



ECONOMICS WORKING PAPERS

**Trends in Labor Force Participation and  
Unemployment, 1976-2024\***

Andreas Hornstein<sup>†</sup> Marianna Kudlyak<sup>‡</sup>

Economics Working Paper 26113

HOOVER INSTITUTION  
434 GALVEZ MALL  
STANFORD UNIVERSITY  
STANFORD, CA 94305-6010

May 28, 2026

Using CPS microdata, 1976-2024, we estimate trend and cyclical components of unemployment and labor force participation for 44 age-gender-education groups. We fit a parsimonious state-space model in which each series is the sum of latent cohort and time-varying age effects and a latent cyclical factor shared across unemployment and participation, without imposing structural covariates. Aggregating group trends with observed population shares, we find that population aging and educational upgrading explain most long-run movements in aggregate trends, while cohort effects drive large gender differences in participation. Combining our estimates with demographic projections and an estimated cohort model of education shares, we forecast that over the next two decades, trend participation declines by about 1.5 pp and trend unemployment falls by about 0.4 pp, remaining historically low.

**Keywords:** Labor Force Participation Rate. Unemployment Rate. Demographic Composition. Age Effects. Cohort Effects. Model of Educational Attainment.

**JEL Code:** E24, J11.

The Hoover Institution Economics Working Paper Series allows authors to distribute research for discussion and comment among other researchers. Working papers reflect the views of the authors and not the views of the Hoover Institution.

---

\* This paper is a substantially revised version of Hornstein and Kudlyak (2019). We thank Vinh Phan and David Ramachandran for excellent research assistance. Any opinions expressed are those of the authors and do not reflect those of the Federal Reserve Bank of Richmond, the Federal Reserve Bank of San Francisco, or the Federal Reserve System.

<sup>†</sup> Federal Reserve Bank of Richmond. E-mail address: and.horn.re@gmail.com.

<sup>‡</sup> Federal Reserve Bank of San Francisco, Hoover Institution at Stanford University, CEPR, IZA. E-mail address: marianna.kudlyak@sf.frb.org.

# 1 Introduction

Long-run movements in unemployment and labor force participation (LFP) shape assessments of labor market slack and potential output. Over the past half-century, U.S. participation rose through the late 1990s and then declined, while unemployment trended down. Because unemployment and participation differ systematically by age, gender, and education—and because the population has become older and more educated—aggregate trends mechanically reflect both *within-group* changes and *changes in demographic composition*. Focusing only on aggregate series can therefore obscure the sources of long-run labor-market change.<sup>1</sup>

This paper quantifies how demographic composition and group-specific trends jointly determine the trend components of aggregate unemployment and participation. We use CPS microdata from 1976 to 2024 to construct annual unemployment and LFP rates for 44 age-gender-education groups. For each gender-education group, we estimate a parsimonious state-space model that jointly decomposes age-group unemployment and participation rates into (i) low-frequency cohort and age components and (ii) a high-frequency cyclical component. We define a group’s trend as the sum of its cohort and age components, construct aggregate trend series by aggregating group trends using observed population shares, and decompose changes in aggregate trends into contributions from (i) changes in age-gender shares, (ii) changes in education shares within age-gender groups, and (iii) changes in group-specific trends.

Our empirical framework is deliberately parsimonious and fully latent. Within each gender-education group, unemployment and participation in each age bracket are defined as the sum of three *unobserved* components: cohort effects, age effects that vary over time, and a cyclical component driven by a common factor shared by unemployment and participation. The model does not have a natural rate interpretation for unemployment and does not contain any structural covariates (e.g., schooling enrollment, disability, or program variables). This avoids committing to a particular set of structural drivers and keeps the decomposition comparable across groups, at the cost that movements in cohort and age effects are reduced-form objects.

We highlight three findings. First, the trend component of aggregate unemployment declines smoothly from about 8% in 1976 to about 5% by 2024, while cyclical deviations from trend are large. Our decomposition attributes this 3 percentage-point (pp) decline, approximately equally, to changes in workforce composition—especially rising educational

---

<sup>1</sup>Section 2 documents these systematic level differences and differential long-run movements across broad demographic groups.

attainment (the largest contribution) and population aging—and to declining group-level unemployment and participation trends; cohort effects account for most of the within-group drift in these group trends.

Second, the trend component of aggregate LFP is hump-shaped, rising from the 1970s and peaking in the late 1990s before gradually declining thereafter. From 1976 to the mid-1990s, the trend LFP rate rose by about 4.1 pp. Educational upgrading explains a little more than half of this rise (+2.2 pp), with rising within-group participation trends contributing an additional +1.2 pp. Conversely, the mid-1990s to 2024 saw a 4.4 pp decline in trend LFP, driven predominantly by population aging (−4.3 pp) and reinforced by declining within-group LFP trends (−2.3 pp), partly offset by continued educational upgrading.

Third, at the group level, unemployment trends drift down only modestly for both men and women, largely through cohort effects. In contrast, participation trends exhibit pronounced cohort-driven differences by gender: successive male cohorts tend to participate less, while successive female cohorts participate more until cohorts entering in the late 1990s and early 2000s.

We also provide twenty-year projections of aggregate trend unemployment and trend participation. Projecting aggregate trends requires forecasts of demographic composition and of group-specific trend rates. We take age-gender population projections from U.S. Census Bureau (2025) and forecast education shares using a cohort state-space model of educational attainment conditional on age and gender estimated on the CPS. Combining projected composition with our estimated group trends, we project that over the next two decades, trend LFP declines by about 1.5 pp from its 2024 level, while trend unemployment declines by about 0.4 pp. These projected trends are mainly driven by population ageing, with minor offsets from cohort effects in educational upgrading, and group participation and unemployment rates.

This paper contributes to the literature in three ways. First, we build aggregate trend unemployment and participation *bottom up* from group-specific trend estimates and quantify the relative roles of age structure, educational upgrading, and within-group trends. Second, relative to standard age-cohort decompositions that treat age profiles as fixed or absorb time variation via observed controls, we allow age effects to evolve stochastically over time in a transparent latent-component framework, while also modeling unemployment and participation jointly through a shared cyclical factor. Third, we incorporate an explicit model of educational attainment dynamics and use it to construct demographic-consistent forecasts of aggregate trend unemployment and participation.

As noted before, our analysis of aggregate trends is built on the aggregation of estimated group-specific trends. In Section 2, we document systematic differences in unemployment

and LFP rates between broad demographic groups and how their composition changes over time. For the more granular decomposition of demographic groups underlying our estimates of aggregate trends in the remainder of the paper, we only discuss broad outlines of group trends and how they differ across groups. In Section 3 we describe the modeling framework. In Section 4 we describe data and estimation. In Section 5 we provide an overview of group-level trend and cycle estimates. In Section 6 we report aggregate trend estimates and their decomposition. In Section 7, we provide twenty-year projections based on a cohort model of educational attainment. Section 8 concludes.

We provide a detailed documentation of the data and estimation methods, and estimated group-specific trends in a detailed online Technical Appendix, Hornstein and Kudlyak (2026).

## 1.1 Related literature

Our analysis relates to work that studies long-run labor force participation through demographic accounting and age-cohort models (e.g., Aaronson et al. (2006), Fallick and Pingle (2007), Kudlyak (2013), Aaronson et al. (2014), Montes (2018)). A common approach in this literature treats age profiles within demographic groups as fixed and attributes remaining time variation to cohort effects and observed controls (such as enrollment, program participation, or health).<sup>2</sup> We take an alternative route: we allow age profiles to vary stochastically over time and discipline the evolution of both age and cohort components within a state-space model. Like Montes (2018), we stratify groups by education rather than treating educational attainment solely as a control.

The paper is also related to the work that estimates the trend unemployment rate and emphasizes the joint determination of unemployment and participation, most prominently Barnichon and Mesters (2018). Their analysis highlights that changes in unemployment trends can reflect changes in participation trends through labor force flows. We similarly model unemployment and participation jointly, but we do so at the level of age-gender-education groups and focus on how changing demographic composition and group-specific trends together shape aggregate trend unemployment. Our results suggest that once education and age composition are accounted for explicitly, population shares play a large role in shaping the long-run decline in the aggregate unemployment trend.<sup>3</sup>

---

<sup>2</sup>Abraham and Kearney (2020) provide a detailed account of various factors behind trends in demographic groups' employment rates.

<sup>3</sup>For other work on demographics and unemployment, see, for example, Shimer (1998), Elsby et al. (2010), and Aaronson et al. (2012).

## 2 Long Run Averages of Demographic Groups

Unemployment and LFP rates differ substantially across gender, age, and education. Aggregate LFP is a population-share weighted average of group-specific LFP rates, while aggregate unemployment is a labor-force-share weighted average of group-specific unemployment rates. If group-specific rates moved largely in parallel over long horizons, long-run movements in aggregate rates would primarily reflect changes in the demographic composition of the population and a common trend. However, group-specific rates trend differently over time, and accounting for the change in aggregate participation rates necessitates a decomposition of the contributions coming from demographic trends and group-specific trends.

Before turning to a formal analysis, we report simple sample averages of unemployment and LFP rates over selected subperiods for a coarse socio-demographic decomposition of the U.S. population, using CPS microdata. We split the population by gender (men versus women), age (young: 16–24; prime-age: 25–54; old: 55+), and education (high school or less versus some college or more).<sup>4</sup> To focus on long-term trends, we compare nine-year averages of unemployment and LFP rates, the employment-to-population ratio (EPOP), and population shares. Table 1 reports these averages for the aggregate economy and for broad sub-aggregates over three subperiods (1976–1984, 1996–2004, and 2016–2024).<sup>56</sup>

Several broad patterns stand out. Men have higher LFP than women, prime-age individuals have higher LFP than younger and older individuals, and more-educated individuals have higher LFP than less educated ones; these rankings are broadly stable across subperiods. Between the first and third subperiods, the population share of individuals aged 55+ increased by about 10 pp, while the share with more than a high school education increased by about 26 pp. Holding group-specific LFP rates fixed, the aging of the population would lower aggregate participation, whereas rising educational attainment would raise it. At the same time, group-specific LFP rates also changed: participation among women and the less-educated increased, whereas participation among the young declined.

Unemployment rates also vary systematically across groups. The unemployment rate among young individuals is higher than that of prime-age individuals, which in turn is higher than that of older individuals, and unemployment is higher among the less educated than among the more educated. Holding group-specific unemployment and LFP rates fixed, shifts toward an older and more-educated population would lower aggregate unemployment. Over time, however, group-specific unemployment rates also changed, with declines among

---

<sup>4</sup>Education is defined for the adult population; see Section 4.1 for details.

<sup>5</sup>The final subperiod includes the pandemic and recovery.

<sup>6</sup>We report averages for a more detailed decomposition by gender, age, and education in the online Technical Appendix, Hornstein and Kudlyak (2026).

Table 1: LFP, Unemployment, and Employment Trends: Aggregate and Broad Sub-Aggregates

	A. Aggregates	
1976-1984	(63.6, 7.7, 58.7, 100.0)	
1996-2004	(66.7, 5.0, 63.4, 100.0)	
2016-2024	(62.6, 4.6, 59.7, 100.0)	
	B. Young: 16-24 years old	
1976-1984	(67.9, 14.6, 58.0, 22.0)	
1996-2004	(64.3, 11.1, 57.2, 16.1)	
2016-2024	(55.3, 9.6, 50.0, 14.7)	
	C. Prime Age: 25-54 years old	
	C.1 HS or less	C.2 More than HS
1976-1984	(73.8, 7.5, 68.2, 29.8)	(84.7, 3.9, 81.4, 21.0)
1996-2004	(78.6, 5.6, 74.2, 24.7)	(87.6, 2.9, 85.0, 31.9)
2016-2024	(74.6, 5.9, 70.2, 16.8)	(86.2, 3.1, 83.5, 31.9)
	D. Old Age: 55 years and older	
	D.1 HS or less	D.2 More than HS
1976-1984	(29.3, 4.6, 27.9, 21.3)	(45.4, 2.6, 44.2, 6.0)
1996-2004	(25.9, 3.7, 24.9, 16.5)	(43.7, 2.8, 42.4, 10.8)
2016-2024	(32.1, 4.1, 30.8, 15.3)	(44.9, 3.3, 43.4, 21.3)

Notes: Sample averages for LFP rate, unemployment rate, employment rate, and population share in percent, (L,U,E,P), for (A) overall economy, population (B) 16-24 years old, (C) 25-54 years old, (D) 55 years and older. Those older than 24 years are split into two groups: one with a high school (HS) education or less, and one with more than a high school education.

the young and prime age individuals, and a modest increase among the more educated older ones.

These simple averages already imply that aggregate trends reflect both compositional changes and group-specific trends. In what follows, we construct trend measures of the unemployment and LFP rates for more granular demographic groups and then attribute changes in aggregate trends to changes in group-specific trends and to demographic shifts.

### 3 Framework for Aggregation of Group-Specific Trends

This section lays out the mapping from group-specific labor-market dynamics to the aggregate trends. The key payoff is twofold: (i) it provides a transparent definition of aggregate trend unemployment and trend participation as population-share-weighted objects built up from group-level trend components; and (ii) it delivers a decomposition that attributes changes in aggregate trends to changes in demographic composition (age-gender and education shares) versus changes in the underlying group-specific trend rates. We proceed in three

steps. First, we define the aggregate rates and their trend counterparts implied by group trends and observed population shares. Second, we specify the state-space model that produces internally consistent trend/cycle decompositions for unemployment and participation within each gender-education group. Third, we derive a first-order (linear) decomposition that we use in the results and in the forecast section to quantify the relative importance of composition shifts versus within-group trend changes.

### 3.1 Aggregate rates and aggregate trends

Let  $i = (s, e)$  index a gender-education group, and let  $g$  index age groups. Denote by  $A_g$  the set of single-year ages contained in age group  $g$  (e.g.,  $A_1 = \{25, \dots, 34\}$  for education groups aged 25+). Let  $\ell_{g,t}^i$  be the labor force participation (LFP) rate of group  $(i, g)$  at time  $t$ ,  $u_{g,t}^i$  its unemployment rate, and  $p_{g,t}^i$  its population share.

The aggregate LFP rate is the population-share weighted average of group LFP rates:

$$\ell_t = \sum_{i,g} p_{g,t}^i \ell_{g,t}^i. \quad (1)$$

The aggregate unemployment rate is the labor-force-share weighted average of group unemployment rates:

$$u_t = \frac{\sum_{i,g} p_{g,t}^i \ell_{g,t}^i u_{g,t}^i}{\sum_{i,g} p_{g,t}^i \ell_{g,t}^i} = \sum_{i,g} \omega_{g,t}^i u_{g,t}^i, \quad (2)$$

where  $\omega_{g,t}^i \equiv p_{g,t}^i \ell_{g,t}^i / \ell_t$  denotes group  $(i, g)$ 's share of the labor force.

Let  $\bar{\ell}_{g,t}^i$  and  $\bar{u}_{g,t}^i$  denote the trend components of the group-specific LFP and unemployment rates estimated using the state-space model below. We define aggregate trends by applying the aggregation formulas (1) and (2) to the trend components, holding population shares at their observed values:

$$\bar{\ell}_t = \sum_{i,g} p_{g,t}^i \bar{\ell}_{g,t}^i, \quad (3)$$

and

$$\bar{u}_t = \frac{\sum_{i,g} p_{g,t}^i \bar{\ell}_{g,t}^i \bar{u}_{g,t}^i}{\sum_{i,g} p_{g,t}^i \bar{\ell}_{g,t}^i} = \frac{\sum_{i,g} p_{g,t}^i \bar{\ell}_{g,t}^i \bar{u}_{g,t}^i}{\bar{\ell}_t}. \quad (4)$$

For the aggregate trends, we take population shares  $\{p_{g,t}^i\}$  as given.

### 3.2 A dynamic factor model for group-specific rates

We estimate a separate dynamic factor model for each gender-education group. Within each group, we model unemployment and LFP jointly by allowing their cyclical components to be driven by a common latent factor. The model decomposes each age-group series into latent

trend and cyclical components: the trend is defined as the sum of the latent cohort and age effects, and the cycle is defined as the latent factor-driven component. We use Kalman filtering and smoothing to estimate these latent components.

For each gender-education group  $i$ , we observe annual outcomes  $q_{g,t}^i$  for  $q \in \{u, \ell\}$  and all age groups  $g = 1, \dots, n_G^i$  over  $t = 1, \dots, T$ , along with population shares  $p_{g,t}^i$ . To simplify notation, we drop the superscript  $i$  in what follows.

### 3.2.1 Measurement equation and time aggregation from cohorts to age groups

Let  $a$  denote single-year age (in years). The cohort index associated with observations at age  $a$  and time  $t$  is  $t - a$  (birth year up to a constant). We model cohort effects at the single-year level and allow the observed age-group outcome to average over the cohorts present in that age group.

For each outcome  $q \in \{u, \ell\}$ , the measurement equation for age group  $g$  is

$$q_{g,t} = \tilde{x}_{qg,t} + y_{qg,t} + c_{qg,t} + e_{qg,t}, \quad \text{with } e_{qg,t} \sim N(0, \sigma_{qg}^2). \quad (5)$$

Here,  $y_{qg,t}$  is a (potentially time-varying) age-group effect,  $c_{qg,t}$  is the cyclical component, and  $\tilde{x}_{qg,t}$  is the average cohort effect of the age group,

$$\tilde{x}_{qg,t} = \frac{1}{\#A_g} \sum_{a \in A_g} x_{qa,t} \quad (6)$$

with fixed cohort effects  $x_{ag,t}$ .<sup>7</sup> Let  $n_A \equiv \sum_{g=1}^{n_G} \#A_g$  denote the number of single-year ages in the sample (equivalently, the number of distinct birth cohorts in any given gender-education group).

### 3.2.2 Trend component

We model two persistent effects of the group-specific rates—cohort and age effect, which together constitute the trend. Age effects represent the influence of an individual’s current age on labor market outcomes (holding cohort constant), while cohort effects capture the lasting influences specific to the birth cohort (holding age constant).

Cohort effects are specified as

$$\begin{aligned} x_{q,a_{\min},t} &= x_{q,a_{\min},t-1} + \varepsilon_{qx,t}, & \text{with } \varepsilon_{qx,t} &\sim N(0, \sigma_{qx}^2), \\ x_{q,a,t} &= x_{q,a-1,t-1}, & \text{for } a &> a_{\min}, \end{aligned}$$

---

<sup>7</sup>For simplicity, we assume equal weights within an age group. For observed positive population growth, younger cohorts in an age group should receive somewhat higher weights, but the deviations from equal weights are small.

where  $a_{\min} = \min A_1$ . The first equation governs the cohort effect for the entering (youngest) cohort at time  $t$ , which evolves as a random walk. The second equation implies that a cohort's effect remains fixed as the cohort ages.

Age-group effects are specified as

$$\begin{aligned} y_{q1,t} &= y_{q1,t-1}, \\ y_{qg,t} &= y_{qg,t-1} + \varepsilon_{qy,g,t}, \quad \text{with } \varepsilon_{qy,g,t} \sim N(0, \sigma_{qy,g}^2) \text{ for } g > 1. \end{aligned}$$

That is, the age effect for the first age group is time-invariant, while the remaining age effects follow random walks.

We define the trend of outcome  $q$  for age group  $g$  as the sum of the cohort and age effects:

$$\bar{q}_{g,t} = \tilde{x}_{qg,t} + y_{qg,t}. \quad (7)$$

The assumption on the first age group implies that the posterior covariance matrix of the Kalman filter remains bounded as more observations become available. In Appendix A, we show that, given the assumptions on the cyclical component below, age and cohort effects are identified once we normalize the level of initial entering cohorts,

$$x_{q1,1} = 0. \quad (8)$$

### 3.2.3 Cyclical component

Cyclical movements are driven by a latent unemployment cycle factor  $z_{u,t}$  that follows an AR(1):

$$z_{u,t} = \rho z_{u,t-1} + \varepsilon_{uz,t}, \quad \text{with } \varepsilon_{uz,t} \sim N(0, \sigma_{uz}^2).$$

We allow the cyclical factor of LFP to respond to the unemployment cycle contemporaneously and with lags,

$$z_{l,t} = \sum_{s=0}^{nz} \phi_s z_{u,t-s} + \varepsilon_{lz,t}, \quad \text{with } \varepsilon_{lz,t} \sim N(0, \sigma_{lz}^2).$$

All innovations  $\varepsilon_{\cdot,t}$  and measurement errors  $e_{\cdot,t}$  are mutually independent across time, across outcomes  $q$ , and across age groups.

The cyclical components in (5) are then

$$c_{ug,t} = \gamma_{ug} z_{u,t},$$

$$c_{lg,t} = \gamma_{lg} z_{l,t}.$$

The loadings  $\gamma_{qg}$  capture systematic differences across age groups in the sensitivity to the common cyclical factor. We normalize the scale of the cyclical factors by setting the loading of the first age group to one:

$$\gamma_{u1} \equiv 1 \quad \text{and} \quad \gamma_{\ell 1} \equiv 1.$$

The first normalization leaves  $\sigma_{uz}^2$  unrestricted, and the second leaves  $\sum_{s=0}^{n_Z} \phi_s$  and  $\sigma_{\ell z}^2$  unrestricted.

Given smoothed estimates of the cohort and age effects, we construct  $\bar{u}_{g,t}^i$  and  $\bar{\ell}_{g,t}^i$  using (7) and then compute aggregate trends using (3) and (4).

### 3.3 Decomposing changes in aggregate rates

We write changes in the aggregate LFP rate and unemployment rate as arising from (i) changes in the age–gender distribution, (ii) changes in educational composition conditional on age and gender, and (iii) changes in group-specific LFP and unemployment rates.

Equations (1) and (2) define mappings from  $(p_{g,t}^i, \ell_{g,t}^i, u_{g,t}^i)$  to the aggregate LFP rate and unemployment rate. For individuals aged 25 and older, we decompose each group’s population share into an age-gender marginal and an education share conditional on age and gender.<sup>8</sup>

Define the age-gender marginal for group  $(s, g)$  as

$$p_{g,t}^s = \sum_e p_{g,t}^{(s,e)}, \tag{9}$$

and the conditional education share within  $(s, g)$  as

$$p_{e|g,t}^s = \frac{p_{g,t}^{(s,e)}}{p_{g,t}^s}. \tag{10}$$

Then  $p_{g,t}^{(s,e)} = p_{g,t}^s p_{e|g,t}^s$ .

Using this factorization, the aggregate LFP rate can be written as

$$\ell_t = \sum_{s,g} p_{g,t}^s \sum_e p_{e|g,t}^s \ell_{g,t}^{(s,e)}. \tag{11}$$

The aggregate unemployment rate can be written as

$$u_t = \frac{\sum_{s,g} p_{g,t}^s \sum_e p_{e|g,t}^s \ell_{g,t}^{(s,e)} u_{g,t}^{(s,e)}}{\sum_{s,g} p_{g,t}^s \sum_e p_{e|g,t}^s \ell_{g,t}^{(s,e)}}. \tag{12}$$

---

<sup>8</sup>Since individuals younger than 25 are not differentiated by education, this procedure is not applied to them.

Let  $\Delta x_t \equiv x_t - x_{t-1}$ . We compute first-order (linear) approximations for the changes in aggregate rates:

$$\Delta \ell_t = \sum_{s,g} \frac{\partial \ell_t}{\partial p_{g,t}^s} \Delta p_{g,t}^s + \sum_{s,g,e} \frac{\partial \ell_t}{\partial p_{e|g,t}^s} \Delta p_{e|g,t}^s + \sum_{s,g,e} \frac{\partial \ell_t}{\partial \ell_{g,t}^{(s,e)}} \Delta \ell_{g,t}^{(s,e)}, \quad (13)$$

$$\Delta u_t = \sum_{s,g} \frac{\partial u_t}{\partial p_{g,t}^s} \Delta p_{g,t}^s + \sum_{s,g,e} \frac{\partial u_t}{\partial p_{e|g,t}^s} \Delta p_{e|g,t}^s + \sum_{s,g,e} \frac{\partial u_t}{\partial \ell_{g,t}^{(s,e)}} \Delta \ell_{g,t}^{(s,e)} + \sum_{s,g,e} \frac{\partial u_t}{\partial u_{g,t}^{(s,e)}} \Delta u_{g,t}^{(s,e)}. \quad (14)$$

The first term captures the contribution of changes in the age–gender distribution, the second the contribution of changes in educational composition conditional on age and gender, the third the contribution of changes in group LFP rates, and the fourth (for unemployment) the contribution of changes in group unemployment rates.<sup>9</sup>

Cumulative changes satisfy

$$\ell_t - \ell_{t_0} = \sum_{\tau=t_0+1}^t \Delta \ell_{\tau}, \quad u_t - u_{t_0} = \sum_{\tau=t_0+1}^t \Delta u_{\tau}, \quad (15)$$

with an analogous decomposition of the sources of change. We apply this decomposition to either the observed aggregates or the trend aggregates.

## 4 Data and Estimation

In this section, we broadly describe our data and estimation method. For more detailed descriptions, see the online Technical Appendix, Hornstein and Kudlyak (2026).

### 4.1 Data

The data come from the monthly Basic CPS files (January 1976 to December 2024), obtained from IPUMS-CPS (see Flood et al. (2022)). Using the CPS labor force status variable, we classify each civilian (age 16+) as employed, unemployed, or out of the labor force (OLF). We aggregate the individual microdata into age–gender–education cells using the composite weights such that group aggregates by age, gender, and education are comparable with the published Bureau of Labor Statistics (BLS) data.<sup>10</sup> For each cell, we construct the unemployment rate, the LFP rate, and the population share, and calculate annual averages of the raw monthly data.

<sup>9</sup>Derivatives are calculated locally, year by year. The linear approximation works well, and unlike in standard shift-share decompositions, second-order terms are negligible.

<sup>10</sup>The composite weights are constructed using a raking procedure such that individual observations add to group totals at the state and national level, Census (2019), page 77, and are available starting 1998, Census (2019), p.35. For years before 1998, we use final weights.

We define 10 gender–education categories: (i) Youth (ages 16–24) for men and women (no education distinction), and (ii) Adults (25+) for men and women, each split into four education levels (less than high school, high school, some college, college or above). Within each category, we further divide by age brackets: 16–19 and 20–24 for youth; 25–34, 35–44, 45–54, 55–64, and 65+ for adults. In total, this yields 44 distinct age–gender–education groups.

#### 4.1.1 Data adjustments

The CPS regularly adjusts the population controls based on information from the Census, especially after a new decennial census has been processed. These population controls are not adjusted backwards in time, which can lead to noticeable discrete jumps in published population numbers for the demographic groups we consider. In the past, the BLS did publish research series which smoothed the population series between these decennial adjustments, e.g. Di Natale (2003). These research series spread out the discrete adjustment of the month when the new population controls are introduced over the months since the last time population controls were adjusted.

Because CPS population controls cause discrete jumps at adjustment points, we apply smoothing adjustments (similar to BLS research series) to avoid artificial jumps in group population counts. We use the same procedure to adjust our monthly population series for decennial adjustments and some large annual adjustments.<sup>11</sup>

Education classification is by grade before 1992 as described in Kominski and Siegel (1993). The 1992 redefinition increases measured educational attainment.<sup>12</sup> For men and women of all age groups, it mainly reduces the share of those with a high school education and increases the share of those with some college. We rescale the pre-1992 monthly shares to match the education shares by degree in 1992. This procedure eliminates discrete jumps in the series, but does not affect trends.

## 4.2 Estimation with Gibbs sampler

We estimate the dynamic factor model described in Section 3.2 using a Bayesian approach. Specifically, we implement a Markov Chain Monte Carlo Gibbs sampling procedure to draw

---

<sup>11</sup>For various age-gender categories, especially males and females less than 25 years old, the CPS increased population controls noticeably in January 2023 and then reduced them again in January 2024. We remove these population ‘bumps’ by smoothing the population series from December 2022 to January 2024. After this pre-filter, we adjust population controls in the following years: 1980, 1990, 2000, 2003, 2012, 2022, and 2025.

<sup>12</sup>We use the IPUMS EDUC variable to classify individuals by education. EDUC is related to HI-GRADES before 1992 and EDUC99 from 1992 on. The description of education variables is available at <https://www.bls.gov/cps/definitions.htm#education>.

the latent state and parameters iteratively. We estimate the model separately for young men and women, not differentiated by education, and for each gender-education group for individuals aged 25 and older.

The state-space model of a demographic group in matrix representation is

$$Y_t = H'X_t + W_t \text{ with } W_t \sim \mathcal{N}(0, R) \quad (16)$$

$$X_t = FX_{t-1} + V_t \text{ with } V_t \sim \mathcal{N}(0, Q) \quad (17)$$

where the observation and state vectors are

$$Y_t = [u_t^T, \ell_t^T]^T \quad (18)$$

$$X_t = [x_{u,t}^T, y_{u,t}^T, x_{\ell,t}^T, y_{\ell,t}^T, z_t^T]^T, \quad (19)$$

and  $(H, F, Q, R)$  are functions of the model parameters. The model parameters are  $\theta = (\gamma_q, \rho, \phi, \sigma_{qq}, \sigma_{qx}, \sigma_{qy,g}, \sigma_{qz})$  for  $q = u, \ell$  and  $g = 1, \dots, n_G$ . We normalize loading factors,  $\gamma_{u1} = \gamma_{\ell 1} = 1$ , and initial entering cohorts,  $x_{q1,1} = 0$ , as noted above. We use three lags for the LFP cyclical factor,  $n_Z = 3$ . We set the prior distribution for the variances as inverse gamma distributions with a diffuse prior.

The data are denoted by  $Y_{1:T} = (Y_1, \dots, Y_T)$ , the latent state by  $\zeta_{1:T} = (\zeta_1, \dots, \zeta_T)$  with  $\zeta_t = \{X_t, V_t, W_t\}$  and we are estimating  $\Theta = [\zeta_{1:T}, \theta]$ . Bayes' rule implies that

$$\Pr(\Theta|Y_{1:T}) = \Pr(Y_{1:T}|\Theta) \Pr(\Theta) / \Pr(Y_{1:T}),$$

where  $\Pr(\Theta)$  is the prior distribution,  $\Pr(Y_{1:T}|\Theta)$  is the likelihood,  $\Pr(Y_{1:T})$  is the marginal likelihood, and  $\Pr(\Theta|Y_{1:T})$  is the posterior distribution we are interested in estimating.

We estimate the posterior distribution

$$\Pr(\zeta_{1:T}, \theta|Y_{1:T}),$$

by iterated Gibbs sampling. For each iteration, the first Gibbs step draws  $\zeta_{1:T}|\theta, Y_{1:T}$ , and the second Gibbs step draws  $\theta|\zeta_{1:T}, Y_{1:T}$ . The initialization of the parameters is specific to the labor force status series. The estimated parameter distributions for each demographic group are in the online Technical Appendix, Hornstein and Kudlyak (2026).

## 5 Group Unemployment and LFP Trends

Over time, unemployment rate trends have declined modestly across groups. This downward trend in group unemployment rates is mostly driven by cohort-specific improvements, whereas age effects have remained fairly stable. LFP rates, however, display pronounced

trends that differ across genders. Men’s LFP trends have mostly declined across almost all groups over the last few decades, whereas women’s LFP trends rose significantly through the 1990s—contributing to the well-known rise in female LFP—before plateauing. These patterns align with cohort effects. For men, the trend LFP declines for successive cohorts, whereas for women, successive cohorts participate more.

In this section, we first briefly summarize the cyclical behavior of unemployment and LFP rates by gender-education group and age, and then describe the trend components and their decomposition into age and cohort effects. A more detailed description is provided in the online Technical Appendix, Hornstein and Kudlyak (2026). Our characterization of trends confirms the broad outlines of Section 2, and it provides additional detail on the role of cohort effects.<sup>13</sup>

## 5.1 Cyclical components

The cyclical components of unemployment in the different demographic groups are highly synchronized across groups. In contrast, within demographic groups, the cyclical component of LFP tends to be weakly negatively correlated with the corresponding unemployment cycle.

For each demographic group, its unemployment cyclical factor is highly persistent and strongly correlated with the aggregate unemployment rate, Figure 1. Across demographic groups, the cyclical volatility declines with education and is lower for women than for men. Within each demographic group, the sensitivity of unemployment to the cyclical factor tends to fall with age.

The LFP cyclical factor tends to be a moving average of the current and lagged cyclical factor for unemployment for all demographic groups, Figure 2. For youths 16–24, unemployment and LFP cyclical components are strongly negatively correlated and roughly coincident; for adults 25+, the negative relation is weaker and often strongest at a one-year lag. For LFP, the youngest group in each cell has the strongest cyclical response, while older groups—especially men 65+ and women 55+ and college-educated women—show much weaker or statistically insignificant responses.

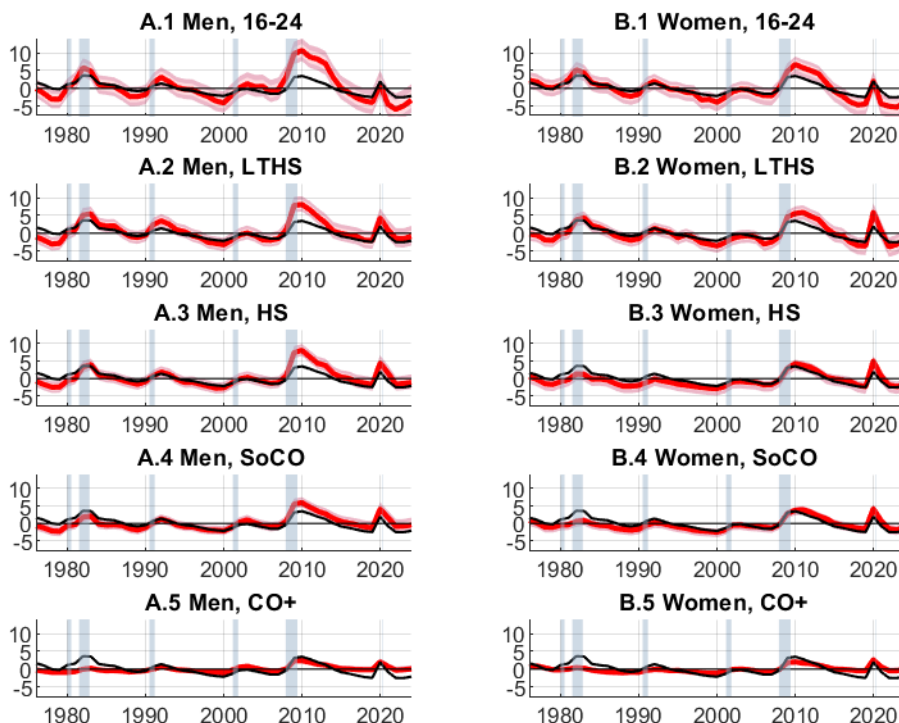
## 5.2 Trend components: age and cohort effects

Trend unemployment rates are lower for older and more educated workers, and they have drifted down only modestly over time, with earlier declines for men and for college-educated women. Estimated age effects of unemployment are relatively stable, so these slow trend

---

<sup>13</sup>We provide a more detailed description of cycle and trend in the online Technical Appendix, Hornstein and Kudlyak (2026).

Figure 1: Cyclical factor for unemployment rates

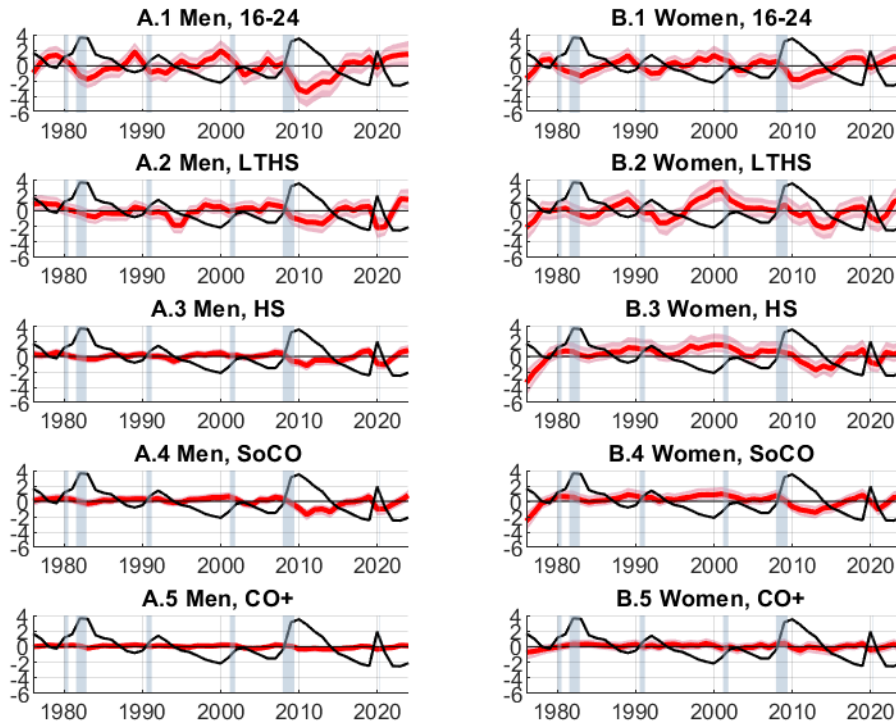


Note: Estimated cyclical effect by gender for the age group 16-24 (not differentiated by education), and for those 25 or older by education groups less than high school (LTHS), high school (HS), some college (SoCO), and college or above (CO+). The left column (A) is for men, and the right column (B) is for women. Solid red lines denote the estimated median cyclical effect, and the light (dark) shaded areas denote the 66% (90%) coverage area. The thin black line is the demeaned aggregate unemployment rate.

declines largely reflect cohort effects: newer cohorts tend to have slightly lower unemployment fixed effects, especially for men and for those without a high school degree.

Trend LFP rates tend to increase with education and with age up to about 35–44, after which LFP declines. Men’s participation trends have fallen over time, whereas women’s rose until the mid-1990s and then declined (except for more educated women whose LFP flattened out later). Exceptions to these trends are men aged 35-44 without a high school degree and men older than 65 without a college degree, whose LFP trends have increased recently. Along these trends, persistent gaps open up between age and trend lines, indicating important cohort effects that propagate across age groups. Cohort effects in LFP are sizable and precisely estimated: they generally decline for men and younger workers, increase, then flatten for older women, and remain high only for women with at least some college.

Figure 2: Cyclical factor for LFP rates



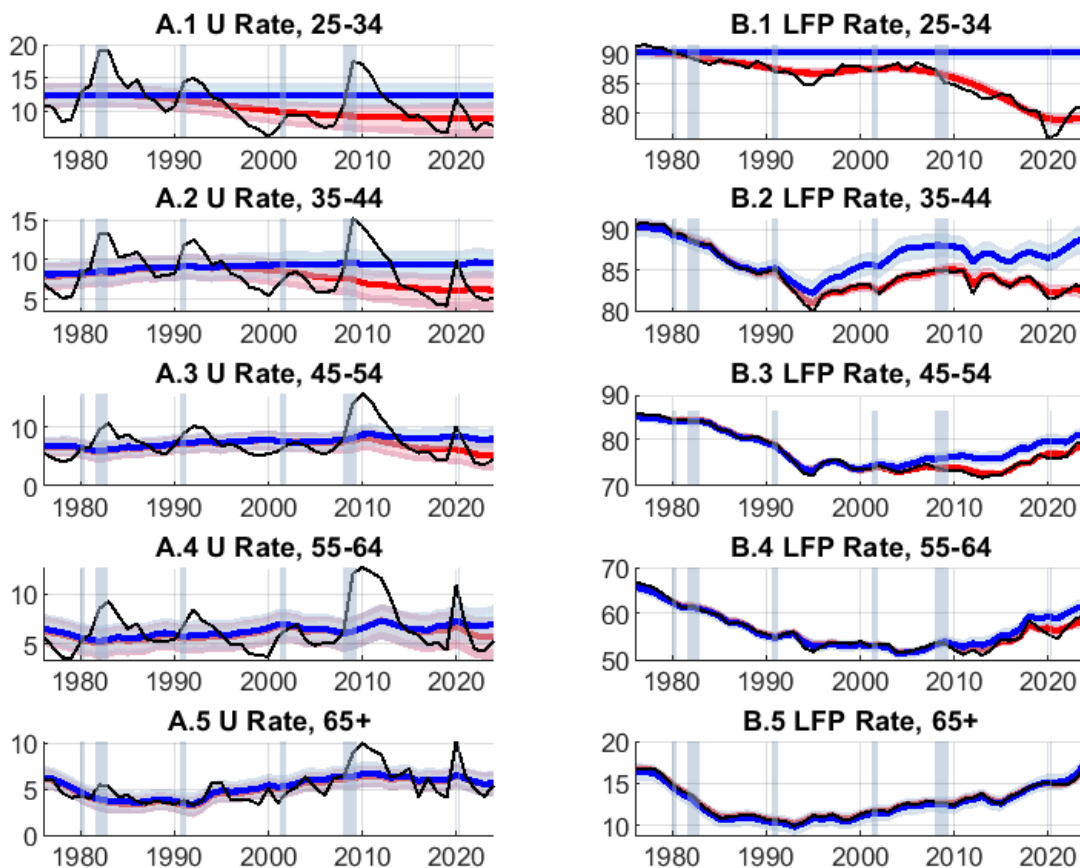
Note: See notes for Figure 1

This suggests that changing cohort behavior, more than changing age profiles, is central for understanding long-run unemployment and LFP trends.

For illustration, we plot the trend decomposition of unemployment and LFP trends for men 25 years and older, with less than a high-school education, in Figure 3. The red and blue lines represent the estimated trend and age effect, respectively. Since, by assumption, the age effect for the youngest group is fixed, the blue lines for those aged 25-34 are flat. Among the 25-34 year olds, there is a gap opening up between the age-effect and trend, starting in 1990 for unemployment, and in 1980 for LFP. This gap reflects the average cohort effect of the entering cohorts. Going down the panels, we see how the gap gets propagated over time in the older age groups with a lag. Unemployment cohort effects for other demographic groups tend to be smaller.

Cohort (entry) effects for unemployment have trended down, for men and women of all education groups, Figure 4. Overall, these declines are gradual, starting later for women

Figure 3: Age effect and trend, Men, 25 years and older, with less than a high-school education

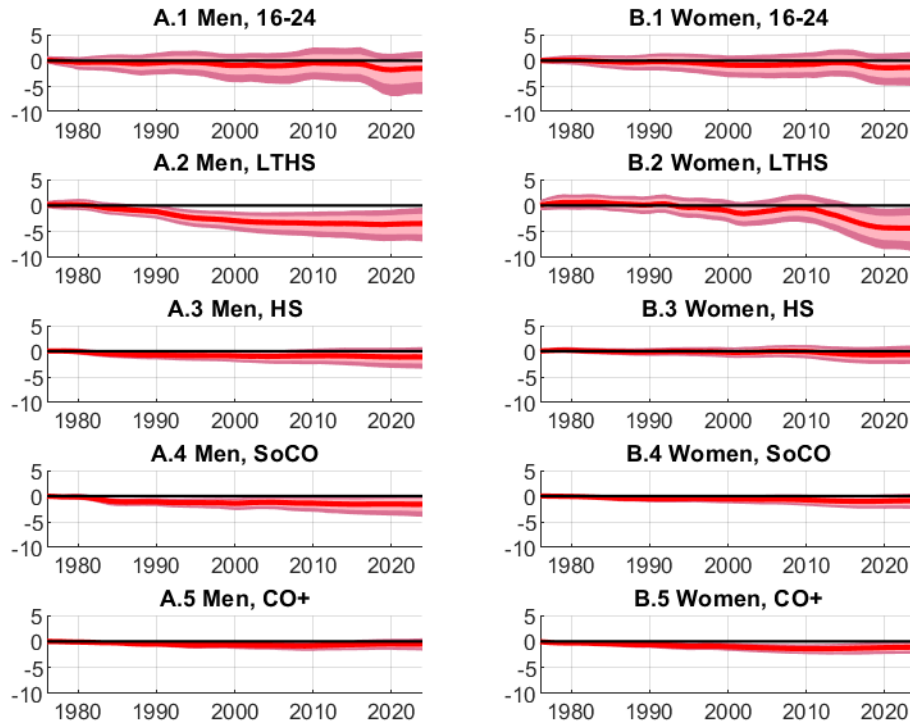


Note: Unemployment rate (A) and LFP rate (B) for age groups 25-34, 35-44, 45-54, 55-64, and 65+. Solid blue lines denote the estimated median age effect, solid red lines denote the estimated median trend, the sum of age and cohort effect, and the corresponding light (dark) shaded areas denote the 66% (90%) coverage area. The thin black lines are the actual unemployment rate and LFP rate, respectively.

(around the 1990s) than for men (1980s), and tend to be modest. The largest decline occurs for those with less than a high school education.

Changes in average LFP cohort effects for different age groups reflect the estimated fixed effects of entering cohorts, Figure 5. For men and women younger than 25 years and all men 25 and older, cohort effects decline over time; the cohort effects for young women decline less than for young men, and they decline less for more educated men. For women older than 25, cohort effects increased until the 1990s, but have remained flat or declined afterward. Cohort effects have remained consistently high only for women with at least some college.

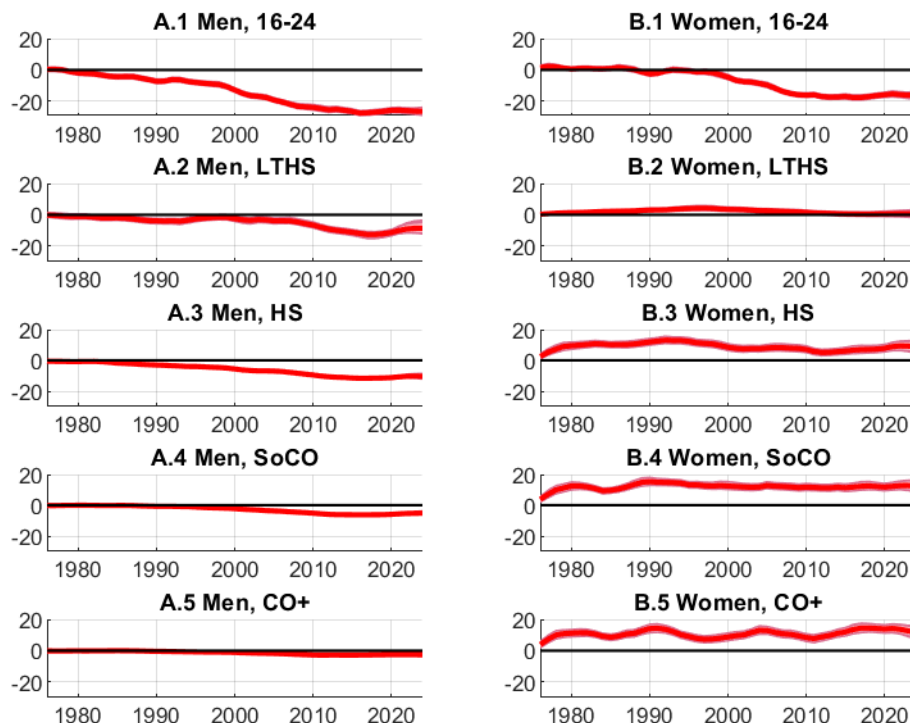
Figure 4: Fixed effect of entering cohorts for unemployment rates



Note: Estimate of median fixed effect of entering cohorts for men (A) and women (B) for the age group 16-24 not differentiated by education and for education groups with less than a high-school degree (LTHS), a high-school degree (HS), some college (SoCO), The light (dark) shaded areas denote the 66% (90%) coverage area.

The cumulative changes in median LFP cohort effects tend to be significant, i.e., the coverage areas do not include zero, unlike the estimated unemployment cohort effects, Figure 4.

Figure 5: Fixed effect of entering cohorts for LFP rates



Note: See notes to Figure 4.

In summary, demographic group analysis reveals that long-run trends differ notably by cohort and gender, even as cyclical fluctuations are broadly shared across groups. We next examine how these group trends aggregate up to the overall labor force and unemployment rates.

## 6 Trends for Aggregate Rates

We now examine aggregate LFP and unemployment trends and how they have been shaped by demographic composition changes in age and education, and changes in group-specific trends in LFP and unemployment. We decompose changes in aggregate rates into contributions from demographics and group-specific trends using the procedure described in Section 3.3.

For the aggregate LFP rate, unemployment rate, and employment-to-population ratio, we plot their estimated trends and decompositions in Figure 6. In the first column, we plot

aggregated data and their estimated trend. In the second column, we plot the cumulative contributions of changes in the age distribution, education shares, and estimated trends in gender-education contingent group trends for the LFP and unemployment rates. Finally, in the third column, we plot the cumulative contributions of changes in cohort and age effects to changes in group trends.

## 6.1 Trend decomposition for aggregate rates

The trend aggregate LFP rate follows a well-known hump-shaped trajectory, rising from the 1970s and peaking in the late 1990s before gradually declining thereafter, Figure 6a. After the mid-1990s, the actual LFP rate deviated notably and persistently from the estimated trend—staying above trend in the late-1990s boom and below trend during the post-2008 recovery.

From 1976 to the mid-1990s, the trend LFP rate rose by about 4.1 pp, Figure 6b. Our decomposition attributes over half of this rise to educational upgrading of the workforce (+2.2 pp), and another portion to within-group LFP trend increases (+1.2 pp), with changes in the age composition accounting for the remainder (+0.7 pp). Conversely, the mid-1990s to 2024 saw a 4.4 pp decline in trend LFP, driven predominantly by population aging (−4.3 pp, as the baby boom moved into older ages) and compounded by a drop in within-group LFP trends (−2.3 pp).

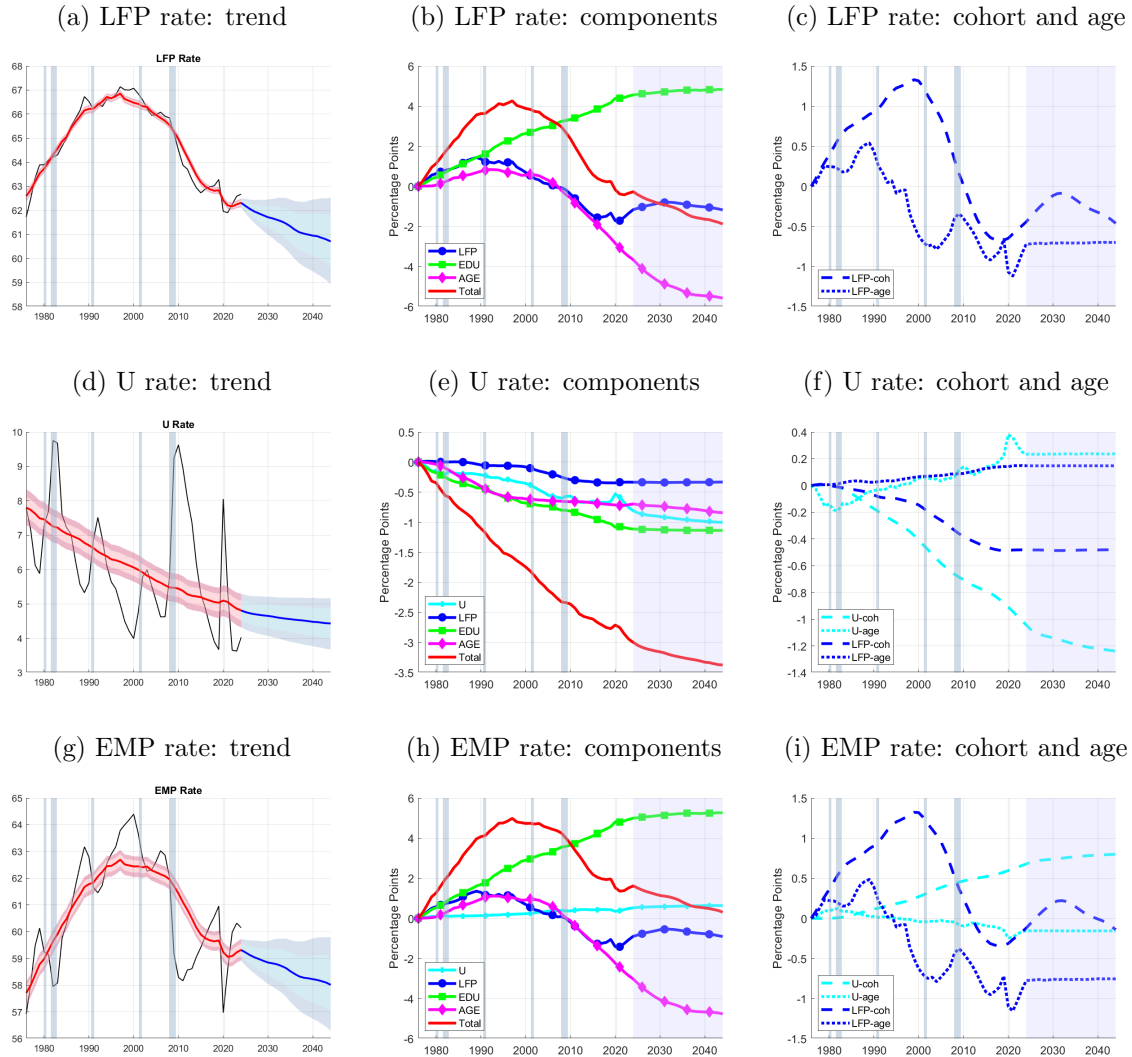
Notably, the hump in the overall LFP trend is primarily due to cohort effects: earlier cohorts raised participation, then later cohorts leveled off—producing a hump-shaped net cohort effect, Figure 6c. Age effects also played a role, as starting in the mid-1990s, some age-specific participation rates (especially at older ages) began to decline, reinforcing the downturn.

The trend aggregate unemployment rate declines smoothly from about 8% in 1976 to 5% by 2024, Figure 6d. Actual unemployment deviates substantially from this trend—noticeably below trend in the late-1990s boom and above trend after 2008 (i.e., slow recovery). Generally, when the LFP rate is above trend, the unemployment rate tends to be below trend, and vice versa.

The declining unemployment trend has been driven in roughly equal measure by changing workforce composition and declining group-level unemployment and participation trends, Figure 6e. Improvements in educational composition provided the largest contribution to lowering the unemployment trend, while changes in groups' own LFP trends contributed the least.

Finally, cohort effects dominate the decline in group trends (for both unemployment and LFP), as shown in 6f. In other words, the downward pull from newer cohorts having

Figure 6: Aggregate trends and projections



Note: The panels in the left column plot the aggregate LFS rates (black lines) and their trends based on the population share-weighted median estimates of the group trends (red lines), with light (dark) red shaded 66% (90%) coverage areas for the sample 1976-2024. The dark blue lines and shaded areas represent the median forecasts and coverage areas for the period 2025-2044. The panels in the middle column plot the cumulative contributions to changes in the trend (Total) coming from changes in population shares by age (AGE), education shares (EDU), and trend LFS rates (LFP and U). The panels in the right column split the contributions of trend LFS rates into age (dotted) and cohort (dashed) effects. The lines in the right shaded areas represent the forecasted contributions.

structurally lower unemployment is the primary factor behind the group trend component of aggregate unemployment declines.

The trend employment-to-population ratio is dominated by the trend LFP rate, whereas the trend deviations are dominated by the unemployment rate deviations, Figure 6g. The

contributions of age and education distributions and group trend LFP rates to changes in the trend employment rate mirror their contributions to the aggregate trend LFP rate, with a minor impact from changes in group trend unemployment rates, Figure 6h. Finally, the hump-shaped path of cohort effects in group LFP trends similarly dominates the shape of the E/P trend (Figure 6i), since employment trend changes are overwhelmingly driven by participation trends.

## 6.2 Trend decomposition for gender aggregates

Overall, our decomposition of the trends in the aggregate LFP rate, unemployment rate, and employment-to-population ratio points to the important role of composition effects arising from changes in the age and education distribution: they account for one-half to three-fourths of trend changes. However, the contribution of the compositional changes in the aggregate LFP rate trend masks diverging trends among men and women. From 1976 to 2024, men’s trend LFP rate plummeted by roughly 10 pp, while women’s trend LFP rate rose by about 10 pp through 2000 and then receded by 2 pp thereafter. This striking difference is illustrated in Figure 7a and 7b.

Interestingly, compositional shifts (age and education) contributed similarly to men’s and women’s LFP trends – roughly half of the change for each. The big difference is in the within-group (cohort) LFP trends: men’s own participation rates fell by 9 pp on average across cohorts, whereas women’s rose by 7 pp. Figures 7c and 7d show that men’s group trends were down, while women’s were up.

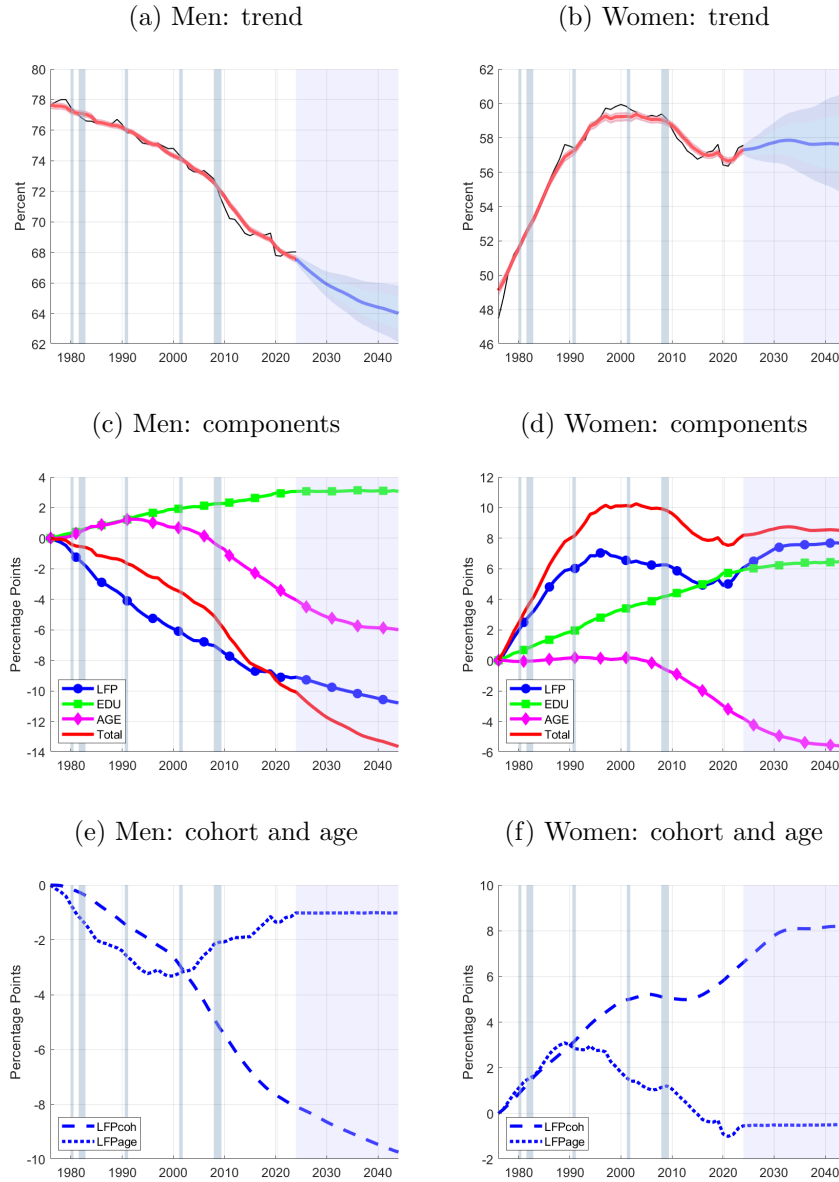
Since age and education effects offset each other, the changes in group trend LFP rates account for almost all of the net change in men’s and women’s trend LFP rates. Finally, we again see the important role of cohort effects in the determination of group trend LFP rates, blue lines in Figures 7e and 7f: the net gender gap in trend LFP changes is almost entirely due to the different cohort trends (declining for men, rising for women).<sup>14</sup>

We see much less disparity in the trend unemployment rates for men and women, Figure 8. For both genders, the trend unemployment rates decline, with qualitatively similar contributions from changes in the age and education distribution, and group trend LFP and unemployment rates. The only difference is the relatively larger decline in group trend unemployment rates for men. Again, cohort effects tend to dominate changes in group trend LFP and unemployment rates.

---

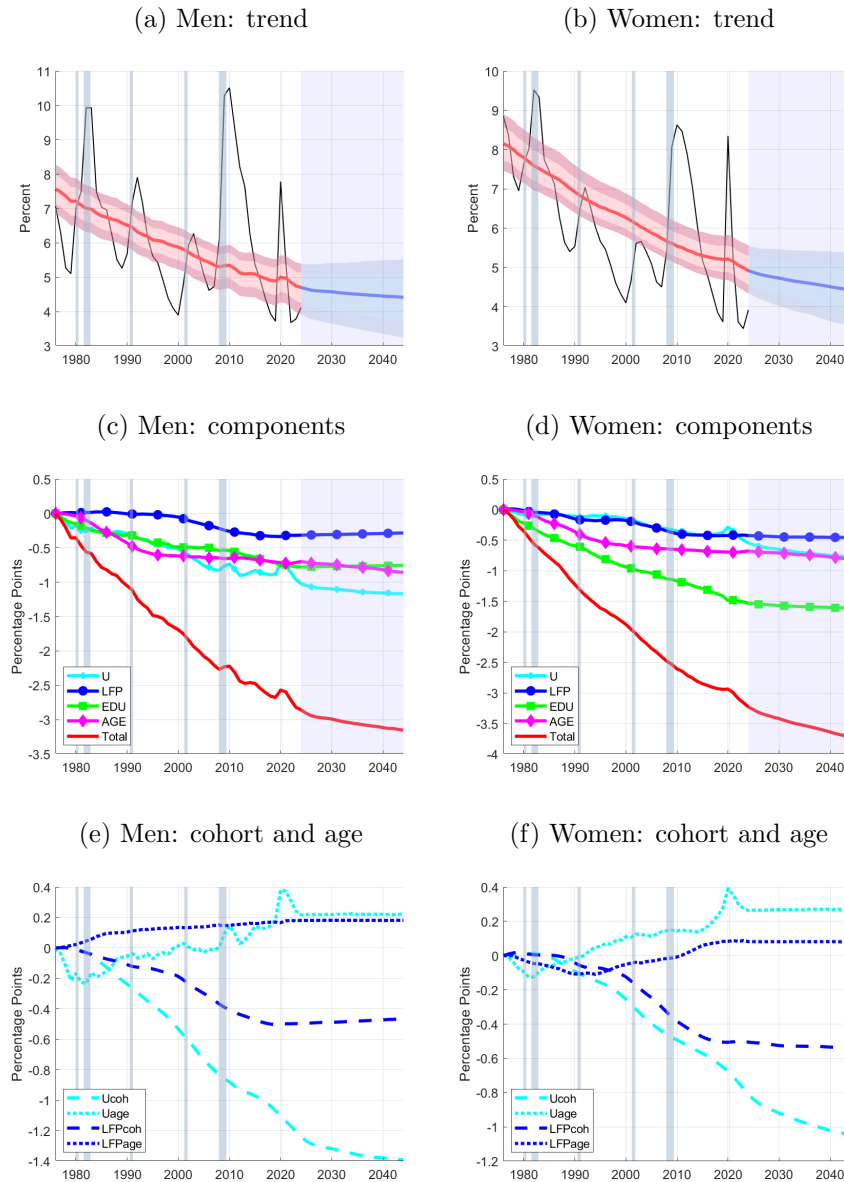
<sup>14</sup>Not shown in figures 7e and 7e are the relative contributions of changes in LFP rates from those 16-24 years old and those 25+. Here, it is worth noting that the declining LFP rates of young males less than 25 years old account for roughly one-third of the declining contribution coming from men’s trend LFP rates, whereas young women have much less of an impact on women’s trend LFP rates.

Figure 7: Aggregate LFP rate trends and projections for men and women



Note: The top panels plot the LFP rates (black lines) and their trends based on the population share-weighted median estimates of the group trends (red lines), with light (dark) shaded 66% (90%) coverage areas for the sample 1976-2024. The dark blue lines and shaded areas represent the median forecasts and coverage areas for the period 2025-2044. The middle panels plot the cumulative contributions to changes in the trend (Total) coming from changes in population shares by age (AGE), education shares (EDU), and trend LFP rates (LFP). The bottom panels split the contributions of trend LFP rates into age (dotted) and cohort (dashed) effects.

Figure 8: Aggregate unemployment rate trends and projections for men and women



Note: The top panels plot the unemployment rates (black lines) and their trends based on the population share-weighted median estimates of the group trends (red lines), with light (dark) shaded 66% (90%) coverage areas for the sample 1976-2024. The dark blue lines and shaded areas represent the median forecasts and coverage areas for the period 2025-2044. The middle panels plot the cumulative contributions to changes in the trend (Total) coming from changes in population shares by age (AGE), education shares (EDU), and trend LFP (LFP) and unemployment (U) rates. The bottom panels split the contributions of trend LFP and U rates into age (dotted) and cohort (dashed) effects.

## 7 Forecasts for Aggregate Rates

In this section, we discuss twenty-year-ahead forecasts for the aggregate LFP and unemployment rate trends. Forecasting the aggregate trends requires (1) a forecast of the group LFP

and unemployment rate trends, and (2) the forecast of the demographic composition of the population. Since the components of the groups are random walks, the best forecast of the future value is the value of the series at the end of the sample. We thus use the cohort and age effects to forecast the group LFP and unemployment rate trends.

We take the population composition projections by gender and age from the U.S. Census Bureau (2025). We forecast future educational attainment using an estimated cohort state-space model for age-gender conditional education shares with time-varying cohort effects.

We project that over the next twenty years, the LFP rate will decline another 1.5 pp from its 2024 value of 62.3 percent, and the unemployment rate will decline by 0.4 pp from its current value of 4.7 percent. These projections are subject to considerable uncertainty.

## 7.1 A factor model of education shares

Educational attainment is defined for individuals aged 25 and older. For each gender  $s$ , education category  $e$ , and age group  $g$  we define the *conditional* education shares

$$m_{g,t}^{se} \equiv \frac{p_{g,t}^{se}}{p_{g,t}^s},$$

that is, the fraction of the age-gender group  $(s, g)$  with education level  $e$ .

We estimate a cohort-based state-space model for  $\{m_{g,t}^{se}\}$  with three key differences relative to the unemployment-LFP model in Section 3.2. First, we do not include separate age effects: differences in education shares across age groups arise from cohort composition and (potentially) within-cohort evolution. Second, we do not impose that cohort effects are fixed over the life cycle. Instead, as cohorts age, their education shares can change, reflecting (for example) later-life educational upgrading or differential mortality across education groups (e.g., Aaronson and Sullivan (2001)). Finally, education shares are empirically very smooth, and we do not detect meaningful cyclical variation, so we omit an explicit cyclical component.

Fix  $(s, e)$  and suppress the superscript for ease of notation. For  $t = 1, \dots, T$  and age groups  $g = 1, \dots, n_G$ , we observe

$$m = \{m_{g,t} : t = 1, \dots, T, g = 1, \dots, n_G\}.$$

We assume that the observed education share for age group  $g$  is the average of latent single-year cohort shares plus measurement error:

$$m_{g,t} = \frac{1}{\#A_g} \sum_{a \in A_g} x_{a,t} + \varepsilon_{g,t}, \quad \varepsilon_{g,t} \sim N(0, \sigma_{mg}^2). \quad (20)$$

The latent state  $x_{a,t}$  can be interpreted as the education share of the cohort observed at single-year age  $a$  in year  $t$ . We assume a random-walk specification for the evolution of cohort shares,  $\{x_{a,t}\}$ , so that the best forecast of future education shares is their current value:

$$x_{1,t} = x_{1,t-1} + \varepsilon_{x1,t}, \text{ with } \varepsilon_{x1,t} \sim N(0, \sigma_{x1}^2), \quad (21)$$

$$x_{a,t} = x_{a-1,t-1} + \varepsilon_{xg,t}, \text{ with } \varepsilon_{xg,t} \sim N(0, \sigma_{xg}^2) \text{ for } a \in A_g \text{ and } g = 2, \dots, n_G. \quad (22)$$

We estimate the model separately for each education category  $e$  (and gender  $s$ ); that is, we do not impose the adding-up constraint that the four education shares sum to one at each  $(g, t)$ . When producing joint forecast paths, we therefore draw forecasts for the four education categories independently and then normalize the four draws so that the shares sum to one. The online Technical Appendix, Hornstein and Kudlyak (2026) provides a more detailed description of the model and its estimation.

## 7.2 Trends and forecasts for education shares

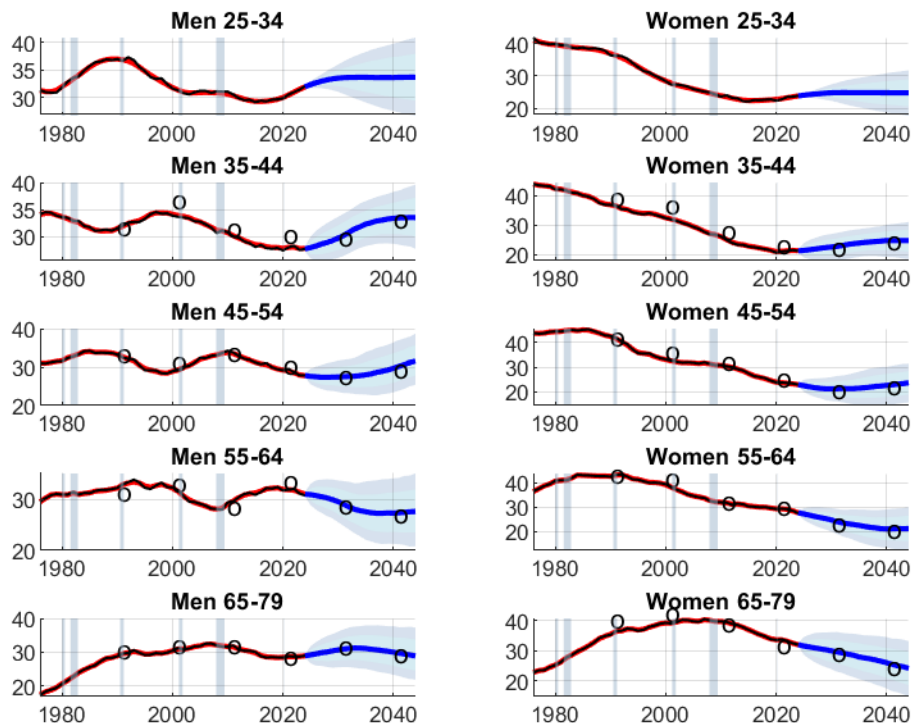
In Section 6, we showed that rising educational attainment has been an important driver of upward trends in the aggregate LFP rate and corresponding declines in the aggregate unemployment rate. Here, we examine in detail how the population’s educational composition has evolved. In particular, we highlight the role of vintage effects – cohort-specific attainment differences – in shaping the path of educational attainment over time. This discussion also sets the stage for Section 7.3, where we combine our models to forecast the aggregate LFP and unemployment trends.

The shares of the four education categories reveal a clear upward shift in educational attainment over the past several decades. Among those aged 25–34, the share with a high-school education or less has steadily decreased, while the share with some college or a college degree has increased. With few exceptions, these changes have been mostly monotonic. Moreover, education shares have evolved very smoothly: thus, the estimated trends closely track the actual series, and the confidence bands are narrow, indicating very little cyclical or random deviation. As an example, we plot the average education shares of men and women 25 years and older with a high school education across age groups. In Figure 9, the black lines are measured shares, the red lines are median estimates of the trend shares, and the red shaded areas are the coverage areas.

This example also illustrates the role of cohort effects. For the youngest age group, those 25–34 years old, the average share of women with a high-school education declines from the start of the sample, and the average share of men starts to decline around 1990. About ten years later, similar declines become evident among men and women aged 35–44, and twenty

years later among those aged 45–54. In other words, the share of each cohort is significantly lower than that of cohorts born a decade earlier, and this pattern “travels” upward through the age brackets over time. We observe analogous cohort-driven patterns for other education groups as well.

Figure 9: Share with a completed high school education, Actual, trend, and forecast

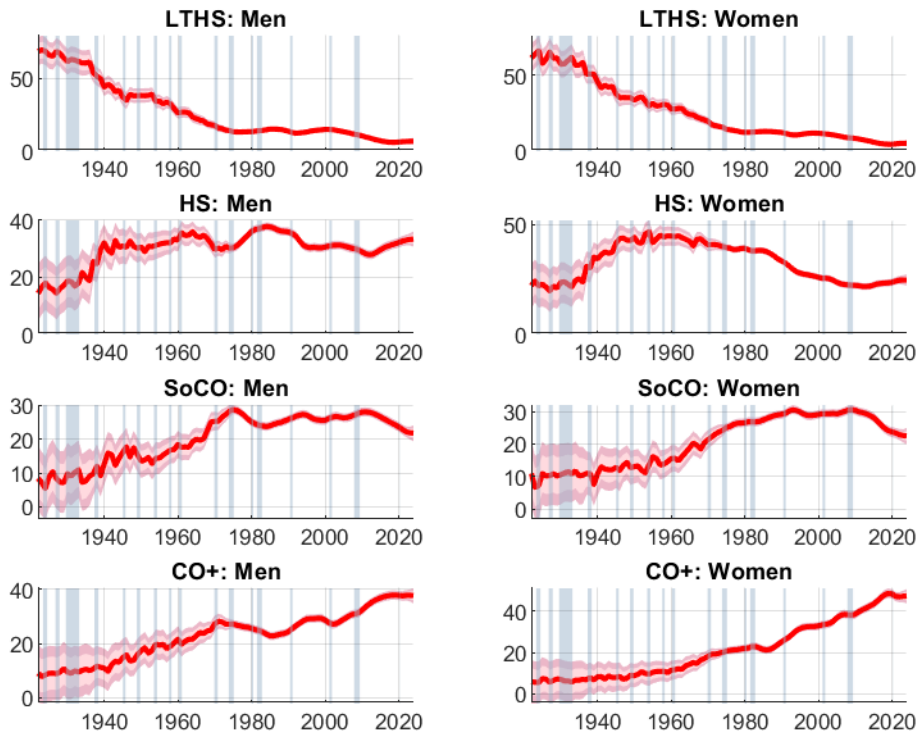


Note: Share of gender-age group with the indicated education. The thin black lines are the actual shares. The solid red lines denote the estimated median trend, and the corresponding light (dark) shaded areas denote the 66% (90%) coverage area. The solid blue lines denote the median forecast, and the corresponding light (dark) shaded areas denote the 66% (90%) coverage area of forecasts. The circles denote the average cohort effect of the preceding age group ten years prior.

From the observed trends across average age groups, the model infers the trends of the entering cohorts’ education shares, plotted in Figure 10. For men, the big gains in educational attainment mostly occurred before our sample period began in 1976—by then, younger male cohorts had already attained much higher levels of education on average than earlier generations. From 1976 onward, the fraction of men in new cohorts with less than a high school (LTHS) or only a high school (HS) education remained roughly constant at a low level. It is only around 2000 that we see another uptick in attainment for men: the shares of

new cohorts completing high school and college rose further, reducing the fraction without a diploma. By contrast, women’s educational attainment keeps improving throughout the sample period. The share of women in entering cohorts with LTHS or HS education steadily declined, while the shares with some college or a completed college degree increased – at least until the most recent few years, where the series plateaued.

Figure 10: Education share of entry cohort



Note: Shares of the entering gender cohort with the indicated education. Solid red lines denote the estimated median trend, and the corresponding light (dark) shaded areas denote the 66% (90%) coverage area. The thin black lines with circle markers are the actual shares.

Our model of education shares implies no predictable changes in education shares for entering and current cohorts. Nevertheless, past changes in educational attainment are embedded in the current cohort structure, implying that future changes in the observed education shares are predictable as those cohorts age. We illustrate this point by including 20-year projections of education shares by age group for men and women with a high-school education in Figure 9.<sup>15</sup> For this group, the cohort dynamics lead to some noticeable shifts.

<sup>15</sup>Our projections are based on Monte Carlo simulations and involve random draws on estimated parameters and the state space model.

For instance, the share of men in the 35–44 age group is expected to increase over the next twenty years. This rise occurs because the cohort of men currently age 25–34 has a higher high-school-educated share than the cohort currently age 35–44; as the younger cohort ages into the 35–44 bracket, it boosts that group’s HS share. On the other hand, for similar reasons, the share of college-educated men in the 35–44 year old age group is projected to decline as time progresses. These offsetting movements, driven entirely by the cohort vintage structure, mean that the overall projected changes in education shares are modest and reflect patterns already observable in recent data. Overall, the projections indicate modest changes in the share of individuals with different education levels

It is worth noting two points regarding our projections and trends. First, our education-share model treats each category’s share as an independent stochastic process, Section 7.1, and, as a consequence, independent simulations for each of the four education categories will not necessarily sum to 100% in a given year. We handle this by drawing the four education shares jointly (one random draw for an entire vector of shares) and then normalizing the outcomes to ensure they add up to one. The two are virtually identical, indicating that the normalization procedure does not alter the results in any meaningful way. We therefore use the normalized joint draws for all reported forecast results.

Second, we have constructed the trend of the aggregate rates (LFP, unemployment, and employment) from measured population share weighted trends of the corresponding rates for gender-age-education-based groups. Once we estimate the trend in education shares, one may ask why not incorporate these estimated trends in the calculation of aggregate trends. But as we have noted before, the deviations of estimated trends from actual education shares are small with narrow confidence intervals. Thus, the constructed aggregate trends and their confidence intervals are essentially the same, and we have decided not to replace actual education shares with their estimated trends in the calculation of aggregate trends.

### **7.3 Forecasts for aggregate LFP and unemployment rates**

Having projected the future demographic composition, we now forecast the *trend* components of the aggregate LFP rate and unemployment rate over the next twenty years. These projections combine (i) projected group-specific trend rates from the state-space model in Section 3.2 and (ii) projected population shares by age, gender, and education (from the U.S. Census Bureau (2025) age projections and our education-share model in Section 7.1). We then aggregate group trends using the formulas in Section 3.1. Note that our projections incorporate forecast and parameter uncertainty; they do not incorporate population forecast uncertainty.

In Figure 6 we plot the historical series (black), the estimated median trend (red), and the median forecast (blue) with uncertainty bands (shaded) for the aggregate unemployment rate, LFP rate, and employment-to-population ratio. In Figures 7 and Figures 8, we plot the same objects separately for men and women for the LFP and unemployment rates.

Over the next two decades, the aggregate trend LFP rate is forecast to decline by roughly 1.5 pp from its 2024 level of 62.3%, while the trend unemployment rate is projected to decline by about 0.4 pp from around 4.7%. The employment-to-population trend largely mirrors the LFP projection, reflecting the relatively small projected movements in the unemployment trend.

The gender split in Figure 7 highlights a source of heterogeneity for projected aggregate LFP rates: while for men the persistent decline of LFP rates across entering cohorts implies a negative projected impact from group-specific LFP rates, for women recent changes of LFP rates across entering cohorts imply a positive projected impact from group-specific LFP rates. Thus, men’s LFP rates are projected to decline, whereas women’s LFP rates are projected to stagnate. On the other hand, trend unemployment for both men and women edges down slightly and then remains near a historically low level, Figure 8.

Finally, interpretation of these projections should keep the trend specification in mind. Under our baseline model, cohort and age components are driftless random walks, so the conditional mean forecast of each cell’s trend is approximately flat. As a result, projected movements in aggregate trends are driven mostly by projected composition changes (aging and education upgrading) and cohort replacement, while the uncertainty bands widen with the forecast horizon.

Overall, the forecast suggests a gradual decline in trend participation and a low, slowly declining trend unemployment rate, consistent with the demographic momentum embedded in the current population structure. The forecast suggests that the U.S. labor force will continue to feel the weight of an aging population, which drags down participation rates, while enjoying the stabilizing benefits of a more highly educated workforce, which blunts some of that decline and keeps unemployment trends low.

## 8 Conclusions

In this paper, we estimate long-run trends for U.S. LFP and unemployment rates from 1976 to 2024, incorporating demographic heterogeneity. We model age-gender-education group unemployment and participation rates with a parsimonious state-space framework that separates a persistent trend component—cohort and age effects—from a cyclical component. The estimated group trends are then aggregated using observed population shares to form

the trends of aggregate LFP and unemployment rates. This structure allows for a transparent decomposition of changes in aggregate trends into three conceptually distinct forces: shifts in the age–gender distribution, shifts in educational attainment conditional on age and gender, and changes in within-group trend rates. Finally, combining the estimated group trends with external demographic projections and a cohort-based model of educational attainment, we generate medium-run forecasts of aggregate trend participation and unemployment.

The familiar hump-shaped pattern estimated for the aggregate LFP rate trend arises from the interaction of increasing educational attainment (though at a diminishing rate), population ageing (at an accelerating rate), and a hump-shaped contribution from estimated group-specific trends. The latter is driven by the interaction of monotonically declining men’s group trends and increasing and then flattening women’s group trends. The modest uniform decline of the estimated aggregate trend unemployment rate arises mainly from increased educational attainment and somewhat less but about equal contributions from population ageing and declining group-specific trend unemployment rates. The statistical model attributes a substantial part of the estimated changes in group trends and educational attainment to cohort effects. These cohort effects generate predictable movements in education shares and group-specific trends for LFP and unemployment, more so for women than for men. The statistical model predicts further declines of the LFP and unemployment rate due to population ageing, which is only partially offset by the predictable movements in educational attainment and group trends.

Overall, our approach offers a transparent accounting framework that can be updated as new data arrive and used in policy analysis of potential labor employment and labor market slack. But our framework also has obvious limitations. Because the estimated cohort and age effects are not linked to measured covariates, the model summarizes shifts in life-cycle patterns without attributing them to particular mechanisms, for example, trade, automation, schooling, health, policy, or family structure. Nevertheless, the framework can provide some perspective on proposed explanations for changes in trend participation and unemployment.

For example, Abraham and Kearney (2020), in their survey of potential explanations for the decline of the employment rate since 2000, point to automation and the ‘China shock’, that is, the substitution of imported goods from China for domestically manufactured goods, as the main suspects. One would expect that if both the China shock and automation lower the employment rate through job displacement, they would be associated with a simultaneous decline in the trend LFP rate and an increase in the trend unemployment rate. Indeed, Autor et al. (2015) find such a relationship for trade but not automation shocks. Yet our work suggests a simultaneous decline in group participation and unemployment rate trends. Furthermore, limiting attention to the overall participation decline after 2000 may also be

misguided, since one of the underlying drivers, group participation rates for men, has been declining since the late 1970s.

## References

- Aaronson, Daniel and Daniel Sullivan, “Growth in Worker Quality,” *Federal Reserve Bank of Chicago Economic Perspectives*, 2001, 25 (4), 53 – 74.
- , Luojia Hu, Arian Seifoddini, and Daniel G. Sullivan, “Changing Labor Force Decomposition and the Natural Rate of Unemployment,” *Federal Reserve Bank of Chicago, Chicago Fed Letter*, 2012, (338).
- Aaronson, Stephanie, Bruce Fallick, Andrew Figura, Jonathan Pingle, and William Wascher, “The Recent Decline in the Labor Force Participation Rate and Its Implications for Potential Labor Supply,” *Brookings Papers on Economic Activity*, 2006, pp. 69–134.
- , Tomaz Cajner, Bruce Fallick, Felix Galbis-Reig, Christopher Smith, and William Wascher, “Labor Force Participation: Recent Developments and Future Prospects,” *Brookings Papers on Economic Activity*, 2014, pp. 197–255.
- Abraham, Katharine G. and Melissa S. Kearney, “Explaining the Decline in the US Employment-to-Population Ratio: A Review of the Evidence,” *Journal of Economic Literature*, 2020, 58 (3), 585 – 643.
- Autor, David H., David Dorn, and Gordon H Hanson, “Untangling Trade and Technology: Evidence from Local Labour Markets,” *The Economic Journal*, 2015, 125 (584), 621–646.
- Barnichon, Regis and Geert Mesters, “On the Demographic Adjustment of Unemployment,” *The Review of Economics and Statistics*, 2018, 100 (2), 219–231.
- Census, Bureau U.S., “Current Population Survey, Design and Methodology,” Technical Report, U.S. Census Bureau October 2019.
- Elsby, Mike, Bart Hobijn, and Aysegul Sahin, “The Labor Market in the Great Recession,” *Brookings Papers on Economic Activity*, 2010, *Spring* (1), 1–48.
- Fallick, Bruce and Jonathan Pingle, “A Cohort-Based Model of Labor Force Participation,” Finance and Economics Discussion Series 2007-09, Board of Governors of the Federal Reserve System 2007.

- Flood, Sarah, Miriam King, Renae Rodgers, Steven Ruggles, J. Robert Warren, and Michael Westberry, “Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset],” Technical Report, IPUMS 2022. Accessed April 30, 2026.
- Hornstein, Andreas and Marianna Kudlyak, “Aggregate Labor Force Participation and Unemployment and Demographic Trends,” Working Paper 19-08, FRB Richmond 2019.
- and —, “Technical Appendix for “Trends in Labor Force Participation and Unemployment, 1976-2024,”” Technical Report 2026. <https://sites.google.com/site/mariannakudlyak/research/lfptrend>.
- Kominski, Robert and Paul M. Siegel, “Measuring Education in the Current Population Survey,” *Monthly Labor Review*, 1993, 116 (9), 34 – 38.
- Kudlyak, Marianna, “A Cohort Model of Labor Force Participation,” *Federal Reserve Bank of Richmond Economic Quarterly*, 2013, 99 (1), 25–43.
- Montes, Joshua, “CBO’s Projection of Labor Force Participation Rates,” Working Paper 2018-04, Congressional Budget Office 2018.
- Natale, Marisa L. Di, “Creating Comparability in CPS Employment Series,” Technical Report, Bureau of Labor Statistics 2003.
- Shimer, Robert, “Why Is the U.S. Unemployment Rate so Much Lower?,” *NBER Macroeconomics Annual 1998*, 1998, 13, 11–61.
- U.S. Census Bureau, “2023 National Population Projections Tables: Main Series,” Data set, U.S. Department of Commerce, U.S. Census Bureau 2025. Accessed 2025-01-27.

## A Identification of age and cohort effects

We illustrate the identification of age and cohort effects for a simplified version of the state-space model that ignores the time aggregation of cohorts and separates unemployment from LFP. In the online Technical Appendix, Hornstein and Kudlyak (2026), we generalize this argument for our model with cohort aggregation and joint estimation of unemployment and LFP rates.

The model is

$$\begin{aligned}
q_{g,t} &= \tilde{x}_{g,t} + y_{g,t} + \gamma_g z_t + e_{g,t} \text{ with } e_{g,t} \sim N(0, \sigma_g^2), \\
\tilde{x}_{1,t} &= \tilde{x}_{1,t-1} + \varepsilon_{1\tilde{x},t} \text{ with } \varepsilon_{1\tilde{x},t} \sim N(0, \sigma_{\tilde{x}}^2), \\
\tilde{x}_{g,t} &= \tilde{x}_{g-1,t-1} \text{ for } g > 1, \\
y_{g,t} &= y_{g,t-1} + \varepsilon_{gy,t} \text{ with } \varepsilon_{gy,t} \sim N(0, \sigma_{gy}^2) \text{ for } g \geq 1, \\
z_t &= \rho z_{t-1} + \varepsilon_{z,t} \text{ with } \varepsilon_{z,t} \sim N(0, \sigma_z^2).
\end{aligned}$$

We have dropped the subindex  $q$ . To emphasize that normalization yields identification, we allow the first age effect to follow a random walk.

Cohort, age, and cycle effects are not identified if there exists a non-zero perturbation  $(\delta\tilde{x}, \delta y, \delta z)$  of the state that is consistent with the dynamics of the model and that does not affect outcomes. We now show that such a perturbation does not exist for the normalization

$$\tilde{x}_{1,1} = 0.$$

The perturbations have to satisfy the following constraints implied by the observation equation, the law of motion for cohorts, the law of motion for age effects, and the law of motion for cycle effects, respectively.

$$\begin{aligned}
0 &= \delta\tilde{x}_{g,t} + \delta y_{g,t} + \gamma_g \delta z_t, \\
\delta\tilde{x}_{1,t} &= \delta\tilde{x}_{1,t-1} \text{ and } \delta\tilde{x}_{g,t} = \delta\tilde{x}_{g-1,t-1}, \\
\delta y_{g,t} &= \delta y_{g,t-1}, \\
\delta z_t &= \rho \delta z_{t-1}.
\end{aligned}$$

Using these laws of motion, we get the following restrictions on perturbations

$$\begin{aligned}
\delta\tilde{x}_{g,t} &= \delta\tilde{x}_{1,1} \text{ for } g \leq t, \\
\delta y_{g,t} &= \delta y_{g,1}, \\
\delta z_t &= \rho^{t-1} \delta z_1.
\end{aligned}$$

And the restriction on the observation equation becomes

$$0 = \delta\tilde{x}_{1,1} + \delta y_{g,1} + \gamma_g \rho^{t-1} \delta z_1 = \delta y_{g,1} + \gamma_g \rho^{t-1} \delta z_1 \text{ for } g \leq t,$$

using the normalization of cohort effects. Since by assumption  $\gamma_1 \equiv 1$  and  $0 < |\rho| < 1$ ,  $\delta z_1 = 0$ . And therefore  $\delta y_{g,1} = 0$ . Therefore,  $\delta\tilde{x}_{g,t} = 0$  for all  $t, g$ . Thus, there is no non-zero perturbation of the state that does not affect outcomes.