The Summer Drop in Female Employment*

Brendan M. Price
Federal Reserve Board

Melanie Wasserman
UCLA

June 1, 2022

Abstract

We provide the first systematic account of summer declines in women's labor market activity. From May to July, the employment-to-population ratio among prime-age US women declines by 1.1 percentage points, whereas male employment rises; women's total hours worked fall by 11 percent, twice the decline among men. School closures for summer break—and corresponding lapses in implicit childcare—provide a unifying explanation for these patterns. The summer drop in female employment aligns with cross-state differences in the timing of school closures, is concentrated among mothers with young school-age children, and coincides with increased time spent engaging in childcare. Decomposing the gender gap in summer work interruptions across job types defined by sector and occupation, we find large contributions from both gender differences in job allocation and gender differences within jobs in the propensity to exit employment over the summer. Summer childcare constraints may contribute to gender gaps in career choice and earnings: women—particularly those with young school-age children—disproportionately work in the education sector, which offers greater summer flexibility but lower compensation relative to comparable jobs outside of education.

Keywords: gender gap, seasonality, labor force participation, childcare, time use, school closure
JEL codes: J13, J16, J22, J24

 Email: `brendan.m.price@frb.gov` and `melanie.wasserman@anderson.ucla.edu`. Price gratefully acknowledges financial support from the W.E. Upjohn Institute's Early Career Research Award program. We are grateful to Monica Rodriguez-Guevara and Matias Giaccobasso for outstanding research assistance. We thank Martha Bailey, John Coglianese, Andrew Garin, Claudia Goldin, Angela Kilby, Colleen Flaherty Manchester, Paola Giuliano, Seth Murray, Christopher Nekarda, Romain Wacziarg, and participants at the 2021 Labor and Employment Relations Association meetings, the Federal Reserve Board, the University of Michigan Population Studies Center, the University of Illinois at Urbana-Champaign, the University of Chicago Inclusive Economy Lab, the 2022 Society of Labor Economists annual meeting, and the 2022 CESifo Area Conference on Labor Economics for helpful feedback. All errors are ours. The views expressed in this paper are those of the authors and do not necessarily represent the views or policies of the Board of Governors of the Federal Reserve System or its staff.

1 Introduction

The COVID-19 pandemic has disproportionately affected women's labor market outcomes. As shown in Figure 1, the labor force participation rate among prime-age US women plummeted by 3.6 percentage points in spring 2020, 0.7 percentage points more than the decline among prime-age men. The gender profile of pandemic job losses marks a departure from previous recessions, which weighed more heavily on male employment. While gender differences in sectoral composition partly explain this deviation from historical norms, recent research points to school and daycare closures as a key driver of the differential declines in female employment and participation.[1]

Although the effects of pandemic school closures on parental labor supply have attracted intense public interest and concern, the labor market ramifications of *yearly* school closures have received much less attention. Figure 1 highlights in gray the months of June, July, and August—when US schools are typically closed for summer breaks—revealing an equally striking *seasonal* pattern. Summer after summer, women's labor force participation drops sharply, whereas men's participation remains comparatively stable.

Figure 1: The summer drop in prime-age female labor force participation

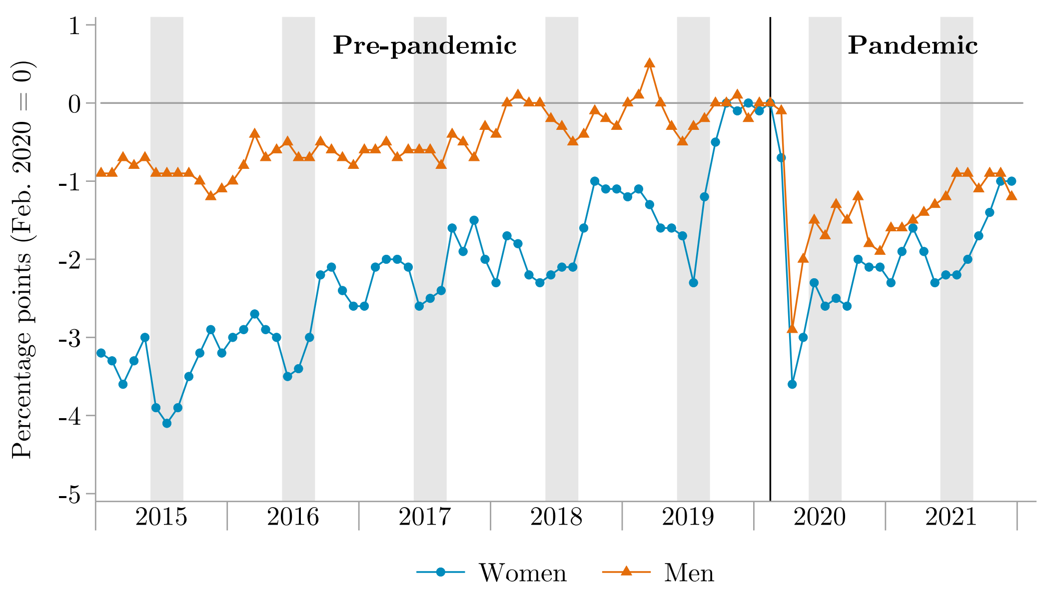

Source: Bureau of Labor Statistics, Labor Force Statistics, series LNU01300061 (men) and LNU01300062 (women). Notes: Non-seasonally adjusted labor force participation rates among individuals ages 25–54, normalized to zero on the eve of the pandemic. Shaded regions correspond to the months of June, July, and August.

[1] See, for example, Amuedo-Dorantes et al. (2021); Couch, Fairlie, and Xu (2022); Garcia and Cowan (2022); Hansen, Sabia, and Schaller (2022); Heggeness (2020); Montes, Smith, and Leigh (2021); and Russell and Sun (2020).

This paper provides the first systematic account of summer declines in female employment. Using Current Population Survey data spanning 1989–2019, we first show that the employment-to-population ratio among prime-age US women falls by an average of 1.1 percentage points between May and July each year, with equal contributions from increased unemployment and diminished participation. The annual reduction in female employment is economically meaningful, amounting to almost one third of the decline in the prime-age female employment rate during the Great Recession. In contrast, employment among prime-age men edges up slightly throughout summer. Declines in female work activity along the intensive margin reinforce those along the extensive margin: conditional on being employed, both women and men work fewer hours in the summer months, but the drop is much larger for women. Combining both margins, women's total hours worked contract by 11.2 percent from May to July, twice the decline among men.

School closures for summer break—and corresponding lapses in implicit childcare—provide a unifying explanation for these patterns. To build intuition, we develop a dynamic model of labor supply and career choice in the face of summer childcare constraints. Because women shoulder a disproportionate share of childcare—as evidenced by observed patterns of parental time use as well as gender differences in single parenthood—their choices are more heavily influenced by seasonal reductions in access to external childcare. Consistent with the model's predictions, we show empirically that the summer drop in female employment (1) aligns with cross-state differences in the timing of schools' summer breaks; (2) is concentrated among mothers, especially those with children old enough to attend school but young enough to require supervision when not in school; (3) is driven primarily by non-participants who cite household or family duties as their main activity while out of the labor force; and (4) coincides with an increase in women's time spent engaging in childcare. These regularities are absent or much less evident among men.

Our model also clarifies the proximate roles of sectoral allocation and within-sector employment flows in accounting for gender differences in summer work. We model two sectors: education, which provides *summer flexibility*, and non-education, which does not. Because education jobs are structured around the school calendar, school staff have the option of taking the summer off without penalties to their school-year compensation. Other jobs, by contrast, penalize workers who deviate from full-year employment. With this device, the model generates both *between-job* and *within-job* gender gaps. On the one hand, because women are more likely to work in education, they have weaker incentives to work over summer, whether by taking on summer employment at school or by seeking temporary employment elsewhere. On the other hand, women in both sectors exhibit summer drops in employment relative to their male colleagues who bear a smaller share of childcare responsibilities. Decomposing the gender gap in summer work interruptions across job types defined by sector and occupation, we find quantitatively significant roles for both channels. Notably, about half of the gender gap in summer work interruptions arises from within-job differences in the propensity of men and women to exit employment over the summer. Within education, female teachers, managers, and bus drivers all work less over the summer than do their male counterparts. Outside education, too, women exit employment each summer at higher rates than men.

Summer childcare constraints may contribute to gender gaps in pay by reducing women's annual hours worked, curbing productivity, impeding human capital accumulation, or limiting occupational choices. We provide suggestive evidence for two such channels. First, anticipated and realized summer childcare costs may induce women to sort into lower-paying jobs that provide summer flexibility. Not only are women much more likely than men to work in education, but their propensity to do so peaks precisely when their children are of school-going age. This sectoral sorting may come at a cost: occupations represented both within and outside the education sector typically exhibit an earnings penalty for working in education, suggesting that women may be trading off compensation for access to summer flexibility. Second, among teachers, women are 19 percentage points less likely than men to engage in paid work during the summer months (inside or outside of schools), leading to a gender gap in summer earnings of over 50 percent.

This paper contributes to the voluminous literature that studies gender disparities in labor market activity along both the extensive and intensive margins. There are well-documented gender differences in part-time work (Wiswall and Zafar, 2018), conventional work schedules (Mas and Pallais, 2017; Cubas, Juhn, and Silos, 2019; Bolotnyy and Emanuel, 2022), long work hours (Cortés and Pan, 2019; Wasserman, 2022), and career interruptions (Bertrand, Goldin, and Katz, 2010). Our paper establishes a new dimension of temporal flexibility—the timing of work throughout the year—and shows that summer childcare constraints both prompt women to gravitate to jobs that provide summer flexibility and reduce their summer employment within a given job.

A closely related literature studies the labor market ramifications of school availability and timing. Expansions in the availability of schooling generally have positive effects on mothers' labor supply (Gelbach, 2002; Cascio, 2009; Fitzpatrick, 2012). With regard to the timing of schooling, Duchini and Van Effenterre (2022) find gains in the continuity of maternal employment when France's school week switched from having Wednesdays off to running Monday through Friday. In a similar vein, Graves (2013) documents that year-round school schedules—which chop up the school year into smaller intervals of schooling—have negative effects on maternal employment. We add evidence on how school closures shape the timing of women's work throughout the year.

Our paper also complements the burgeoning literature on the gendered labor market effects of the COVID-19 pandemic.[2] Despite clear parallels, the school closures that occur each summer differ in important respects from those caused by the pandemic. While pandemic school closures were unprecedented and unanticipated, summer school closures are longstanding, predictable events, to which career choices have ample time to respond. In addition, pandemic closures shifted many education jobs online, whereas summer closures entail a multi-month reduction in labor demand across the education sector. Lastly, when the pandemic subsides, the summer work interruptions resulting from regularly scheduled school closures will likely remain a fixture of the US labor market.

We also contribute to a body of research analyzing seasonal regularities both in the macroeconomy (Barsky and Miron, 1989; Miron and Beaulieu, 1996; Olivei and Tenreyro, 2007; Ngai and Tenreyro, 2014; Geremew and Gourio, 2018) and among individual workers and households (Moretti, 2000; Del Bono and Weber, 2008; Coglianese and Price, 2020). A recurring theme in these papers is that seasonal phenomena—though routinely regarded as statistical nuisances to be adjusted away—can have important real-world consequences that go unnoticed in annualized or adjusted data. Sounding the same theme, we demonstrate how seasonal lapses in publicly provided implicit childcare shape the timing and continuity of women's labor market activity.

Section 2 describes our sample and regression specifications. Section 3 documents summer declines in female employment and hours. Section 4 develops a model of life-cycle labor supply under summer childcare constraints. Section 5 provides evidence that school closures for summer break are central to the summer drop in female employment. Section 6 decomposes the gender gap in summer work interruptions between and within jobs. Section 7 discusses ramifications for the gender gap in pay. Section 8 concludes.

[2]See, among others, Albanesi and Kim (2021); Alon et al. (2021); Amuedo-Dorantes et al. (2021); Couch, Fairlie, and Xu (2022); Furman, Kearney, and Powell (2021); Garcia and Cowan (2022); Goldin (2022); Hansen, Sabia, and Schaller (2022); Heggeness (2020); Montes, Smith, and Leigh (2021); and Russell and Sun (2020).

2 Data and Methodology

We trace seasonal shifts in labor market activity using the Current Population Survey (CPS), with auxiliary analyses drawing on the American Time Use Survey (ATUS) and the Schools and Staffing Survey (SASS). We describe the CPS here, with further details in Appendix B.1. We defer discussion of the ATUS and SASS until later in the paper.

2.1 Sample construction

The CPS is a representative survey of US households conducted monthly by the US Census Bureau on behalf of the Bureau of Labor Statistics. From basic CPS extracts provided by the Integrated Public Use Microdata Series (IPUMS, Flood et al., 2021), we assemble a person $\times$ year-month panel of civilian individuals ages 25–49 spanning the years 1989–2019. We focus on prime-age adults to abstract from seasonality in labor supply linked to an individual's own school enrollment and retirement decisions; we restrict to civilians because key labor market questions are not asked of members of the armed forces. Our analysis period begins in 1989, when the CPS first reports actual hours worked—allowing us to examine the intensive as well as the extensive margin of labor input—and ends on the eve of the COVID-19 pandemic, which upended typical seasonal patterns. Appendix Table A.1 reports summary statistics for our CPS sample.

CPS households are in-sample for four consecutive months, out-of-sample for eight months, and then back in-sample for a final four months. We use the cross-sectional dimension of the CPS to trace seasonality in labor market stocks, and we use the longitudinal dimension to track labor market flows both month-to-month and in back-to-back years (Rivera Drew, Flood, and Warren, 2014). For cross-sectional analyses, we use IPUMS sampling weights to ensure that our estimates are representative of the prime-age US population. For longitudinal analyses, we use iterative proportional fitting to construct sex-specific raked sampling weights that ensure consistency between labor market stocks and flows throughout our analysis period (Frazis et al., 2005). Following Madrian and Lefgren (2000), we validate cross-period individual linkages on the basis of sex, age, and race, and we exclude probable mismatches from our longitudinal analyses.

We observe household characteristics and labor market activity as of the survey reference week, which usually straddles the 12th day of the month. We partition individuals into those employed, those unemployed, and those not participating in the labor force. To account for vacation/leave-taking during the summer months, we separately analyze whether individuals are employed and at work or employed but absent from work. Since education-sector contracts can span 9 or 12 months, focusing on whether or not individuals are employed and at work also sidesteps the subtleties of how school staff report spells of non-work during the summer months. We also leverage CPS data on industry, occupation, and hours worked during the reference week, as well as reasons for non-participation or unemployment among those not employed.

We code individuals as "married" if their spouse is present, absent, or separated; we code single, divorced, and widowed individuals as "unmarried". We define parental status based on the presence or absence in the household of one or more own children under age 18. This definition encompasses adopted children and step-children as well as biological children, but it excludes other children residing in the household (such as nieces and nephews) as well as biological children who have already moved out.

2.2 Main specifications

We employ simple regression specifications that recover the typical seasonal movements in a given time series. Because the variation of interest is cross-month, we aggregate our data to the year-month level for each population we consider. To trace seasonal shifts in labor market activity within a given population, we then estimate time-series specifications of the form

$$y_t = \alpha + \sum_{m \neq 5} \beta_m \cdot \mathbb{1}\{M(t) = m\} + f(t) + weeks_t + \varepsilon_t \tag{1}$$

where y_t is an outcome in year-month t, $M(t) \in \{1, 2, \ldots, 12\}$ returns the calendar month for period t, $f(t)$ controls for lower-frequency trends, and $weeks_t$ is the number of weeks elapsed since the previous month's reference week. Because our focus is on summer work interruptions, we normalize β_5 to zero, so that the coefficients of interest β_m capture average differences in an outcome relative to the month of May.

To account flexibly but parsimoniously for secular trends and business-cycle dynamics that might otherwise bias estimation of seasonal patterns, we specify $f(t)$ as a linear spline in calendar time, with knots at roughly five-year intervals corresponding to turning points in the prime-age employment and participation rates. Appendix B.2 details our knot-selection procedure, which we adapt from the algorithm used by Dupraz, Nakamura, and Steinsson (2019) to locate turning points in the unemployment rate. Our spline function flexibly captures low-frequency dynamics in our core outcomes of interest and, more generally, allows for non-parametric time trends in all of our specifications. We additionally control for the number of weeks elapsed between successive months' reference weeks, since these time intervals are correlated with month length and holiday timing. We estimate Equation (1) separately for each of the demographic groups we consider, since trend and cyclical movements in labor market outcomes vary strongly with sex and household structure (Juhn and Potter, 2006; Albanesi and Şahin, 2018; Bardóczy, 2022).

Equation (1) is designed for use with stock variables, such as employment rates. When examining labor market flows, we estimate the first-differenced analogue of Equation (1):

$$\Delta y_t = \sum_{m \neq 5} \delta_m \cdot \mathbb{1}\{M(t) = m\} + \Delta f(t) + \Delta weeks_t + \Delta \varepsilon_t \tag{2}$$

where Δy_t represents gross inflows, gross outflows, or net flows into employment as a share of the relevant population. In this formulation, the coefficients of interest δ_m capture the magnitude of flows between months $m-1$ and m relative to April–May flows, and the differenced spline terms morph into indicator variables that allow for structural breaks in flow rates at the knot dates.

In both stock and flow specifications, we allow for heteroskedastic and autocorrelation-consistent standard errors (Newey and West, 1987) correlated up to a maximum lag of 26 months, a horizon suggested by the automatic lag selector of Newey and West (1994).[3] When our interest lies in functions of the estimated coefficients (rather than $\widehat{\beta}_m$ and $\widehat{\delta}_m$ directly), we construct confidence intervals via the delta method.[4]

[3]To choose an appropriate lag structure, we ran our main specification separately by sex and by sex $\times$ household structure for several key outcome variables (employment, participation, hours worked, and gross employment flows). Across these specifications, the optimal bandwidth often equaled (and never exceeded) 27 months, corresponding to a maximal lag of 26 months. For consistency and simplicity, we impose this same bandwidth throughout the paper.

[4]We estimate all models in Stata using the command `ivreg2` and construct confidence intervals using `nlcom`.

3 Summer Declines in Female Employment and Hours

This section establishes that women's labor market activity contracts each summer—along both extensive and intensive margins—in ways much less evident among men.

3.1 Women's employment drops in the summer

We start with the extensive margin. Figure 2 plots coefficients $\widehat{\beta}_m$ from estimating Equation (1) for three outcomes—employment, unemployment, and non-participation—separately by sex, with each measure expressed as a percentage of the corresponding population. As shown in the left panel, the prime-age female employment-to-population ratio (EPOP) declines by 1.1 percentage points (p.p.) between May and July, then rebounds strongly in the fall. Unemployment and non-participation contribute equally to the summer reduction in employment, with each rising 55 basis points from May to July.[5] In contrast, prime-age male employment actually *rises* slightly over the summer months. The summer decline in female employment is sizable, equaling almost one third of the decline in prime-age female EPOP in the wake of the Great Recession.[6]

Employment also contracts sharply with the onset of winter, especially for men. Because the main drivers of winter work interruptions—adverse weather, which triggers layoffs in male-dominated sectors like construction, and a post-holiday retreat in consumer spending—are not operative in the summer months, we confine our analysis to summer work interruptions, though we continue to show year-round seasonal movements to place the summer in context.

3.2 The employment drop mostly stems from increased outflows

The summer drop in female employment could reflect weak inflows to employment, strong outflows from employment, or both. In Appendix C.1, we show how the flow coefficients $\widehat{\delta}_m$ from estimating Equation (2) for gross per-capita inflows and for gross per-capita outflows can be transformed to express seasonal changes in employment rates as *excess inflows* minus *excess outflows*. Intuitively,

[5] While we pool all CPS survey years for our main analysis, in Appendix Figure A.1 we explore how the summer drop in female employment changes over our sample period. The summer drop appears relatively stable over time, with no obvious trend or cyclical variation in its magnitude. In addition, Appendix Figure A.2 shows that the female drop in summer employment appears consistently across age, education, and racial and ethnic groups.

[6] Prime-age female EPOP fell 3.7 percentage points from the start of the Great Recession in December 2007 (72.4 percent) to its nadir in September 2011 (68.7 percent). BLS Labor Force Statistics, series LNS12300062.

Figure 2: Seasonal shifts in per capita employment, unemployment, and non-participation

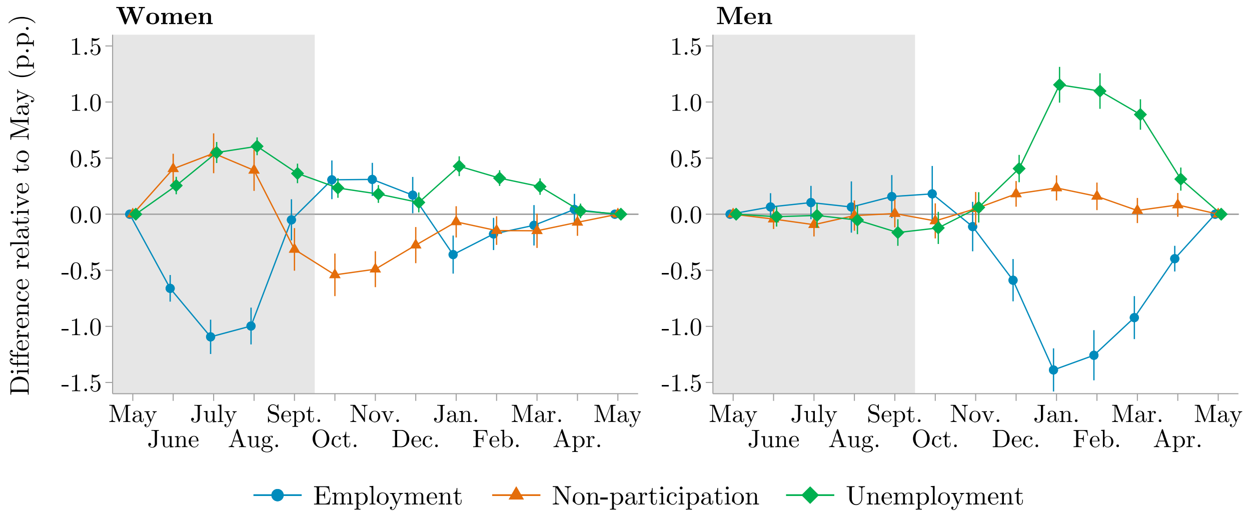

Source: Current Population Survey.
Notes: Coefficients $\widehat{\beta}_m$ from estimating Equation (1), separately by sex, for CPS respondents ages 25–49 grouped to the year-month level. Each measure is expressed as a share of the population. Bars show 95 percent confidence intervals based on Newey-West standard errors. In this and many subsequent figures, coefficients for May are normalized to zero, and plotted points are offset horizontally for visual clarity.

Figure 3: Decomposition of seasonal changes in EPOP into excess inflows vs. depressed outflows

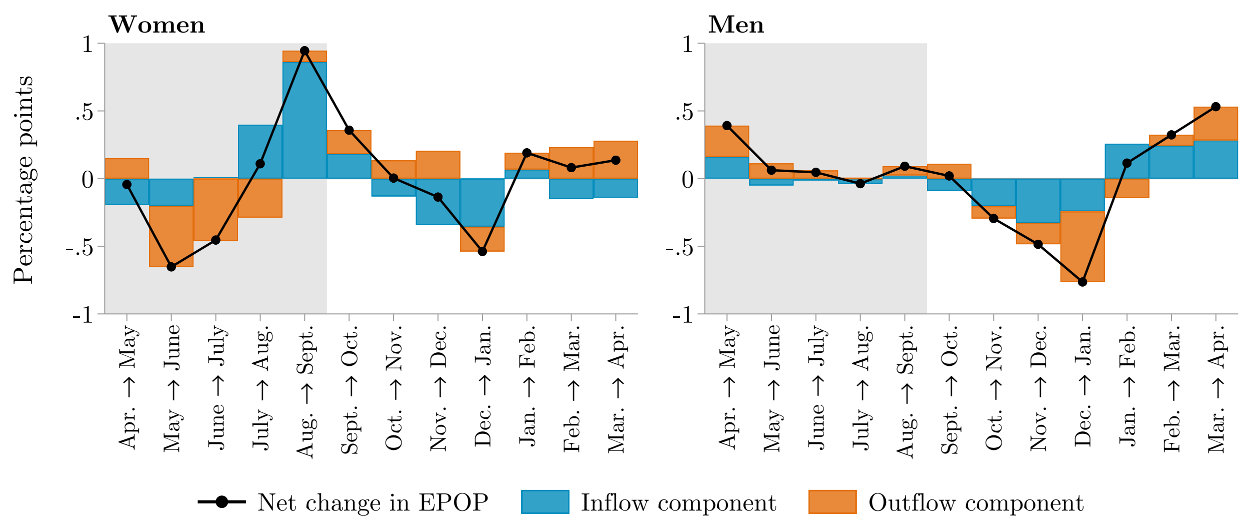

Source: Current Population Survey.
Notes: Additive decomposition of month-to-month changes in prime-age EPOP into contributions from above-average inflows and below-average outflows. The net change in EPOP reports coefficients $\widehat{\delta}_m$ estimated from the first-difference specification in Equation (2). Positive bar segments (respectively, negative segments) indicate that a given margin boosts (lowers) EPOP in a given month. See Appendix C for details on the inflow and outflow terms, which are transformations of the coefficients $\widehat{\delta}_m$ obtained using gross monthly transitions into or out of employment.

employment rises between two consecutive months if monthly inflows exceed their annual average and/or if monthly outflows fall short of their annual average.

Figure 3 decomposes month-to-month changes in EPOP into these respective margins. The left panel shows that the summer drop in female employment is primarily a story of summer exits: elevated May–June and June–July outflows drive female employment rates down by a combined 0.9 p.p. from May to July, with depressed inflows contributing an additional 0.2 p.p. The summer decline reverses in the autumn months, when employment is buoyed by a wave of entries. Among men, the relative stability in employment stocks is echoed in gross employment flows, which hover near their annual averages throughout the summer months.

For many women, summer exits are a recurring phenomenon, rather than one-time occurrences. As shown in Appendix Figure A.3, the same women who exit employment at the start of a given summer tend to do so again in the summer of the following year.[7]

3.3 Summer hours contract more for women than for men

The summer drop in female employment is also evident in total hours worked, an omnibus measure that encompasses shifts in labor market activity along both extensive and intensive margins. In Table 1, row (1) reports the average May–July change in hours worked during the reference week among prime-age individuals observed in both months. Alongside reductions in women's employment and participation rates, their hours fall by an average of 3.0 per week (11.2 percent) from May to July. Men's hours also decrease, but by a more modest 2.0 hours per week (5.2 percent).

Table 1 further decomposes the May–July decline in aggregate hours into extensive and intensive margin changes. For this decomposition, we define three groups: (i) those who are employed and at work with positive hours during the reference week (*present at work*); (ii) those who are unemployed or out of the labor force (*non-employed*); and (iii) those who are employed but absent from work for the entire reference week (*absent from work*). Row (2) quantifies the net change in hours along the extensive margin by tallying up positive and negative changes in hours among individuals who transition between non-employment and presence at work. Row (3) quantifies the intensive margin change in hours associated with transitions of employed workers between presence at work and absence from work. Row (4) quantifies the intensive margin change

[7] Specifically, women who exit employment between the May and June reference weeks in a given year are 4.9 percentage points more likely to experience another such separation exactly 12 months later than would be expected based on separation rates 11 and 13 months after baseline. July separations are also unusually likely to be repeated in back-to-back years. See Appendix E.1 for details.

Table 1: Decomposition of summer hours decline along the extensive and intensive margins

	Women		Men	
Change in hours worked during reference week	Δ	$\%\Delta$	Δ	$\%\Delta$
Total change from May to July:				
(1)	-3.0	-11.2	-2.0	-5.2
Contribution from extensive margin:				
(2) Employed, at work $\longleftrightarrow$ not employed	-0.3	-1.3	-0.0	-0.1
Contribution from intensive margin:				
(3) Employed, at work $\longleftrightarrow$ employed, absent	-2.1	-7.9	-1.3	-3.4
(4) Δ among those employed, at work	-0.5	-1.9	-0.6	-1.7

Source: Current Population Survey.
Notes: Row (1) reports the per capita change in hours from May to July among prime-age CPS respondents observed in both months. Rows (2)–(4) decompose this change by tabulating net hours changes among workers in the indicated categories. "Employed, at work" are employed individuals with positive hours worked in the previous week; "employed, absent" are employed individuals who worked zero hours in the previous week; and "not employed" are those unemployed or out of the labor force. "Δ among those employed, at work" is the change in hours worked among those employed with positive hours in both the May and July reference weeks. Percent changes are relative to May.

in hours from shifts in hours worked among those present at work in both May and July.

Consistent with the decline in women's employment during the summer months, women experience a 1.3 percent reduction in hours along the extensive margin (row 2). This extensive margin change is reinforced by much larger reductions on the intensive margin, due primarily to increased absences from work (row 3). By examining seasonality in individuals' stated reasons for being absent from work, we find that the increase in summer absences is concentrated among individuals taking vacation or personal days. A small portion of the intensive margin change stems from decreases in hours worked among those at work in both May and July (row 4). For men, the entirety of the 5.2 percent decline in summer hours comes via intensive margin changes, primarily in the form of increased absences from taking vacation.

In Figure 4 we further document that women experience prolonged summer work interruptions at higher rates than men. Let W denote being employed and at work and NW denote being either non-employed or absent from work. During the summer months, women and men experience similarly sharp upticks in the frequency of $W \rightarrow NW \rightarrow W$ spells, whereas women experience a much larger uptick than men in the frequency of $W \rightarrow NW \rightarrow NW$ spells. These results suggest that men's intensive-margin changes in summer hours are almost entirely due to brief vacations, whereas women's reflect both vacations and longer periods of non-work.

Figure 4: The prevalence of briefer versus longer summer work interruptions

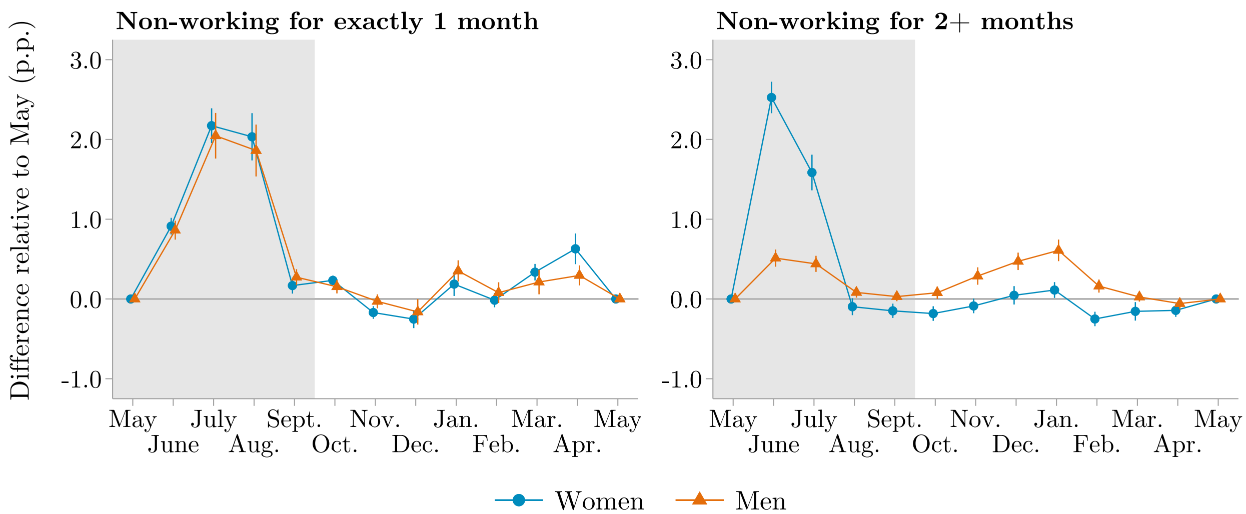

Source: Current Population Survey.
Notes: Coefficients $\widehat{\beta}_m$ from estimating Equation (1) separately by sex in a sample of respondents observed for at least three consecutive months. Let W denote being both employed and at work and NW denote being either non-employed or absent from work. In the left panel, the dependent variable is an indicator for having a $W \to NW \to W$ spell, with non-work beginning in the indicated month. In the right panel, the dependent variable is an indicator for having a $W \to NW \to NW$ spell. Bars show 95 percent confidence intervals based on Newey-West standard errors.

Summary. To recap, we document that prime-age women disproportionately experience declines in employment rates, elevated outflows from employment, and reductions in hours during the summer months. While men also experience a modest reduction in summer hours, their labor force participation and employment rates remain stable throughout the summer.

4 Conceptual Framework: the Central Role of School Closures

Why is there a summer drop in female employment? To frame our subsequent analysis, we describe a two-period model with career choices and summer childcare constraints that can rationalize the summer drop in female employment as a byproduct of the traditional school calendar. We provide one formalization of the model in Appendix D but note that other formulations (such as a model with frictional job search) would yield similar predictions.

Model setup. We consider a two-period partial equilibrium model in which individuals decide whether and in which sector to work at different points throughout the year and throughout their lives. Each period represents a distinct phase of the life cycle—pre-parenthood or parenthood—

and is subdivided into two distinct seasons, the summer and the school year. In each season, an individual may choose to (i) work in the education sector, (ii) work in the non-education sector, or (iii) not work, by being either jobless or on leave from a job.

Jobs differ in the extent to which they reward continuous employment or (equivalently) penalize interrupted employment. Because jobs in the education sector have a built-in summer recess, we assume that education workers may choose whether or not to work over the summer without influencing their earnings during the school year.[8] By contrast, non-education jobs offer a continuity bonus for full-year employment.[9] Together, these assumptions imply that the earnings penalty for summer work interruptions is smaller in the education sector. Put differently, education jobs provide a form of summer flexibility that other jobs do not.

While the sectors differ in their treatment of *within-year continuity*, we assume they both reward *career continuity*: individuals who stay in the same sector throughout their careers receive an earnings premium for doing so. This premium is meant to capture a range of real-world returns to sectoral or job tenure, such as specific human capital (Parent, 2000), backloaded salary scales (Lazear, 1981), or the vesting of pension benefits (Allen, Clark, and McDermed, 1993).

Each individual derives disutility from working, comprised of a sector-specific distaste for work as well as costs associated with childcare. During the school year, both working and non-working parents incur zero childcare costs, since schools provide implicit childcare when in session. During summer, by contrast, working parents incur childcare costs whereas other agents do not. In keeping with observed patterns of parental time use, we assume that mothers shoulder a disproportionate share of childcare costs (Handwerker and Mason, 2017), since they are more likely to be single parents and, if married, are less likely to have a non-working spouse available to cover childcare. In two-earner households, gender gaps in earnings within couples could create an incentive for women, rather than men, to curtail their summer employment if parental childcare is needed.

We allow individuals to vary in their earnings potential, sectoral choices, and childcare costs

[8]Education-sector workers often have the option to keep working over the summer, at least part-time: a teacher might teach summer school or coach a sports team, while a bus driver might drive a limited number of summer routes. Education workers without such an option may instead seek temporary employment in another sector. In any event, these individuals' summer choices are unlikely to impact the earnings they receive outside of summer.

[9]The continuity bonus is a stand-in for many real-world work configurations. For example, some employers (such as consulting or law firms) might only hire workers who commit to full-year employment, whereas others might offer a lower-paying career track for workers who seek fewer hours or weeks worked per year. See Podgursky (2011) for a discussion of the differences in the pecuniary and non-pecuniary attributes of education versus non-education jobs.

by virtue of underlying differences in productivity, comparative advantage, leisure preferences, societal norms, or household structures. By permitting such heterogeneity, the model can generate a rich set of different employment strategies pursued by different individuals.[10]

Model implications. This analytically tractable framework yields several predictions regarding how summer childcare costs—stemming from summertime lapses in school-provided implicit childcare—shape individual employment patterns throughout the year and over the life cycle. We highlight three key implications.

1. **Summer drop in female employment:** Individuals who face larger childcare costs—in particular, women relative to men and mothers relative to non-mothers—will exhibit summer declines in employment.

2. **Within-sector gender disparities:** Conditional on working in a given sector during the school year, women are less likely to work over the summer than are their male counterparts, since they face larger summer childcare costs. These gender differences arise in both the education sector and the non-education sector.

3. **Job-sorting effects:** Larger childcare costs will cause some individuals to sort into education jobs, so as to avail themselves of the education sector's summer flexibility. Such sorting takes two distinct forms. First, there is *contemporaneous sorting*: some individuals work in non-education early in their careers, then switch to education jobs once they have school-age children and therefore face summer childcare costs. Second, there is *anticipatory sorting*: due to the returns to career continuity, women are more likely than men to sort into education jobs earlier in their careers in anticipation of future childcare costs.

We explore these predictions empirically throughout the rest of the paper.

[10]For example, some individuals will find it optimal to work year-round in non-education jobs, while others will work in education during the school year and not work over the summer. Some individuals will work year-round prior to having children, then switch to taking summers off or opt out of working entirely after having children.

5 Evidence: Timing and Incidence of the Summer Drop

In this section, we provide a constellation of evidence that the summer drop in female employment stems from school closures. First, we show that the timing of the summer drop lines up with cross-state differences in the timing of schools' summer breaks. Second, we document that mothers of school-age children are especially likely to experience summer declines in employment. Third, these declines are accompanied by an uptick in time spent on childcare during the summer months.

5.1 The summer drop in female employment tracks school calendars

We exploit cross-state variation in the timing of school closures to establish that the summer drop in female employment is inextricably tied to schools' closure for summer break.[11] To determine when schools typically close in each state, we leverage information about how many 16-year-old CPS respondents report being enrolled in high school during the May, June, and July reference weeks. For each state, we compute the average decline in school enrollment rates from May to July during our analysis period. We then classify as "early-closure states" those in which at least two thirds of the total May–July decline occurs between May and June; we classify as "late-closure states" those in which less than one third of the decline occurs between June and July.[12]

Applying this classification, the right panel of Figure 5 plots the summer drop in female employment, separately for respondents in early-closure versus late-closure states. The data speak clearly: in states where the large majority of K–12 schools have closed by the June reference week, female employment also starts its summer decline in June. By contrast, in states where most closures occur between the June and July reference weeks, female employment instead holds steady in June and starts its decline in July. The tight synchronization between the onset of school summer breaks and declines in female employment points to school closures as the underlying cause.

Figure 5: Cross-state synchronization of school closures with the summer drop in female EPOP

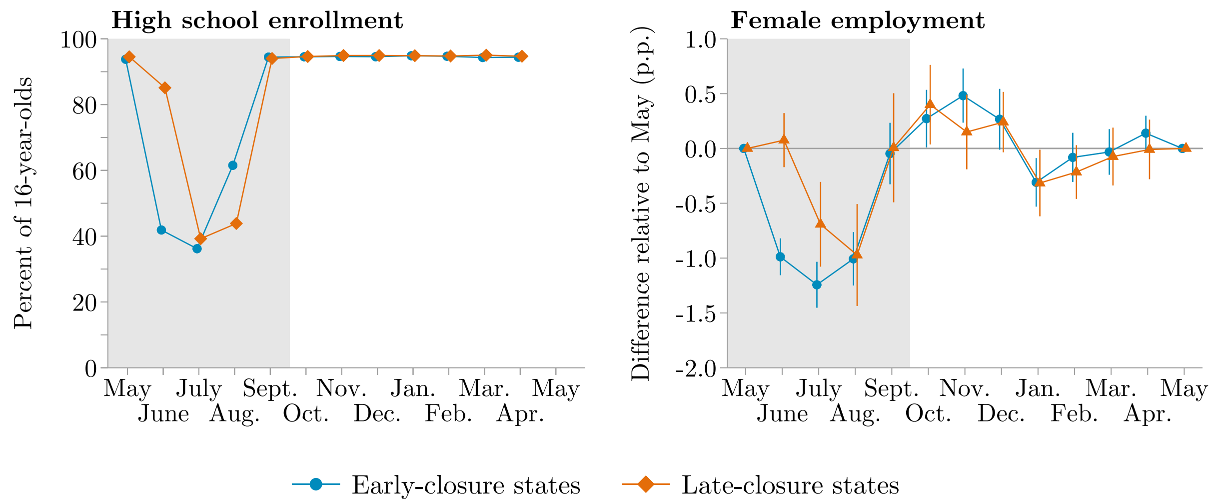

Source: Current Population Survey.
Notes: Left panel shows the percentage of 16-year-olds who report being enrolled in high school in the indicated month in states with early school closures (mostly in effect by the June reference week) or late school closures (mostly in effect only as of July). Right panel shows coefficients $\widehat{\beta}_m$ from estimating Equation (1) for female EPOP separately in early and late closure states. Bars show 95 percent confidence intervals based on Newey-West standard errors.

Figure 6: Seasonal patterns in female employment by household structure and parental status

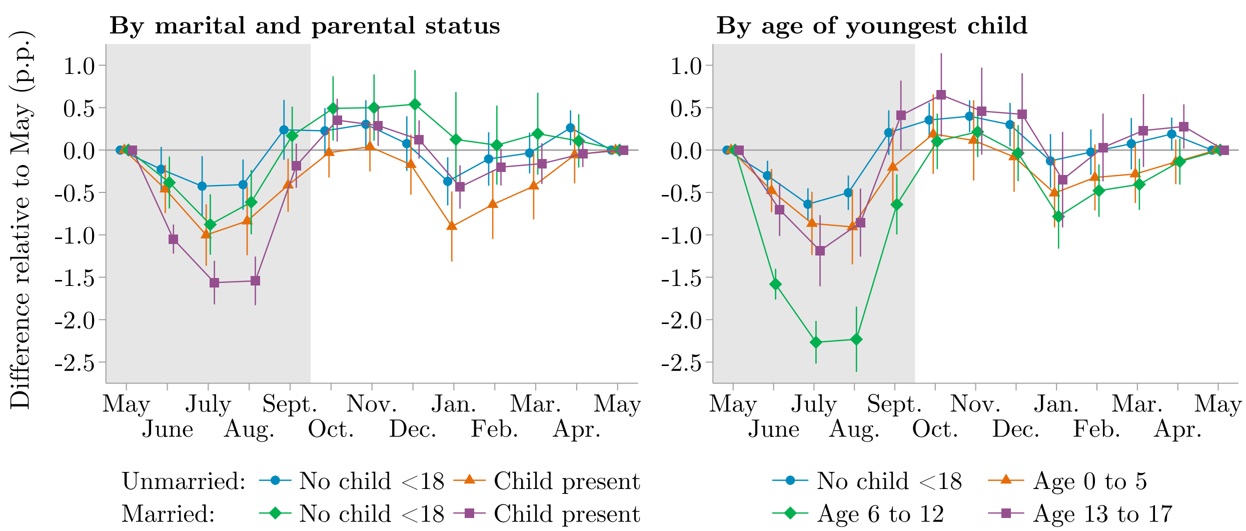

Source: Current Population Survey.
Notes: Coefficients $\widehat{\beta}_m$ from estimating Equation (1) for female EPOP separately by sex, marital, and parental status. Separated individuals are coded as married; parental status is defined in reference to an individual's own children, including adoptees and step-children but excluding younger siblings and other children not one's own. Bars show 95 percent confidence intervals based on Newey-West standard errors.

5.2 The summer drop is largest for women with school-age children

Our conceptual framework predicts that summer declines in employment will be most pronounced among women who experience lapses in externally provided childcare during the summer months. To test this prediction, the left panel of Figure 6 examines heterogeneity in women's seasonal employment patterns by marital status interacted with the presence or absence of a child under 18 in the household. The presence of children—within both the unmarried and married groups—amplifies the summer drop in female employment. The decline is steepest, at 1.6 percentage points, among married mothers residing with their children.[13]

Childcare needs are most likely to constrain summer employment when children are old enough to attend school from fall through spring, but too young to be left unattended for extended periods of time. Since childcare constraints are likely to be determined by a mother's youngest child, the right panel of Figure 6 stratifies mothers (of any marital status) by the age of that child: children under 6 years old, who have yet to enter the K–12 education system; those aged 6–12, who attend school and require supervision when not in school; and those aged 13–17, who attend school and require less supervision when not in school. Mothers of children aged 6–12 experience the largest drop in employment, of 2.3 percentage points.[14]

5.3 Women spend more time on childcare in the summer months

Our assertion that childcare responsibilities account for women's reduced summer employment is consistent with their self-reported summer activities. Beginning in 1994, the CPS reports each non-participant's major activity while not in the labor force. As shown in the top panel of Figure 7, both for prime-age women as a whole and for mothers of school-age children in particular, the

[11] Year-round schooling—in which schools replace the long summer vacation with a series of shorter breaks spread throughout the year—commands only a small share of the market. According to the National Center for Education Statistics, 4.1 percent of public schools used year-round calendars in 2011–2012, the latest school year for which this tabulation is available (Digest of Education Statistics, Table 234.12).

[12] Appendix Figure A.5 plots the distribution of this statistic across states. As shown in Appendix Figure A.6, most states in the American interior and the South Atlantic are classified as having early school closures, while much of the Northeast and Washington state have late school closures. A number of states in the Northeast, Midwest, and West Coast exhibit mixed patterns that defy neat classification.

[13] From a life-cycle standpoint, these patterns imply that the summer drop in female employment should widen during women's prime child-rearing years. Appendix Figure A.7 confirms this implication by plotting May–July changes in employment from estimating Equation (1) separately for each sex × one-year age bin.

[14] Appendix Figure A.8 shows analogous plots for men. No subgroup of men experiences a decline in summer employment. Men with children younger than age 13 experience a slight increase in summer employment.

Figure 7: Decomposition of summer changes in non-participation and unemployment

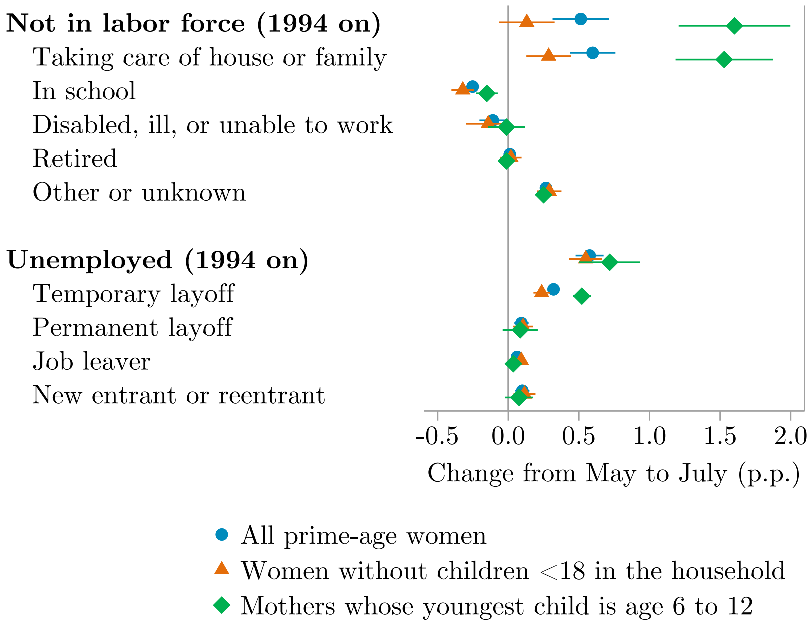

Source: Current Population Survey.
Notes: Coefficients $\widehat{\beta}_7$ (representing May–July changes) from estimating Equation (1) for the indicated outcomes in grouped data spanning 1994–2019. Subcategories of individuals not in the labor force denote respondents' major activity during the reference week. Subcategories of unemployed individuals denote respondents' reason for being unemployed. Bars show 95 percent confidence intervals based on Newey-West standard errors.

net increase in non-participation is fully accounted for by an increase in the share of women who report that they are "taking care of house or family". In contrast, women without children in the household exhibit no change in their labor force participation during the summer months (and only a slight increase in their propensity to cite family duties in the event of non-participation).

Some of the summer increase in unemployment may also reflect women providing childcare while awaiting recall. As shown in the bottom panel of Figure 7, the uptick is driven—especially for mothers—by a jump in the share of respondents who are job losers on temporary layoff, meaning that they expect to be called back to work within the next six months (Katz and Meyer, 1990). Since temporary summer layoffs are concentrated in the education sector, they are likely to align closely with the span of time for which a laid-off worker's children are on summer break.

To further probe the role of childcare in women's time allocation during the summer months, we turn to the American Time Use Survey. As detailed in Appendix B.3, we compute total childcare time by summing time spent on primary childcare (childcare as one's main activity) and secondary childcare (childcare while doing other tasks). We decompose secondary childcare according to

Figure 8: Decomposition of total childcare time among parents of school-age children

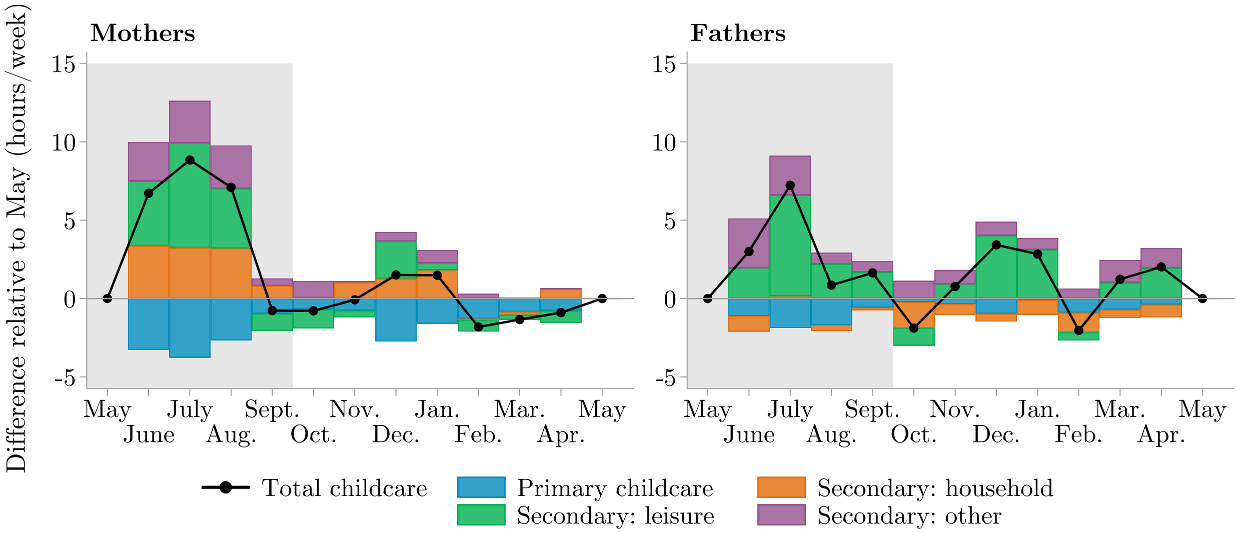

Source: American Time Use Survey.
Notes: Coefficients $\widehat{\beta}_m$ from estimating an individual-level version of Equation (1) on 2004–2019 ATUS respondents aged 25–49 who reside with a youngest child aged 6–12. We control for a linear spline in calendar time and for day-of-week fixed effects. See Appendix B.3 for definitions of each childcare category.

the primary tasks that accompany it: leisure activities, household activities, or other activities. Motivated by our earlier results, we focus on parents whose youngest child is aged 6–12.

Consistent with summer school closures prompting women to shift their time use from employment to childcare, the left panel of Figure 8 documents that mothers' total time spent on childcare rises by 8.9 hours per week from May to July. The increase in total childcare time embeds a sharp rise in secondary childcare partly offset by a reduction in primary childcare.[15] Consistent with the more modest drop in men's hours worked associated with summer vacations (Table 1), the right panel of Figure 8 shows that fathers experience a smaller rise in total time spent on childcare, owing mainly to increased secondary childcare while engaged in leisure activities.

6 Evidence: Job Sorting and Within-Job Gender Differences

Our conceptual framework in Section 4 generates predictions about the sectoral allocation and within-sector employment patterns of individuals who face greater summer childcare costs. First, we provide evidence that sorting across jobs as well as within-job gender differences in the propensity

[15]This pattern is consistent with prior research finding a summer decline in primary childcare involving educational activities, such as helping children with homework or driving them to school events (Handwerker and Mason, 2017).

to exit employment contribute to the summer drop in female employment. Next, we implement a formal decomposition of the gender gap in summer work interruptions into between- and within-job components and find large contributions from both. Notably, about half of the gap arises from within-job gender differences in the propensity to exit employment during the summer months.

6.1 Job sorting contributes to the summer drop

Our model predicts that women—particularly those who face summer lapses in school-provided childcare—will be more likely than men to work in the education sector, in part because they value jobs that provide summer flexibility. Empirically, women are indeed disproportionately represented in education: 13.2 percent of female workers are employed in educational services in May, compared with just 4.7 percent of male workers. But women might gravitate toward the education sector for a variety of reasons unrelated to the alignment of work with children's school schedules: tastes for working in education, comparative advantage, historical path dependence in occupational choice, or norms. To test whether women's propensity to work in education tracks childcare demands, we analyze the sorting of parents based on the age of their youngest child.

In Figure 9, the left panel presents raw shares of employed men and women working in education, according to the age of their youngest child. The right panel presents regression-adjusted probabilities, which control for the age of the parent and for secular time trends. Individuals with children less than one year old are the omitted group, and their propensity to work in education is normalized to zero.[16] Relative to mothers with an infant, the share of working mothers employed in education first declines with child age, shoots up sharply as the youngest child reaches school age, peaks for mothers whose youngest child is 10 years old, and then declines as the youngest child progresses through adolescence. In contrast, men's propensity to work in education is invariant to the age of their youngest child.

[16] For this exercise, we expand our sample to include ages 20–64 in order to better capture the tails of the child-age distribution. We drop the months of June, July, and August to avoid conflating differences in sectoral choice with differences in summer behavior. For each sex, we then estimate individual-level regressions of the form

$$e_{it} = \sum_{a=0}^{17} \beta_a \cdot \mathbb{1}\{a_{it} = a\} + \beta_n n_{it} + o_{it} + f(t) + \varepsilon_{it} \tag{3}$$

where e_{it} is an indicator for working in education, a_{it} is the age of individual i's youngest child, n_{it} is an indicator for having no child under 18 residing in the household, o_{it} is a full set of own-age fixed effects, and $f(t)$ is our standard linear spline. The coefficients of interest β_a capture working parents' propensity to work in education as a function of child age. We two-way cluster on individual and time period to allow for serial correlation and common shocks.

Figure 9: Share of those employed who work in the education sector as a function of child age

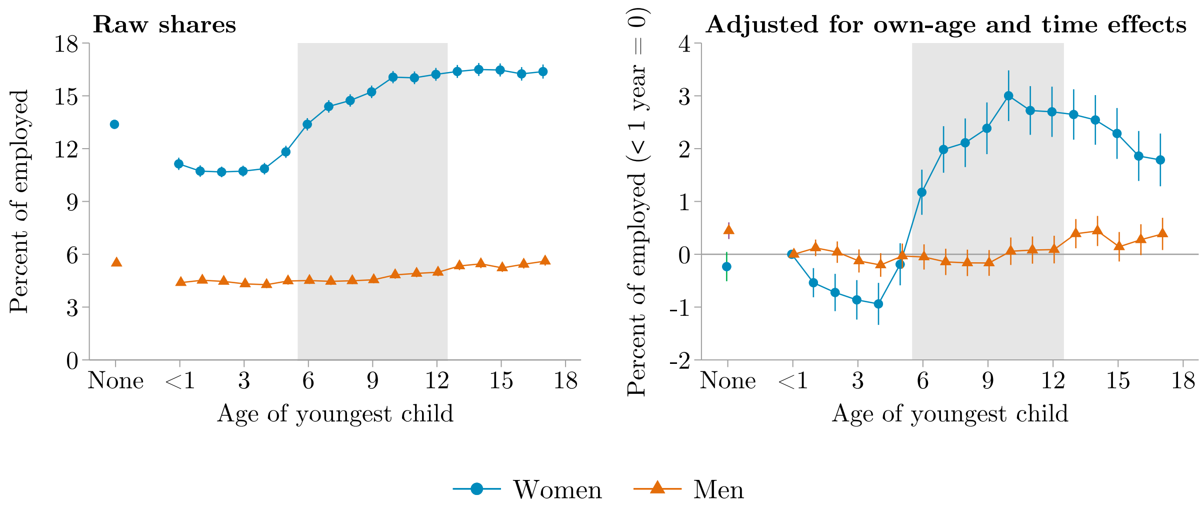

Source: Current Population Survey.
Notes: Coefficients $\widehat{\beta}_n$ (no child under 18) and $\widehat{\beta}_a$ (youngest child of age a) from estimating Equation (3) in a sample of employed individuals aged 20–64 and observed during the non-summer months. The left panel reports raw shares employed in educational services; the right panel reports shares adjusted for a full set of one-year own-age effects and for a linear spline in calendar time, with the coefficient for parents with a newborn (β_0) normalized to zero. Bars show 95 percent confidence intervals, based on standard errors two-way clustered on individual and year-month.

While our model focuses on the choice of working in either education or non-education, childcare costs should also induce women to seek out jobs that provide summer flexibility *within* the education sector. We find that women who work in educational services are more likely to work in occupations that shed more workers over the summer months. Appendix Table A.3 reports female and male employment shares as well as summer separation hazards for select occupations within the education sector. For example, the share of women employed in education who work as primary school teachers is nearly double that of men, while the reverse pattern holds for secondary school teachers. Consistent with the logic of our model, primary school teachers experience higher separation rates during the summer: averaging the male and female hazard rates (to neutralize the effect of differences in gender composition), primary school teachers are 1.7 p.p. more likely than secondary school teachers to exit employment from May to July.

6.2 Within-job differences contribute to the summer drop

Our model predicts that women in particular jobs—within and outside of the education sector—will be less likely to work during the summer months than men in the same job. We provide evidence of gender differences in the propensity to exit employment during the summer months, even within

Figure 10: Hazard rate of exiting employment among workers in educational services

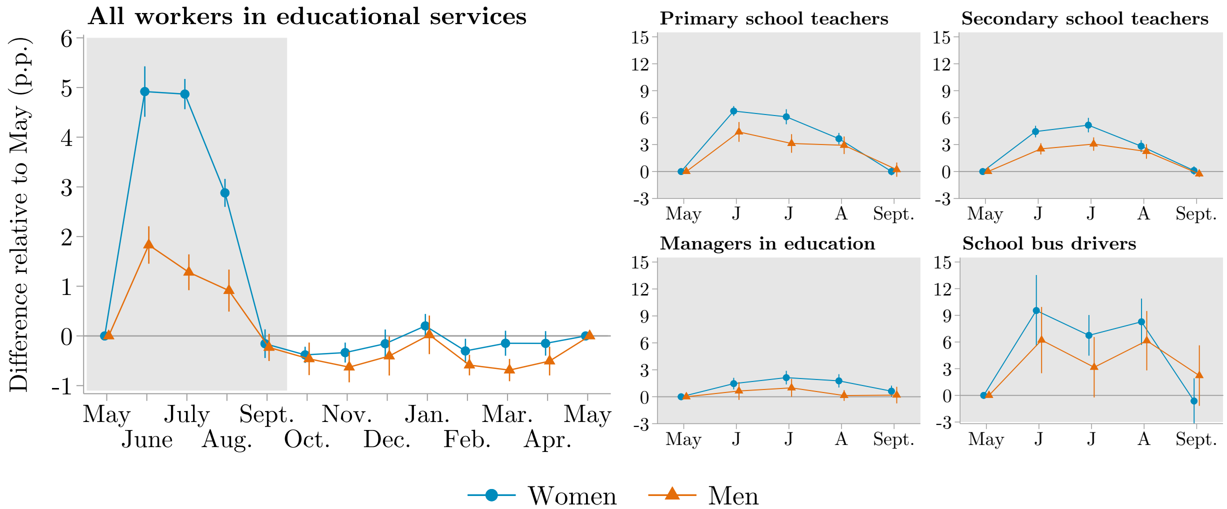

Source: Current Population Survey.
Notes: Coefficients $\widehat{\beta}_m$ from estimating Equation (2) in a sample of CPS respondents ages 25–49 with valid longitudinal links. The sample is restricted to individuals employed in educational services in the previous month, either in any occupation (left panel) or in the indicated occupation (right panel); the outcome is the percentage of these individuals who exited into non-employment in the current month. Bars show 95 percent confidence intervals based on Newey-West standard errors. In the right panel, coefficients for October–April are estimated but not shown.

narrowly defined occupations in education. As shown in Figure 10, female primary school teachers, secondary school teachers, managers in education, and school bus drivers are all more likely to exit employment each summer than their male counterparts.

A potential concern with these comparisons is that, conditional on objective circumstances, women and men who work in education may differ in their likelihood of self-reporting their summer status as non-employed versus employed but absent from work. To rule out this concern, Appendix Figure A.9 presents an alternative version of Figure 10 that measures the hazard rate of transitioning from positive hours worked in the reference week to zero hours worked, without distinguishing between absence and non-employment. For all four occupations, the same qualitative picture emerges, and the gender gaps are generally even larger in absolute terms.

Outside of the education sector, women are also more likely than men to exit employment during the summer. The left panel of Figure 11 shows that in each summer month, the hazard rate of exiting employment for workers outside of educational services is higher for women than for men. Furthermore, gender differences in the propensity to exit employment during the summer within non-education sectors are not simply incidental: the right panels document that women's

Figure 11: Hazard rate of exiting employment among workers outside of educational services

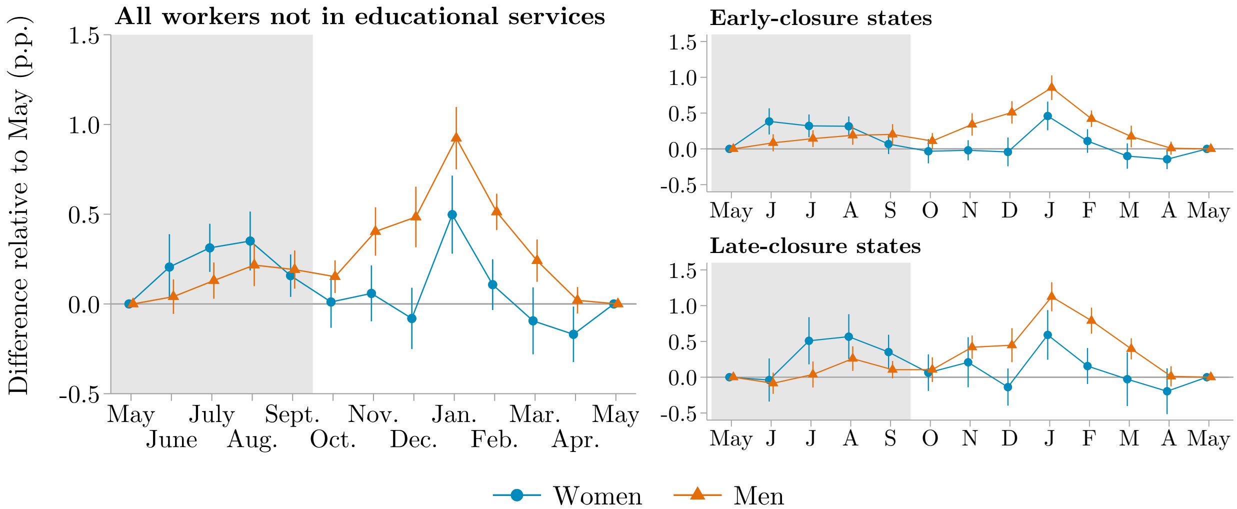

Source: Current Population Survey.
Notes: Coefficients $\widehat{\beta}_m$ from estimating Equation (2) in a sample of CPS respondents ages 25–49 with valid longitudinal links. The sample is restricted to individuals employed outside of educational services in the previous month; the right panel further restricts the sample to respondents residing in states with early or later school closures, as classified in the text. The outcome is the percentage of these individuals who exited into non-employment in the current month. Bars show 95 percent confidence intervals based on Newey-West standard errors.

hazard rates of exiting employment are tightly connected to school calendars. Using our earlier classification of early- and late-closure states, we observe that outside of the education sector, women experience an uptick in exits from employment precisely when schools in their state close for summer break.

6.3 Quantifying the roles of job-sorting and within-job effects

What share of the gender gap in summer work interruptions reflects gender differences in job sorting, and what share reflects gender differences conditional on job type? To answer this question, we develop a nested Oaxaca-Blinder decomposition that quantifies contributions from six distinct channels. We describe the decomposition verbally here and formalize it in Appendix C.

Consider the May–July change in women's EPOP minus the same change among men. Men and women differ in their allocation across job types, which in turn differ in their propensity to generate net outflows from employment between May and July. Conditional on job allocation, men and women also differ in their propensity to exit employment. By the standard Oaxaca-Blinder logic, we can therefore decompose the gender gap in summer work interruptions as

$$\text{overall gender gap} = \text{between jobs} + \text{within jobs} \tag{4}$$

We define "jobs" on the basis of both sector and occupation. Within educational services, we distinguish five job types: (i) pre-K, kindergarten, and primary school teachers; (ii) secondary school teachers; (iii) postsecondary teachers; (iv) other staff in elementary and secondary schools; and (v) other staff in educational services. Outside of education, we distinguish 13 job types corresponding to "one-digit" sectors, such as construction, manufacturing, and retail trade. Using these job groupings, we can subdecompose the "between" component as

$$\begin{aligned} \text{between jobs} = \; & \text{sorting into education vs. non-education} \\ & + \text{sorting across jobs within education} \\ & + \text{sorting among non-education sectors} \\ & + \text{baseline differences in EPOP} \end{aligned} \tag{5}$$

The first of these terms captures gender differences in sorting into education, coupled with the fact that education contracts each summer relative to non-education. The second term captures gender differences in sorting among education jobs, which likewise differ in their seasonal patterns; for example, primary school teachers are more likely to exit employment each summer than are secondary school teachers. The third term captures gender differences in the rest of the labor market; for example, men are disproportionately employed in the construction sector, which expands every summer, relative to health care, which is comparatively stable through the summer months. The final term, a scaling component that adjusts for gender differences in baseline EPOP, is of little economic interest and will be quantitatively small in practice.

The within-job component, in turn, can be expressed as a share-weighted average of the gender difference in employment seasonality observed within each job type. Summing these differences across education and non-education jobs, we obtain

$$\text{within jobs} = \text{within education jobs} + \text{within non-education jobs} \tag{6}$$

Figure 12: Decomposition of female-male differences in the seasonality of EPOP

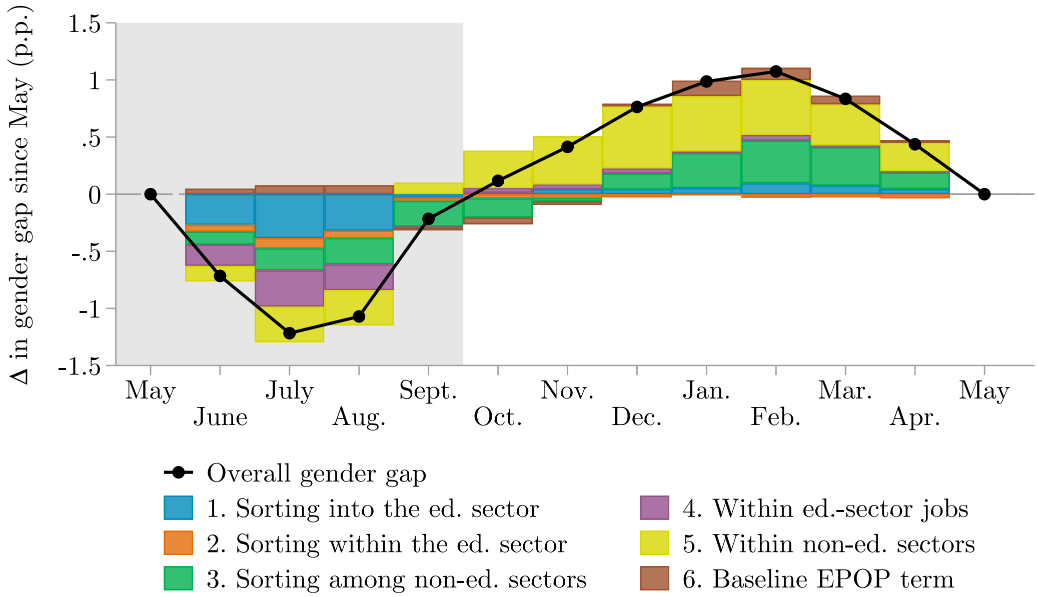

Source: Current Population Survey.
Notes: Additive decomposition of the gender gap in cumulative changes in EPOP between May and the indicated month. See text and Appendix C for details on the decomposition methodology. See Appendix Table A.2 for accompanying point estimates and standard errors.

Differences in the propensities of male and female secondary school teachers to exit employment during the summer months will be credited to the within-education term. Likewise, differences between male and female construction workers will be credited to the within-non-education term.

Figure 12 implements this decomposition, with the methodology extended to span the full calendar year. Consistent with the evidence presented in Section 6.1 and Section 6.2, each of the five components emphasized above contributes to the overall 1.2 percentage point gender gap in May–July employment changes. Notably, gender differences in job sorting and gender differences conditional on job type each explain about half of the gender gap in summer work interruptions. Sorting into education explains just over 30 percent of the overall change, while sorting across jobs within the education sector contributes an additional 7 percent. Sorting outside of education—such as between construction and health care—explains 16 percent of the total. Finally, gender differences within education jobs and gender differences within non-education jobs each account for 26 percent of the total.[17]

[17] The shares attributed to these five components sum to a little over 100 percent, owing to the small baseline EPOP scaling term acting in the opposite direction.

Summary. The evidence presented in Section 5 and Section 6 draws tight empirical links between summer school closures and the summer drop in female employment. Summer childcare constraints, operating both between and within job types, provide a unifying explanation for these links.

7 Implications for the Gender Pay Gap

This section provides an overview of various channels through which summer childcare constraints may affect female earnings and consequently contribute to the gender pay gap. We document empirical evidence for two channels. First, women—particularly those with young school-age children—disproportionately work in education-sector jobs, which offer greater summer flexibility but lower compensation relative to comparable jobs outside of education. Second, using survey data on teachers, we find that women earn less over the summer months than their male counterparts, in part due to their lower likelihood of taking on supplemental summer work.

7.1 Potential channels for summer childcare constraints to affect earnings

Summer childcare constraints may decrease women's earnings through several channels. First, anticipation of these constraints could dissuade some women from participating in the labor force if there are steep costs associated with summer childcare or substantial penalties associated with taking time off to care for one's children.[18] Second, conditional on working, women might disproportionately seek out employment in the education sector, which offers summer flexibility but lower compensation; likewise, they may seek work in other sectors (such as retail) that permit intermittent employment but offer few opportunities for career advancement. Third, within education, women may be less likely than men to take advantage of opportunities for supplemental work during the summer months, such as teaching summer school or coaching a sports team. Finally, outside of education, women may exit employment during the summer months in order to care for their children. In addition to the direct earnings consequences of a reduction in weeks worked, these work interruptions could indirectly affect contemporaneous or future earnings through lower productivity or suspended human capital accumulation.

[18] Using a longitudinal sample of young MBA professionals, Bertrand, Goldin, and Katz (2010) document a rise in female non-participation post-childbirth, particularly due to family reasons. Incompatibility between childcare demands and the long and inflexible hours of many corporate jobs may explain the decline in participation.

Figure 13: Fraction female in occupations present both within and outside of educational services

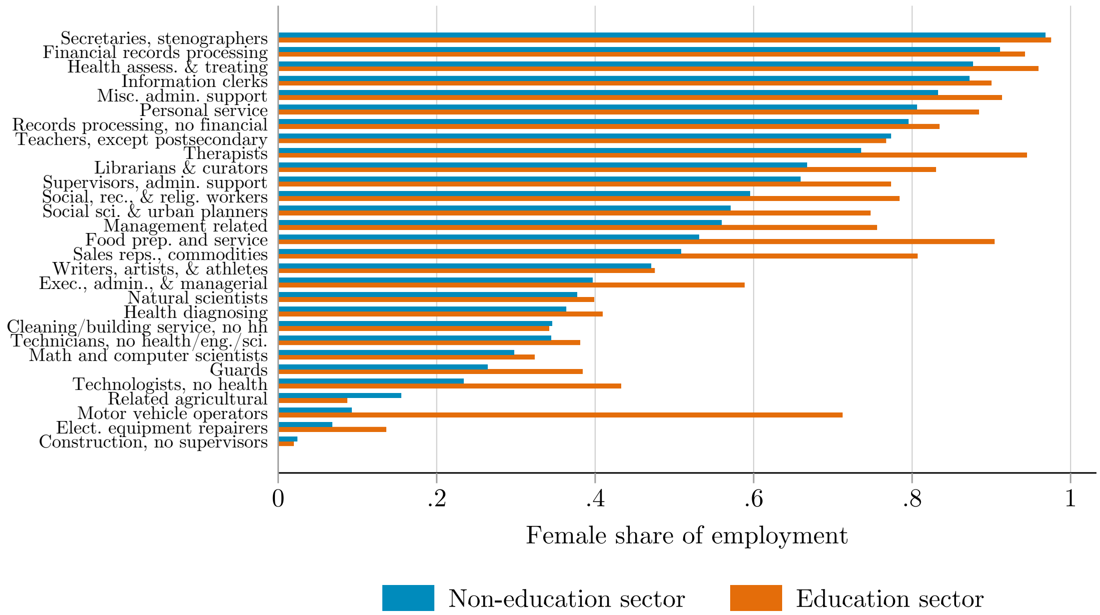

Source: Current Population Survey.
Notes: Average female employment as a share of each occupation over 1989–2019, computed separately for the education and non-education sectors. The 29 listed occupations, drawn from a set of 73 two-digit Census occupations, are those for which average employment exceeds 20,000 in each of the two sectors.

7.2 Earnings differentials between education and non-education jobs

Women's disproportionate representation in education services, in part due to its provision of summer flexibility, may contribute to gender gaps in pay. In Figure 13 we select 29 occupations present in both the education and non-education sectors, then compute the female share of each occupation, by sector. Consistent with women actively seeking out work in the education sector, the female share is higher in the education sector for 25 out of 29 occupations, often by a wide margin. Using these same occupations, we then estimate the education-sector earnings premium or penalty in each occupation by estimating a Mincer regression on annual *male* earnings in the Annual Social and Economic Supplement (ASEC) to the March CPS. As shown in Figure 14, a large majority of occupations display an earnings penalty associated with working in the education sector, suggesting that women may be trading off compensation for access to work within the education sector.

Figure 14: Education-sector earnings penalties and premia within occupations present both within and outside of educational services

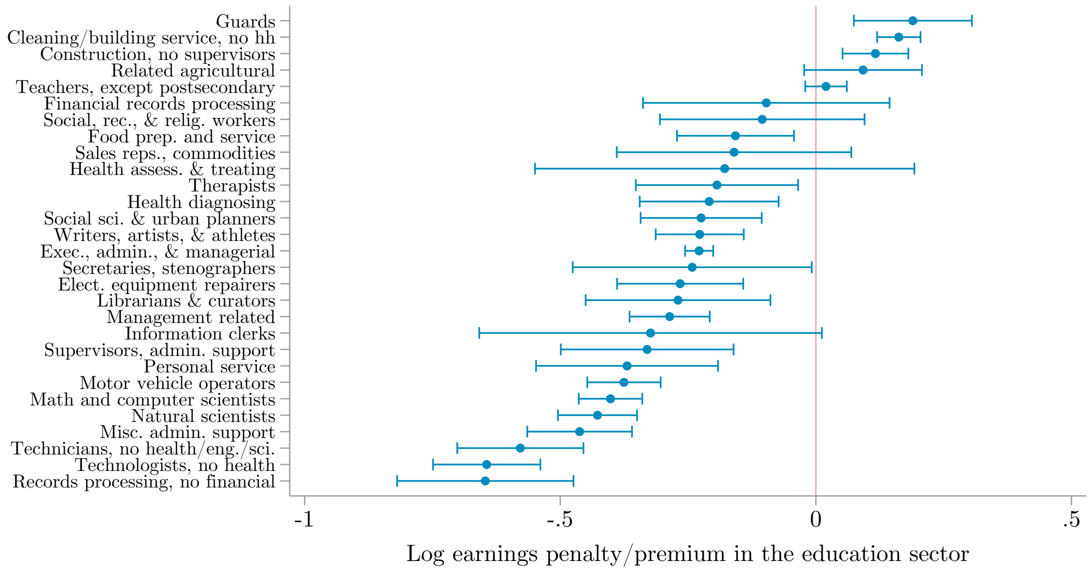

Source: Current Population Survey, March ASEC.
Notes: Coefficients on the interaction of occupation fixed effects with an education-sector dummy, from an individual-level regression of log annual earnings among men employed in the listed occupations (selected as described in Figure 13). The regression also controls for occupation main effects, educational attainment, a quadratic in age, and calendar year. Bars show 95 percent confidence intervals, with standard errors clustered at the household level.

7.3 Within-education differences in summer earnings

Women are less likely than men to work during the summer months, even within jobs. Although data limitations in the CPS preclude a comprehensive analysis of gender differences in summer earnings, the 1999–2000 Schools and Staffing Survey from the National Center for Education Statistics allows us to document summer earnings losses in the teaching profession.[19] Female teachers are 18.8 percentage points less likely than male teachers to engage in any type of paid work during the summer months (see Appendix E.2 for further details).[20] These gender disparities in summer work have implications for earnings. Figure 15 shows that male teachers earn, on average, $2600 during the summer months from teaching summer school and non-teaching summer jobs (both inside and outside the school). Women, by contrast, earn less than half that amount, controlling for individual, job, and district characteristics. The gender gap in earnings reflects women's lower likelihood

[19]The CPS Outgoing Rotation Groups are asked about their "usual" weekly earnings rather than what they actually earned during the reference week, as would be needed to estimate seasonal fluctuations in earnings.

[20]The gender gap in paid summer work among teachers further corroborates that women's relative decline in summer employment is not an artifact of how education workers report vacation/leave.

Figure 15: Gender gaps in supplemental earnings among teachers

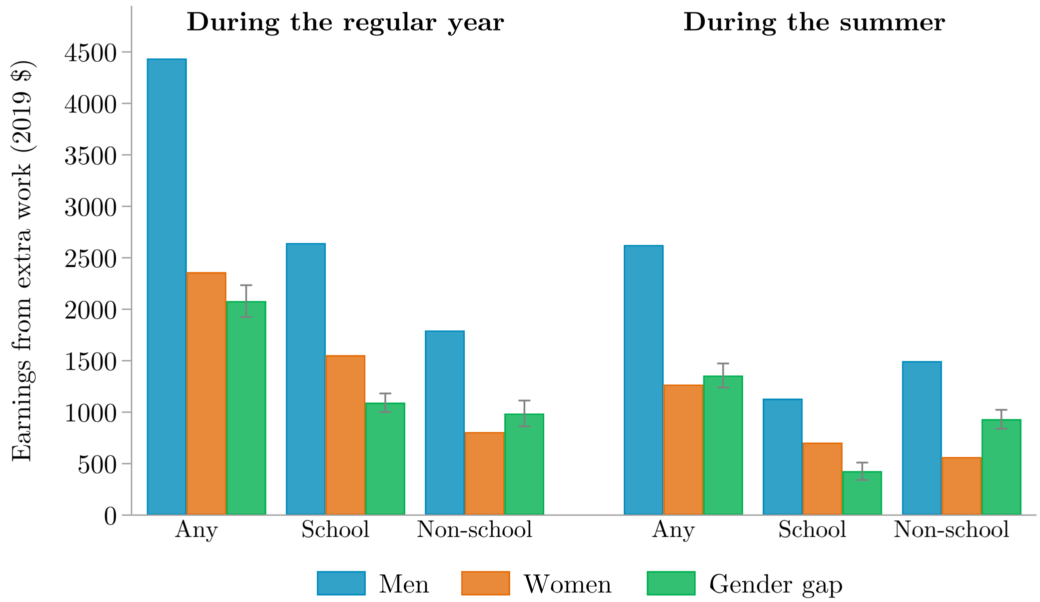

Source: Schools and Staffing Survey.
Notes: Supplemental earnings of full-time teachers during the regular school year and the summer months, by gender, in 2019 dollars. School-based work entails participation in extracurricular activities, coaching, or summer/evening teaching. The regression-adjusted gender gaps control for age categories, teaching experience, race, ethnicity, master's degree, school type (primary, secondary), subject taught, urban status of school, and Census region.

of working at all during the summer as well as the fact that, conditional on working during the summer, women earn less than men.

8 Conclusion

This paper documents pervasive summer declines in women's labor market activity. Extending prior research into the causes and consequences of interruptions to women's careers, we document that the summer season brings with it significant reductions in female employment and a steep reduction in women's total hours worked, especially along the intensive margin. In contrast, men's employment increases slightly during the summer months, and their hours fall half as much.

We establish the central role of school closures in driving gender gaps in summer employment patterns. The summer drop in female employment aligns with cross-state differences in the timing of summer breaks, is concentrated among women with school-age children, and coincides with an uptick in time spent on childcare. A decomposition of the gender gap in summer work interruptions reveals substantial contributions from both gender differences in sorting across jobs with varying degrees of summer flexibility and gender differences within jobs in the propensity to exit employment over the summer. Even within narrowly defined education-related occupations—such as primary school teachers, secondary-school teachers, and school bus drivers—the hazard rate of exiting employment rises more in summer for women than for men. Outside of education, too, women experience a spate of summer work interruptions that find little parallel among men.

While the phenomena we document are longstanding, the school closures induced by the COVID-19 pandemic cast them in a new light. A flurry of recent research has shown how the challenges of juggling work and childcare have weighed heavily on female labor force participation, with potentially severe and long-lasting ramifications for gender disparities in career progression and earnings potential. Pandemic school closures differ in important respects from those that occur on an annual basis: pandemic closures were unanticipated and intermittent, and working parents were limited in their ability to rely on alternative sources of childcare provided by extended family members or the private market. When the pandemic subsides, however, the summer work interruptions wrought by regularly scheduled school closures will likely remain a fixture of the US labor market. The heavy imprint of schools' summer break on female labor force participation, employment, and hours worked raises important questions about the potential need for policy solutions to alleviate the remaining barriers to women's equal participation in the labor market.

References

Albanesi, Stefania and Ayşegül Şahin. 2018. "The Gender Unemployment Gap." *Review of Economic Dynamics*, 30: 47–67.

Albanesi, Stefania and Jiyeon Kim. 2021. "Effects of the COVID-19 Recession on the US Labor Market: Occupation, Family, and Gender." *Journal of Economic Perspectives*, 35(3): 3–24.

Allen, Steven G., Robert L. Clark, and Ann A. McDermed. 1993. "Pensions, Bonding, and Lifetime Jobs." *Journal of Human Resources*, 28(3): 463–481.

Alon, Titan, Sena Coskun, Matthias Doepke, David Koll, and Michèle Tertilt. 2021. "From Mancession to Shecession: Women's Employment in Regular and Pandemic Recessions." NBER Working Paper No. 286322.

Amuedo-Dorantes, Catalina, Miriam Marcén, Marina Morales, and Almudena Sevilla. 2021. "COVID-19 School Closures and Parental Labor Supply in the United States." IZA Discussion Paper Series.

Bardóczy, Bence. 2022. "Spousal Insurance and the Amplification of Business Cycles." Working paper.

Barsky, Robert B. and Jeffrey A. Miron. 1989. "The Seasonal Cycle and the Business Cycle." *Journal of Political Economy*, 97(3): 503–534.

Bertrand, Marianne, Claudia Goldin, and Lawrence F. Katz. 2010. "Dynamics of the Gender Gap for Young Professionals in the Financial and Corporate Sectors." *American Economic Journal: Applied Economics*, 2(3): 228–255.

Bolotnyy, Valentin and Natalia Emanuel. 2022. "Why Do Women Earn Less than Men? Evidence from Bus and Train Operators." *Journal of Labor Economics*, 40(2): 283–323.

Cajner, Tomaz, John Coglianese, and Joshua Montes. 2021. "The Long-Lived Cyclicality of the Labor Force Participation Rate." Federal Reserve Board Finance and Economics Discussion Series No. 2021-047.

Cascio, Elizabeth U. 2009. "Maternal Labor Supply and the Introduction of Kindergartens into American Public Schools." *Journal of Human Resources*, 44(1): 140–170.

Coglianese, John and Brendan M. Price. 2020. "Income in the Off Season: Household Adaptation to Yearly Work Interruptions." Federal Reserve Board Finance and Economics Discussion Series No. 2020-084.

Cortés, Patricia and Jessica Pan. 2019. "When Time Binds: Substitutes for Household Production, Returns to Working Long Hours, and the Skilled Gender Wage Gap." *Journal of Labor Economics*, 37(2): 351–398.

Couch, Kenneth A., Robert W. Fairlie, and Huanan Xu. 2022. "The Evolving Impacts of the COVID-19 Pandemic on Gender Inequality in the US Labor Market: The COVID Motherhood Penalty." *Economic Inquiry*.

Cubas, German, Chinhui Juhn, and Pedro Silos. 2019. "Coordinated Work Schedules and the Gender Wage Gap." NBER Working Paper No. 26548.

Del Bono, Emilia and Andrea Weber. 2008. "Do Wages Compensate for Anticipated Working Time Restrictions? Evidence from Seasonal Employment in Austria." *Journal of Labor Economics*, 26(1): 181–221.

Duchini, Emma and Clémentine Van Effenterre. 2022. "School Schedule and the Gender Pay Gap." *Journal of Human Resources*, 0121-11431R2.

Dupraz, Stéphane, Emi Nakamura, and Jón Steinsson. 2019. "A Plucking Model of Business Cycles." NBER Working Paper No. 26351.

Fitzpatrick, Maria Donovan. 2012. "Revising Our Thinking About the Relationship Between Maternal Labor Supply and Preschool." *Journal of Human Resources*, 47(3): 583–612.

Flood, Sarah, Miriam King, Renae Rodgers, Steven Ruggles, J. Robert Warren, and Michael Westberry. 2021. "Integrated Public Use Microdata Series, Current Population Survey: Version 9.0 [dataset]." Minneapolis, MN: IPUMS. https://doi.org/10.18128/D030.V9.0.

Frazis, Harley J., Edwin L. Robison, Thomas D. Evans, and Martha A. Duff. 2005. "Estimating Gross Flows Consistent with Stocks in the CPS." *Monthly Labor Review*, 128: 3–9.

Furman, Jason, Melissa Schettini Kearney, and Wilson Powell. 2021. "The Role of Childcare Challenges in the US Jobs Market Recovery During the COVID-19 Pandemic." NBER Working Paper No. 28934.

Garcia, Kairon Shayne D. and Benjamin W. Cowan. 2022. "The Impact of School and Childcare Closures on Labor Market Outcomes during the COVID-19 Pandemic." NBER Working Paper No. 29641.

Gelbach, Jonah. 2002. "Public Schooling for Young Children and Maternal Labor Supply." *American Economic Review*, 92(1): 307–322.

Geremew, Menelik and François Gourio. 2018. "Seasonal and Business Cycles of US Employment." *Economic Perspectives*, 42(3): 1–28.

Goldin, Claudia. 2022. "Understanding the Economic Impact of COVID-19 on Women." NBER Working Paper No. 29974.

Graves, Jennifer. 2013. "School Calendars, Child Care Availability, and Maternal Employment." *Journal of Urban Economics*, 78: 57–70.

Guryan, Jonathan, Erik Hurst, and Melissa Kearney. 2008. "Parental Education and Parental Time with Children." *Journal of Economic Perspectives*, 22(3): 23–46.

Hamermesh, Daniel S., Harley Frazis, and Jay Stewart. 2005. "Data Watch: The American Time Use Survey." *Journal of Economic Perspectives*, 19(1): 221–232.

Handwerker, Elizabeth Weber and Lowell Mason. 2017. "BLS Spotlight on Statistics: Differences in Parents' Time Use Between the Summer and the School Year." Available at https://www.bls.gov/spotlight/2017/differences-in-parents-time-use-between-the-summer-and-the-school-year/pdf/differences-in-parents-time-use-between-the-summer-and-the-school-year.pdf.

Hansen, Benjamin, Joseph J. Sabia, and Jessamyn Schaller. 2022. "Schools, Job Flexibility, and Married Women's Labor Supply: Evidence From the COVID-19 Pandemic." NBER Working Paper No. 29660.

Heggeness, Misty L. 2020. "Estimating the Immediate Impact of the COVID-19 Shock on Parental Attachment to the Labor Market and the Double Bind of Mothers." *Review of the Economics of the Household*, 18(4): 1053–1078.

Hobijn, Bart and Ayşegül Şahin. 2021. "Maximum Employment and the Participation Cycle." NBER Working Paper No. 29222.

Hofferth, Sandra L., Sarah M. Flood, and Matthew Sobek. 2020. "American Time Use Survey Data Extract Builder: Version 2.8 [dataset]." College Park, MD: University of Maryland and Minneapolis, MN: IPUMS. https://doi.org/10.18128/D060.V2.8.

Juhn, Chinhui and Simon Potter. 2006. "Changes in Labor Force Participation in the United States." *Journal of Economic Perspectives*, 20(3): 27–46.

Katz, Lawrence F. and Bruce D. Meyer. 1990. "Unemployment Insurance, Recall Expectations, and Unemployment Outcomes." *Quarterly Journal of Economics*, 105(4): 973–1002.

Lazear, Edward P. 1981. "Agency, Earnings Profiles, Productivity, and Hours Restrictions." *American Economic Review*, 71(3): 606–620.

Madrian, Brigitte C. and Lars John Lefgren. 2000. "An Approach to Longitudinally Matching Current Population Survey (CPS) Respondents." *Journal of Economic and Social Measurement*, 26(1): 31–62.

Mas, Alexandre and Amanda Pallais. 2017. "Valuing Alternative Work Arrangements." *American Economic Review*, 107(12): 3722–3759.

Miron, Jeffrey A. and J. Joseph Beaulieu. 1996. "What Have Macroeconomists Learned about Business Cycles from the Study of Seasonal Cycles?" *Review of Economics and Statistics*, 78(1): 54–66.

Montes, Joshua, Christopher Smith, and Isabel Leigh. 2021. "Caregiving for Children and Parental Labor Force Participation During the Pandemic." FEDS Notes. Washington: Board of Governors of the Federal Reserve System.

Moretti, Enrico. 2000. "Do Wages Compensate for Risk of Unemployment? Parametric and Semiparametric Evidence from Seasonal Jobs." *Journal of Risk and Uncertainty*, 20(1): 45–66.

Newey, Whitney K. and Kenneth D. West. 1987. "A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica*, 55(3): 703–708.

Newey, Whitney K. and Kenneth D. West. 1994. "Automatic Lag Selection in Covariance Matrix Estimation." *Review of Economic Studies*, 61(4): 631–653.

Ngai, L. Rachel and Silvana Tenreyro. 2014. "Hot and Cold Seasons in the Housing Market." *American Economic Review*, 104(12): 3991–4026.

Olivei, Giovanni and Silvana Tenreyro. 2007. "The Timing of Monetary Policy Shocks." *American Economic Review*, 97(3): 636–663.

Parent, Daniel. 2000. "Industry-Specific Capital and the Wage Profile: Evidence from the National Longitudinal Survey of Youth and the Panel Study of Income Dynamics." *Journal of Labor Economics*, 18(2): 306–323.

Podgursky, Michael. 2011. "Teacher Compensation and Collective Bargaining." *Handbook of the Economics of Education*, , ed. Eric A. Hanushek, Stephen Machin, and Ludger Woessmann Vol. 3, 279–313. Elsevier.

Rivera Drew, Julia A., Sarah Flood, and John Robert Warren. 2014. "Making Full Use of the Longitudinal Design of the Current Population Survey: Methods for Linking Records Across 16 Months." *Journal of Economic and Social Measurement*, 39(3): 121–144.

Russell, Lauren and Chuxuan Sun. 2020. "The Effect of Mandatory Child Care Center Closures on Women's Labor Market Outcomes During the COVID-19 Pandemic." *Covid Economics: Vetted and Real-Time Papers*, 62: 124–154.

Wasserman, Melanie. 2022. "Hours Constraints, Occupational Choice, and Gender: Evidence from Medical Residents." *Review of Economic Studies*, forthcoming.

Wiswall, Matthew and Basit Zafar. 2018. "Preference for the Workplace, Investment in Human Capital, and Gender." *Quarterly Journal of Economics*, 133(1): 457–507.

A Additional Figures and Tables

Appendix Figure A.1: Summer changes in employment-to-population ratios, 1989–2019

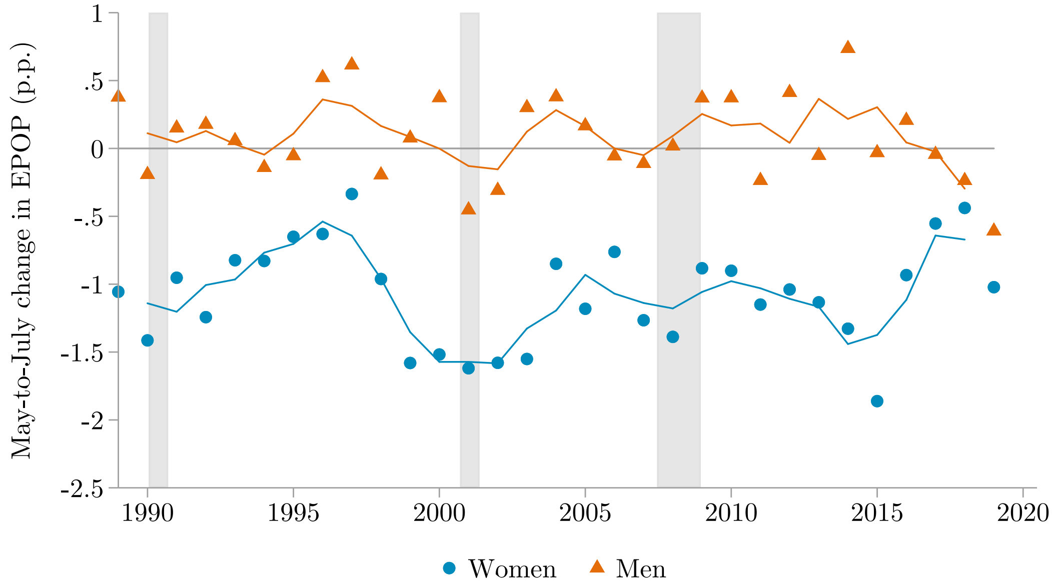

Source: Current Population Survey.
Notes: Plotted points show unadjusted May–July changes in EPOP in our sample of prime-age individuals. Smoothed curves show three-year centered moving averages. Shading denotes recessions, as dated by the NBER.

Appendix Figure A.2: Demographic heterogeneity in the May–July change in EPOP

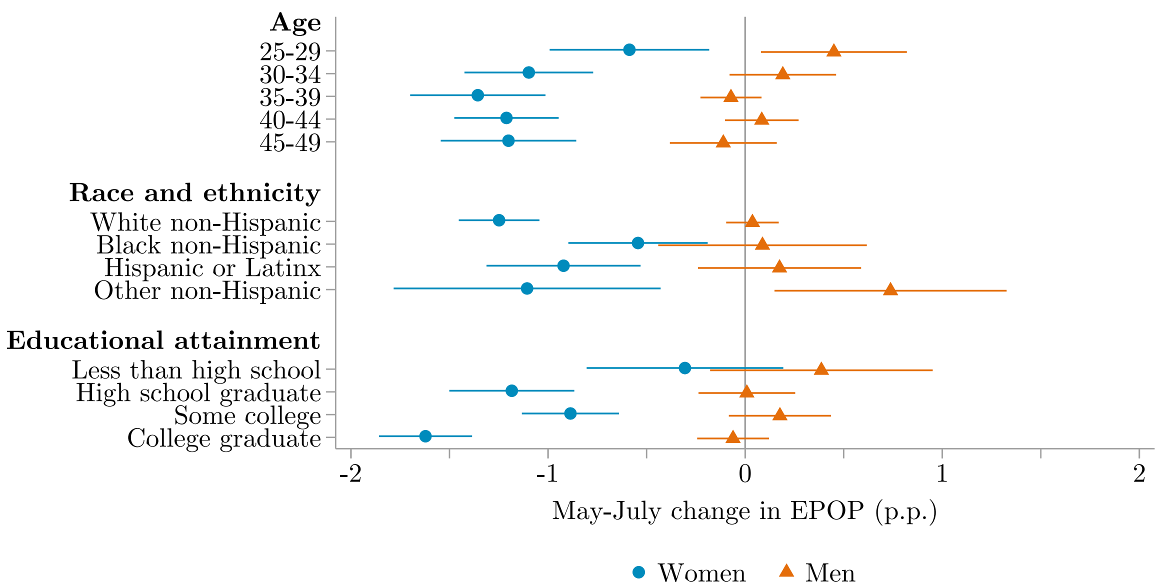

Source: Current Population Survey.
Notes: Coefficients $\widehat{\beta}_7$ (representing May–July changes) from estimating Equation (1) separately by sex × the indicated characteristic. See Appendix B.1 for details on our coding of race, ethnicity, and educational attainment. Bars show 95 percent confidence intervals based on Newey-West standard errors.

Appendix Figure A.3: Excess recurrence of work interruptions 12 months after an initial one

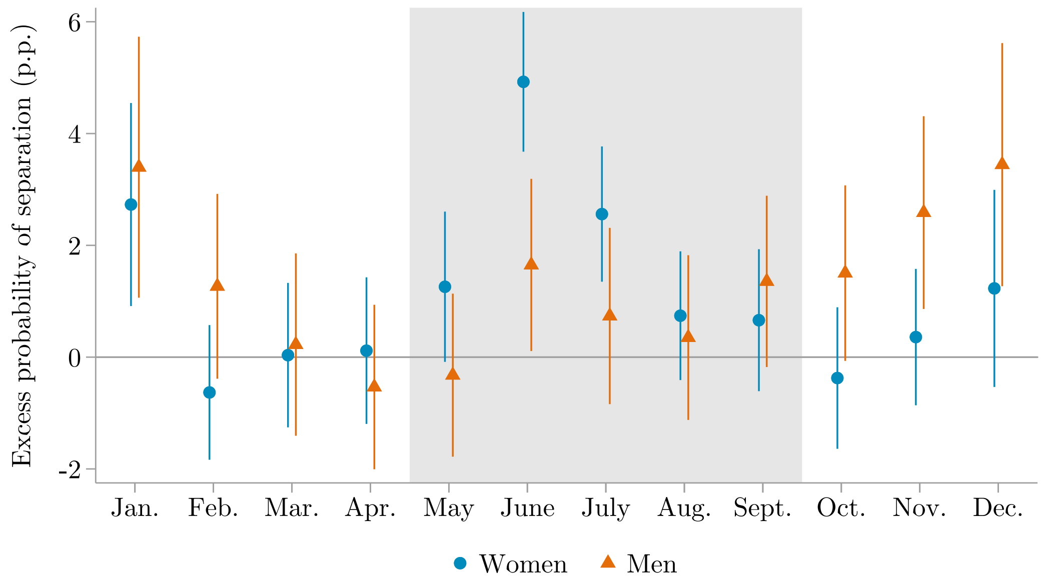

Source: Current Population Survey.
Notes: Excess recurrence of work interruptions at annual intervals, as defined by Coglianese and Price (2020) and obtained by estimating $\widehat{\rho}_{12} - \frac{1}{2}(\widehat{\rho}_{11} + \widehat{\rho}_{13})$ in Equation (21) using a sample of CPS respondents ages 25–49. Bars show 95 percent confidence intervals, with standard errors clustered at the household level. See Appendix E.1 for details.

Appendix Figure A.4: Validation of linear splines fitted to prime-age EPOP and LFPR

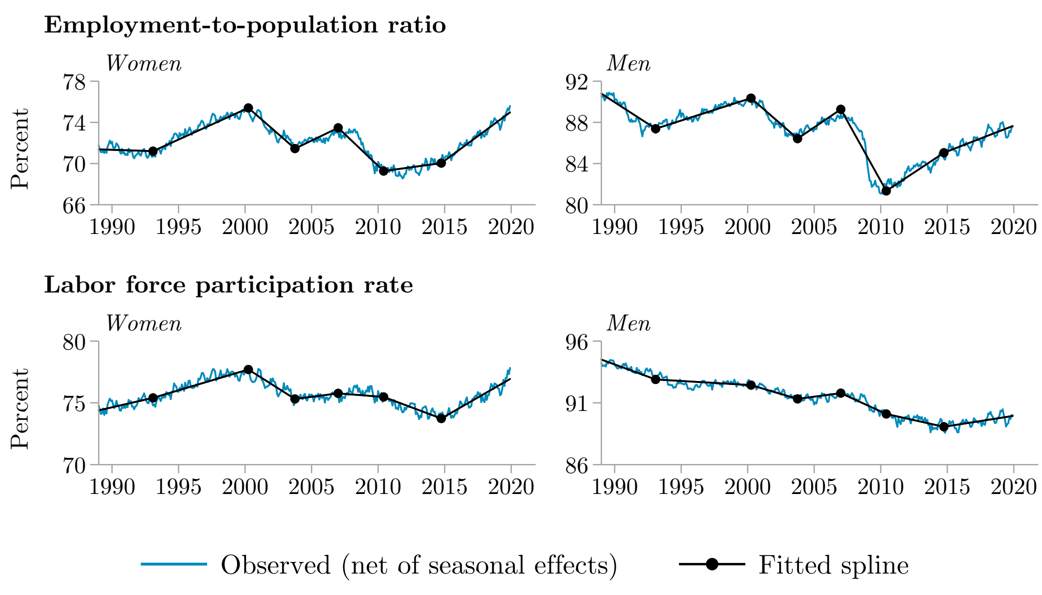

Source: Current Population Survey.
Notes: Series labeled "observed" plot the indicated measure net of estimated month effects and weeks-between-reference-weeks effects. Series labeled "fitted spline" plot the model predictions using only the linear spline, which places knots at select turning points in prime-age EPOP and LFPR. See Appendix B.2 for details.

Appendix Figure A.5: Cross-state distribution of the share of the total May–July drop in high school enrollment observed by the June reference week

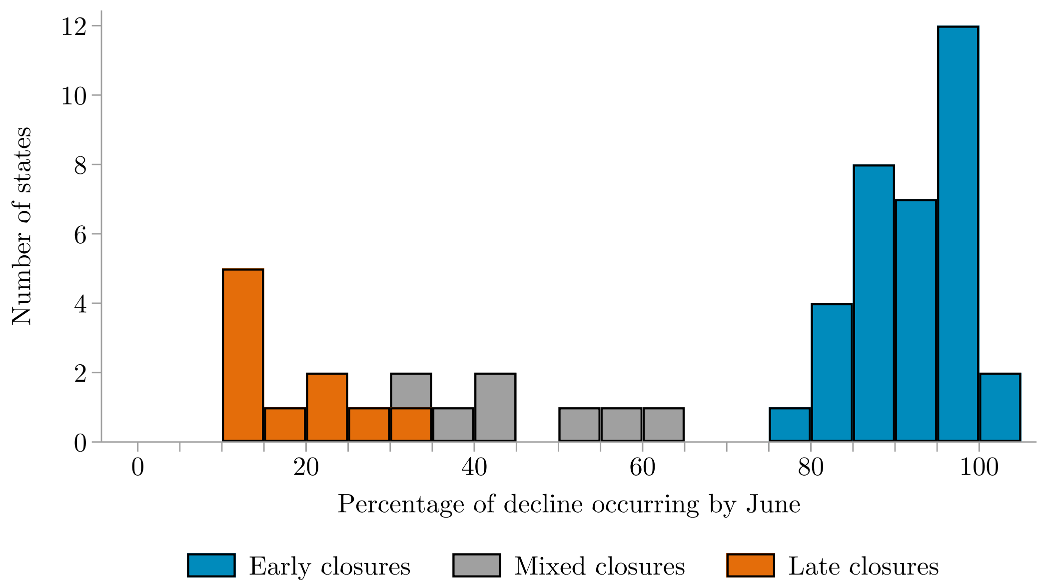

Source: Current Population Survey.
Notes: The statistic shown represents the share of the total May–July drop in high school enrollment among 16-year-old CPS respondents that occurs by the June CPS reference week. "Early-closure" states are those in which this statistic exceeds two thirds; "late-closure" states are those in which it falls short of one third. The remaining states are classified as "mixed-closure" states.

Appendix Figure A.6: Classification of US states by the timing of K–12 school closures

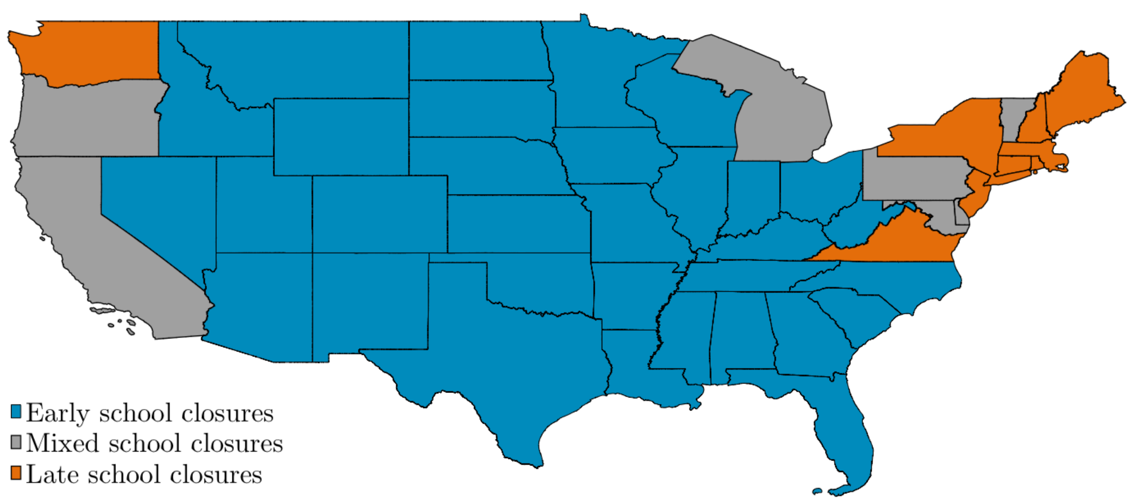

Source: Current Population Survey; US Census Bureau shapefiles.
Notes: See notes to Appendix Figure A.5. Alaska and Hawaii (not shown) are classified as early-closure states.

Appendix Figure A.7: Evolution of May–July employment gaps over the life cycle

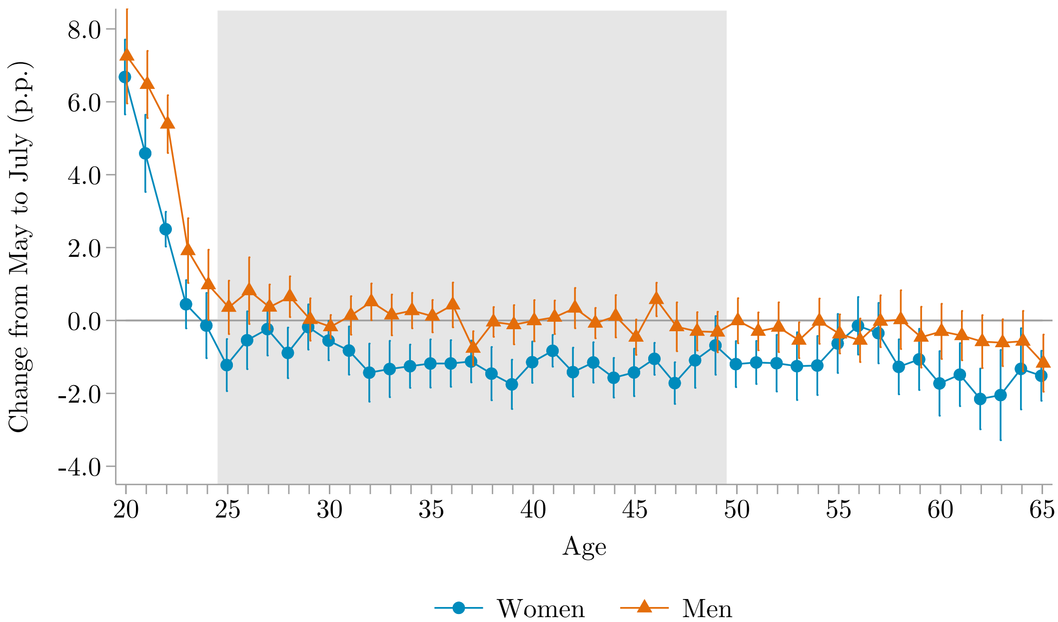

Source: Current Population Survey.
Notes: Coefficients $\widehat{\beta}_7$ (representing May–July changes) from estimating Equation (1) separately by sex × one-year age bins. The shaded region denotes the age range used in our main estimation sample. Bars show 95 percent confidence intervals based on Newey-West standard errors.

Appendix Figure A.8: Seasonal patterns in male employment by household structure

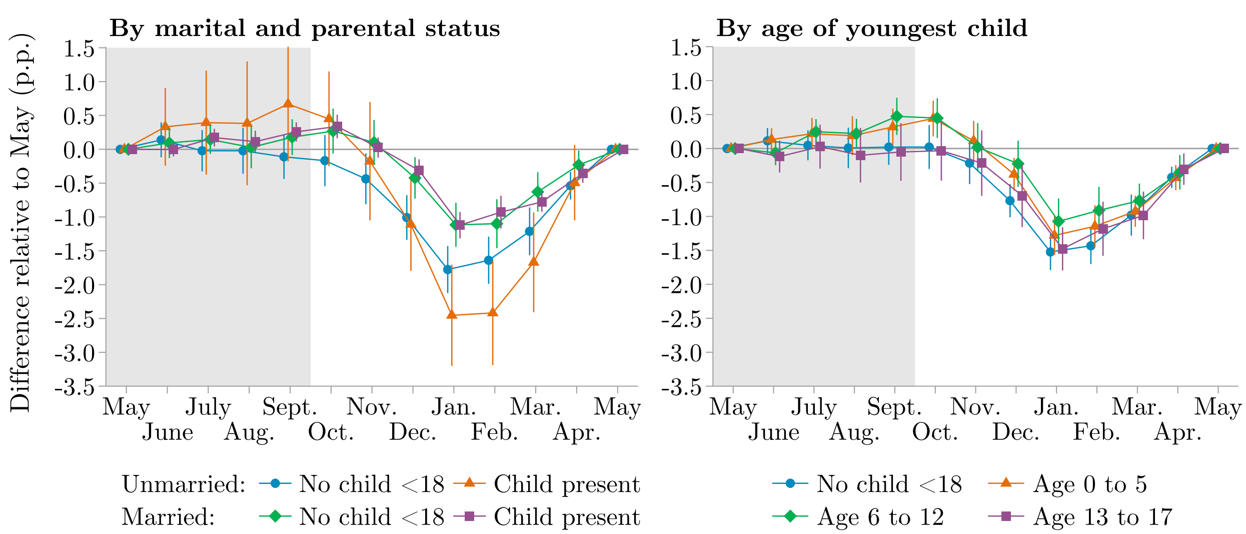

Source: Current Population Survey.
Notes: Coefficients $\widehat{\beta}_m$ from estimating Equation (1) for male EPOP separately by sex, marital, and parental status. Separated individuals are coded as married; parental status is defined in reference to an individual's own children, including adoptees and step-children but excluding younger siblings and other children not one's own. Bars show 95 percent confidence intervals based on Newey-West standard errors.

Appendix Figure A.9: Hazard rate of switching to zero hours worked among education workers

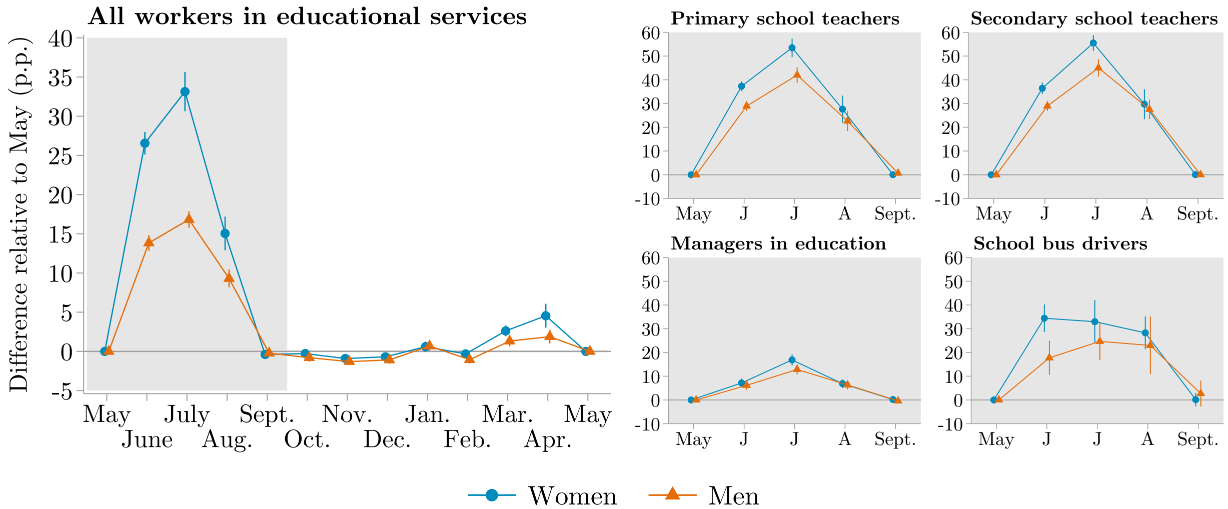

Source: Current Population Survey.

Notes: Coefficients $\widehat{\beta}_m$ from estimating Equation (2) in a sample of CPS respondents ages 25–49 with valid longitudinal links. The sample is restricted to individuals who worked positive hours in educational services during the previous month's reference week, either in any occupation (left panel) or in the indicated occupation (right panel); the outcome is the percentage of these individuals who worked zero hours in the current month's reference week (whether non-employed or absent). Bars show 95 percent confidence intervals based on Newey-West standard errors. In the right panel, coefficients for October–April are estimated but not shown.

Appendix Table A.1: Summary statistics for the CPS estimation sample

	All prime-age		Parents (child 6–12)	
	Women	**Men**	**Mothers**	**Fathers**
	(1)	(2)	(3)	(4)
Demographics				
Age	37.0	36.9	38.2	39.9
	(7.1)	(7.1)	(5.7)	(5.5)
Married	65.6	63.3	75.4	89.9
Own child < 18 in household	59.6	49.0	100.0	100.0
Youngest < 6 years old	26.2	24.2	0.0	0.0
Youngest 6–12 years old	22.0	16.9	100.0	100.0
Youngest 13–17 years old	11.3	7.9	0.0	0.0
Labor market activity				
Employed	71.9	86.8	71.9	91.3
At work during reference week	68.4	84.2	68.6	88.6
Absent during reference week	3.5	2.6	3.4	2.7
Unemployed	3.8	4.6	4.1	3.7
Temporary layoff	0.4	0.8	0.5	0.8
Other reason unemployed	3.4	3.8	3.6	2.9
Not in labor force	24.3	8.6	24.0	5.0
Not in labor force (1994 or later)	24.2	9.0	23.7	5.2
Taking care of house or family	16.5	1.3	18.2	1.2
Other major activity	7.7	7.7	5.5	4.0
Hours worked in reference week	25.6	36.6	24.7	39.4
	(20.0)	(19.4)	(19.5)	(18.2)
Observations	9,033,776	8,351,163	2,005,503	1,443,127

Source: Current Population Survey.
Notes: The sample consists of individuals aged 25–49. All statistics are sample means, with standard deviations reported in parentheses for non-binary variables. All statistics other than age and hours worked are expressed as percentages. Columns (3) and (4) restrict to parents whose youngest child residing in the household is aged 6–12. Observations are weighted to obtain representative estimates for the prime-age US population.

Appendix Table A.2: Decomposition of female-male differences in the seasonality of EPOP

Component	May	June	July	Aug.	Sept.	Oct.	Nov.	Dec.	Jan.	Feb.	Mar.	Apr.	May
Overall change in gender gap	0.000	-0.716	-1.217	-1.071	-0.216	0.117	0.414	0.764	0.987	1.074	0.834	0.437	0.000
	(0.000)	(0.087)	(0.098)	(0.135)	(0.151)	(0.140)	(0.105)	(0.104)	(0.114)	(0.107)	(0.132)	(0.091)	(0.000)
Sorting into the education sector	0.000	-0.269	-0.385	-0.319	-0.031	0.012	0.041	0.043	0.054	0.095	0.075	0.045	0.000
	(0.000)	(0.014)	(0.016)	(0.015)	(0.013)	(0.010)	(0.011)	(0.011)	(0.014)	(0.013)	(0.008)	(0.007)	(0.000)
Sorting across non-education sectors	0.000	-0.114	-0.193	-0.229	-0.225	-0.168	-0.031	0.137	0.304	0.373	0.335	0.145	0.000
	(0.000)	(0.027)	(0.039)	(0.045)	(0.042)	(0.031)	(0.039)	(0.050)	(0.060)	(0.063)	(0.052)	(0.032)	(0.000)
Sorting across education-sector jobs	0.000	-0.058	-0.089	-0.066	-0.028	-0.037	-0.034	-0.029	-0.008	-0.032	-0.029	-0.035	0.000
	(0.000)	(0.006)	(0.008)	(0.009)	(0.008)	(0.008)	(0.007)	(0.009)	(0.008)	(0.007)	(0.006)	(0.005)	(0.000)
Within education-sector jobs	0.000	-0.186	-0.316	-0.227	-0.009	0.041	0.047	0.045	0.017	0.051	0.019	0.011	0.000
	(0.000)	(0.020)	(0.021)	(0.032)	(0.028)	(0.024)	(0.025)	(0.024)	(0.025)	(0.022)	(0.020)	(0.017)	(0.000)
Within sectors other than education	0.000	-0.136	-0.312	-0.308	0.100	0.327	0.421	0.544	0.484	0.480	0.359	0.249	0.000
	(0.000)	(0.093)	(0.119)	(0.137)	(0.138)	(0.124)	(0.102)	(0.111)	(0.103)	(0.098)	(0.122)	(0.088)	(0.000)
Baseline EPOP component	0.000	0.047	0.078	0.078	-0.023	-0.058	-0.030	0.022	0.135	0.107	0.076	0.022	0.000
	(0.000)	(0.008)	(0.010)	(0.013)	(0.011)	(0.013)	(0.012)	(0.011)	(0.011)	(0.014)	(0.011)	(0.008)	(0.000)

Source: Current Population Survey.
Notes: Point estimates and standard errors associated with the decomposition presented in Figure 12. Decomposition components are derived from a set of job-specific flow specifications (as in Equation (2)), estimated separately by sex. See Appendix C.2 for details on the decomposition procedure, which employs Oaxaca-Blinder techniques, and calculation of standard errors, which we obtain by stacking regression models in the manner of seemingly unrelated regression.

Appendix Table A.3: Gender differences in sectoral and occupational sorting

		% in sector/occ Women	Men	Pr(E → N)
Sector:				
(1)	Education	13.5	4.8	7.2
(2)	Non-education	86.5	95.2	3.7
	Total	100.0	100.0	4.0
Occupation in education sector:				
(3)	Primary school teacher	27.8	15.3	8.7
(4)	Secondary school teacher	8.9	17.4	7.0
(5)	Other non-college teacher	17.8	8.2	9.6
(6)	College teacher	5.4	14.2	6.5
(7)	Administrative staff	13.6	3.1	7.6
(8)	Managers	7.9	11.3	2.5
(9)	Food/trans./cleaning services	7.3	9.2	9.2
(10)	Other	11.3	21.3	5.4
	Total education	100.0	100.0	7.2

Source: Current Population Survey.
Notes: Employment shares and separation hazards calculated among individuals employed as of May and observed in July of the same year. The last column reports the percentage of these individuals who were non-employed as of July (computed as the average of the female and male shares non-employed). All statistics are averaged over the 1989–2019 analysis period.

B Details on Data Preparation

B.1 Current Population Survey

Our analysis draws primarily on basic monthly CPS extracts provided by IPUMS (Flood et al., 2021). We also use the Annual Social and Economic Supplement (ASEC), which accompanies the March CPS.

Sample restrictions. We limit our analysis to individuals aged 25–49. We further exclude members of the armed forces, who are not counted towards the official unemployment rate and for whom we do not observe key labor market variables (such as hours worked).

Longitudinal linkages. We link CPS observations across individuals over time and across individuals in the same household using the IPUMS variables `cpsidp` and `cpsid`, respectively. We lack reliable linkages in mid-1995, owing to changes in the CPS household identifiers.

IPUMS cautions that `cpsidp` sometimes yields erroneous links stemming from errors in data collection and advises researchers to validate individual linkages using age, sex, and race. For month-to-month analyses, we exclude individuals whose observed sex, race, or ethnicity differs between consecutive months, as well as those whose age differs by more than two years.[1] For the analysis of annually recurrent work interruptions in Appendix E.1, we exclude individuals for whom we observe inconsistencies in any of these variables at any point in their tenure in the survey.

Sampling weights. For cross-sectional analyses, we weight observations using the variables `wtfinl` (for analyses using basic monthly CPS extracts) and `asecwt` (for analyses using the March ASEC). For longitudinal analyses, we weight observations using raked weights we compute ourselves. Although IPUMS provides a set of raked longitudinal weights that align gross labor market flows with stocks in the full set of adult CPS respondents, this equivalence holds only in the aggregate and breaks down once we restrict to prime-age individuals. Adapting replication files supplied by IPUMS, we use the Stata package `ipfraking` to construct raked weights via iterative proportional fitting, separately by sex and separately for each pair of consecutive months. Applying these raked weights to the set of longitudinally linkable individuals in our sample yields gross flows among employment, unemployment, and non-participation consistent with observed changes in the stock of individuals in each status in our full cross-sectional sample.

Demographic characteristics.

- **Marital status:** We code individuals as married or unmarried using the variable `marst`. We treat separated but non-divorced individuals as married.
- **Parental status:** We link children in each household to their parents—whether biological, adoptive, or step-parents—using the variables `momloc`, `poploc`, `momloc2`, and `poploc2`, which encompass both same- and opposite-sex couples. We take particular note of the age of each adult respondent's youngest child in the household and bin parents into four groups: youngest child is under 6 years old, youngest child is 6–12 years old, youngest child is 13–17 years old,

[1] Although gender and racial identity can evolve over time, changes in these variables in the brief periods between CPS survey rounds are more likely to reflect distinct respondents than changes in self-identification. We allow for slight inconsistencies in the reporting of age because, among observations exhibiting logically impossible combinations of lagged and current age, very slight discrepancies are disproportionately common, suggesting that in many cases the same individual is in fact being observed on both occasions.

or there is no child under 18 in the household. These age cutoffs mirror similar groupings used by IPUMS in its preparation of the ATUS data.

- **Educational attainment:** We classify individuals into four educational categories—“less than a high school degree”, “high school degree”, “some college”, and “college degree or higher”—using the variable `educ`, which is populated both before and after changes to the underlying CPS questions in 1992.

- **Race and ethnicity:** We classify individuals as “White non-Hispanic”, “Black non-Hispanic”, “Hispanic”, or “other non-Hispanic” using the variables `race` and `hispan`.

Sectors. We define the educational services sector using the variable `ind1990`, which bridges changes over time in the CPS industry codes. Educational services encompasses five industry codes:

- 842: Elementary and secondary schools
- 850: Colleges and universities
- 851: Vocational schools
- 852: Libraries
- 860: Educational services, n.e.c.

The decomposition in Figure 12 partitions non-education jobs into 13 sectors, delineated by capitalized headers in the IPUMS codebook. These are: “Agriculture, forestry, and fishing”; “Mining”; “Construction”; “Manufacturing”; “Transportation, communication, and utilities”; “Wholesale trade”; “Retail trade”; “Finance, insurance, and real estate”; “Business and repair services”; “Personal services”; “Entertainment and recreational services”; “Professional services” (excluding educational services); and “Public administration”.

Occupations. We use the variable `occ1990`, which harmonizes CPS occupation codes over time.

Reference week timing. The CPS reference week usually, but not always, straddles the 12th day of the month. We calculate the number of weeks elapsed between successive CPS reference weeks by following BLS guidance:

1. Define the reference week as the 7-day calendar week (Sunday to Saturday) that includes the 12th day of the month.

2. Shift the December reference week one week earlier if the calendar week that includes December 5 would otherwise be contained entirely within the month of December.

3. Shift the November reference week one week earlier if Thanksgiving falls during the week containing November 19.[2]

In our CPS sample, reference weeks are spaced four weeks apart in 63.7 percent of observations, five weeks apart in 33.9 percent of observations, and three or six weeks apart in the remainder.

[2]According to the BLS, the Census Bureau sometimes advances the November reference week by one week in other years as well, when it determines that there is not enough time to process the data before December interviews begin. We do not observe these judgmental deviations and thus do not adjust for them.

Annual Social and Economic Supplement (ASEC). For the analysis of the education sector earnings premium/penalty, we use the CPS Annual Social and Economic Supplement (ASEC), 1989–2019, which is administered in March of each year. The supplement includes respondents' annual income derived from wage and salary income (variable `incwage`). We trim extremely low values of annual income, equivalent to earning less than the nadir of the minimum wage over our sample period, working 10 hours per week, and working 20 weeks per year. We use this income information alongside the respondent's industry and occupation during the previous year to compute the regression-adjusted education-sector earnings premium or penalty in each occupation. The regression controls for educational attainment, a quadratic in age, and calendar-year fixed effects, in addition to occupation fixed effects and their interactions with the education sector.

B.2 Choosing knots for the linear spline

Our workhorse specifications in Equations (1) and (2) control for a linear spline in calendar time. To motivate this approach, suppose first that a given outcome variable (such as female EPOP) contains a linear time trend. Because our analysis period runs from January 1989 through December 2019, later months in the year tend to occur slightly later in calendar time, so that a naïve regression on month dummies alone would be biased in proportion to the degree of secular drift. In addition, one might worry that turning points in the business cycle happen to occur at particular points in the calendar year. To address these potential biases, and to improve the precision of our estimates, we use a flexible spline function with knots at key turning points in the business cycle.

Our choice of knots is inspired by recent research on the cyclical properties of unemployment and labor force participation. Dupraz, Nakamura, and Steinsson (2019) note that turning points in the unemployment rate do not align perfectly with official business cycle dates from the National Bureau of Economic Research, while Cajner, Coglianese, and Montes (2021) and Hobijn and Şahin (2021) document the sluggish response of the labor force participation rate (LFPR) to cyclical conditions, especially in the wake of the Great Recession. Motivated by these observations, we adopt a data-driven approach that locates knots tailored to prime-age EPOP and LFPR:

1. We start with an algorithm from Dupraz, Nakamura, and Steinsson (2019, hereafter DNS) that locates turning points in the US unemployment rate by searching for local extrema while ignoring small fluctuations within a tolerance band. Adapting their replication code, we locate turning points in seasonally adjusted EPOP for ages 25–54, as published by the Bureau of Labor Statistics (Labor Force Statistics series LNS12300060).[3]

2. The DNS procedure yields six turning points that fall within our 1989–2019 analysis period: January 1990, February 1993, April 2000, October 2003, January 2007, and June 2010. We discard the January 1990 turning point since it falls near the edge of our period.

3. Although a linear spline with these knots can effectively capture broad movements in prime-age EPOP, it imposes a linear trend for the last decade of our analysis period, and it misses an important turning point in prime-age participation during the mid-2010s. To remedy these defects, we rerun the DNS algorithm using the seasonally adjusted LFPR for ages 25–54 (Labor Force Statistics series LNS11300060) and retain the turning point in October 2014.

We end up with six knots: February 1993, April 2000, October 2003, January 2007, June 2010, and October 2014. Besides corresponding to notable inflection points in prime-age labor market

[3]The DNS algorithm deals with the possibility of "ties" by selecting the earliest peak or trough within a given expansion or contraction. We depart slightly from their procedure by instead taking the midpoint between the earliest and latest candidate inflection points (rounding up to the nearest month when needed).

conditions, these knots are situated roughly five years apart and hence serve as a flexible means to model trend movements in other outcomes we examine as well.

To verify that our chosen knots satisfy their intended function, Appendix Figure A.4 plots our fitted splines from estimation of Equation (1) for prime-age EPOP and LFPR, separately by sex, against the observed time series, net of our estimated month effects and weeks-between-reference-weeks effects. The close correspondence between these series indicates that the model residuals are systematically modest through the ups and downs of the business cycle.

B.3 American Time Use Survey (ATUS)

Launched in 2003, the ATUS (another BLS product) surveys a random subset of outgoing CPS respondents a few months after their final CPS interview (Hamermesh, Frazis, and Stewart, 2005; Guryan, Hurst, and Kearney, 2008). One randomly selected adult member of the household is asked to provide a detailed, minute-by-minute accounting of their activities throughout the previous day. Paralleling our CPS sample, we assemble IPUMS ATUS data on individuals ages 25–49 over the period 2004–2019 (Hofferth, Flood, and Sobek, 2020); we discard 2003 because of issues with data completeness. We exclude respondents with incomplete time diaries, so that time allocations sum to 24 hours. We multiply minutes spent on each activity by $7/60$, so that our measures are expressed in terms of hours per week.

We examine both narrow and broad measures of time allocated to childcare activities. First, we follow Guryan, Hurst, and Kearney (2008) in constructing a measure of "primary" childcare, defined as intervals of time in which the respondent was mainly engaged in childcare activities. Second, we compute "total" childcare by adding in the ATUS measure of secondary childcare, defined as time spent engaging in childcare concurrently with some other primary activity. We exploit the granular structure of the ATUS to decompose secondary childcare according to whether it accompanies household activities, leisure activities, or other activities.

The ATUS diary dates are distributed evenly throughout the year, but weekends are deliberately oversampled. We employ IPUMS sampling weights that adjust for both cross-household and day-to-day differences in sampling probability, so that our estimates are representative of prime-age adults' time allocation throughout the week as well as the year.

Primary childcare time. Our definition of primary childcare time follows Guryan, Hurst, and Kearney (2008), who write:

> *We define "total child care" as the sum of four primary time use components. "Basic" child care is time spent on the basic needs of children, including breast-feeding, rocking a child to sleep, general feeding, changing diapers, providing medical care (either directly or indirectly), grooming, and so on. However, time spent preparing a child's meal is included in general "meal preparation," a component of nonmarket production. "Educational" child care is time spent reading to children, teaching children, helping children with homework, attending meetings at a child's school, and similar activities. "Recreational" child care involves playing games with children, playing outdoors with children, attending a child's sporting event or dance recital, going to the zoo with children, and taking walks with children. "Travel" child care is any travel related to any of the three other categories of child care. For example, driving a child to school, to a doctor, or to dance practice are all included in "travel" child care.*

We identify the ATUS activities matching these verbal descriptions and use them to construct measures of basic, educational, recreational, and travel childcare, then sum these measures to obtain primary childcare.

Secondary childcare time. Alongside each person × primary activity observation, the ATUS reports whether the respondent had a child under age 13 in their care while engaging in that activity. Following our definition of parental status, we use a measure of secondary childcare that counts only instances when the child under an adult's care is the parent's own child. We define total childcare as the sum of primary and secondary childcare. To shed additional light on seasonal changes in time use, we also partition time spent on secondary childcare according to the primary activity it accompanies:

1. Household activities, a category reported directly in the ATUS;
2. Leisure activities, which we define as the union of the ATUS categories "socializing, relaxing, and leisure", "sports, exercise, and recreation", and "traveling"; and
3. All other activities.

Data quality and completeness. We exclude observations with data quality flags (which note, for example, cases in which a respondent intentionally provided a wrong answer or could not remember their activities), as well as those with incomplete time diaries (cases in which total time usage sums to less than 24 hours).

C Decomposition Details

In this appendix, we derive two key decompositions used in the main text. First, we show how seasonal changes in employment rates can be decomposed into contributions from inflows versus outflows (Figure 3). Second, we show how gender differences in employment seasonality can be decomposed into gender differences in job sorting as well as gender differences conditional on job type (Figure 12).

Notation. We begin by introducing notation common to both decompositions.

- Let $g \in \{♀ \text{ (female)}, ♂ \text{ (male)}\}$ index gender.
- Let $m \in \{0, 1, \ldots, 12\}$ index calendar months relative to the base month 0, which we take to be May. We sometimes use $m = 12$ as an alternative label for the base month.
- Let e_{gm} denote group g's EPOP in month m. Let f_{gm} and s_{gm} denote the shares of each population *finding* or *separating from* employment in month m, and let $n_{gm} \equiv f_{gm} - s_{gm}$.

  We refer to these shares as *inflows*, *outflows*, and *net inflows*, respectively. Since our empirical implementation implicitly averages across years after netting out low-frequency time trends, monthly changes in (e, f, s, n) represent the typical seasonal pattern in each outcome.

- For any variable x, we define the operators

$$\begin{aligned}
\Delta_g(x) &\equiv x_{♀} - x_{♂} && \text{(gender gap)} \\
\Delta_m(x) &\equiv x_m - x_{m-1} && \text{(month-to-month change)} \\
\mathbb{E}_g(x) &\equiv \frac{1}{2}(x_{♀} + x_{♂}) && \text{(cross-gender average)} \\
\mathbb{E}_y(x) &\equiv \frac{1}{12}\sum_{m=1}^{12} x_m && \text{(within-year average)}
\end{aligned}$$

These operators may be nested: for example, $\Delta_g(\Delta_m(x)) = (x_{♀m} - x_{♀,m-1}) - (x_{♂m} - x_{♂,m-1})$.

C.1 Stock-flow decomposition

We begin with the decomposition shown in Figure 3, which expresses changes in each group's EPOP between months $m-1$ and m as the sum of an inflow component and an outflow component.

Stock-flow identity. Since month-to-month changes in EPOP equal net inflows, we have the law of motion

$$e_{gm} = e_{g,m-1} + f_{gm} - s_{gm} \qquad \text{for } m > 0 \tag{7}$$

By recursive substitution, $e_{g12} = e_{g0} + \sum_{m=1}^{12}(f_{gm} - s_{gm})$. But since e_{gm} represents a seasonal cycle, we know that $e_{g0} = e_{g12}$: net of low-frequency trends and idiosyncratic shocks, EPOP evolves from May through April and then returns to its May level. It follows that

$$\sum_{m=1}^{12} f_{gm} = \sum_{m=1}^{12} s_{gm} \tag{8}$$

Intuitively, EPOP can remain stable over a 12-month cycle only if total inflows exactly counterbalance total outflows over that period.

Dividing Equation (8) by 12 yields $\mathbb{E}_y(f_g) = \mathbb{E}_y(s_g)$: average inflows equal average outflows over the seasonal cycle. Adding and subtracting these (equal) terms to Equation (7), we obtain

$$\Delta_m(e_{gm}) \equiv \underbrace{(f_{gm} - \mathbb{E}_y(f_g))}_{\text{excess inflows}} - \underbrace{(s_{gm} - \mathbb{E}_y(s_g))}_{\text{excess outflows}} \tag{9}$$

Intuitively, EPOP rises between two consecutive months to the extent that inflows exceed their average monthly rate and/or outflows fall short of their average monthly rate.

Estimation. Equation (9) is estimable. Let β^f_{gm} and β^s_{gm} denote the parameters of interest in our inflow and outflow specification, respectively. Start with inflows. Since these parameters represent differences in flows between month m and the base month, we have $f_{gm} = f_{g0} + \beta^f_{gm}$, so that

$$\mathbb{E}_y(f_g) = \frac{1}{12}\sum_{m=1}^{12} f_{gm} = \frac{1}{12}\sum_{m=1}^{12}(f_{g0} + \beta^f_{gm}) = f_{g0} + \mathbb{E}_y(\beta^f_g) \tag{10}$$

We can then rewrite excess inflows as

$$f_{gm} - \mathbb{E}_y(f_g) = (f_{g0} + \beta^f_{gm}) - (f_{g0} + \mathbb{E}_y(\beta^f_g)) = \beta^f_{gm} - \mathbb{E}_y(\beta^f_g) \tag{11}$$

Rewriting excess outflows in the same fashion, and replacing each parameter with its empirical estimate, we obtain our stock-flow decomposition:

$$\Delta_m(e_{gm}) \equiv \underbrace{(\widehat{\beta}^f_{gm} - \mathbb{E}_y(\widehat{\beta}^f_g))}_{\text{excess inflows}} - \underbrace{(\widehat{\beta}^s_{gm} - \mathbb{E}_y(\widehat{\beta}^s_g))}_{\text{excess outflows}} \tag{12}$$

Although Figure 3 is expressed in terms of one-month changes, one could cumulate these decomposition terms across months to estimate the contributions of inflows versus outflows to changes in EPOP between any pair of months m and m'. In addition, confidence intervals can be readily constructed via the delta method.

C.2 Job decomposition

We now turn to the decomposition shown in Figure 12. Our goal is to decompose $\Delta_g(\Delta_m(e_g))$, which represents gender differences in the evolution of EPOP between months $m-1$ and m, into a set of terms representing gender differences in sorting across job types and gender differences in seasonality conditional on job type. Having done so, we can then cumulate the decomposition terms across months to characterize gender differences over the full seasonal cycle.

Step 1: Partition employment into jobs and sectors. We partition employment into a finite set $\mathbb{J}$ of job types, indexed by j. These jobs are nested within S sectors, so that $\mathbb{J} = J_A \cup J_B \ldots \cup J_S$. In our empirical implementation, we label sector A as "educational services" and distinguish five job types within that sector, whereas we treat all other sectors as singletons, each comprised of a single undifferentiated job type. For the moment, however, we keep the notation general, allowing for the possibility that sectors $B, \ldots, S$ are each subdivided into multiple job types.

Step 2: Express seasonal changes in EPOP in shift-share form. To leverage standard decomposition techniques, we first write $\Delta_m(e_g)$ as a share-weighted average of job-level flow rates.

Let e_{gjm} denote the share of population g employed in job j in month m, so that $e_{gm} = \sum_{j\in\mathbb{J}} e_{gjm}$. Let f_{gjm} denote the share of population g moving from non-employment into job j, and let s_{gjm} denote the share moving from job j into non-employment. We define $n_{gjm} \equiv f_{gjm} - s_{gjm}$ as net inflows from non-employment into job j. Note that these flows exclude job-to-job transitions, which cancel out in the aggregate and hence leave no imprint on overall EPOP.

Next, we express seasonal changes in EPOP as the sum of net inflows across job types:

$$\Delta_m(e_g) = n_{gm} = \sum_{j\in\mathbb{J}} n_{gjm} \tag{13}$$

As in the aggregate case, these seasonal movements must cumulate to zero over a full 12-month cycle, so that $\mathbb{E}_y(f_{gj}) = \mathbb{E}_y(s_{gj})$ and hence $\mathbb{E}_y(n_{gj}) = 0$. Subtracting this expression, we obtain

$$\Delta_m(e_g) = \sum_{j\in\mathbb{J}} (n_{gjm} - \mathbb{E}_y(n_{gj})) \tag{14}$$

Now, multiply and divide the summand by e_{gj0}, the share of population g employed in job j in the base month:

$$\Delta_m(e_g) = \sum_{j\in\mathbb{J}} e_{gj0} \underbrace{\left(\frac{n_{gjm}}{e_{gj0}} - \frac{\mathbb{E}_y(n_{gj})}{e_{gj0}}\right)}_{\equiv \lambda_{gjm}} = \sum_{j\in\mathbb{J}} e_{gj0}\lambda_{gjm}, \tag{15}$$

where the newly defined term λ_{gjm} represents *group g's excess net flows from non-employment into job j in month m as a share of baseline employment.*

Step 3: Decompose the gender gap between and within job types. With the shift-share formulation in hand, we can express the gender gap in employment seasonality as

$$\Delta_g(\Delta_m(e_g)) = \Delta_g\left(\sum_{j\in\mathbb{J}} e_{gj0}\lambda_{gjm}\right) \tag{16}$$

We are now in the realm of familiar decomposition techniques. Using the standard trick of adding and subtracting cross-terms, we can decompose the righthand side as[4]

$$\Delta_g\left(\sum_{j\in\mathbb{J}} e_{gj0}\lambda_{gjm}\right) = \underbrace{\sum_{j\in\mathbb{J}} \Delta_g(e_{gj0})\mathbb{E}_g(\lambda_{gjm})}_{\text{between jobs}} + \underbrace{\sum_{j\in\mathbb{J}} \mathbb{E}_g(e_{gj0})\Delta_g(\lambda_{gjm})}_{\text{within jobs}} \tag{17}$$

Intuitively, the *between-job* component captures gender differences in seasonality arising from differences in the share of each group employed at various jobs that differ in their propensity to generate employment inflows/outflows throughout the year. The *within-job* component captures gender differences in employment flows conditional on a given allocation across job types.

Step 4: Separate the job-sorting and baseline EPOP effects. Whereas the within-job component in Equation (17) has a straightforward economic interpretation, the between-job component does not, as it confounds gender differences in sorting with gender differences in employment rates. With a little more algebra, however, we can separate these effects:

$$\sum_{j\in\mathbb{J}} \Delta_g(e_{gj0})\mathbb{E}_g(\lambda_{gjm}) = \underbrace{\sum_{j\in\mathbb{J}} \Delta_g\left(\frac{e_{gj0}}{e_{g0}}\right)\mathbb{E}_g(e_{g0})\mathbb{E}_g(\lambda_{gjm})}_{\text{job-sorting effect}} + \underbrace{\Delta_g(e_{g0})\sum_{j\in\mathbb{J}} \mathbb{E}_g\left(\frac{e_{gj0}}{e_{g0}}\right)\mathbb{E}_g(\lambda_{gjm})}_{\text{baseline EPOP effect}} \tag{18}$$

The *job-sorting effect* captures the extent to which—conditional on being employed—male and female workers differ in their propensity to work in jobs with different seasonal patterns. The *baseline EPOP effect* is a scaling term that accounts for gender differences in employment rates: because male EPOP exceeds female EPOP, a seasonal shift that has the same *proportional* impact on male and female employment rates will have a bigger *absolute* impact on men than on women. By splitting out the baseline EPOP effect (which we regard as a nuisance term), we can better assess how job sorting contributes to the gender gap in summer work interruptions.

Step 5: Distinguish sorting across sectors from sorting within sectors. We can further unpack the job-sorting effect to distinguish sectoral sorting from sorting across jobs within a given sector. To condense notation:

- Let $\phi_{gj} \equiv \frac{e_{gj0}}{e_{g0}}$ denote group g's employment in job j as a fraction of its total employment.

[4] As with any Oaxaca-Blinder-style decomposition, we face the question of which gender to use as the base group in each term. Equation (17) uses cross-gender averages in each term to avoid making an arbitrary choice.

- Let $\widetilde{\lambda}_{jm} \equiv \mathbb{E}_g(e_{g0})\mathbb{E}_g(\lambda_{gjm})$ denote excess net flows in job j, averaged across genders and then scaled by aggregate EPOP.

- Define analogous terms for sector k: $\Phi_{gk} \equiv \sum_{j \in J_k} \phi_{gj}$ and $\Lambda_{km} \equiv \sum_{j \in J_k} \mathbb{E}_g\left(\frac{\phi_{gj}}{\Phi_{gk}}\right)\widetilde{\lambda}_{jm}$.

- Let $\Phi_{g\not{A}} \equiv \sum_{k \neq A} \Phi_{gk}$ and $\Lambda_{\not{A}m} \equiv \sum_{k \neq A} \mathbb{E}_g\left({}^{\Phi_{gk}}/{}_{\Phi_{g\not{A}}}\right)\Lambda_{km}$ describe non-education as a whole.

The job-sorting effect then becomes simply $\sum_j \Delta_g(\phi_{gj})\widetilde{\lambda}_{jm}$, which we subdecompose as follows:

$$
\begin{aligned}
\sum_{j \in \mathbb{J}} \Delta_g(\phi_{gj})\widetilde{\lambda}_{jm} = {} & \underbrace{\Delta_g\left(\Phi_{gA}\Lambda_{Am} + \Phi_{g\not{A}}\Lambda_{\not{A}m}\right)}_{\text{sorting into the education sector}} + \underbrace{\Delta_g\left(\sum_{k \neq A} \frac{\Phi_{gk}}{\Phi_{g\not{A}}}(\Lambda_{km} - \Lambda_{\not{A}m})\right)}_{\text{sorting across non-education sectors}} \\
& + \underbrace{\sum_{j \in J_A} \Delta_g(\phi_{gj})(\widetilde{\lambda}_{jm} - \Lambda_{Am})}_{\text{sorting across ed.-sector jobs}} + \underbrace{\sum_{k \neq A}\sum_{j \in J_k} \Delta_g(\phi_{gj})(\widetilde{\lambda}_{jm} - \Lambda_{km})}_{\text{sorting across jobs within non-ed. sectors}}
\end{aligned}
\tag{19}
$$

Intuitively:

- The first term captures gender differences in sorting into educational services, "priced" using average seasonal patterns in the education sector versus non-education as a whole.

- The second term captures gender differences in sorting into sectors with different seasonal patterns (such as construction versus health care), conditional on sorting into non-education.

- The third term captures gender differences in sorting across jobs within educational services (such as primary school teaching versus secondary school teaching). The final term captures analogous sorting patterns across jobs within each non-education sector.

In our empirical implementation, we treat each non-education sector as consisting of a single undifferentiated job, so this final term vanishes.

Step 6: Isolate gender differences within jobs in each sector. In a similar (but simpler) fashion, we can also subdecompose the within-job component from Equation (17) into two terms:

$$
\sum_{j \in \mathbb{J}} \mathbb{E}_g(e_{gj0})\Delta_g(\lambda_{gjm}) = \underbrace{\sum_{j \in J_A} \mathbb{E}_g(e_{gj0})\Delta_g(\lambda_{gjm})}_{\text{within education-sector jobs}} + \underbrace{\sum_{k \neq A}\sum_{j \in J_k} \mathbb{E}_g(e_{gj0})\Delta_g(\lambda_{gjm})}_{\text{within jobs outside of the education sector}} \tag{20}
$$

Note that, in Equation (20), the second subcomponent can be viewed as a single entity representing the non-education sector as a whole, or one could separately examine within-job contributions from the construction sector, manufacturing sector, retail sector, and so on. By the same token, one could go further and examine the contribution made by gender differences within specific jobs, such as differences among primary school teachers or differences among lawyers.

Empirical implementation. Equations (17) and (18) give us a three-way decomposition of the gender gap in employment seasonality into within-job, job-sorting, and baseline EPOP components. Equations (18) to (20) unpack these further into as many as seven components. To implement these decompositions, we need (1) a partition $\mathbb{J}$ of industry-occupation pairings into job types, (2) estimates of the share of women/men employed in aggregate and in each job, and (3) estimates of the λ terms capturing net excess flows.

We start by distinguish 14 "one-digit" sectors, such as construction, manufacturing, and retail trade, on the basis of the variable `ind1990`. These follow headers used by IPUMS in its codebook entries for that variable, except that we split "professional and business services" into "educational services" and "other professional and business services". Next, we partition jobs in the educational services sector into five categories:

- Pre-K, kindergarten, and primary school teachers
- Secondary school teachers
- Postsecondary teachers
- Other staff in elementary and secondary schools
- Other staff in educational services

We code all other sectors as singletons: e.g., we recognize only a single manufacturing "job". Doing so eliminates the term representing sorting across jobs within non-education sectors, so we are left with the six-way decomposition presented in Figure 12.

We then compute baseline employment shares as simple average employment shares across all May observations in our analysis period. We estimate the λ terms by estimating our standard seasonal specification on grouped data, with one observation per sex $\times$ job type.

Confidence intervals. Each term in our decomposition combines parameter estimates from a subset of 2 $\times$ J notionally independent regressions. To construct confidence intervals for each decomposition term, we stack a copy of the group-level data for each constituent regression, then estimate a single stacked model in the manner of seemingly unrelated regression. We cluster errors at the year $\times$ month level, so that the error terms can be arbitrarily correlated across outcomes in each stack. Using the stacked covariance matrix, we can then construct confidence intervals via the delta method, using the Stata command `nlcom` as we do throughout the paper.[5]

D A Model of Sectoral Choice with Summer Childcare Constraints

We use a two-period model to illustrate how summer childcare constraints—which are likely to differentially affect women's labor supply—may contribute to gender differences in employment over both the seasonal cycle and the life cycle. In our model, period 1 represents a typical year during the early part of an agent's working life, before she (or he) has children. Period 2 represents years later in life when children are old enough to attend school but young enough to require supervision during the summer months, when schools are not in session. We abstract from other portions of the life cycle so as to focus attention on the most pertinent theoretical issues.

We proceed in four steps. In step I, we develop a static variant of the model that is isomorphic to period 2 in the full dynamic model. In step II, we determine which of the available strategies are "admissible" in the sense of being optimal for some possible parameter values. In step III, we perform comparative statics showing how optimal behavior responds to the increased implicit cost of childcare associated with working during the summer months. We interpret these comparative statics as a reduced-form representation of comparisons between agents who differ in parental status,

[5]As a check on the logic of this procedure, we compared the confidence intervals for the overall gender gap in EPOP obtained via this method with those obtained from a direct regression using aggregate flows. These match up to the slight numerical errors one would expect from repeated application of the delta method.

child age, the availability of spousal childcare, or access to market-provided childcare. In step IV, we extend the static model into a two-period dynamic model and derive additional implications about life-cycle career choices.

Step I: static setup

We consider a single agent deciding whether and in which sector to work at different points throughout the year. Here and throughout, the model is in partial equilibrium in the sense that we do not endogenize employment opportunities or wages. We also abstract from fertility decisions and take the presence or absence of children as exogenous.

Time periods. Each period, or "year", is divided into two subperiods, which we call "seasons" and index by $\tau \in \{A, B\}$. Season A, which we sometimes call "winter" for concreteness, represents the school year, whereas season B represents the summer.[6] Since our initial focus is on a single year, we omit year subscripts until step IV.

Work status. The agent chooses whether to supply one unit of labor and, if so, in which sector to work. In a given season, the agent's work status ("job") is $j \in \{E, N, O\}$, where:

- E represents being employed (and at work) in the *education* sector;
- N represents being employed (and at work) in the *non-education* sector; and
- O represents being non-employed (or employed but absent), which we call the *outside* sector.

We abstract from both job search and leave-taking. First, we assume that the agent can obtain a job in either sector at zero cost and at any time. As a result, there is no meaningful distinction between unemployment and non-participation in our model.[7] Second, as detailed below, there is also no meaningful distinction between unemployment and vacation. Under these assumptions, the single status O suffices to capture all forms of non-work during a given season.

Strategies. A *strategy*, denoted by s, is an ordered pair $(j_A,\ j_B)$ representing the agent's employment status during both winter and summer. Since each status can assume three different values, there are nine available strategies. For brevity, we often write $s = EE$ or $s = NO$ in place of $s = (E, E)$ or $s = (N, O)$. We write $s^*(\theta)$ for the optimal strategy under parameter vector θ, which we define explicitly below.

Utility. Utility $u(s|\theta)$ from choosing strategy s given parameters θ equals earnings net of distaste for labor and childcare costs, with each component summed across seasons.

Earnings. Let $w_{j\tau}$ denote base wages from working in job j during season τ. Let b_j be a bonus awarded for working year-round in job j. We make four assumptions about earnings in each sector:

[6] Although (for simplicity) we model the two seasons as being of equal length, it would be straightforward to modify the model to allow for "winter" to be three times as long as "summer".

[7] Although unemployment is a first-order consideration at high frequencies, it is of secondary importance relative to participation decisions over longer time horizons. We focus on a single year for purposes of exposition, but we interpret our model as capturing employment dynamics over longer periods each lasting for at least several years.

Assumption A1. $w_{OA} = w_{OB} = 0$.

Non-work yields zero earnings. Intuitively, this assumption abstracts from unemployment benefits, cash welfare, or any other forms of non-labor income.

Assumption A2. $w_{EA} > w_{EB} > 0$.

Education jobs pay less over the summer. This assumption captures the idea that the demand for education workers is greater during the school year but remains positive during the summer.[8]

Assumption A3. $w_{NA} = w_{NB} \equiv w_N > 0$.

Non-education jobs pay the same base wage in each season. This assumption abstracts from seasonal differences in labor demand in sectors like agriculture, construction, and retail.

Assumption A4. $b_N > 0$, $b_E = b_O = 0$.

Non-education jobs offer a continuity bonus, whereas education jobs do not. This assumption captures the idea that, in many industries, full-year employment offers premium earnings (or, equivalently, interrupted employment carries an earnings penalty).

These assumptions amount to a parsimonious way of modeling the key idea that education-sector jobs are more flexible than non-education jobs, especially as pertains to summer work. Although many of our theoretical results would obtain under weaker assumptions, the sharp parameter restrictions assumed above simplify the exposition and streamline the proofs.

Distaste for labor and childcare costs. Working in sector j in season τ entails flow disutility $\Phi_{j\tau}$. We decompose this disutility as $\Phi_{j\tau} \equiv \phi_{j\tau} + \Delta_{j\tau}$, where $\phi_{j\tau}$ is distaste for working in job j and $\Delta_{j\tau}$ represents utility costs associated with covering childcare. We assume as follows:

Assumption B1. $\phi_{OA} = \phi_{OB} = 0$.

We normalize the intrinsic distaste for nonemployment to zero in both seasons.

Assumption B2. $\phi_{EA} = \phi_{EB} \equiv \phi_E, \quad \phi_{NA} = \phi_{NB} \equiv \phi_N$.

In both sectors, the distaste for labor is the same across seasons. We thus abstract from the possibility that some jobs are more pleasant or unpleasant to perform at certain times of year. Together with our previous normalization, we also abstract from the possibility that agents may have different leisure preferences at different times of year for reasons other than childcare constraints—say, a taste for recreation opportunities available in warm weather.

Assumption B3. *(i)* $\Delta_{OA} = \Delta_{EA} = \Delta_{NA} = 0$, *(ii)* $\Delta_{OB} = 0$, *(iii)* $\Delta_{EB} = \Delta_{NB} \equiv \Delta > 0$.

Childcare costs are (i) zero during the school year, (ii) zero for the non-employed during the summer, and (iii) positive for the employed during the summer.[9] Intuitively, we assume that schools

[8] School-year employment and summer employment need not involve the same employer: for example, some education workers may work for different school districts in different seasons, or switch between K–12 and college.

[9] Although we intend our model to encompass both parents and childless individuals, it is convenient to assume that Δ is strictly positive for all agents, as doing so rules out certain uninteresting cases below and thereby simplifies the exposition. Our model is also isomorphic to one in which $\Delta = \Delta_C + \Delta_R$, where $\Delta_C \geq 0$ represents childcare costs and $\Delta_R > 0$ represents a taste for summer leisure. In this alternative formulation, our interest would lie in comparative statics as Δ_C increases.

provide implicit childcare when in session and that non-employed parents can provide childcare at zero cost during summer recess, whereas working parents must pay for market-provided childcare when schools are in recess.

As with our assumptions about the earnings process, our main results would continue to obtain under weaker assumptions about leisure preferences and childcare costs.

Parameters. Let $\theta \equiv (w_{EA}, w_{EB}, w_N, b_N, \phi_E, \phi_N, \Delta)$ be a vector of exogenous parameters. Apart from the restrictions made above, we have in mind that these parameters vary freely across agents with different productivities, comparative advantages, leisure preferences, and household structures. We will be chiefly interested in comparative statics with respect to Δ, which captures childcare costs associated with summer employment.

Step II: admissible strategies

We can write out the utility associated with each available strategy as follows:

Strategy	Earnings	Distaste	Childcare	Utility
OO	0	0	0	0
EO	w_{EA}	ϕ_E	0	$w_{EA} - \phi_E$
OE	w_{EB}	ϕ_E	Δ	$w_{EB} - \phi_E - \Delta$
EE	$w_{EA} + w_{EB}$	$2\phi_E$	Δ	$w_{EA} + w_{EB} - 2\phi_E - \Delta$
NO	w_N	ϕ_N	0	$w_N - \phi_N$
ON	w_N	ϕ_N	Δ	$w_N - \phi_N - \Delta$
NN	$2w_N + b_N$	$2\phi_N$	Δ	$2w_N + b_N - 2\phi_N - \Delta$
EN	$w_{EA} + w_N$	$\phi_E + \phi_N$	Δ	$w_{EA} + w_N - \phi_E - \phi_N - \Delta$
NE	$w_N + w_{EB}$	$\phi_N + \phi_E$	Δ	$w_{EB} + w_N - \phi_E - \phi_N - \Delta$

By inspection, three strategies can be immediately ruled out:

- *OE* is strictly dominated by *EO* since $w_{EA} > w_{EB}$ and $\Delta > 0$.
- *ON* is strictly dominated by *NO* since $\Delta > 0$.
- *NE* is strictly dominated by *EN* since $w_{EA} > w_{EB}$.

The dominated strategies represent work configurations that, though of course present to some extent in the real world, are of secondary importance for our analysis.[10] Each of the remaining six strategies is *admissible* in the sense of being the optimal strategy for some parameter vector θ.

Lemma 1. *For each strategy $s_k \in \{OO, EO, EE, NO, NN, EN\}$, there exists a parameter vector θ_k such that $s^*(\theta_k) = s_k$. Moreover, θ_k can be chosen such that s_k is the unique optimum.*

Proof: Fix an initial vector θ_0 satisfying the assumptions stated previously. By taking certain parameter values to the limit while keeping all other parameters fixed, we can make each of the six strategies uniquely optimal:

- *OO*: take $\phi_E \to \infty$ and $\phi_N \to \infty$.

[10] Pure summer employment (*OE* or *ON*) is common among young adults but less common among prime-age adults. Although employment rates among prime-age men are significantly higher in summer than in winter, seasonal patterns among men primarily track the timing of adverse winter weather rather than the timing of summer recess.

- *EO*: take $w_{EA} \to \infty$ and $\Delta \to \infty$.
- *EE*: take $w_{EA} \to \infty$ and $w_{EB} \to \infty$.
- *NO*: take $w_N \to \infty$ and $\Delta \to \infty$, with $\Delta > w_N + b_N - \phi_N$.
- *NN*: take $w_N \to \infty$.
- *EN*: take $w_{EA} \to \infty$ and $w_N \to \infty$, with $w_{EA} - \phi_E > w_N + b_N - \phi_N$.

The surviving strategies mirror employment patterns that commonly arise in the data.

Step III: comparative statics

We now consider comparative statics as Δ increases to $\Delta' = \Delta + \delta$, with all other parameters held fixed. To illustrate how summer childcare costs may shape employment decisions, we show how agents pursuing each admissible strategy under the original parameter vector θ reoptimize under the new vector θ'. Under each admissible strategy, utility changes as follows:

Strategy (s)	$u(s\|\theta)$	$u(s\|\theta') - u(s\|\theta)$
OO	0	0
EO	$w_{EA} - \phi_E$	0
EE	$w_{EA} + w_{EB} - 2\phi_E - \Delta$	$-\delta$
NO	$w_N - \phi_N$	0
NN	$2w_N + b_N - 2\phi_N - \Delta$	$-\delta$
EN	$w_{EA} + w_N - \phi_E - \phi_N - \Delta$	$-\delta$

To streamline the exposition, we ignore the edge cases where the agent is initially indifferent between two or more strategies.

Theorem 1. *Consider agents whose optimal strategy $s^*(\theta)$ is initially inframarginal, so that for $\delta \approx 0$ the new optimum $s^*(\theta')$ coincides with the original one. For sufficiently large values of δ, we observe the following changes in optimal behavior:*

(i) If $s^*(\theta) \in \{OO, EO, NO\}$, then $s^*(\theta') = s^*(\theta)$.

(ii) If $s^*(\theta) \in \{EE, EN\}$, then $s^*(\theta') = EO$.

(iii) If $s^*(\theta) = NN$, then each of $s^*(\theta') \in \{NO, EO, OO\}$ is potentially optimal.

Proof: Strategies *EE*, *NN*, and *EN* are clearly subobtimal when δ is large, so it suffices to consider whether *OO*, *EO*, or *NO* yields the most utility in each case.

(i) If $s^*(\theta) \in \{OO, EO, NO\}$, then $u(s^*(\theta)|\theta') = u(s^*(\theta)|\theta)$, whereas $u(s|\theta') \leq u(s|\theta)$ for all $s \neq s^*(\theta)$. It follows that $s^*(\theta)$ remains optimal under θ'.

(ii) By revealed preference, it must be that $w_{EA} - \phi_E > 0$, since otherwise the agent could have profitably deviated to strategy *OE* (in the case $s^*(\theta) = EE$) or *ON* (if $s^*(\theta) = EN$). Therefore $u(EO|\theta') > u(OO|\theta')$, so that *EO* is preferred to *OO*.

Likewise, it must be that $w_{EA} - \phi_E > w_N - \phi_N$ (in the case $s^*(\theta) = EE$) or $w_{EA} - \phi_E > w_N + b_N - \phi_N$ (in the case $s^* = EN$), since otherwise the agent could have profitably deviated to *NE* or *NN*, respectively. Thus *EO* is preferred to *NO*, as well.

(iii) Let θ_{-b} denote all parameters other than b. For any given choice of θ_{-b}, there exists some threshold b^* such that strategy NN is optimal for $b > b^*$. Fix such a value of b, then take $\delta \to \infty$, so that strategy NN is dominated and the new optimum is either NO, EO, or OO. Among these possibilities:
- NO dominates if $w_N - \phi_N > \max\{0, w_{EA} - \phi_E\}$.
- EO dominates if $w_{EA} - \phi_E > \max\{0, w_N - \phi_N\}$.
- OO dominates if $0 > \max\{w_{EA} - \phi_E, w_N - \phi_N\}$.

Intuitively, as summer childcare costs rise, (i) agents who counterfactually would have been non-employed over the summer are simply reinforced in their original decisions; (ii) agents whose primary job is in education choose to engage in home production over the summer; and (iii) agents who would otherwise have worked year-round outside of education either take the summer off, switch to education, or withdraw from employment altogether.

Step IV: two-period model

Now suppose the agent lives for two periods, indexed by $t \in \{1, 2\}$, each with seasons $\tau \in \{A, B\}$, and chooses a strategy s_t in each period to maximize lifetime utility.

Parameter vectors. Let $\theta \equiv (\theta_1, \theta_2, \beta)$, where θ_t is defined as in the static model and β is defined below. Let $\theta_{-\Delta,t}$ be a list of all period t parameters other than summer childcare costs Δ_t. We maintain assumptions *A1–A4* and *B1–B3* from the static model and additionally assume:

Assumption C1. $\theta_{-\Delta,1} \equiv \theta_{-\Delta,2}, \quad \Delta_1 = 0, \quad \Delta_2 > 0.$

The earnings and distaste parameters are identical across periods, so we omit t subscripts. Summer childcare costs arise only in period 2, when agents have school-age children.

Career premium. Potential earnings are linked across periods because of returns to career continuity. If the agent is employed in job $j \in \{E, N\}$ during the *winter* (season A) of period 1, we assume she receives supplemental income β_j in the event she remains employed in that same job during the winter of period 2. For simplicity, we assume that this supplemental income—which we call the *career premium*—is the same across sectors, though this assumption is inessential.

Assumption C2. $\beta_E = \beta_N \equiv \beta > 0, \quad \beta_O = 0.$

We regard β as a reduced-form representation of sector-specific human capital, seniority provisions, defined-benefit pensions, and other mechanisms that reward agents who remain in the same line of work throughout their careers. Because (in the real world) the school year lasts much longer than the summer, we assume that receipt or non-receipt of the career premium depends only on employment status in the winter season.

Utility. We assume that utility is additively separable across periods and can be written as

$$v(s_1, s_2|\theta) = u(s_1|\theta_1) + u(s_2|\theta_2) + \beta(s_1, s_2)$$

where $u(\cdot)$ is defined as in the static model. The function $\beta(s_1, s_2)$ equals β if the agent receives a career premium and zero otherwise. The bonus for year-round work, if received, is embedded in $u(\cdot)$. Since we consider only two periods, we ignore discounting to avoid cluttering the notation.

Strategies. The full strategy space consists of $9 \times 9 = 81$ ordered pairs $s \equiv (s_1, s_2)$ corresponding to actions taken in each of the two years, but—as in the static model—strategies OE, ON, and NE are dominated within each year, leaving $6 \times 6 = 36$ remaining possibilities.

Of these, only 11 strategies are admissible (potentially optimal) under our assumptions. Although a full characterization of the model solution would proceed by backward induction, we can establish the results of interest more directly by exploiting the fact that only the career premium links choices across years: decisions are otherwise separable between the two periods.

Lemma 2. *In the first period, each of the strategies $s_1^*(\theta) \in \{OO, EO, EE, NO, NN, EN\}$ is optimal for some set of parameter values. In the second period:*

(i) If $s_1^*(\theta) \in \{OO, EO, NO\}$, then $s_2^*(\theta) = s_1^*(\theta)$.

(ii) If $s_1^*(\theta) = EE$, then each of $s_2^*(\theta) \in \{EE, EO\}$ is potentially optimal.

(iii) If $s_1^*(\theta) = EN$, then each of $s_2^*(\theta) \in \{EN, EO\}$ is potentially optimal.

(iv) If $s_1^*(\theta) = NN$, then each of $s_2^*(\theta) \in \{NN, NO, EO, OO\}$ is potentially optimal.

Proof: To show that all six strategies may be optimal in the first period, it suffices to consider the case $\beta \approx 0$ and appeal to the corresponding arguments in the static setup of step II. Next:

(i) Suppose $s_1^*(\theta) \in \{OO, EO, NO\}$. Since the increment to utility from summer work is lower in period 2 than in period 1 (reflecting increased childcare costs), revealed preference ensures that the agent won't work in the summer of period 2, or equivalently $s_2^*(\theta) \in \{OO, EO, NO\}$. Furthermore, revealed preference—reinforced by the career premium, which discourages sectoral switching across years—ensures that these agents will make the *same* choice in both years. It follows that $s_2^*(\theta) = s_1^*(\theta)$.

(ii) If $s_1^*(\theta) = EE$, then revealed preference—again reinforced by the career premium—ensures that the agent will continue to work in the education sector in the winter of period 2. Revealed preference also ensures that, in the summer of period 2, the agent will either work in education (if Δ_2 is small) or refrain from working (if Δ_2 is large), so that $s_2^*(\theta) \in \{EE, EO\}$.

(iii) If $s_1^*(\theta) = EN$, the argument is analogous to that for $s_1^*(\theta) = EE$.

(iv) Suppose that $s_1^*(\theta) = NN$, and consider the limiting case $\beta \to 0$ so that the problem becomes separable across periods. Then, since the two periods are identical except that $\Delta_1 = 0$ and $\Delta_2 > 0$, the same arguments used in the proof of Theorem 1 establish the potential optimality of NN (if $\Delta_2 \approx 0$) and NO, EO, OO (if $\Delta_2 \gg 0$).

The lemma characterizes the distinct kinds of life-cycle career patterns that arise in our model. First, many agents make the same labor supply decisions throughout their working lives. Second, some agents work in the education sector, engage in summer work early in their careers, and then refrain from summer work once they have school-age children. Third, some agents work in the non-education sector early in their careers, then switch to the more flexible education sector once they face summer childcare costs. Finally, some agents work in non-education early in their careers, then withdraw from the labor force altogether when raising children.[11]

Our final result extends Theorem 1 from our static model to the two-period setting.

[11] In our model, decisions to quit the labor force are driven solely by *summer* childcare costs. Although we abstract from the costs of caring for pre-school-age children year-round and those of caring for young school-age children during the school year, accounting for these costs would provide additional incentives for agents to leave employment while raising children.

Theorem 2. *Consider comparative statics as Δ_2 increases to $\Delta_2' = \Delta_2 + \delta$, with all other parameters held fixed; let θ and θ' describe the original and perturbed parameter vectors. Conditional on choices made in period 1 (s_1^*), choices made in period 2 (s_2^*) respond as in the static model. Additionally, however, some agents for whom $s^*(\theta) = (NN, NN)$ will instead choose $s^*(\theta') = (EE, EO)$ or $s^*(\theta') = (EO, EO)$ when summer childcare costs rise. All other choices made in period 1 are unaffected by changes in Δ_2.*

Proof: By backward induction, period 2 in our two-period model is isomorphic to the single period considered in our static model, with potential earnings modified where appropriate for agents eligible for a career premium. As a result, all of our earlier comparative statics pass through unaltered in the second period of our dynamic setup.

We now show that some agents switch from NN to EE or EO in period 1 in response to *future* summer childcare costs. To see how this can arise, consider the special case in which $w_{EA} > 0$, $w_{EB} = 0$, $w_N = 0$, $\phi_E = \phi_N = 0$, $\Delta \approx 0$, and $b > w_{EA}$. In this special case, the agent initially chooses (NN, NN) under baseline parameters θ because earnings from doing so—which come exclusively in the form of the year-round continuity bonus b—exceed earnings available in the education sector.

Now increase summer childcare costs ($\delta \to \infty$) to the point that the agent no longer finds it optimal to work in the summer of period 2. Because the agent's earnings from non-education employment were predicated on year-round employment, the agent will switch from $s_2^*(\theta) = NN$ to $s_2^*(\theta') = EO$, thereby taking advantage of the more flexible earnings opportunities afforded by education employment. But if, in addition, $b < w_{EA} + \beta$, the agent will also switch from $s_1^*(\theta) = NN$ to $s_1^*(\theta') = EO$ because doing so secures receipt of the education sector's career premium in period 2.

If we modify this example so that $w_{EB} = \varepsilon > 0$, the agent will instead switch to (EE, EO).

With this last result, we can see that the model generates two kinds of sectoral sorting in response to summer childcare constraints. First, there is *contemporaneous sorting*: some agents switch from non-education into education upon experiencing summer childcare costs. Second, there is *anticipatory sorting*: some agents switch from non-education into education earlier in their careers. Intuitively, agents who know they will eventually want to make such a change may seek education employment from the beginning because of the returns to career continuity.

Our model also formalizes two distinct ways in which agents may be penalized for interrupted employment: agents who refrain from summer work miss out on the returns to continuous *year-round* employment, while those who switch sectors mid-career upon encountering summer childcare costs miss out on the returns to continuous *life-cycle* employment.

E Supplemental analyses

We close this appendix with two additional analyses: first, an examination of the tendency for a given individual to experience summer work interruptions in back-to-back years; second, a look at supplemental earnings among teachers during the summer versus the school year.

E.1 Recurrent summer work interruptions in consecutive years

Coglianese and Price (2020) introduce a method for identifying seasonal work interruptions at the individual level on the basis of patterns of recurrent transitions from employment into non-employment spaced exactly 12 months apart. Exploiting the limited longitudinal dimension of the

CPS, we apply that method to determine the extent to which individuals experiencing summer work interruptions tend to do so in back-to-back years.

Let y_{it} be an indicator variable equal to 1 if individual i was employed in period $t-1$ but not in period t. Using our sample of prime-age CPS respondents, we first identify all such work interruptions that occur during an individual's first four months in the sample, such that—barring attrition—we can observe that individual's employment status one year later. Letting t_0 denote the period in which the base separation occurred, we stack all available observations 10–14 months after baseline and estimate regressions of the form

$$y_{it} = \sum_{\tau=10}^{14} \rho_\tau \mathbb{1}\{t - t_0 = \tau\} + \beta weeks_t + \varepsilon_{it} \tag{21}$$

Thus $\rho_{10}, \ldots, \rho_{14}$ capture the relative probability of a recurrent work interruption occurring 10, 11, 12, 13, or 14 months after the initial one, adjusting for the fact that more separations tend to be observed when successive reference weeks are further apart. We cluster standard errors at the household level to allow for within-person serial correlation in the outcome variable as well as cross-sectional dependence among members of the same household.

Following Coglianese and Price (2020), we define the *excess recurrence* of work interruptions at annual intervals as $\rho_{12} - \frac{1}{2}(\rho_{11} + \rho_{13})$. Intuitively, excess recurrence tells us to what extent a given group of workers exhibit repeated exits from employment spaced exactly 12 months apart, net of the background rate of exit observed at similarly distant (but non-annual) horizons. Coglianese and Price demonstrate that excess recurrence aligns well with the demographic, sectoral, and temporal hallmarks of seasonal fluctuations in US employment.

Appendix Figure A.3 plots estimates of excess recurrence obtained by stratifying our CPS sample by sex and by the calendar month in which the base separation occurred. For women, work interruptions occurring between the May and June reference weeks are 4.9 percentage points more likely to be repeated 12 months later than 11 or 13 months later. Excess recurrence is also elevated in July—echoing the continued outflows of women from employment we see in that month (Figure 3)—as well as in January, when many businesses are trimming payrolls after the holiday shopping season. As a point of comparison, Coglianese and Price (2020) estimate an excess recurrence of 1.4 p.p. among all prime-age CPS respondents. By this measure, then, women show a pronounced tendency not only to exit employment at the start of summer, but to do so in (at least) two consecutive years.[12] The figure also provides another illustration of the very different seasonal work patterns we observe among men, for whom separations at the onset of winter—rather than the onset of summer—are most likely to recur 12 months later.

E.2 Schools and Staffing Survey (SASS)

The 1999–2000 Schools and Staffing Survey (SASS) from the National Center for Education Statistics provides a nationally representative snapshot of US public school teachers.

Variable definitions: The survey asks teachers about their supplemental earnings—i.e., earnings in addition to their base salary—during the summer months and, separately, during the regular school year. The survey additionally delineates between school-based and non-school-based supplemental work, where school-based work entails participation in extracurricular activities, coaching,

[12] Although annually recurrent separations are quite common among those employed in the education sector (consistent with patterns reported elsewhere in the paper), we observe qualitatively similar—albeit quantitatively muted—patterns if we exclude baseline departures from educational employment into non-employment.

and summer/evening teaching. From these earnings variables, we create indicator variables for supplemental work (school- or non-school-based) during the regular school year and summer months. We use these variables in our analysis of gender differences in the propensity to engage in supplemental work and earnings from supplemental work during the school year and summer months. SASS provides earnings categories for each type of supplemental work. To construct numeric earnings, we take the midpoint of each category. We multiply the top-coded earnings category by a constant factor of 1.5. We assign zero earnings when the individual does not engage in that type of supplemental work. We then deflate earnings to December 2019 dollars using the Personal Consumption Expenditures price index.

We also define the following regression controls:

- Teacher total experience: total years of teaching experience, in years
- Teacher age category: <30 years, 30–39 years, 40–49 years, 50+ years
- Teacher race/ethnicity: White non-Hispanic, Black non-Hispanic, Hispanic, other non-Hispanic
- Teacher educational attainment: indicator for whether the teacher has a master's degree
- School urban/rural status: large/mid-size city, urban fringe, small town/rural
- School region: Northeast, Midwest, South, West
- School level: elementary, secondary, combined
- Teacher field of assignment: pre-K, kindergarten, general elementary; math/science; English/language arts; social science; special education; foreign languages; bilingual/ESL; vocational/technical education; all others

Sample restrictions: We limit our sample to regular full-time teachers.

Regression-adjusted gender gaps: We regress the earnings from each type of supplemental work on a female indicator, age categories, teaching experience, race/ethnicity, master's degree, school type (primary, secondary), subject taught, urban status of school, and Census region. Each regression is weighted by the SASS sampling weights.

Figure 15 plots the regression-adjusted gender gaps in earnings from supplemental work among full-time public school teachers, throughout the summer months and the regular school year, controlling for demographic, job, and school characteristics.

We also explored gender differences in the propensity to engage in each type of supplemental work. Conditional on observables, female teachers are 18.8 percentage points less likely than male teachers to engage in any type of paid summer work. Furthermore, the gender gap in supplemental work is 3.8 percentage points larger during the summer months than during the regular school year, with the growth stemming from a differential uptick in men working outside of schools during summer. Overall, these results echo our above findings that, within granular educational occupations, women's work hours fall during the summer months, relative to men's and relative to the regular school year.