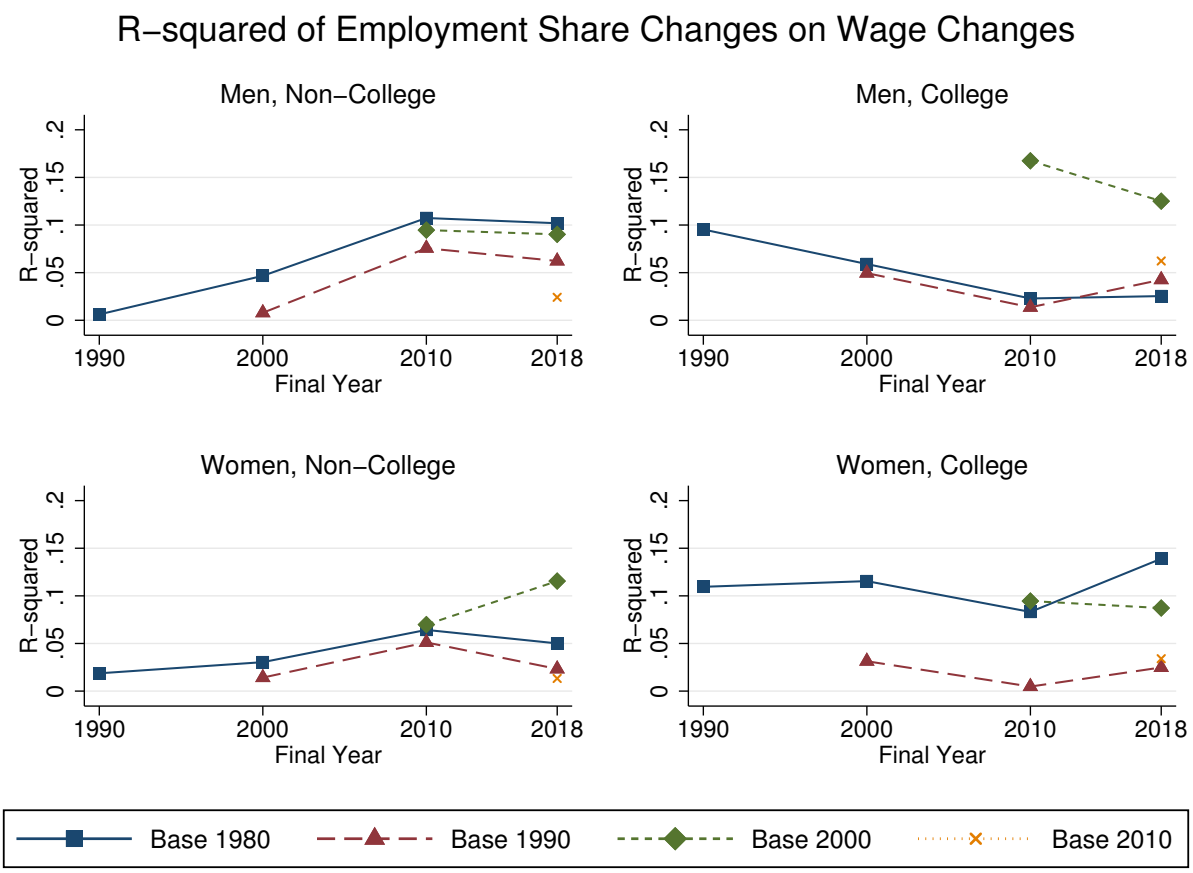


Figure 2: Explanatory Power of Wage Changes on Employment Share Changes Across Time Horizons. This figure plots the share of explained variation from regressions of $\Delta \log(\hat{s}_{jt_0})$ on a fourth-degree polynomial in $\Delta\tilde{w}^n_{jt_0}$, where $\hat{s}_{jt}$ is the employment share of occupation j relative to occupation zero and $\tilde{w}^n_{jt} = w^n_{jt} - w^n_{0t}$ denotes wage terms demeaned by the baseline. The regressions are estimated separately for four demographic groups: Non-College Men, College Men, Non-College Women, and College Women. Within each panel, the x-axis denotes the final year of the time difference (1990, 2000, 2010, or 2018). The distinct connected lines represent different base years (1980, 1990, 2000, and 2010). The persistently low explained variation across most horizons suggests that wage changes alone account for a minimal fraction of the observed shifts in employment shares.

Workers choose whether to work and, if so, which occupation $j \in \{1, \ldots, J\}$; the outside option $j = 0$ corresponds to non-employment. When working, workers supply labor inelastically. Conditional on occupation j at time t, a worker of type i with ability a earns income $y_{ijt}(a) = w_{ijt} \cdot a$, where w_{ijt} is the per-efficiency-unit wage for group i in occupation j. Utility is:

$$U_{ijt}(w_{ijt}, a) = u(w_{ijt}\, a) + b_{ijt} \tag{3}$$

where $u(c) = (c^{1-\gamma} - 1)/(1 - \gamma)$ is CRRA consumption utility and b_{ijt} denotes the latent match value for a type i worker in occupation j at time t. The latent value absorbs all non-wage job attributes, including hours arrangements, working conditions, and access costs; it is measured relative to non-employment, with $b_{i0t} = 0$ by normalization. Occupation choice combines systematic utility with an idiosyncratic preference shock θ^{ι}_{jt}, drawn i.i.d. from a Type I Extreme Value distribution with scale σ_θ:

$$\max_{j=0,1,\ldots,J} U_{ijt}(w_{ijt}, a) + \theta^{\iota}_{jt} \tag{4}$$

By the properties of the Extreme Value distribution, the fraction of workers in group i with ability a employed in occupation j at time t is:

$$s_{ijt}(a) = \frac{\exp(U_{ijt}(w_{ijt}, a)/\sigma_\theta)}{\sum_{j'=0}^{J} \exp(U_{ij't}(w_{ij't}, a)/\sigma_\theta)} \tag{5}$$

The total measure of group i workers in occupation j is $\mu_{ijt} = \mu_i \int s_{ijt}(a)\, dF(a)$.

Firms and production. On the production side, a representative final good producer aggregates a continuum of intermediate goods, each produced by a separate firm employing a single occupation. Occupations are grouped into four task clusters defined by two binary dimensions: whether tasks are routine or non-routine, and whether they are cognitive or manual (Autor et al., 2003; Autor and Dorn, 2013). The four resulting clusters are non-routine cognitive (NRC), routine cognitive (RC), routine manual (RM), and non-routine manual (NRM). Within each occupation j, the effective labor input aggregates across demographic groups and the distribution of individual abilities:

$$\tilde{L}_{jt} = \sum_i \beta_{ijt} \int a\, s_{ijt}(a)\, \mu_i\, dF(a) \tag{6}$$

where β_{ijt} is the match-specific productivity of group i in occupation j. The cluster composite T_{kt} aggregates effective labor across the occupations in cluster k through an inner CES

(Krusell et al., 2000):

$$T_{kt} = \left[\sum_{j \in k} \alpha_{jt} \, \tilde{L}_{jt}^{\rho_k} \right]^{1/\rho_k} \tag{7}$$

Final output combines the four cluster composites through an outer CES:

$$Y_t = A_t \left[\sum_{k} \nu_k \, T_{kt}^{\eta} \right]^{1/\eta} \tag{8}$$

The parameter ρ_k governs substitutability among occupations within cluster k; η governs substitutability across clusters. Firm optimization implies the equilibrium wage for match (i, j) with $j \in k$:

$$w_{ijt} = A_t^{\eta} \, Y_t^{1-\eta} \, \nu_k \, \alpha_{jt} \, \beta_{ijt} \, \tilde{L}_{jt}^{\rho_k - 1} \, T_{kt}^{\eta - \rho_k} \tag{9}$$

When $\rho_k = \eta$ for all k, the nested structure collapses to a flat CES. The grouping of occupations into task clusters embeds the assumption that productive substitutability is higher within clusters than across them, which translates into the restriction $\rho_k > \eta$ for all k, or equivalently that the within-cluster elasticity of substitution $1/(1-\rho_k)$ exceeds the across-cluster elasticity $1/(1-\eta)$. This restriction is testable; if the estimates violate it for some cluster k, the task-based grouping of occupations within that cluster is at odds with their behavior in production. A further restriction embedded in the nested CES is that η is common to all pairs of clusters: routine cognitive and non-routine cognitive workers are assumed to be equally substitutable across the routine/non-routine divide as routine manual and non-routine manual workers are across the manual/cognitive divide. This is a maintained simplification; with four clusters there are in principle six pairwise elasticities, but limited cross-cluster wage variation makes them difficult to identify separately. **Equilibrium.** A competitive equilibrium is a set of per-efficiency-unit wages $\{w_{ijt}\}$ and occupational allocations $\{\mu_{ijt}(a)\}$ such that: (i) workers solve problem (4); (ii) intermediate and final good producers optimize; (iii) labor markets clear, so effective labor supply matches firm demand in each occupation.

Search frictions. The framework accommodates search frictions implicitly rather than setting them aside. More persistent matching frictions are embedded in the latent match values b_{ijt}: following the logic of discrete choice theory (McFadden, 1974), the revealed-preference approach treats the lower present value of jobs with higher access costs as part of the worker's revealed valuation of the match. Frictions that make an occupation harder to enter reduce its attractiveness, and this is captured in the estimated b_{ijt} without imposing additional structural assumptions on the underlying search process (e.g. Moscarini, 2005). Time variation in friction intensity, which may generate polarization (Martellini and Menzio,

2021), is similarly absorbed through the time-varying b_{ijt}. Shorter-run frictions are less of a concern given the long-run focus of the analysis. Kambourov and Manovskii (2008) estimate annual occupation switching rates of 16–20 percent in the U.S. labor market, implying average occupational tenures of 5–6 years; monthly CPS switching rates imply shorter tenures still (Moscarini and Thomsson, 2007). Combined with the high frequency of early-career transitions documented by Topel and Ward (1992), workers encounter repeated reallocation opportunities over the decadal horizons we study (Sullivan and To, 2025), and short-run matching frictions become progressively less binding.

4 Identification and Estimation

The model is estimated by GMM, following a nested fixed-point structure. We distinguish between deep structural parameters $\Theta_{deep} = \{\gamma, \sigma_a, \sigma_\theta, \rho_k, \eta\}$ and time-varying parameters $\Theta_t = \{b_{ijt}, \beta_{ijt}, \alpha_{jt}, \nu_k\}$. For any candidate Θ_{deep}, the time-varying parameters are recovered analytically in an inner loop by inverting observed wages and employment shares; only Θ_{deep} requires numerical search.

Supply parameters. The deep supply parameters $(\gamma, \sigma_a, \sigma_\theta)$ are identified from the second moments of the income distribution. For a worker in group i employed in occupation j at time t, log income equals $\ln w_{ijt} + \ln a$. The population variance of log income can be broken down in within and between components:

$$\mathrm{Var}(\ln y_t) = \underbrace{\mathrm{Var}(E[\ln y \mid i, j, t])}_{\text{between-group}} + \underbrace{E[\mathrm{Var}(\ln y \mid i, j, t)]}_{\text{within-group}} \tag{10}$$

The parameter σ_a is proportional to the total variance $\mathrm{Var}(\ln y_t)$: greater ability dispersion raises income variation for any given wage structure.

The parameter σ_θ is identified from the ratio of between-group to within-group income variance. As $\sigma_\theta \to \infty$, idiosyncratic preference shocks dominate and workers sort randomly across occupations; within-group ability variance equals the population variance and the between-group component reflects only wage variation. As $\sigma_\theta \to 0$, sorting becomes deterministic and within-group ability variance is minimized while between-group variance is maximized. The ratio is monotone in σ_θ, so the observed between/within ratio uniquely pins it down, provided $\gamma \neq 1$.

The curvature parameter γ is identified from the gradient of within-group income variance with respect to average occupation wages. When $\gamma < 1$, utility exhibits increasing differences in wages and ability, generating positive sorting: high-ability workers gravitate toward

high-wage occupations (Groes et al., 2015). Because the right tail of the lognormal ability distribution is unbounded, within-group ability variance is larger in high-wage occupations, and a lower γ steepens this gradient. Logarithmic utility ($\gamma = 1$) makes $\ln a$ additive and cancels from the selection probabilities, rendering occupational sorting independent of ability and the gradient uninformative. Matching the empirical gradient identifies γ separately from σ_θ.

Given $(\gamma, \sigma_a, \sigma_\theta)$, the latent values b_{ijt} are recovered by inverting the observed employment shares μ_{ijt}. For any trial $b_{ijt}^{(n)}$, the update

$$b_{ijt}^{(n+1)} = \sigma_\theta \ln \mu_{ijt} - \sigma_\theta \ln\left(\mu_i \int s_{i0t}^{(n)}(a)\, \exp(u(w_{ijt}a)/\sigma_\theta)\, dF(a)\right) \tag{11}$$

converges to the unique b_{ijt} that matches observed shares, by the market-share inversion of Berry (1994) extended to integrate over ability.

Technology parameters. The match-specific productivity terms β_{ijt} are recovered from within-occupation, cross-group wage ratios. The wage equation (9) implies $w_{ijt}/w_{i'jt} = \beta_{ijt}/\beta_{i'jt}$ for any two groups in the same occupation, so:

$$\hat{\beta}_{ijt} = \frac{w_{ijt}}{w_{1jt}} \tag{12}$$

normalizing $\beta_{1jt} = 1$ for the reference group.

Given $\{\hat{\beta}_{ijt}\}$, the within-cluster elasticity ρ_k is identified from within-cluster wage ratios. For occupations j, j' in the same cluster k:

$$\log \frac{w_{ijt}}{w_{ij't}} = \log \frac{\alpha_{jt}\beta_{ijt}}{\alpha_{j't}\beta_{ij't}} + (\rho_k - 1) \log \frac{\tilde{L}_{jt}}{\tilde{L}_{j't}} \tag{13}$$

Occupation-pair fixed effects absorb the first term; the slope on relative labor inputs identifies $\rho_k - 1$. We estimate this separately for each cluster, instrumenting relative labor inputs with demographic shift instruments (Appendix A).

Given $\{\hat{\rho}_k\}$, we construct cluster aggregates $\hat{T}_{kt} = \left[\sum_{j \in k} \hat{\alpha}_{jt} \hat{\tilde{L}}_{jt}^{\hat{\rho}_k}\right]^{1/\hat{\rho}_k}$ and identify η from between-cluster wage ratios. For $j \in k$ and $j'' \in k''$:

$$\log \frac{w_{ijt}}{w_{ij''t}} = \mathrm{FE}_{jj''} + (\rho_k - 1) \log \tilde{L}_{jt} - (\rho_{k''} - 1) \log \tilde{L}_{j''t} + (\eta - \rho_k) \log \hat{T}_{kt} - (\eta - \rho_{k''}) \log \hat{T}_{k''t} \tag{14}$$

Conditional on $\{\hat{\rho}_k, \hat{T}_{kt}\}$, the equation is linear in η, identified from the differential sensitivity of between-cluster wage ratios to movements in cluster-level labor aggregates. The restriction

$\rho_k > \eta$ for all k is testable. As a preliminary check, the flat-CES estimate $\hat{\rho} \approx 1.65$ from Alonzo and Gallipoli (2023) averages over within- and between-cluster pairs; the nested structure should recover $\rho_k > \hat{\rho} > \eta$ if the task-based grouping is empirically meaningful.

5 The Role of Technology versus Latent Values

We use four counterfactual experiments to assess the respective contributions of technology and latent employment values to the evolution of employment and wages. In the first, latent values are held at their 1980 levels and the model is re-solved; in the second, technology parameters are held fixed. Two additional experiments separate partial from general equilibrium effects: one holds wages fixed, isolating employment responses to latent value changes absent price adjustment; the other holds employment shares fixed, isolating the partial equilibrium wage impact of technology alone.

Tables 1 and 2 summarize the results by occupation cluster. Wage changes are measured in constant 2000 dollars per hour. The experiments suggest that shifts in latent employment values are the primary force behind worker reallocation across occupations, while technology is the primary driver of changes in the wage structure. The two channels interact: the supply contraction induced by deteriorating latent values in routine jobs partially offsets the wage declines that technology would otherwise have produced, and the influx of workers into cognitive occupations has compressed wage growth there. Neither margin can be read in isolation.

Table 1: Cumulative changes in employment shares by occupation cluster, 1980 to 2018 (percentage points). Counterfactual columns hold the indicated variable at its 1980 level and re-solve the model. [Preliminary - To Be Updated]

	Actual	Fixed latent	Fixed tech	Fixed wages
Non-routine cognitive (NRC)	+10.0	+9.0	+7.5	+9.0
Routine cognitive (RC)	−1.0	−1.5	−0.5	−1.0
Non-routine manual (NRM)	+5.0	+0.1	+4.5	+5.5
Routine manual (RM)	−5.0	−2.5	−3.0	−4.5

5.1 Latent values drive occupational reallocation

The clearest finding is for non-routine manual (NRM) jobs. Employment in that cluster rose by about 5 percentage points; fixing latent values at 1980 reduces this increase to zero, while fixing technology has almost no effect. NRM wages were roughly unchanged over the period (Table 2), so a model driven solely by wages would predict no reallocation there at all.

Table 2: Cumulative changes in average hourly wages by occupation cluster, 1980–2018 (constant 2000 dollars per hour). Counterfactual columns hold the indicated variable at its 1980 level and re-solve the model. [Preliminary - To Be Updated]

	Actual	Fixed empl.	Fixed latent	Fixed tech
Non-routine cognitive (NRC)	+4.3	+7.5	+4.5	−0.8
Routine cognitive (RC)	+1.0	−0.5	+1.0	+1.5
Non-routine manual (NRM)	+0.0	+1.2	+2.8	−1.5
Routine manual (RM)	−2.0	−6.0	−3.0	+2.5

Non-routine cognitive (NRC) employment also grew substantially, by about 10 percentage points, but the attribution is more mixed. Both latent values and technology contributed: fixing either one reduces the increase, to roughly 9 and 7.5 percentage points respectively. The technology channel is more important in NRC than anywhere else, consistent with skill-biased demand shifts raising productivity in cognitive work.

At the bottom, routine manual (RM) employment fell by about 5 percentage points, with technology and latent values contributing in roughly equal measure. Each accounts for about 2 to 3 percentage points of the decline. Routine cognitive (RC) employment changed little, making attribution less clear.

A pattern runs through all four clusters: the partial-equilibrium (fixed-wages) and general equilibrium (baseline) counterfactuals are close to each other, indicating that wage adjustment played a secondary role in the employment reallocation. Employment responds primarily to the total value of the match, of which wages are only a part.

Finding. Latent value changes account for essentially all of the employment growth in non-routine manual occupations and a substantial share of the decline in routine manual occupations. Technology matters mainly for the shift toward non-routine cognitive work. Wage adjustment is secondary throughout.

5.2 The female employment surge

The demographic findings in Figure 3 are most striking for women. Employment rose by about 15 percentage points for both college and non-college women. Fixing latent values at 1980 levels collapses nearly all of that increase: the counterfactual shows a gain of roughly 0.5 percentage points for non-college women and 2.5 for college women. Fixing technology, by contrast, leaves most of the increase intact. The implication is that essentially the entire female employment expansion over this period is associated with changes in non-wage job

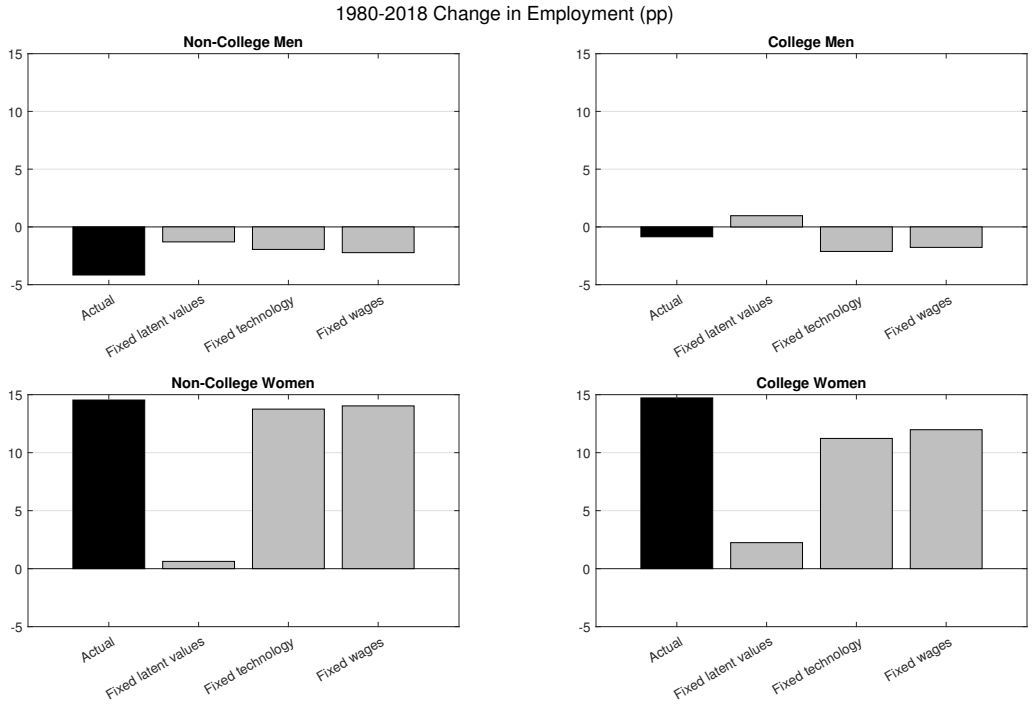


Figure 3: Changes in employment rates by demographic group, 1980 to 2018. Black bars show historical percentage-point changes. Counterfactual bars hold, respectively, latent values, technology parameters, and wages at their 1980 levels.

attributes.[3]

Men's employment fell for both education groups, with a different balance between the two forces. For non-college men, the roughly 4 percentage point decline largely reflects latent value deterioration: fixing latent values at 1980 reduces the counterfactual decline to about 1 percentage point. For college men, the actual change is small (about -0.5 percentage points), but the two forces are larger and oppose each other: latent values pulled employment down, while technology raised it (the fixed-technology counterfactual shows a sharper decline than the actual). Both channels matter; they partially cancel.

Finding. The 14-15 percentage point rise in women's employment between 1980 and 2018 is mostly attributable to improved latent job values; fixing them at 1980 levels reduces the increase to near zero for both education groups. For men, both channels operate but are smaller and partially cancel for college graduates.

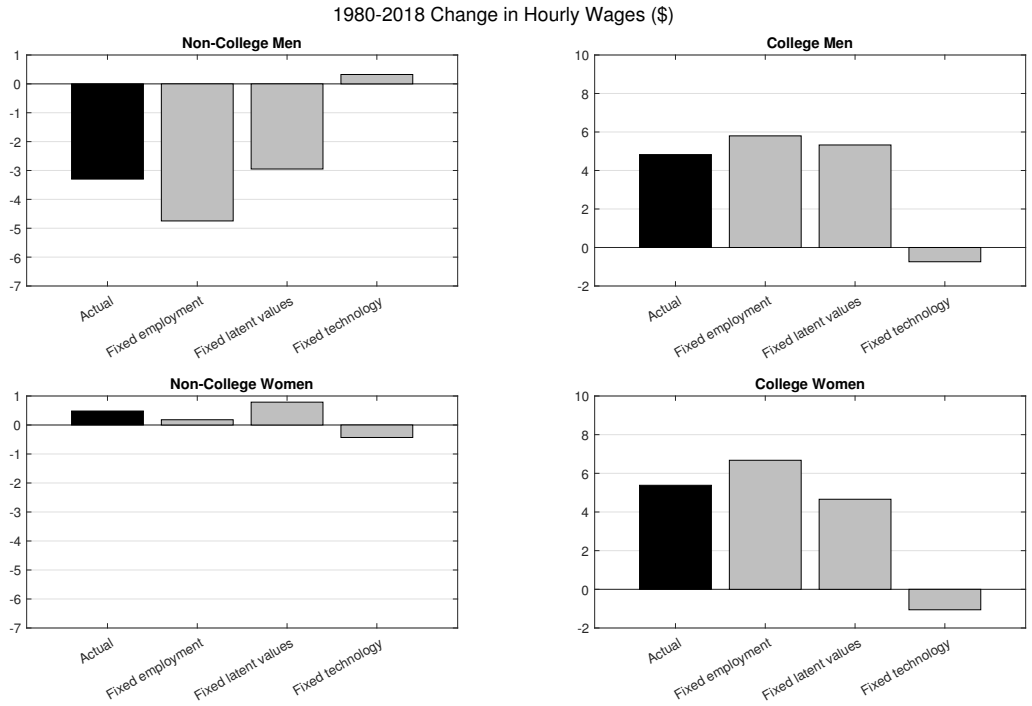


Figure 4: Changes in average hourly wages by demographic group. Baseline and counterfactual scenarios, 1980–2018.

5.3 Technology drives wages, supply responses dampen it

For all four demographic groups (Figure 4), fixing technology at 1980 values reverses or eliminates most of the actual wage change. Fixing latent values, by contrast, has little effect on wages. This is consistent with technological change being the dominant force behind the wage structure, with latent value shifts operating mainly through occupational sorting rather than through wages directly.

The occupation-level results (Table 2) add an important qualification: equilibrium feedbacks from employment changes do have an influence on wage outcomes, in opposite directions at the two ends of the distribution.

In NRC occupations, wages rose by about \$4.30. Holding employment fixed at 1980 levels, wages would have risen by roughly \$7.50: the large supply inflow into cognitive jobs, driven mainly by improving latent values among college workers, compressed wages in those occupations by about \$3. In RM occupations the logic runs the other way. Wages fell by about \$2.00. Without any employment adjustment, the wage decline would have been closer to \$6.00; the supply contraction, again driven largely by deteriorating latent values, raised

[3]Declining occupational barriers enter the model through latent match values b_{ijt} rather than technology parameters; improvements in occupational access (Hsieh et al., 2019) are therefore reflected in the latent value channel.

the marginal product of remaining routine workers and recovered roughly $4.00 of that loss. Technology shock and supply response nearly cancel in RM, leaving a small observed wage decline that understates the magnitude of what happened in that market.

NRM wages were essentially flat, reflecting two roughly offsetting forces: technology raised marginal products modestly, while the large labor inflow driven by improving latent values depressed them. Routine cognitive wages changed little and the counterfactuals are uninformative.

Finding. Technology is the dominant driver of wages across all groups, but equilibrium supply responses substantially modified the outcomes. The surge of workers into cognitive occupations compressed wages there by roughly $3 relative to a fixed-employment benchmark; the supply contraction out of routine manual jobs offset about $4 of a $6 technology-driven wage decline. Latent value shifts are largely transmitted through employment, not wages.

6 Missing Wage Polarization: an Equilibrium Outcome

The counterfactual results raise a question: why do shifts in latent values move employment without moving wages, while technology shifts do the opposite? The coexistence of sharp employment polarization and muted wage polarization follows from two properties of the estimated model that interact through equilibrium. This section examines the mechanics of these responses.

Two elasticities recur throughout the analysis. The *extensive margin elasticity to wages* $\varepsilon^{ext}_{w,ij}$ measures the percentage change in the share of group i workers choosing occupation j per percentage change in its wage. The *extensive margin elasticity to latent values* $\varepsilon^{ext}_{b,ij}$ is the corresponding response to a percentage change in $|b_{ij}|$; the absolute value accommodates matches with negative latent values. Closed-form expressions under the parametric assumptions of Section 3 are derived in Appendix B.

6.1 Differential elasticities and equilibrium dampening

First, we consider a shock to latent employment values, say an improvement in the amenities of non-routine manual jobs. The match becomes more attractive, drawing workers in from outside and from competing occupations. Employment inflows can be large because the extensive margin elasticity to latent values, ε^{ext}_{b}, ranges from 1.51 to 1.71 for routine manual occupations and is comparably elevated elsewhere.[4] The inflow raises labor supply

[4]Estimates from Alonzo and Gallipoli (2023); see Appendix B for definitions and closed-form expressions. Re-estimation with the finer occupation classification used here may shift these values.

in the occupation and depresses wages; under a pure amenity shock, the wage declines in equilibrium.

A technology shock works differently. Higher productivity raises wages and labor demand. Labor supply responds through the extensive margin, but with a smaller elasticity: ε_w^{ext} lies between 0.50 and 0.57 for routine occupations. The modest employment response means that the composition of workers in the match changes little and the marginal worker's characteristics remain approximately stable and the technology improvement shows up in observed wages with little dilution.

These two mechanisms together generate the polarization pattern. In routine manual occupations, both productivity and latent values deteriorated. The combined employment outflow is large because the amenity channel induces a larger outflow than the wage channel. That outflow is itself an equilibrium force: as workers exit, the marginal product of those remaining rises, partially supporting wages. Formally, let $\Delta \log \mu_j$ denote the equilibrium employment share change in occupation j and $\Delta \log w_j$ the wage change following a joint shock. The employment response decomposes as:

$$\Delta \log \mu_j \approx \varepsilon_w^{ext} \cdot \Delta \log w_j + \varepsilon_b^{ext} \cdot \Delta \log |b_j| \tag{15}$$

The wage is an equilibrium object. Under finite substitutability across occupations, the employment change feeds back into the wage through the inverse labor demand elasticity η_j:

$$\Delta \log w_j \approx \underbrace{\Delta \log A_j}_{\text{direct technology}} - \underbrace{\eta_j \cdot \Delta \log \mu_j}_{\text{equilibrium dampening}} \tag{16}$$

Substituting (16) into (15) and solving gives the reduced-form responses:

$$\Delta \log \mu_j \approx \frac{\varepsilon_w^{ext} \Delta \log A_j + \varepsilon_b^{ext} \Delta \log |b_j|}{1 + \varepsilon_w^{ext} \eta_j} \tag{17}$$

$$\Delta \log w_j \approx \frac{\Delta \log A_j - \eta_j \varepsilon_b^{ext} \Delta \log |b_j|}{1 + \varepsilon_w^{ext} \eta_j} \tag{18}$$

Under a pure latent value shock ($\Delta A_j = 0$), the wage response in (18) equals $-\eta_j$ times the employment response in (17): an inflow raises labor supply and depresses wages, so employment and wages move in opposite directions. An amenity improvement draws workers in without raising wages; a deterioration pushes workers out while cushioning wages through the resulting rise in marginal products. Under a pure technology shock ($\Delta b_j = 0$), wages respond with unit weight while employment responds with weight ε_w^{ext} only. When $\varepsilon_b^{ext}/\varepsilon_w^{ext}$ is large, latent value shocks drive most employment variation while their associated wage

effects are dampened or reversed. The 'bodies without wages' pattern is the result.

6.2 The elasticity ratio as a sufficient statistic

The ratio $\varepsilon_b^{ext}/\varepsilon_w^{ext}$ varies systematically across clusters in a way that predicts where the employment-wage disconnect should be sharpest. Alonzo and Gallipoli (2023) find that the elasticity to latent values is roughly two to three times larger than the elasticity to wages, with the gap most pronounced in routine manual occupations and narrower in non-routine cognitive ones, where pecuniary returns play a greater role in attracting workers. Both employment and wages moved substantially in the latter, as the counterfactuals confirm.

The NRM cluster offers the clearest illustration of the amenity mechanism. Employment grew by 5 percentage points and wages barely moved. Fixing latent values eliminates essentially the entire employment increase; fixing technology leaves it intact. This is consistent with an amenity improvement that drew workers in through the latent value channel, with the employment inflow dampening wages through equilibrium labor demand and leaving them near their 1980 level.

The NRC wage result in Table 2 confirms the same equilibrium logic from the opposite direction. Without the employment inflow, NRC wages would have risen by roughly \$7.50; the actual increase was \$4.30. The difference reflects compression by new entrants, predominantly college-educated women whose participation rose sharply. Technology drove the wage growth; the supply response attenuated it by changing the composition of marginal workers.

6.3 Wage resilience without downward rigidity

Institutional rigidities, such as downward wage stickiness, offer one standard account of why routine wages did not fall further. The framework here requires no such rigidities: wages are fully flexible. The dampened wage decline in routine occupations arises because the employment outflow raises the marginal product of remaining workers, not because wages are stuck.[5]

More broadly, the magnitude of employment polarization relative to wage polarization is informative about the structure of labor supply preferences, specifically the ratio $\varepsilon_b^{ext}/\varepsilon_w^{ext}$. In a labor market where workers respond strongly to non-wage attributes, shocks to amenities generate sharp employment reallocation with little wage movement. That is what the estimates suggest for the U.S. between 1980 and 2018: wages are an incomplete signal of

[5]Hunt and Nunn (2022) attribute stable routine wages to selective exit of low-skill workers, which mechanically raises average wages in those jobs. We emphasize an equilibrium mechanism where wages are sustained by rising marginal products, independent of unobserved worker composition.

what happened to the total value of worker-occupation matches over this period.

7 Gender and the Polarization of Bodies

Figure 1 documents that men and women participate in polarization differently. Women's employment growth is concentrated in the upper-middle portion of the skill distribution, while men show a sharper peak at the bottom. The wage panel offers little guidance: as Figure 2 shows, wage changes account for a small share of employment reallocation for both genders, though the share is somewhat higher for college-educated women. The gender asymmetry in employment reallocation is difficult to account for with wages alone.

The counterfactual analysis in Section 5 provides a quantitative assessment of these patterns. The 14–15 percentage point rise in women's employment between 1980 and 2018 collapses to near zero when latent values are held at their 1980 levels; fixing technology leaves most of the increase intact. For men, the two channels are smaller and partially offset. Latent value changes are the primary driver of female employment growth across both education groups.

The mechanism is the differential elasticity documented in Alonzo and Gallipoli (2023): women's labor supply is more responsive to non-wage job attributes than men's. Faberman et al. (2026) recover the same asymmetry from job search data, finding that women place greater weight on non-wage amenities when evaluating job offers; Faberman et al. (2022) document systematic gender differences in job search behavior among both employed and non-employed workers, with women placing higher value on non-pecuniary job characteristics; Le Barbanchon et al. (2021) find that women accept lower wages for shorter commutes and Wiswall and Zafar (2018) show that women have a higher willingness to pay for workplace flexibility and stability (Goldin, 2014), consistent with higher amenity elasticity across multiple settings. When the amenity content of non-routine service and cognitive occupations improved relative to routine ones, women reallocated more strongly. Bagga et al. (2023) provide independent evidence for this channel, showing that amenity shocks generate substantial labor reallocation; the estimates here suggest gender asymmetry amplifies that response considerably. Cortés and Pan (2018) document that women are systematically overrepresented in occupations offering greater flexibility, higher social contribution, and lower physical demands, and that these non-pecuniary characteristics are significant predictors of the female employment share across occupations.

If workers partly exchange lower wages for better non-wage attributes, a portion of the gender wage gap reflects amenity differences rather than discrimination or human capital gaps (Morchio and Moser, 2024). The rise in female labor force participation between 1980

and 2000 reflects in part an improved allocation of talent as barriers to high-skill occupations declined (Hsieh et al., 2019). Our analysis lends support to that account. Declining barriers, whether formal or informal, enter the model through b_{ijt}. A worker who anticipates resistance or limited access in an occupation values that match less, so improvements in occupational access show up as rising latent values and drive the employment growth we attribute to that channel (Hurst et al., 2024).

8 Conclusion

The reallocation of workers across occupations since 1980 has been far more pronounced than the concurrent changes in the wage structure. This paper argues that the disconnect follows from how equilibrium interactions shape both employment and wages jointly: when labor supply responds more strongly to non-wage job attributes than to wages, shifts in latent match values generate large employment movements with little accompanying wage movement. The relative responsiveness of employment to latent values versus wages governs the magnitude of the employment-wage disconnect. Alonzo and Gallipoli (2023) find that the elasticity to latent values is two to three times larger than the elasticity to wages, implying that amenity shocks carry several times the employment weight of equivalent wage shocks while producing a fraction of the wage response. The counterfactuals confirm this: fixing latent values at their 1980 levels collapses most of the observed employment polarization, while wage polarization survives.

A second lesson is about how equilibrium adjustments influence wages. The muted wage changes visible in Figure 1 do not require institutional rigidities. They are an equilibrium consequence of the employment flows themselves, operating in opposite directions across the skill distribution. Occupations in the middle that lost workers, routine jobs where both technology and latent values deteriorated, saw the marginal product of remaining workers rise as the outflow proceeded, partially supporting wages from below. The result is the near-flat wage profile observed between the 5th and 85th skill percentiles. Only at the top, where technology gains were large enough to outpace the supply inflow, did the equilibrium dampening fail to absorb the shock, producing the spike visible for both genders and more pronounced for women. No assumption of downward wage rigidity is needed; employment polarization and wage resilience are joint equilibrium outcomes of the same flexible-wage model.

The gender dimension provides the sharpest single illustration of these forces. Women's employment rose by roughly 15 percentage points between 1980 and 2018; fixing latent values at their 1980 levels collapses that increase to near zero for both education groups,

while fixing technology leaves most of it intact. For men, the two channels operate in opposite directions and largely cancel for college graduates. The asymmetry is consistent with women's labor supply being more responsive to non-wage job attributes, a finding corroborated by independent methods in the amenity and job search literatures (Mas and Pallais, 2017; Maestas et al., 2023; Faberman et al., 2026). Declining occupational barriers, formal or informal, enter the framework through improvements in b_{ijt}, so the estimated latent value shifts partly reflect the improved allocation of talent documented by Hsieh et al. (2019).

These results suggest that market wages are an incomplete signal of total labor market returns over this period. Workers in routine occupations saw the non-wage component of their returns deteriorate in ways that drove large employment outflows; the observed wage decline understates the fall in total returns. Workers who shifted into non-routine occupations may have experienced the reverse, with wage gains that understate the improvement in total match value. The wage record understates what happened to worker welfare over this period. Losses among routine workers exceeded what wages reveal; for those who reallocated into better-matched occupations, the gains likely did too.

References

Acemoglu, D. and Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. *Handbook of Labor Economics*, 4:1043–1171.

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

Alonzo, D. and Gallipoli, G. (2023). The changing value of employment. Discussion Paper 17943, Centre for Economic Policy Research.

Autor, D. H. and Dorn, D. (2013). The growth of low skill service jobs and the polarization of the U.S. labor market. *American Economic Review*, 103(5):1553–1597.

Autor, D. H., Katz, L. F., and Kearney, M. S. (2006). The polarization of the US labor market. *The American Economic Review*, 96(2):189–194.

Autor, D. H., Levy, F., and Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. *Quarterly Journal of Economics*, 118(4):1279–1333.

Bagga, S., Mann, L., Sahin, A., and Violante, G. (2023). Job amenity shocks and labor reallocation. Working paper, Princeton University.

Bárány, Z. L. and Siegel, C. (2018). Job polarization and structural change. *American Economic Journal: Macroeconomics*, 10(1):57–89.

Berry, S. T. (1994). Estimating discrete-choice models of product differentiation. *RAND Journal of Economics*, 25(2):242–262.

Charles, K. K., Hurst, E., and Schwartz, M. (2019). The transformation of manufacturing and the decline in US employment. In *NBER Macroeconomics Annual 2018, Volume 33*, pages 307–372. University of Chicago Press.

Cortes, G. M., Jaimovich, N., and Siu, H. E. (2017). Disappearing routine jobs: Who, how, and why? *Journal of Monetary Economics*, 91:69–87.

Cortés, P. and Pan, J. (2018). Occupation and gender. *Oxford Handbook of Women and the Economy*, pages 425–452.

Faberman, J. R., Mueller, A. I., and Sahin, A. (2026). Job search, job amenities and the gender pay gap. Working paper, NBER. Previously circulated as "Job Search and the Gender Wage Gap". NBER Working Paper 34877.

Faberman, R. J., Mueller, A. I., Şahin, A., and Topa, G. (2022). Job search behavior among the employed and non-employed. *Econometrica*, 90(4):1743–1779.

Goldin, C. (2014). A grand gender convergence: Its last chapter. *American Economic Review*, 104(4):1091–1119.

Groes, F., Kircher, P., and Manovskii, I. (2015). The U-shapes of occupational mobility. *The Review of Economic Studies*, 82(2):659–692.

Hsieh, C.-T., Hurst, E., Jones, C. I., and Klenow, P. J. (2019). The allocation of talent and U.S. economic growth. *Econometrica*, 87(5):1439–1474.

Hunt, J. and Nunn, R. (2022). Has U.S. employment really polarized? A critical reappraisal. *Labour Economics*, 75:102117.

Hurst, E., Rubinstein, Y., and Shimizu, K. (2024). Task-Based Discrimination. *American Economic Review*, 114(6):1723–1768.

Kambourov, G. and Manovskii, I. (2008). Rising occupational and industry mobility in the United States: 1968-97. *International Economic Review*, 49(1):41–79.

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

Lamadon, T., Mogstad, M., and Setzler, B. (2022). Imperfect Competition, Compensating Differentials, and Rent Sharing in the US Labor Market. *American Economic Review*, 112(1):169–212.

Le Barbanchon, T., Rathelot, R., and Roulet, A. (2021). Gender differences in job search: Trading off commute against wage. *The Quarterly Journal of Economics*, 136(1):381–426.

Maestas, N., Mullen, K. J., Powell, D., von Wachter, T., and Wenger, J. B. (2023). The value of working conditions in the united states and the implications for the structure of wages. *American Economic Review*, 113(7):2007–47.

Martellini, P. and Menzio, G. (2021). Jacks of All Trades and Masters of One: Declining Search Frictions and Unequal Growth. *American Economic Review: Insights*, 3(3):339–352.

Mas, A. and Pallais, A. (2017). Valuing alternative work arrangements. *American Economic Review*, 107:3722–3759.

McFadden, D. (1974). Conditional logit analysis of qualitative choice behavior. In Zarembka, P., editor, *Frontiers in Econometrics*, pages 105–142. Academic Press, New York.

Morchio, I. and Moser, C. (2024). The gender pay gap: Micro sources and macro consequences. Working Paper 32408, National Bureau of Economic Research.

Moscarini, G. (2005). Job Matching and the Wage Distribution. *Econometrica*, 73(2):481–516.

Moscarini, G. and Thomsson, K. (2007). Occupational and job mobility in the US. *Scandinavian Journal of Economics*, 109(4):807–836.

Sorkin, I. (2018). Ranking firms using revealed preference. *The Quarterly Journal of Economics*, 133:1331–1393.

Sullivan, P. and To, T. (2025). Job dispersion and compensating wage differentials. *Industrial and Labor Relations Review*. Forthcoming.

Taber, C. and Vejlin, R. (2020). Estimation of a roy/search/compensating differential model of the labor market. *Econometrica*, 88:1031–1069.

Topel, R. H. and Ward, M. P. (1992). Job mobility and the careers of young men. *The Quarterly Journal of Economics*, 107(2):439–479.

Wiswall, M. and Zafar, B. (2018). Preference for the workplace, investment in human capital, and gender. *The Quarterly Journal of Economics*, 133:457–507.

A Estimation Details

Inner loop: analytical inversion

For a given candidate Θ_{deep}, the time-varying parameters are recovered analytically for each period t:

- b_{ijt}: inverted from observed employment shares μ_{ijt} via the market-share inversion in equation (11).
- β_{ijt}: recovered from within-occupation wage ratios (equation (12)), normalizing the reference group to unity.
- α_{jt}, ν_k: recovered from within- and across-cluster wage ratios given β_{ijt} and $\tilde{L}_{jt}$, with normalizations $\sum_{j \in k} \alpha_{jt} = 1$ and $\sum_k \nu_k = 1$.

Outer loop: GMM objective

The outer loop searches over $\Theta_{deep} = \{\gamma, \sigma_a, \sigma_\theta, \rho_k, \eta\}$ to minimize the weighted distance between model-implied and empirical moments. The targeted moments are:

- Total log-income variance $\text{Var}(\ln y_t)$, identifying σ_a.
- Ratio of between-group to within-group income variance, identifying σ_θ.
- Gradient of within-group income variance with respect to average occupation wages, identifying γ.
- Co-movement of within-cluster relative wages and relative effective labor over time, identifying ρ_k (equation (13)).
- Co-movement of across-cluster relative wages and cluster-level labor aggregates over time, identifying η (equation (14)).

All moments are stacked and minimized using an identity weight matrix to reduce small-sample bias.

B Labor Supply Elasticities

We define two elasticities from the model of Section 3. Throughout, $\bar{s}_{ij} \equiv \mu_{ij}/\mu_i = \int s_{ij}(a)\, dF(a)$ denotes the aggregate conditional employment share of group i in occupation j.

Equivalence of Elasticity of Measure and Share

We are interested in the elasticity of the measure of workers μ_{ij} with respect to the wage w_{ij}. By the properties of log-derivatives:

$$\varepsilon^{ext}_{w,ij} = \frac{\partial \ln \mu_{ij}}{\partial \ln w_{ij}} = \frac{w_{ij}}{\mu_{ij}} \frac{\partial \mu_{ij}}{\partial w_{ij}} \tag{19}$$

Since $\mu_{ij} = \mu_i \bar{s}_{ij}$ where μ_i is the fixed population mass, we have $\frac{\partial \mu_{ij}}{\partial w_{ij}} = \mu_i \frac{\partial \bar{s}_{ij}}{\partial w_{ij}}$. Substituting this back:

$$\varepsilon^{ext}_{w,ij} = \frac{w_{ij}}{\mu_i \bar{s}_{ij}} \left(\mu_i \frac{\partial \bar{s}_{ij}}{\partial w_{ij}} \right) = \frac{w_{ij}}{\bar{s}_{ij}} \frac{\partial \bar{s}_{ij}}{\partial w_{ij}} = \frac{\partial \ln \bar{s}_{ij}}{\partial \ln w_{ij}} \tag{20}$$

Thus, the elasticity of the measure is identical to the elasticity of the share because the population mass μ_i is constant.

Extensive margin elasticity to wages

The percentage change in the employment share of match (i, j) per percentage change in its wage is $\varepsilon^{ext}_{w,ij}$. To derive its closed form, we follow three steps:

Step 1: Derivative of Utility. First, we find the change in utility with respect to the wage:

$$\frac{\partial U_{ij}(a)}{\partial w_{ij}} = \frac{\partial}{\partial w_{ij}} \left[\frac{w_{ij}^{1-\gamma} a^{1-\gamma}}{1-\gamma} \right] = w_{ij}^{-\gamma} a^{1-\gamma}$$

Step 2: Derivative of Choice Probability. Using the property of the Logit function $\frac{\partial s_j}{\partial U_j} = \frac{1}{\sigma_\theta} s_j(1 - s_j)$, we have:

$$\frac{\partial s_{ij}(a)}{\partial w_{ij}} = \frac{1}{\sigma_\theta} s_{ij}(a)(1 - s_{ij}(a)) \cdot w_{ij}^{-\gamma} a^{1-\gamma}$$

Step 3: Aggregate Elasticity. Integrating over ability a and converting to elasticity:

$$\varepsilon_{w,ij}^{ext} \;=\; \frac{w_{ij}^{1-\gamma}}{\sigma_\theta\, \bar{s}_{ij}} \int a^{1-\gamma}\, s_{ij}(a)\big(1 - s_{ij}(a)\big)\, dF(a) \tag{21}$$

Extensive margin elasticity to latent values

The percentage change in the employment share of match (i, j) per percentage change in $|b_{ij}|$ is $\varepsilon_{b,ij}^{ext}$.

Step 1: Derivative of Utility. Since b_{ij} enters utility linearly, $\frac{\partial U_{ij}(a)}{\partial b_{ij}} = 1$.

Step 2: Derivative of Choice Probability. The individual response is:

$$\frac{\partial s_{ij}(a)}{\partial b_{ij}} = \frac{1}{\sigma_\theta} s_{ij}(a)(1 - s_{ij}(a))$$

Step 3: Aggregate Elasticity. Integrating over ability and multiplying by $\frac{|b_{ij}|}{\bar{s}_{ij}}$:

$$\varepsilon_{b,ij}^{ext} \;=\; \frac{|b_{ij}|}{\sigma_\theta\, \bar{s}_{ij}} \int s_{ij}(a)\big(1 - s_{ij}(a)\big)\, dF(a) \tag{22}$$

The Elasticity Ratio

Dividing (22) by (21) cancels σ_θ and $\bar{s}_{ij}$, yielding:

$$\frac{\varepsilon_{b,ij}^{ext}}{\varepsilon_{w,ij}^{ext}} \;=\; \frac{|b_{ij}|}{w_{ij}^{1-\gamma}} \cdot \frac{\displaystyle\int s_{ij}(a)\big(1 - s_{ij}(a)\big)\, dF(a)}{\displaystyle\int a^{1-\gamma}\, s_{ij}(a)\big(1 - s_{ij}(a)\big)\, dF(a)} \tag{23}$$

The first term represents the relative weight of latent values vs. wages in utility. The second factor is a selection-variance-weighted ratio of ability moments; when ability dispersion is small or $\gamma \to 1$, it approaches unity and the ratio collapses to $|b_{ij}|/w_{ij}^{1-\gamma}$. When this ratio exceeds one, amenity shocks generate larger employment movements than equivalent-sized wage shocks; Alonzo and Gallipoli (2023) estimate it to be roughly two to three across occupation clusters.