

The Lost Generation? Labor Market Outcomes for Post Great Recession Entrants

Jesse Rothstein¹

UC Berkeley, NBER, and IZA

May 2020

Abstract:

I study cohort patterns in the labor market outcomes of recent college graduates, examining changes surrounding the Great Recession. Recession entrants have lower wages and employment than those of earlier cohorts; more recent cohorts' wages have recovered but their employment rates continue to fall. I relate these changes to "scarring" effects of initial conditions. I demonstrate that adverse early conditions permanently reduce new entrants' employment probabilities. I also replicate earlier results of medium-term scarring effects on wages that fade out by the early 30s. But scarring cannot account for the recent employment collapse. There was a dramatic negative structural break in college graduates' employment rates, beginning around the 2005 entry cohort, that shows no sign of abating.

¹ rothstein@berkeley.edu. An earlier version of this paper was titled "The Lost Generation? Scarring After the Great Recession." I am grateful to Audrey Tiew, Rachel Young, Nathaniel Ruby, and Nicolas Ghio for excellent research assistance; to Avi Feller, Sara Hinkley, Rob Valletta, Till von Wachter, and Danny Yagan for helpful discussions and comments; and to conference and seminar participants at Berkeley, Georgetown, the University of Illinois, LSE, Michigan, NBER, UC Berkeley, and UC Merced.

I. Introduction

I study the early career economic outcomes of college graduates in the period between the Great Recession and the Covid-19 economic collapse. Anecdotes suggest that the cohorts that entered the labor market during the long period of economic weakness following the Great Recession had trouble finding footholds on job ladders and generally have had poor outcomes. My analysis largely confirms these anecdotes. I relate average outcomes for recent graduates to predictions from existing models of cohort outcomes, emphasizing “scarring” models that capture persistent effects of initial conditions (e.g., Kahn 2010; Oreopoulos, von Wachter, and Heisz 2012; Borgschulte and Martorell 2018; Schwandt and von Wachter 2019). I investigate whether models fit to pre-Great Recession entrants can account for recent cohorts’ experiences.

I document several facts regarding the economic outcomes of recent college graduates.

First, employment rates fell precipitously. The cohort that entered the labor market in 2010 has had an employment rate that, averaged over its experience to date, is two percentage points lower than what would have been expected based on prior cohorts’ age profiles and the state of the economy. The decline began with cohorts that started their careers before the Great Recession, and has only continued for more recent cohorts. Employment rates for 2015 entrants, who faced strong labor markets for the first four years of their careers, are three percentage points lower than projections based on pre-recession trends. Although more recent cohorts have only a few years of data, each is thus far posting lower employment rates than the one before.

Second, wages for those who are employed show a quite different pattern. While those who entered the labor market in 2009 have lower wages than those who entered earlier, by

about two percent, post-recession entrants have fully recovered. The most recent entrants are earning higher wages than did pre-recession cohorts at similar ages.

Third, both employment and wages are subject to scarring effects of early career conditions. The form of scarring is quite different for the two outcomes, however. I replicate past findings that those who face high unemployment rates at the start of their careers have lower wages through their first several years in the labor market, but that the scars gradually fade and by the tenth year the effects of initial conditions have faded. Scarring effects on employment, by contrast, are permanent: Cohorts that face high unemployment when they enter the labor market have statistically and substantively lower employment rates throughout their careers, not just at the beginning.

Fourth, these scarring effects, while quantitatively important, are not large enough to account for the employment outcomes of the Great Recession entrants. Projections based on earlier business cycles account for only about half of the observed downturn in employment rates for the 2009-2012 cohorts. Moreover, these cannot help explain poor outcomes for the most recent entrants, who encountered historically *low* unemployment rates at the outset of their careers (at least until the coronavirus pandemic in 2020, which is outside the scope of my data). Again, the evidence for wages is different: Early career scarring effects account for all of the decline in wages for Great Recession entrants, and subsequent cohorts have earned *higher* wages than past patterns would suggest.

Last, there was a dramatic structural break in the employment rate of young college graduates, beginning with the cohorts that entered the labor market around 2005, well before the Great Recession. Although employment rates rose dramatically following the Great Recession

for incumbent workers, new entrants did not share in this improvement. The new, dramatically negative trend shows no sign of letting up in the most recent data. The most recent cohorts have relative employment rates three to four percentage points lower than what one would have anticipated based on the pre-2005 trend.

These results provide reason for great concern about the future of the Great Recession and post-recession cohorts. Something seems to have changed in the labor market with notable negative effects on new generations' employment rates, and the long, steady recovery of the 2010s did not fix it. This has long-run implications for prosperity. The coronavirus crisis provides an additional cause for concern, with potentially enormous medium- and longer-term consequences for the class of 2020 and, if it lasts, for subsequent cohorts as well.

My analysis relates to several distinct literatures. First, several studies have found evidence for labor market changes during and after the Great Recession. Hershbein and Kahn (2018) find that employers raised education requirements when the labor market was weak, though Modestino et al (2016) find that this trend reversed during the subsequent recovery. Jaimovich and Siu (2020) find that demand for middle-skill, routine occupations fell sharply in the recession and did not recover thereafter. By focusing on college graduates – approximately the best-educated third of each cohort – I hope to avoid concerns about technological change and skill obsolescence. If new graduates are not hired, it is unlikely to be because they lack sufficient skills for the currently available jobs.

A second related literature explores so-called “jobless recoveries” following recent recessions (see, e.g., Daly, Hobijn, and Kwok 2009; Jaimovich and Siu 2020). Following the Great Recession in particular, the employment-population ratio recovered quite slowly, even as the

unemployment rate fell to historically low levels (Cunningham 2018). My analysis indicates that employment rates of younger workers have been particularly slow to recover. Again, technological change and skill obsolescence are unlikely to generate this pattern, which appears more consistent with reduced dynamism (Decker et al. 2014) and bigger insider-outsider distinctions in the post-recession labor market (Davis and Haltiwanger 2014).

Finally, my paper is most closely related to the aforementioned literature on labor market “scarring” (e.g., Kahn 2010; Oreopoulos et al. 2012; Borgschulte and Martorell 2018; Schwandt and von Wachter 2019). The scarring studies are a small component of an immense literature examining hysteresis more generally. Clark and Summers (1982) conclude participation is fairly persistent: “Workers drawn into the labour force by cyclical upturns tend to remain even after the boom has ended. The converse is true for shocks which reduce employment,” (p. 842). Yagan (2019) finds that workers in areas that experienced larger Great Recession shocks had worse outcomes 6-8 years later, not attenuated by mobility to healthier labor markets. My analysis complements Yagan’s by examining young workers – his estimates are based on workers already over 30 in 2007 – and relating outcomes following the Great Recession to patterns observed in past business cycles.

The paper proceeds as follows. Section 2 describes the Current Population Survey (CPS) data that I rely on for my analysis. Section 3 reviews the evolution of the labor market surrounding the Great Recession, with particular attention to the experience of young college graduates. Section 4 presents decompositions of employment and wages into age, time, and cohort components. I show that the cohort components turn down sharply for those graduating around the Great Recession, with the wage series but not the employment series recovering

quickly for subsequent cohorts. Section 5 augments the basic decompositions to allow for the type of lasting but impermanent scarring effects estimated in past studies. These scarring effects fully account for the downturn in wages around the recession, but do not explain the apparent break in employment outcomes. Section 6 examines permanent effects of initial conditions, by relating estimated cohort effects on employment and wages to the unemployment rate when the cohort entered the labor market. I show that cohort employment rates are systematically lower for those who enter in recessions, but that cohort wages are not cyclical. I also compare predictions based on past cyclical patterns to the observed experiences of the post-Great Recession entrants. This comparison reveals a clear structural break in employment before the Great Recession that is not accounted for by the usual cyclical sensitivity.

II. Data

I use repeated cross section data from the monthly Current Population Survey (CPS) to examine individuals born between 1948 and 1997 and observed at ages 22 to 40 between 1979 and 2019. My primary outcome is an indicator for weekly employment, measured in the monthly CPS. I also analyze log real hourly wages, observed for those who are employed in one-quarter of the CPS sample each month (the Outgoing Rotation Groups, or ORG). Data processing and definitions are discussed in the Appendix.

Following the scarring literature, I focus on those with bachelor's degrees or above. Recent graduates are unlikely to have obsolete skills, so a focus on this subgroup limits the potential scope for skill-biased technical change to produce negative labor demand shocks.

For each outcome and sample, I aggregate to the state-year-birth cohort cell, using CPS sampling weights. I merge onto each cell the contemporaneous state unemployment rate and the unemployment rate in the year that the cohort was 22.² I index cohorts by this year, so that c=2009 corresponds to the cohort that was born in 1987 and entered the labor market in 2009, at the nadir of the Great Recession.

Table 1 shows summary statistics for my main CPS sample and for the ORG subsample used for measuring wages.

There are three important limitations to the CPS data for my purposes. First, though the scarring literature emphasizes the importance of economic conditions when an individual enters the labor market, I do not observe the state where the individual lived at the start of his or her career. In my main specifications, I assume that all respondents entered the labor market in the state where they currently live.³

Second, the CPS data contain measures of current but not eventual educational attainment. This creates sample selection in a synthetic cohort analysis. The sample of college graduates from a particular birth cohort observed at age 22 represents a subset of the population represented by a sample of graduates from the same birth cohort taken later: Those who earn degrees after 22 will be represented in the later sample but not in the earlier sample. Insofar as those who take longer to finish college have different outcomes than those from the same birth

² State unemployment rates are available only from 1976; for cohorts that were 22 before this date, I use the national rate.

³ Schwandt and von Wachter (2019) explore the bias in similar specifications from failing to measure initial location, and conclude that it is “unlikely to be very large.” In the Appendix, I present results from a strategy that they develop that uses long-run inter-state mobility patterns to construct an instrument that is purged of endogenous post-college mobility.

cohort who finished earlier, this may generate spurious changes in cohort employment rates with age. My specifications all include age controls, but changes in the age pattern over time could create bias. A particular concern for my analysis is that the Great Recession may have induced many people to finish college who would not have had conditions been better. Consistent with this, college enrollments surged during the downturn (Barr and Turner 2015; Snyder, De Brey, and Dillow 2019, Table 302.10).

Fortunately for my purposes, the recession surge in enrollments does not seem to have led to a comparable surge in bachelor's degree awards, perhaps because much of the new enrollment was at community colleges (Digest of Education Statistics 2019, Table 302.10). Thus, the age pattern in BA attainment does not change dramatically for the relevant cohorts. Figure 1 shows the graduate share by year and age, with dashed lines connecting successive observations on birth cohorts as they age. The degree attainment rate for 22-year-olds is consistently less than half of what will be seen for the same cohort at age 30. For cohorts born in the 1960s, the great majority of degrees were earned by age 24. There was a substantial increase in post-24 degree attainment around the 1972 birth cohort (turning 24 in 1996; see Bound, Lovenheim, and Turner 2012), but still the great majority of degrees are awarded by 26. Since the mid-1990s, final attainment has grown substantially, with most of this growth in degrees awarded by age 24.⁴ There is no indication either of a surge in BA attainment among Great Recession cohorts or of changes in the share of degrees earned at older ages.⁵ Together, these facts make it plausible

⁴ Post-24 degrees may be growing for the post-1988 birth cohorts (aged 24 after 2012). It will be several years until we can fully measure this, and in the meantime it has little impact on my CPS samples.

⁵ Appendix Figure 1 shows a similar figure for the share of the population with some college or more. Here, we do see important increases in later attainment that are plausibly attributable to

that age effects will absorb most of the bias in cohort mean outcomes averaged across various ages.

I take several approaches to minimizing remaining composition bias. First, I include in all specifications a selection term computed as the inverse Mills ratio of cohort-by-age attainment. Because my specifications include both cohort and age effects, the coefficient on this control is identified from changes in the attainment rate over age that differ across cohorts. In a simple bivariate normal selection model, this will absorb any bias from changing selection (Gronau 1974, Card and Rothstein 2007). In practice, results are little affected by the inclusion of this control. Results are also insensitive to alternative, less parametric control functions (e.g., a quadratic in the cohort-by-age attainment rate). Finally, I have reestimated my primary specifications excluding observations from ages under 25. Results are again largely unchanged.

A last limitation, related to but distinct from the sample selection issue, is that when I observe someone who has a college degree, I do not know when that degree was obtained, so cannot measure labor market conditions that prevailed at that time. Following the scarring literature, I use conditions when the individual was age 22 as my measure of initial conditions. This can be seen as a reduced-form specification that uses age-22 conditions as an instrument for conditions on the (potentially endogenous) date when an individual actually graduated.

the Great Recession. This could create more serious sample selection biases in analyses of the population with any college.

III. The Great Recession in the labor market

In this section, I describe some stylized facts about the Great Recession labor market, to help set expectations for the cohort analyses below. The headline unemployment rate rose by 5.6 percentage points between mid-2007 and late 2009, while the prime-age (25-54) non-employment rate rose by over 5 percentage points (Figure 2). Young people fared particularly poorly. Over the same period, the employment rate among 26-30-year-olds fell by more than 7 percentage points. This is consistent with past experience: Young workers have always been more cyclically sensitive.

Unemployment began to recover in mid 2010 and declined roughly linearly, at a rate of about 0.9 percentage points per year, thereafter. The unemployment rate was below 6% from the third quarter of 2014 and below its pre-recession level from late 2017. Employment, however, was much slower to recover. Only half the decline in prime-age employment had been erased by the end of 2015; the employment rate did not recover its level prior to the recession until late 2019. Employment of young workers was particularly slow to recover: by mid-2014, the 26-30-year-old employment rate remained 5.5 percentage points below its pre-recession peak (as compared with 4.2 percentage points for all prime-age workers).

Average real hourly wages did not fall during the recession, due to changes in the composition of workers (Daly, Hobijn, and Wiles 2012). But they began to fall after the recession ended, with a larger decline for younger workers, then recovered in the later part of the recovery.

Following the scarring literature, I focus particularly on college graduates in their early- and middle careers. College graduates typically have much higher employment rates than non-graduates, but younger people are less likely to be employed than older. Figure 3 shows non-

employment and unemployment for college graduates aged 22-40. The unemployment rate rose by 150% between early 2007 and late 2009, while the non-employment rate rose from 14% in January 2007 to 18% in December 2012.⁶ Employment recovered somewhat more quickly for young graduates than for the prime-age labor force as a whole, but nevertheless did not achieve its level on the eve of the recession until mid 2018. Figure 4 shows the graduate employment series separately for two age groups. The decline in employment was about twice as large for the younger group, and was much more persistent. On the eve of the recession young graduates had similar employment rates to older graduates, as they did at the previous business cycle peak, but have been persistently lower since the recession's onset. Even in the most recent data, the younger group's employment rate is about three percentage points lower than that of the older group. This is a manifestation of the decline in cohort effects that I document below.

IV. Cohort effects on employment and wages

The time series evidence suggests that the experience of the recession and its aftermath may have been different for younger workers than for their older peers. Here, I document this more carefully, focusing on cohorts as the relevant unit of analysis. A complication is that more recent entrants are observed only at young ages, when employment rates tend to be lower. I use a formal decomposition to distinguish cohort differences from age and time effects.

⁶ These calculations are based on 7-month moving averages of seasonally adjusted rates. Non-employment is measured as a share of the population, while unemployment is measured as a share of the labor force.

Initial results are shown in Figure 5. The first series, with short dashes, shows employment rates by cohort, averaged over all years in which each cohort is observed in my sample.⁷ This plummets for the most recent cohorts, largely because recent cohorts are observed only at young ages.

The second series in Figure 5 implements a simple age adjustment. Specifically, let Y_{satc} represent the employment rate of college graduates in state s at age a in time t from birth cohort c (indexed by the year that the group entered the labor market, so $c=t-a+22$). I estimate a regression of the form:

$$Y_{satc} = \alpha + \gamma_a + \delta_c + \zeta_s + \rho\lambda(p_{satc}) + \epsilon_{satc}. \quad (1)$$

Here, γ_a represent a set of fixed effects for age and δ_c are effects for birth cohorts, each in single years. ζ_s represents state effects. As noted earlier, I include a selection correction to capture spurious age patterns deriving from changes in the composition of college graduates within cohorts over ages. This is the $\lambda(\cdot)$ term, the inverse Mills function applied to the state-cohort-age attainment rate, p_{satc} .⁸

Of interest are the cohort effects, δ_c . They represent age-adjusted employment rates (normalized relative to the excluded cohort, those who turned 22 in 1984), and are graphed with long dashes in Figure 5. Age-adjusted employment rates fell gradually across cohorts entering from 1975 through around 2004, with the total decline amounting to around 2.5 percentage

⁷ In this plot and in those that follow, I exclude the cohort that was 22 in 2019, on which I have only a single year of data.

⁸ The inverse Mills ratio is $\lambda(p) = \phi(\Phi^{-1}(p))/p$. In my samples, the 10th percentile of p_{satc} is 0.20 and the 90th percentile is 0.43. $\lambda(p)$ is approximately linear over this range. As noted earlier, all results are robust to replacing $\lambda(p)$ with a quadratic in p .

points. There was then an additional 1.8 percentage point decline between the 2004 and 2010 entrants, with stability thereafter.

These estimates address the fact that recent entrants are seen only at young ages, but not the specific economic conditions in the years in which they were observed – very bad after the Great Recession, then unusually good in the more recent data. Intuitively, the stable age-adjusted employment rates from 2010 onward are disappointing in the context of a steadily strengthening economy and labor market.

To capture the role of time-varying business conditions, I add to the specification a full set of calendar year controls. The new specification is:

$$Y_{satc} = \alpha + \beta_t + \gamma_a + \delta_c + \zeta_s + \rho\lambda(p_{satc}) + \epsilon_{satc}. \quad (2)$$

New here are the β_t coefficients, fixed effects for calendar years. These capture both aggregate demand and any supply-side factors that are common across age groups, potentially including the hysteresis effects documented by Yagan (2019). The cohort effects from (2) are net of the effects of overall economic conditions, so long as they have common effects across ages. They represent permanent differences between one birth cohort and another in the same labor market, beyond those reflected in the age profile γ_a .

The full set of age, time, and cohort effects in (2) is not identified, even after omitting one of each, due to the linear dependency among age, time, and cohort (see, e.g., Schulhofer-Wohl, 2018). A single additional normalization, effectively pinning down the linear trend in one of the three series, is required. I normalize the cohort effects for the 1984 and 2000 entry cohorts (born in 1962 and 1978, respectively) to be equal. This forces the estimated trend to be zero across these cohorts; all estimates can be seen as relative to the true trend over this period. I focus on

changes in trends in the estimated cohort effects, which are not affected by the normalization choice. Below, I fit simple regression models explaining the estimated cohort effects; these always include linear trends to absorb the normalization.

Cohort effects from (2) are plotted as the solid line in Figure 5. This series shows a stable trend from 1970 through 2000, which appears flat due to my normalization. It also shows a clear trend break in the 2000-2004 period.⁹ Across the 18 cohorts since the 2000 entrants, cohort effects have fallen nearly five percentage points (relative to the 1984-2000 trend), with no sign that the decline stabilized after the Great Recession. If the additive decomposition in (2) is correct, the most recent entrants will have employment rates that are five percentage points lower over the course of their careers (through age 40) than the 2000 entrants did, even holding labor market conditions constant.

Age and time effects (γ_a and β_t) are plotted in Appendix Figures 2 and 3, respectively. The time effects in Appendix Figure 3 show a strong upward trajectory from 2011 through 2019, of roughly the same magnitude as the downward trajectory in the cohort effects in Figure 5: Conditions were improving rapidly during this period for college graduates who had already entered the labor market, but new entrants were excluded from this trend, producing the flat overall age-adjusted profile in Figure 5.

Figure 6 presents additional estimates that address two potential concerns with the simple age-time-cohort decomposition. Each is motivated by the idea that the additive separability of age, time, and cohort effects in (2) may be too strong, particularly when used to

⁹ A model with a linear trend and single trend break fits best when the break is in 2003, but breaks placed in any year from 2000 to 2006 all fit nearly as well.

characterize the large changes that occurred during the Great Recession. Because post-recession cohorts are observed only when young, their cohort indicators are strongly collinear with both indicators for young ages and indicators for recent years, making it difficult to disentangle these three. To assess sensitivity to this, I re-estimated the model using only the cohorts that entered the labor market by 2000. I then used the estimated state, year, and age effects and sample selection coefficient ρ from this model to forecast outcomes for subsequent cohorts, and measured cohort effects as average deviations from this forecast. This is shown as the short dashed line in Figure 6. It slightly reduces the magnitude of the post-2004 downturn, but overall makes little difference – age and year effects estimated from the pre-2000 entrants are very similar to those estimated from the full sample, giving confidence in their stability and in the decomposition.

Second, as noted earlier, there may be selection in samples of young college graduates, arising from the fact that not all who will graduate college have done so by age 22. The sample selection term in (1) and (2) is intended to absorb this but may not be successful. I re-estimate model (2), including all cohorts but dropping observations before age 25, when educational attainment is rising most rapidly. This again has no meaningful effect on the employment trends.

Figure 7 repeats the cohort effects analyses (1) and (2), this time for hourly wages among those who are employed. (I do not show a series without age adjustments; because wages rise rapidly with age in the early career.) In the age-adjusted series (1) we see increases in wages across cohorts, interrupted by a period of erosion from 1997-2009. When I add year controls and normalize the 1984-2000 trend to zero, the stagnant period lengthens: The trend is stable between 1979 and 2009, with growth relative to that trend before and afterward. Wages did

decline sharply for the 2009 cohort, and to a lesser extent in 2007 and 2008, but recovered for the 2010 and subsequent cohorts. Moreover, there seems to be an upward trend break starting in 2009; for entrants in 2011 and thereafter, wages are comparable to or higher than the pre-recession trend. This recovery is only strengthened when I estimate age and year effects only from pre-recession cohorts or when I limit the sample to those ages 25 and older (Appendix Figure 4) – these specifications indicate that the most recent cohorts are earning about three percent more than would be expected based on the earlier trend.

Comparing Figures 6 and 7, my analysis of cohort patterns indicates strikingly different patterns for employment and wages. Employment rates for recent cohorts are notably lower than those for earlier entrants, with an apparent trend break around the 2004 cohort. The most recent cohorts have estimated employment rates three to four percentage points lower than would be expected based on the experience of pre-2000 entrants. By contrast, wages for those who are employed dipped for the 2009 entrants, and to a lesser extent for the two preceding cohorts, but have more than recovered for more recent entrants.

V. Scarring and other age-specific effects

In this section and the one that follows, I explore the potential role of the Great Recession in accounting for the negative employment trends in Figure 6. I augment my simple decomposition to include the scarring effects of initial conditions that have been documented in past research, and assess whether estimates of cohort effects that are adjusted for this show less evidence of a downturn following the recession. In Section VI, I further explore whether cohort effects, with or without adjustment for age-varying scarring effects, are systematically related to

entry conditions, and whether this relationship can account for any further component of the pattern.

A. Specifications

The decomposition (2) assumes that age, time, and cohort effects are additively separable. There are two sources of violations of this assumption, beyond those explored earlier, that could lead estimates from (2) to understate the performance of recent cohorts. First, scarring effects – an effect of the entry unemployment rate on later outcomes – would depress employment of the Great Recession cohorts. Since I observe those cohorts primarily in their early career years, any scarring effect that is concentrated in the early career would lead me to understate the permanent component of their outcomes. Second, younger workers may simply be more sensitive to current economic conditions. This too would depress outcomes through the recession and post-recession years for the recent entrants, pulling down my estimated cohort effects.

I begin with scarring. Let UR_{st} represent the unemployment rate in state s in year t , and recall that I index cohorts by the year that they entered the market. The entry unemployment rate for cohort c is thus UR_{sc} . To incorporate early-career scarring effects, I create a sequence of age indicators, D_a^j , for $j=1, \dots, 5$. D_a^1 is an indicator for ages 22-23 (i.e., for observations with 0 or 1 year of potential experience following presumed college graduation at 22), D_a^2 is an indicator for ages 24-25, D_a^3 for ages 26-27, D_a^4 for ages 28-29, and D_a^5 for ages 30-31. I include in my decomposition interactions between these D_a^j indicators and the state unemployment rate in the year that the cohort entered, UR_{sc} .

The augmented specification is:

$$Y_{satc} = \alpha + \beta_t + \gamma_a + \delta_c + \zeta_s + \rho\lambda(p_{satc}) + \sum_{j=1}^5 D_a^j * UR_{sc} * \phi_j + (UR_{st} - UR_t)\kappa + \epsilon_{satc}. \quad (3)$$

The first six terms are as in (2). As before, the cohort effects δ_c capture any permanent differences across cohorts, including any permanent effects of entry conditions. The summation on the second line captures additional effects of entry conditions on early career outcomes, over and above any effects that persist throughout the career.

The term following the summation in (3) requires additional discussion. The unemployment rate is measured at the state level, so scarring effects could in principle be identified solely from across-state variation within cohorts and times. However, most of the variation in unemployment is in the time series rather than the cross section. My enriched model that includes scarring effects thus uses both time series variation and variation across states in the amplitude of shocks to identify the importance of these effects. However, to obtain adequate precision for the cohort effects, I pool members of the cohort across states. The π coefficient is included to bridge between the state-level scarring effects and the national-level cohort effects. If cohort effects vary across states, π will absorb any component predicted by the state unemployment rate, without affecting the estimate of the national cohort effect. That is, the implied intercept for cohort c in state s is $\delta_c + (UR_{sc} - UR_c)\pi$, where UR_c represents the average of UR_{sc} across states, and the average of this across states is simply δ_c . Moreover, the inclusion of the π term ensures that scarring coefficients ϕ_j are identified solely from differential cohort outcomes early in the career for cohorts affected by varying entry unemployment rates, not from other unmodeled state-by-cohort variation.

I also examine a closely related phenomenon using a similar specification: It may be that younger workers are simply more sensitive than are older workers to contemporaneous economic conditions. That is, the impact of a downturn may not be constant across age, but larger for younger workers. I model this as an interaction between the contemporaneous unemployment rate (a calendar time effect) and age. This leads to a specification similar to (3), except that here the cohort's entry unemployment rate, UR_{sc} , is replaced by the contemporaneous rate, UR_{st} :

$$Y_{satc} = \alpha + \beta_t + \gamma_a + \delta_c + \zeta_s + \rho\lambda(p_{satc}) + \sum_{j=1}^5 D_a^j * UR_{st} * \theta_j + (UR_{st} - UR_t)\kappa + \epsilon_{satc}. \quad (4)$$

As before, the κ term bridges the state-level variation in the unemployment rate with the calendar time effects, which are modeled flexibly at the national level via β_t but constrained to vary in proportion to UR_{st} across states.

Past work has generally not distinguished scarring from excess sensitivity effects – they are quite difficult to disentangle empirically, as the unemployment rate experienced by a cohort early in its career is strongly correlated with the rate when the cohort entered – and has often treated the scarring coefficients ϕ_j as the reduced form for the combination of the two (see discussions in Oreopoulos et al. 2012 and Schwandt and von Wachter 2019). The implications are different, however, for assessing whether recessions have effects that persist beyond their ends, and thus potentially for the estimated cohort effects δ_c .

I also estimate a third specification that includes both the scarring and excess sensitivity controls. Across specifications, my interest is the extent to which the inclusion of age-time and/or

age-cohort interactions alters the pattern of estimated cohort effects during the period surrounding the Great Recession.

B. Estimates of scarring and excess sensitivity

Figure 8, panel A shows the estimated scarring coefficients from (3), ϕ_j , in green, along with 95% confidence intervals (allowing for clustering at the state level). The estimated scarring effects on employment are substantial: Cohorts that enter the labor market when the state's unemployment rate is elevated by 1% have employment probabilities that are reduced by 0.7 percentage points at ages 22 and 23, 0.5 percentage points at 24 and 25, and about 0.2 percentage points at 26 and 27, after which the effect fades away. These are consistent with the result from the scarring literature that early conditions have persistent effects, and are quantitatively similar to estimates of scarring effects on employment reported by Schwandt and von Wachter (2019). Importantly, however, these are not the only way that initial conditions can affect outcomes in (3). There are also permanent effects operating through the cohort coefficients δ_c , which I explore below.¹⁰

Panel B shows excess sensitivity coefficients from (4). These are very similar in sign and magnitude to the scarring coefficients, and indicate that early-career workers have lower employment when their state has a higher unemployment rate in that year, relative to older workers.

In both panels, I also show (in orange) results from a combined specification that includes both the scarring terms from (3) and the excess sensitivity effects from (4). These are often hard

¹⁰ There is also a permanent effect operating through π , though I estimate this at -0.03 (SE 0.07).

to distinguish, though the sharp changes in the Great Recession make it easier – excess sensitivity effects would imply declines in employment rates of 2005 entrants in 2008, relative to those of older workers, while scarring effects would not. Perhaps for this reason, I obtain similar precision in the combined model as when I include just one set at a time. In the combined model, the ϕ_j coefficients are much closer to zero than when they were included alone, while the θ_j coefficients are basically unchanged. In other words, the variation loads more strongly onto the interactions of age with the current unemployment rate (as in equation (4)) than onto the interactions of age with the entry unemployment rate (as in (3)) – excess sensitivity explains the data better than does scarring.

Panels C and D of Figure 8 show the same sets of coefficients in models for log wages. In Panel C, we again see evidence of substantial scarring effects of initial unemployment on subsequent wages, with a similar pattern: A 1 percentage point higher unemployment rate in the year of entry reduces wages by about 1.1% at age 22-23, 1% at 24-25, 0.4% at 26-29, and 0.1% (not significant) at 30-31. Panel D shows effects of contemporaneous unemployment that show a similar pattern, though smaller magnitudes. In contrast to the employment results in Panels A and B, here we see that it is the scarring coefficients that are robust to including both sets of interactions, while the excess sensitivity effects disappear when both are included.

A natural concern is that the Great Recession may be playing an outside role in driving these estimates. In Appendix Figure 5, I reestimate the models using only cohorts who turned 22 in 2000 and earlier. With this restriction, only cycles before 2000 contribute to the scarring and excess sensitivity coefficient estimates. Across all specifications, results are similar, though a bit less precise, to those in Figure 8.

C. Adjusted cohort effects

Figure 9 shows estimates of cohort effects δ_c for employment in specifications that allow for scarring and excess sensitivity interactions. The base estimates from specification (2) are shown as a solid line. The long-dashed and short-dashed lines show estimates from the scarring and excess sensitivity specifications (3) and (4), respectively, while the dashed red line shows estimates from the specification that includes both scarring and excess sensitivity controls. In these estimates, the cohort effects can be interpreted as differences among cohorts, beyond those captured by the early career age-unemployment interactions, that persist throughout the observed portion of the career. Early career and permanent cohort differences could not be distinguished for the most recent cohorts alone; in these specifications the scarring and excess sensitivity coefficients in Figures 7 and 8 are identified primarily from earlier cohorts (as in Appendix Figure 5), and the recent cohort effects reflect differences in the early-career experiences of these cohorts from what would have been projected based on past business cycles.

The scarring specification shows a notably smaller decline in cohort employment rates for the 2009-2013 cohorts than in the baseline decomposition. Evidently, this model attributes much of the low employment in the early years of these cohorts' careers to temporary scarring effects rather than permanent differences. However, we still see a trend break in 2004, and that is sufficient to generate a substantial decline in employment rates by the 2010 cohort. Moreover, there have been quite rapid declines for the most recent entrants.

Recall that the scarring specification is not robust. Specifications that allow for excess sensitivity of younger workers, in short dashes without scarring controls or mid-length with them, are much less optimistic about the Great Recession cohorts. These estimates diverge a bit from the unadjusted baseline estimates in the early 2000s, but generally move in parallel thereafter. As a result, the downward trend during the Great Recession is quite steep.

Perhaps the most notable aspect of Figure 9 is that the sharp decline in cohort employment rates for the most recent cohorts is largely robust to the choice of controls. These cohorts entered the labor market when the unemployment rate was relatively low, so should not be scarred by entry conditions, but are nevertheless doing quite poorly. Across all four specifications, employment rates for the 2017 entrants are between 4.4 and 5.0 percentage points below the 1990s trend.

Figure 10 shows the same cohort effect estimates, this time for log wages. Results are quite different here. In the baseline decomposition we see a sharp drop in wages, about 2 percent, for the 2009 entrants, with smaller reductions in 2007, 2008, and 2010. This persists in the model that allows for excess sensitivity, but completely disappears in the models that allow for scarring effects. Evidently, this dip was entirely consistent with historical scarring patterns that, in past cycles, have faded away over the first decade of workers' careers. Moreover, we see sharp increases in cohort intercepts following the recession. The 2011 and subsequent cohorts have wages about 2% higher, on average, than earlier cohorts, after adjusting for normal early career scarring effects.

Figure 11 shows estimated cohort effects from specifications run separately for men and women. The broad patterns are similar, though the decline in employment was steeper, earlier, and larger for women than for men, while post-2010 wage increases are seen only among men.

VI. Cyclical and cohort effects

The coefficients ϕ_j in specification (3) capture scarring effects of initial conditions on graduates' outcomes in the first ten years of their potential careers. The existing literature on scarring effects of initial conditions has focused on these early career effects (e.g., Borgschulte and Martorell 2018; Kahn 2010; Oreopoulos et al. 2012). But initial conditions may have effects that persist beyond that period. These would be captured by cohort effects, included in the specifications in each of the papers above but generally treated as nuisance parameters (e.g., Oreopoulos et al., 2012). Any permanent effects of the entry unemployment rate would manifest as cyclical variation of these cohort effects, with reduced values of δ_c for cohorts who were 22 at times of elevated unemployment. Because these effects would last throughout a worker's career, even a small amount of cyclical variation would be quantitatively important relative to the short-term effects that have been the subject of most previous study.

Figure 9 offers visual evidence for the potential importance of permanent scarring, as cohort effects seem to dip in recessions. Interestingly, the amplitude of the cyclical variation is somewhat lower for the estimates that adjust for early-career scarring than for others, suggesting that failure to account for this may lead to overestimation of the permanent effect of initial conditions.

Table 2 investigates this more carefully. I regress the estimated cohort effects from my various specifications on the entry unemployment rate UR_c . Each column draws the cohort effects from a different specification: Column 1 uses cohort effects from the simple age-time-cohort decomposition (2); column 2 uses effects adjusted for early-career scarring, as in (3); column 3 uses effects adjusted for excess sensitivity of younger workers, as in (4); and column 4 uses effects from a specification that includes both scarring and excess sensitivity controls. In each column, I use only the δ_c s for cohorts entering the labor market before 2005, to capture pre-Great Recession dynamics. Each regression controls for a linear time trend (which absorbs the normalization of the cohort effects), and I report Newey-West standard errors that allow for heteroskedasticity and for autocorrelation of the error terms at up to two lags.

The unadjusted cohort effects from decomposition (2) are very strongly related to the entry unemployment rate: Each 1 percentage point increase in the national unemployment rate is associated with a 0.16 percentage point reduction in the δ_c associated with the cohort turning 22 in that year. Because each cohort is included in my sample for 19 years, until age 40, and is in the labor market for at least 20 years beyond that time, this cumulates to a loss of 0.03 employment-years for each member of the cohort of new entrants, over the course of their careers (through age 40).

Columns 2-4 use cohort effects from decompositions that are modified to allow for age interactions with the entry unemployment rate (scarring), the contemporaneous unemployment rate (excess sensitivity), or both. Any of these modifications reduces the cyclicity of the estimated cohort effects by about half, with entry unemployment rate coefficients between -0.06 and -0.08. That is, each percentage point increase in the national unemployment rate

permanently reduces the entering cohort's employment rate by about 0.07 percentage point in each year of that cohort's careers. This is on top of larger effects early in the career, shown in Figure 8: In the scarring specification, the cohort's employment rate is reduced by an additional 0.7 percentage points at ages 22 and 23 and by 0.5 percentage points at ages 24 and 25.

A useful comparison for the cyclical sensitivity of the cohort effects δ_c is the sensitivity of the time effects β_t , which affect all college graduates in the labor market in a given year but do not persist. The second panel of Table 2 presents regressions of the estimated time effects β_t on the contemporaneous unemployment rate, using only the β_t estimates for $t < 2005$ and again controlling for a linear trend. These are much larger in magnitude than the coefficients for cohort effects – not surprisingly, the immediate effects of unemployment rate fluctuations are much larger than their persistent effects many years later. The amplitude of the cyclical fluctuations is much reduced in models that include additional terms for excess sensitivity of younger workers, but remains about five times as large as that of the cohort effects.

The third and fourth panels of Table 2 repeat the exercise for effects on log wages of those who are employed. In the cohort effect models in Panel C, I add two additional controls, an intercept shift and a trend break in 1978, to capture the clear change then seen in Figure 10; without these, the fit is extremely poor.

Neither cohort effects or time effects on log wages are significantly cyclical, consistent with longstanding evidence that wages do not fall in recessions (e.g., Bewley 1999). Evidently scarring effects on wages conditional on employment are relatively short-lived, justifying the scarring literature's emphasis on early-career effects, but effects on employment itself are more permanent.

C. Relative importance of hangover effects

An important question is the relative magnitude of the immediate and delayed effects of business cycle downturns. To approach this, I use the estimated scarring and sensitivity coefficients from Figure 9 and the cyclical estimates in Table 2 to measure the various effects of a one percentage point increase in the unemployment rate, distributed equally across the country. I present results in Table 3 for each of my main specifications, though my discussion focuses on the model that allows for both scarring and excess sensitivity effects, as in the final series in Figure 9 and column 4 of Table 2. To abstract from purely demographic considerations, I simulate the effect of a transitory one percentage point increase in the unemployment rate in a hypothetical stable population, with identically sized cohorts at each age from 22 through 40. I scale all estimates by presenting them as the total number of employment-years lost in a population with one worker per cohort.

I distinguish four components of the effect of the initial shock, corresponding to different terms in my decomposition. The first component operates through the time effects, β_t in (2)-(4). These are common to everyone in the labor market at the time of the transitory shock, and by Table 2 a shock to the unemployment rate reduces β_t by 0.41, translating to a decline in the employment rate of college graduates by 0.41 percentage points in that year. Because the time effects apply to all 19 cohorts in the labor market (recall that I focus only on those under 40), the total effect is a loss of $0.0041 \times 19 = 0.078$ employment-years.

A second component operates through the excess sensitivity coefficients θ_j . From Figure 8, $\theta_1 = -0.75$, meaning that the employment rate of 22 and 23-year-old graduates falls not by 0.41

percentage points but by $0.41+0.75=1.16$ percentage points for each one point increase in the unemployment rate. The effect is smaller but still substantial for slightly earlier cohorts: 24 and 25- year-olds see employment rate declines 0.34 percentage points larger than do established workers; 26- and 27-year-olds see declines of 0.31 percentage points, again relative to established workers; and so on. Cumulating these effects over all workers under age 31 in the labor market adds up to 0.034 total employment-years lost, on top of the 0.078 coming from common time effects.

A third component comes through medium-term scarring effects ϕ_j . These also apply only to young workers, but rather than affecting all young workers at a particular time, they affect a particular cohort through the beginning of its career. In the specification that includes both scarring and excess sensitivity effects, the scarring effects are small and account for only 0.003 lost employment-years, though in the simpler specification (2) they are larger, 0.025 lost employment-years.

The final component of the effect of a transitory downturn operates through the permanent cohort effect of the cohort that enters the labor market during the downturn. Based on Table 2, a one percentage point increase in the entry unemployment rate reduces the cohort effect for the new entrants by 0.071. This is relatively small, but it accumulates over the 19 years that the cohort is in the labor market, so adds up to 0.014 lost employment years.

In the lower panel of Table 3, the first row reports the share of the total effect of a downturn that derives from the hangover – effects that occur in later years – rather than from the downturn’s immediate effects. This is about one-quarter for the first two specifications, but falls to about one-eighth in specifications that allow for excess sensitivity effects. The final row

of the table shows the share of these delayed effects that is captured by the ϕ_j coefficients that have been the focus of the previous literature. In the specification most similar to those typically used, this share is 69%. However, it falls to just 17% in the specification that allows for excess sensitivity of young workers. Permanent scarring effects, operating through δ_c , are an important component of the persistent effects of transitory downturns.

C. Are post-GR outcomes consistent with pre-GR cyclicalities?

The analysis above demonstrates that recessions prior to 2005 had persistent effects on the employment rates and wages of those who enter the labor market during them. I return now to my investigation of the experience of the Great Recession cohorts and their successors. How much of the decline in these cohorts' employment rates can be attributed to past scarring patterns?

Figure 12 shows my estimated cohort effects along with predicted cohort effects from the specifications in Table 2. Panel A shows the employment estimates from the simple decomposition (2), while panel B shows employment estimates from an augmented decomposition that allows for both scarring and excess sensitivity effects on younger workers. Both predicted series show sharp downturns in the Great Recession, about twice as large in the baseline specification. But these downturns are dwarfed by the downturns in the actual series, which by 2010 are at least twice as large as expected based on past business cycles. Moreover, the predicted series turn up after 2010, reflecting the gradual decline in the unemployment rate after that, while as discussed earlier the actual series continue to decline. The cohort that entered

the labor market in 2014 has had employment rates 3 percentage points lower than predicted in each model, while the shortfall has grown to 5 or 6 percentage points by the 2018 entrants.

Thus, while patterns in prior business cycles would have predicted both early-career and permanent consequences for the Great Recession entrants, the actual trend has been far more negative than past experience would have suggested. Evidently, there has been an additional change to the employment prospects of recent entrants, over and above what is consistent with normal recession effects. Moreover, this change has persisted into the most recent entrants, who were in middle school during the Great Recession.

The lower panels of Figure 11 repeat the exercise for wages. Here, I modify the cyclical regression: I allow for an intercept shift and trend break with the 1978 cohort, as the series break at this point is so clear and models without this do not fit the data well at all. As seen in Figure 10 and Table 2, cohort effects on wages are only modestly cyclical, and become almost completely acyclical when I allow for scarring effects. There is no substantial divergence from the predictions for the Great Recession entrants, but wages have risen faster than predicted among more recent entrants.

VII. Conclusion

Something dramatic has happened to the employment prospects of recent college graduates. Starting with those who entered the labor market around 2005, each successive cohort has had lower employment rates, relative to older workers in the same labor market, than those before.

These cohort effects layer on to other changes in the labor market that are common to workers of different ages. During the Great Recession, overall employment rates were falling, but the employment rates of new entrants fell by more than did those of cohorts already established in the labor market. This is perhaps unsurprising; 100% of new graduates had to find jobs in a historically weak labor market, where most older workers were able to remain in their old jobs.

What is more surprising is that the decline in employment rates continued after the recession ended, not just for those who entered during it but for successive cohorts as well. The labor market grew steadily stronger through the 2011-2019 period, and overall employment rates rose (see Appendix Figure 3). But the employment rates of new graduates were stagnant, despite the rising tide. Even as the secular trend has increased the average employment rate of college graduates by a bit over four percentage points, the cohort effects of new entrants have fallen by more than that, with a net negative effect on (age-adjusted) employment rates.

The existing literature (e.g., Oreopoulos et al. 2012) provides several mechanisms by which weak economic conditions can have persistent effects on workers entering the labor market. For example, these workers may have trouble getting their feet on the first rung of career ladders, with long-term consequences: A worker who might have found a traineeship with opportunities for advancement in good times might instead settle for a dead-end job in bad times, reducing wages immediately but also reducing future career prospects. I confirm earlier evidence (from, e.g., Oreopoulos et al. 2012) that initial conditions have persistent effects on wages through the first several years of graduates' careers. I also document a new stylized fact: Initial conditions also have persistent effects on employment rates that, rather than fading away, last throughout a workers' career (at least through age 40). This represents an important long-

term cost of business cycles, and may be a mechanism behind the long-term health consequences documented by Schwandt and von Wachter (2019).

Nevertheless, while scarring effects are important, the magnitude of the relationship between the entry unemployment rate and cohort employment seen for pre-Great Recession entrants is simply too small to account for more than a fraction of the decline in employment rates more recently. Extrapolation of the past relationship implies that cohort employment rates for the 2010 entrants are 1.6 percentage points lower than would have been predicted, while those for the 2017 entrants are (in the three years they have been observed) 5.0 percentage points lower. Fortunately, the same is not true of wages. While wages for the Great Recession entrants are slightly lower than predicted based on past models, this did not persist for more recent entrants, whose wages are as much as 2.9 log points above the prior trend.


The employment patterns indicate that we were not well prepared for the pandemic-induced downturn of 2020. The 2020 graduates, and likely those of 2021 and 2022, will enter a labor market that is historically weak. The results here indicate that this will permanently scar them, and moreover that this will reduce employment below what would already be a dramatically reduced level compared to pre-2005 cohorts. A deeper understanding of the secular decline is urgently needed to support potential policy responses.

References

- Barr, Andrew, and Sarah Turner (2015). "Out of work and into school: Labor market policies and college enrollment during the Great Recession." *Journal of Public Economics*, 124: 63-73.
- Bewley, Truman F. (1999). *Why Wages Don't Fall during a Recession*. Harvard University Press: Cambridge.
- Borgschulte, Mark, and Paco Martorell (2018). "Paying to avoid recession: Using reenlistment to estimate the cost of unemployment." *American Economic Journal: Applied Economics*, 10.3: 101-127.
- Bound, John, Michael F. Lovenheim, and Sarah Turner (2012). "Increasing time to baccalaureate degree in the United States." *Education Finance and Policy*, 7(4), 375-424.
- Card, David, and Jesse Rothstein. "Racial segregation and the black–white test score gap." *Journal of Public Economics* 91.11-12 (2007): 2158-2184.
- Center for Economic and Policy Research. 2018. CEPR ORG Version 2.3, Apr 10, 2018, http://ceprdata.org/wp-content/cps/programs/org/cepr_org_master.do.
- Clark, Kim B., and Lawrence H. Summers (1982). "Labour force participation: timing and persistence." *The Review of Economic Studies*, 49(5): 825-844.
- Cunningham, Evan (2018). "Great Recession, great recovery? Trends from the Current Population Survey." *Monthly Labor Review*, U.S. Bureau of Labor Statistics, April 2018
- Daly, Mary C., Bart Hobijn and Joyce Kwok (2009). "Jobless recovery redux?" *Federal Reserve Bank of San Francisco Economic Letter*, June.
- Daly, Mary C., Bart Hobijn, and Theodore S. Wiles (2012). "Dissecting aggregate real wage fluctuations: Individual wage growth and the composition effect." Federal Reserve Bank of San Francisco working paper 2011-23, May.
- Davis, Steven J., and John Haltiwanger (2014). "Labor Market Fluidity and Economic Performance." NBER Working Papers 20479, National Bureau of Economic Research, Inc.
- Decker, Ryan A., John Haltiwanger, Ron S. Jarmin, and Javier Miranda (2014). "The Secular Decline in Business Dynamism in the U.S." Working paper.
- Gronau, Reuben (1974). "Wage comparisons--A selectivity bias." *Journal of Political Economy* 82.6: 1119-1143.
- Hershbein, Brad, and Lisa B. Kahn (2018). "Do Recessions Accelerate Routine-Biased Technological Change? Evidence from Vacancy Postings." *American Economic Review*, 108 (7): 1737-72.
- Jaimovich and Henry Siu (2020). "Job polarization and jobless recoveries." *Review of Economics and Statistics*, 102 (1): 129-147.
- Kahn, Lisa B. (2010). "The long-term labor market consequences of graduating from college in a bad economy." *Labour Economics*, 17.2: 303-316.


- Modestino, A., Daniel Shoag, and Joshua Ballance (2016). "Downskilling: Changes in Employer Skill Requirements Over the Business Cycle." *Labour Economics*, 41: 333-347.
- Oreopoulos, Philip, Till von Wachter, and Andrew Heisz (2012). "The short-and long-term career effects of graduating in a recession." *American Economic Journal: Applied Economics* 4.1: 1-29.
- Schmitt, John (2003). "Creating a consistent hourly wage series from the Current Population Survey's Outgoing Rotation Group, 1979-2002." Center for Economic Policy Research technical paper, August.
- Schulhofer-Wohl, Sam (2018). "The age-time-cohort problem and the identification of structural parameters in life-cycle models." *Quantitative Economics*, 9.2: 643-658.
- Schwandt, Hannes, and Till von Wachter (2019). "Unlucky Cohorts: Estimating the long-term effects of entering the labor market in a recession in large cross-sectional data sets." *Journal of Labor Economics*, 37(1), S161-S198.
- Snyder, Thomas D., Cristobal De Brey, and Sally A. Dillow (2019). *Digest of Education Statistics 2018* (NCES 2020-009). National Center for Education Statistics, Institute for Education Sciences, U.S. Department of Education: Washington DC., Table 302.10
- Yagan, Danny (2019). "Employment hysteresis from the Great Recession." *Journal of Political Economy*, 127.5: 2505-2558.

Figure 1. College graduate share by year and age




Notes: Each grey dashed line represents repeated cross sectional averages for a single birth cohort, followed from 22 through 30 across successive CPS surveys. For example, the leftmost series includes those who were 22 in 1980, observed in the 1980-1988 CPS surveys. Only cohorts that are age 22 in even-numbered years are shown. Solid lines connect observations across cohorts at ages 22, 24, and 30.

Figure 2. Employment and unemployment




Notes: The unemployment rate is the headline, seasonally adjusted rate reported by the Bureau of Labor Statistics. The non-employment rate is computed from CPS microdata and pertains to those ages 25-54. It is seasonally adjusted via a regression on year and month indicators, then smoothed using a 7-month centered triangle kernel. Recessions are indicated by shaded bars.

Figure 3. Non-employment and unemployment among college graduates aged 22-40




Notes: Non-employment and unemployment rates are computed from CPS microdata, using the subsample of college graduates aged 22-40. Both are seasonally adjusted via regressions on year and month indicators, then smoothed using a 7-month centered triangle kernel. Recessions are indicated by shaded bars.

Figure 4. Employment of young college graduates, by age




Notes: Employment rates are computed from CPS microdata, using the subsamples of college graduates aged 22-30 and 31-40. Both series are seasonally adjusted via regressions on year and month indicators, then smoothed using a 7-month centered triangle kernel. Recessions are indicated by shaded bars.

Figure 5. Cohort employment rates




Notes: Figure shows cohort effects on employment from three specifications, computed from CPS samples of college graduates aged 22-40. The “raw mean” series uses the simple mean employment rate for each cohort, averaged over all years that the cohort is observed without adjustment for age differences and normalized relative to the 1984 entry cohort. The “age adjusted” series derives from an estimate of equation (1), and is also normalized relative to the 1984 entry cohort. The “age-time-cohort decomposition” series derives from an estimate of equation (2), and both the 1984 and 2000 entry cohorts are normalized to zero. The 2019 entrants are included in the estimation but their coefficients, which are very noisily estimated, are not plotted. Recessions are indicated by shaded bars.

Figure 6. Cohort effects on employment of college graduates



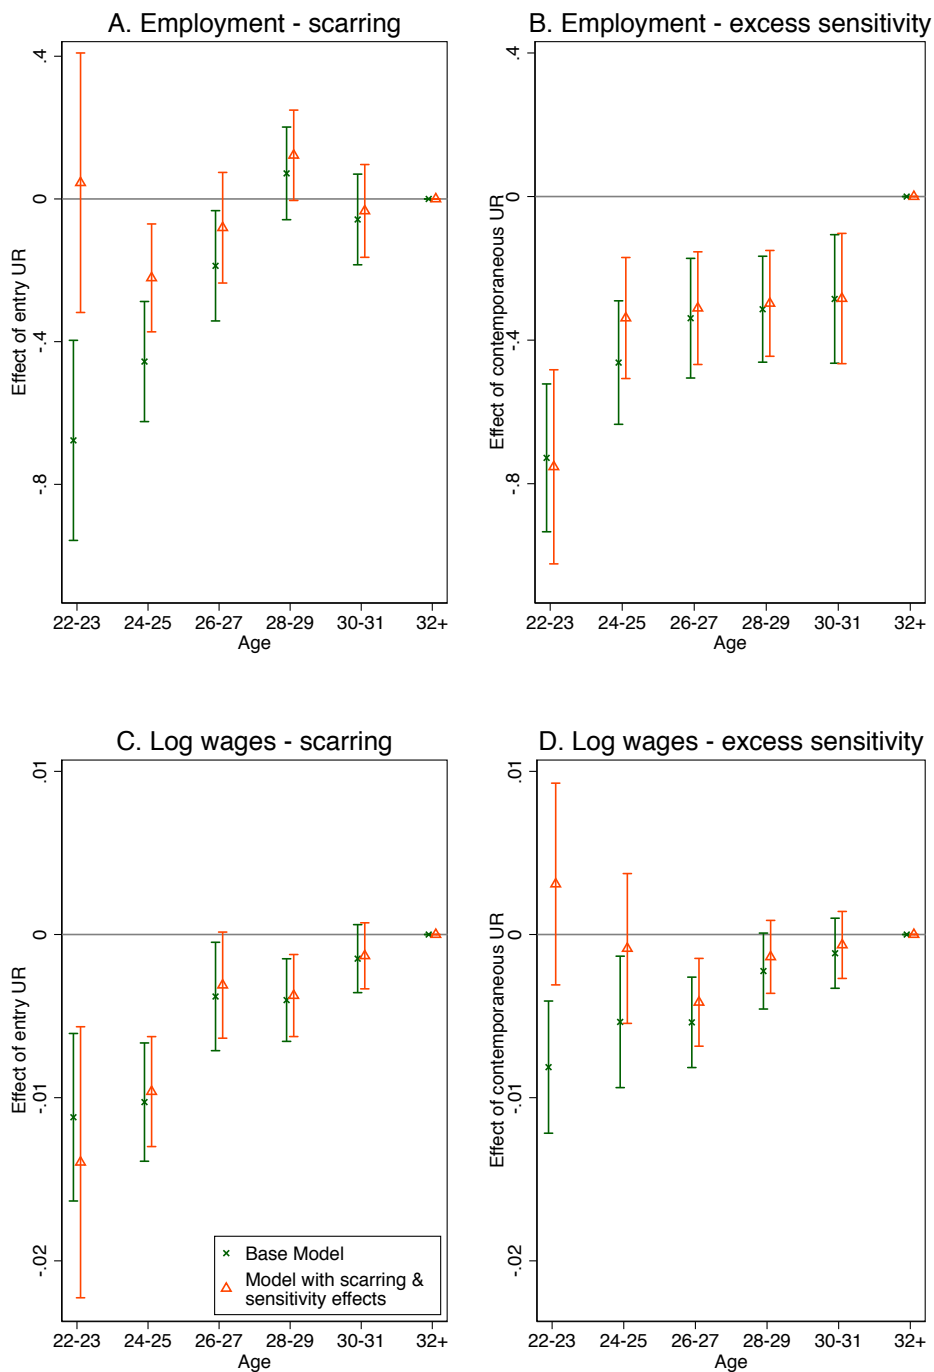
Notes: Figure shows cohort effects on employment from three specifications, computed from CPS samples of college graduates aged 22-40 and normalized to zero in 1984 and 2000. The “baseline” series repeats the “age-time-cohort decomposition” series from Figure 5. The “pre-recession fit” series estimates equation (2) using just cohorts that turned 22 in 2000 or earlier. Cohort effects for subsequent cohorts are estimated as the mean difference between observed employment rates and predictions based on the estimated model. The “age 25+” model includes all cohorts (except the most recent entrants) in the sample, but removes observations when the cohort is less than 25 years old. The 2019 entrants are included in the estimation but their coefficients, which are very noisily estimated, are not plotted. Recessions are indicated by shaded bars.

Figure 7. Cohort wages




Notes: Figure shows cohort effects on log real hourly wages from two specifications, computed from CPS samples of college graduates aged 22-40. The “age adjusted” series derives from an estimate of equation (1), and is normalized relative to the 1984 entry cohort. The “age-time-cohort decomposition” series derives from an estimate of equation (2), and both the 1984 and 2000 entry cohorts are normalized to zero. The 2019 entrants are included in the estimation but their coefficients, which are very noisily estimated, are not plotted. Recessions are indicated by shaded bars.

Figure 8. Scarring and excess sensitivity effects on employment and log wages




Notes: “Base model” derives from estimates of equations (3) (panels A and C) and (4) (panels B and D), and shows coefficients and confidence intervals, clustered at the state level, for ϕ_j and θ_j , respectively. Expanded model shows estimates from a model that includes both terms.

Figure 9. Cohort effects on employment with adjustments for excess sensitivity and scarring



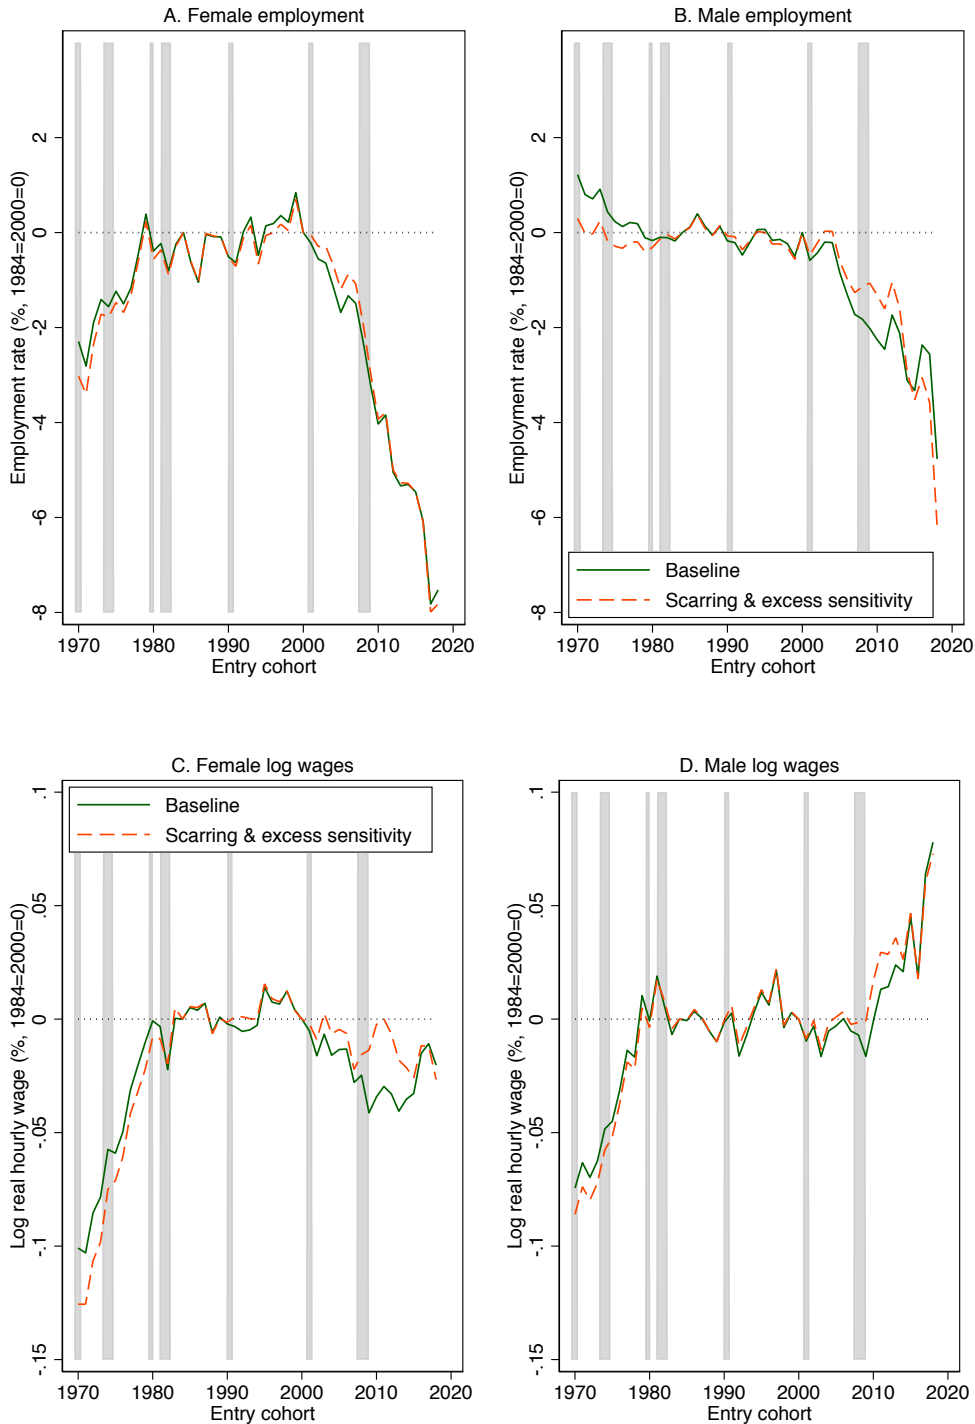
Notes: Figure shows cohort effects on employment from equations (2), (3), and (4), and from an expanded specification that includes all terms from both (3) and (4). All are estimated on CPS samples of college graduates aged 22-40. All cohort effects series are normalized to zero in 1984 and 2000. The 2019 entrants are included in the estimation but their coefficients, which are very noisily estimated, are not plotted. Recessions are indicated by shaded bars.

Figure 10. Cohort effects on log wages with adjustments for excess sensitivity and scarring



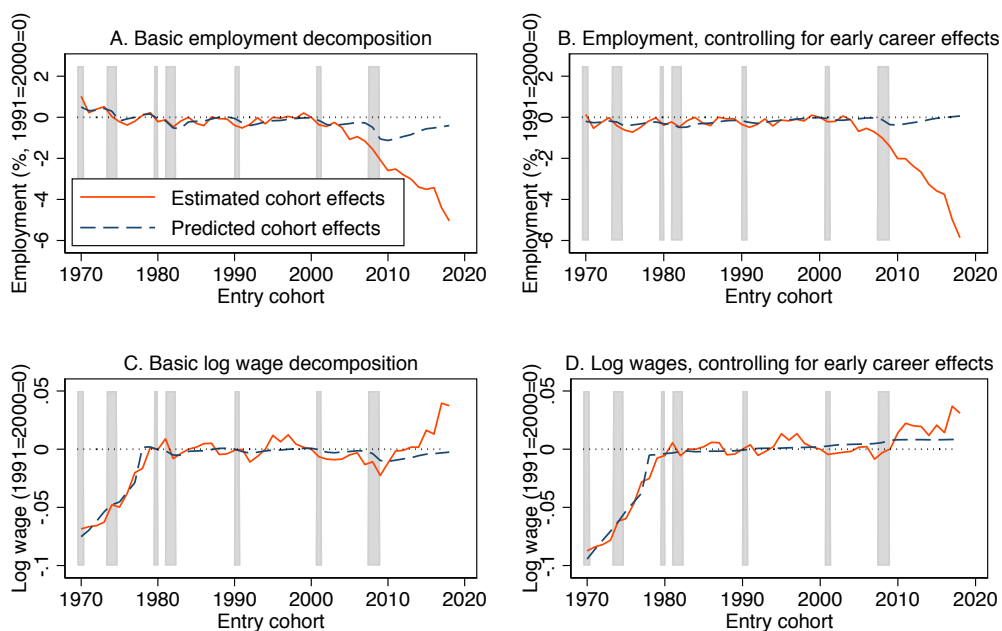
Notes: Figure shows cohort effects on log real hourly wages from equations (2), (3), and (4), and from an expanded specification that includes all terms from both (3) and (4). All are estimated on CPS samples of employed college graduates aged 22-40. All cohort effects series are normalized to zero in 1984 and 2000. The 2019 entrants are included in the estimation but their coefficients, which are very noisily estimated, are not plotted. Recessions are indicated by shaded bars.

Figure 11. Cohort effects on employment and wages, by gender



Notes: Figure repeats the first and last series from Figures 9 and 10, estimating the specifications separately for men and women. No cross-equation restriction are imposed, and the 1984 and 2000 entry cohort effects are separately normalized to zero for each group.

Figure 12. Predicted cohort effects based on pre-2005 patterns



Notes: The solid lines show the estimated cohort effects series from the first and last specifications in Figures 9 and 10. The 2019 entrants are included in the estimations but their coefficients, which are very noisily estimated, are not plotted. The dashed lines show fitted values from regressions of these cohort effects on time trends and the entry-year unemployment rate. These regressions use just the cohorts entering prior to 2005; the models for log wages also include an indicator for pre-1978 entrants and a separate time trend for these cohorts.

Table 1. Summary statistics

	Mean	SD	p10	p90
	(1)	(2)	(3)	(4)
Birth year	1970.8	12.0	1954	1987
Entry year	1992.8	12.0	1976	2009
Age	31.3	5.2	24	39
Unemployment rate (state, %)	6.0	2.0	3.8	8.8
Unemployment rate at entry (state, %)	6.2	1.9	4.2	8.7
Fraction of cohort with college degree	0.31	0.09	0.20	0.43
Graduate degree (0/1)	0.29	0.45		
Employment (0/1)	0.85	0.35		
Log real hourly wage	3.06	0.53	2.38	3.73

Notes: Summary statistics relate to the main analysis sample, consisting of college graduates aged 22-40. Statistics are weighted using CPS sampling weights.

Table 2. Relationship between unemployment rate and cohort and time effects

	Age-time-cohort decomposition specification			
	Baseline (1)	Scarring (2)	Excess sensitivity (3)	Both (4)
<u>A. Cohort effects on employment</u>				
Unemployment rate in entry year	-0.16 (0.04)	-0.06 (0.04)	-0.08 (0.03)	-0.07 (0.03)
Trend	-0.022 (0.004)	0.007 (0.005)	0.001 (0.003)	0.004 (0.003)
<u>B. Time effects on employment</u>				
Unemployment rate	-0.62 (0.30)	-0.58 (0.29)	-0.38 (0.28)	-0.41 (0.28)
Trend	-0.05 (0.09)	-0.09 (0.09)	-0.07 (0.08)	-0.08 (0.08)
<u>C. Cohort effects on log wages</u>				
Unemployment rate in entry year	-0.0017 (0.0011)	0.0006 (0.0013)	-0.0014 (0.0011)	0.0004 (0.0012)
Trend	0.0072 (0.0010)	0.0076 (0.0010)	0.0079 (0.0010)	0.0080 (0.0010)
<u>D. Time effects on log wages</u>				
Unemployment rate	-0.0084 (0.0075)	-0.0078 (0.0076)	-0.0060 (0.0077)	-0.0069 (0.0077)
Trend	0.0006 (0.0022)	-0.0004 (0.0022)	0.0004 (0.0022)	-0.0003 (0.0022)

Notes: Panels A and C show regressions of estimated cohort effects on employment (panel A) and log wages (panel C) on the unemployment rate in the cohort's first year in the labor market. Panels B and D show regressions of estimated calendar year effects on the contemporaneous unemployment rate. Unemployment rates and employment rates are measured in percentage points. The specification in Panel C also includes an indicator for pre-1978 entrants and a separate pre-1978 linear trend. Cohort and year effects in columns 1-4 come, respectively, from specifications (2), (3), (4), or expanded specification that includes both scarring and excess sensitivity effects. Only effects from cohorts entering prior to 2005 (panels A and C) or from years before 2005 (panels B and D) are used in the second-stage regressions. Newey-West standard errors allow for autocorrelation at two lags, and are not adjusted for estimation of the dependent variable.

Table 3. Decomposing the effect of a transitory increase in unemployment

	Employment years lost			
	Baseline (1)	Scarring (2)	Excess sensitivity (3)	Scarring & excess sensitivity (4)
Effect of a one p.p. increase in unemployment operating through:				
Immediate effects				
Time effects	-0.118	-0.110	-0.072	-0.078
Excess sensitivity	--	--	-0.037	-0.034
Delayed effects				
Early-career scarring	--	-0.025	--	-0.003
Cohort effects	-0.030	-0.011	-0.015	-0.014
Total	-0.148	-0.146	-0.124	-0.128
Delayed effects as % of total	20%	25%	12%	13%
Early-career scarring effects as % of total delayed effects	--	69%	--	17%

Notes: Table shows the total effect of a transitory 1 percentage point increase in the unemployment rate, measured as employment-years divided by the size of an entering cohort. Effects are decomposed based on the terms of equations (2)-(4), which are separated into those that manifest in the year of the downturn (immediate effects) and those that manifest primarily in subsequent years (delayed effects). Thus, in the specification in column (4), a one percentage point increase in unemployment reduces employment of all 19 cohorts of college graduates in the labor market in that year by 0.41 percentage points, cumulating to 0.078 employment-years per member of a single cohort. It also has additional effects on younger workers that cumulate to 0.034 employment-years, and hangover effects on the cohort that begins its career in the year of the downturn that cumulate to 0.017 employment-years. The final rows show the delayed effects as a share of the total and the effects coming from early-career scarring (ϕ_j in (3)) as a share of all delayed effects.

Appendices

Appendix A. Data processing and definitions

The primary analysis concerns the cohort-age-year-state employment rate and mean log real wages for college graduates. Employment rates are estimated from the monthly Current Population Survey (CPS). I construct a synthetic panel from data from each month from 1979 through 2019, aggregated to the calendar year. Log wages are computed from the the 1979-2017 Merged Outgoing Rotation Group (ORG) files from the monthly CPS, using an algorithm adapted from Center for Economic and Policy Research (2018; see also Schmitt 2003). For those paid hourly, I use the higher of the reported hourly wage (without overtime, tips, or commissions) and the ratio of weekly earnings (which in principle include overtime, tips, and commissions) to usual weekly hours. For non-hourly workers, I use the latter. When necessary, usual hours are imputed using actual hours last week or the mean by gender and full time/part time status. Hourly wages are converted to real 2015 dollars using the CPI, set to missing if below \$1 or above \$200, and logged.

My sample for age-time-cohort decompositions consists of respondents born 1948-1997 (who turned 22 between 1970 and 2019) who are surveyed at ages 22-40 and who report being college graduates.

Individual respondents are assigned the age (and implied birth cohort) that they are on the survey date, so someone born on July 1, 1980 and surveyed in 2010 will be treated as age 29 in 2010 and in the 1981 birth cohort, if surveyed in the first half of the calendar year, and as 30 and in the 1980 cohort if surveyed in the second half. Before 1992, graduates are those with 16 or more years of completed schooling; afterward, it is those with Bachelors or graduate degrees.

In some specifications (e.g., in Figure 6), I identify the age, time, and auxiliary control coefficients only from pre-Recession cohorts. For these specifications, I estimate the decomposition using the subsample excluding post-1978 birth cohorts, who turned 22 after 2000. I use the resulting coefficients to form predictions for the subsequent cohorts if the relevant cohort effects were zero, then estimate these cohort effects as the average difference between observed outcomes and the predictions.

My unemployment measures are the state unemployment rates, averaged over the calendar years. When assigning the entry unemployment rate, I use the rate in the state of current residence in the year when the respondent was 22 years old. State unemployment rates are not available prior to 1976; for earlier entrants, I use the national unemployment rate.

All analyses are weighted by the sum of CPS sample weights in the cell.

Appendix B. Additional specifications

I present a range of additional specifications and models in the appendix. I discuss them in turn.

Educational attainment

Figure 1 shows the evolution of degree attainment across ages for synthetic cohorts created from the CPS samples. It shows that the share of the cohort with a college degree rises rapidly from age 22 through age 24 and more slowly after that. The share of people with a college degree by age 30 has risen substantially over time. The share earning degrees between age 24 and 30 rose dramatically beginning around the cohort that was 22 in 1990, but has not grown much since then.

Appendix Figure 1 repeats this exercise for the share of each cohort that has some college or more. The pattern is strikingly different. The share with some college has grown dramatically across cohorts. Within each cohort, it generally grows at a similar rate across ages. To see this, note that the point for each cohort at age 30 is six years to the right of the point for the same cohort at age 24. The slope between these points is generally similar to the slope of each age-specific pattern across cohorts. It is a bit flatter before 2000 and a bit steeper after that, but there is no indication of a meaningful change for the Great Recession entrants.

Main decompositions

Appendix Figures 2 and 3 show estimated age and time effects, respectively, from the main decomposition (2) for employment of college graduates. Age effects are normalized to zero at age 32, while year effects are normalized to zero in 2007. (As discussed in the main text, the cohort effects from this decomposition, shown as the solid line in Figure 5, are normalized to zero for both the 1984 and 2000 entering cohorts.)

Appendix Figure 4 shows cohort effects on log wages from three different specifications. Cohort effects on employment from these same specifications are shown in Figure 6, and the specifications are described in the main text.

Appendix Figure 5 shows estimates of the scarring and excess sensitivity coefficients (θ_j and ϕ_j , respectively) from versions of my decompositions that are estimated on samples limited to pre-2000 entering cohorts. These are quite similar to the full sample estimates in Figure 8.

Sensitivity to mobility

In equation (3), I use the unemployment rate when the individual was 22, in the state of current residence, as a measure of labor market conditions when that individual entered the market. Interstate mobility between age 22 and the date of the CPS interview will make this a noisy measure of initial conditions. Further, any endogeneity of this mobility to local conditions could bias my results.

To address this, I use an instrumental variables strategy adapted from Schwandt and von Wachter (2019).¹¹ Construction of the instrument has three steps.

First, using pooled data from the 1980, 1990, and 2000 decennial census and from the 2001-2016 American Community Survey (one-year samples) public use microdata files, I estimate the number of college graduates born in each state b living in state s at age a ($22 \leq a \leq 40$), N_{bsa} . This is a long-run average, pooled across nearly forty years of data, so is not influenced by economic conditions relevant for any single cohort.

Second, I construct the average age-22 unemployment rate to which college graduates born in each state b and each entry cohort c would have been exposed. This average rate is then:

$$\overline{UR22}_{bc} = \frac{\sum_s UR_{sc} N_{bs,22}}{\sum_s N_{bs,22}}.$$


Third, for each state s , cohort c , and age a , I average the average birth-state exposure across all birth states, weighting by the share of age- a graduates in state s born in each state b :

$$\overline{UR22}_{sca} = \frac{\sum_b \overline{UR22}_{bc} N_{bsa}}{\sum_b N_{bsa}}.$$

Last, I interact this measure with age indicators $g_j(a)$ ($j=1, \dots, 5$) and use the results as instruments for the $D_a^j * UR_{sc}$ interactions in equation (3). Appendix Figure 6 shows the estimated scarring coefficients from the base OLS specification and using IV, first for equation (3) (in Panel A) and second for an augmented specification that also includes excess sensitivity controls (in Panel B). Coefficients are similar in IV as in OLS specifications.


¹¹ Schwandt and von Wachter (2019) use a double-weighting estimator to abstract from both endogenous mobility and endogenous attainment rates. I treat attainment rates as exogenous, using the selection correction discussed in the text to address changes in cohort composition with age.

Appendix Figure 1. Some college share by year and age




Notes: Each grey dashed line represents repeated cross-sectional averages for a single birth cohort, followed from 22 through 30 across successive CPS surveys, and shows the share with some college or more at each age. For example, the leftmost series includes those who were 22 in 1980, observed in the 1980-1988 CPS surveys. Only cohorts that are age 22 in even-numbered years are shown. Solid lines connect observations across cohorts at ages 22, 24, and 30. Crossings of these lines indicate that the cross-sectional ranking of these age groups by attainment is not stable over time.

Appendix Figure 2. Age effects from age-cohort and age-time-cohort decompositions of employment




Notes: Figure shows estimated age effects γ_a on employment from equation (2), estimated on cohort-age-state-year cells constructed from CPS microdata samples of college graduates aged 22-40. The age-32 coefficient is normalized to zero; the model also normalizes the cohort effects to zero for both the 1984 and 2000 entering cohorts.

Appendix Figure 3. Time effects from age-time-cohort decomposition of employment



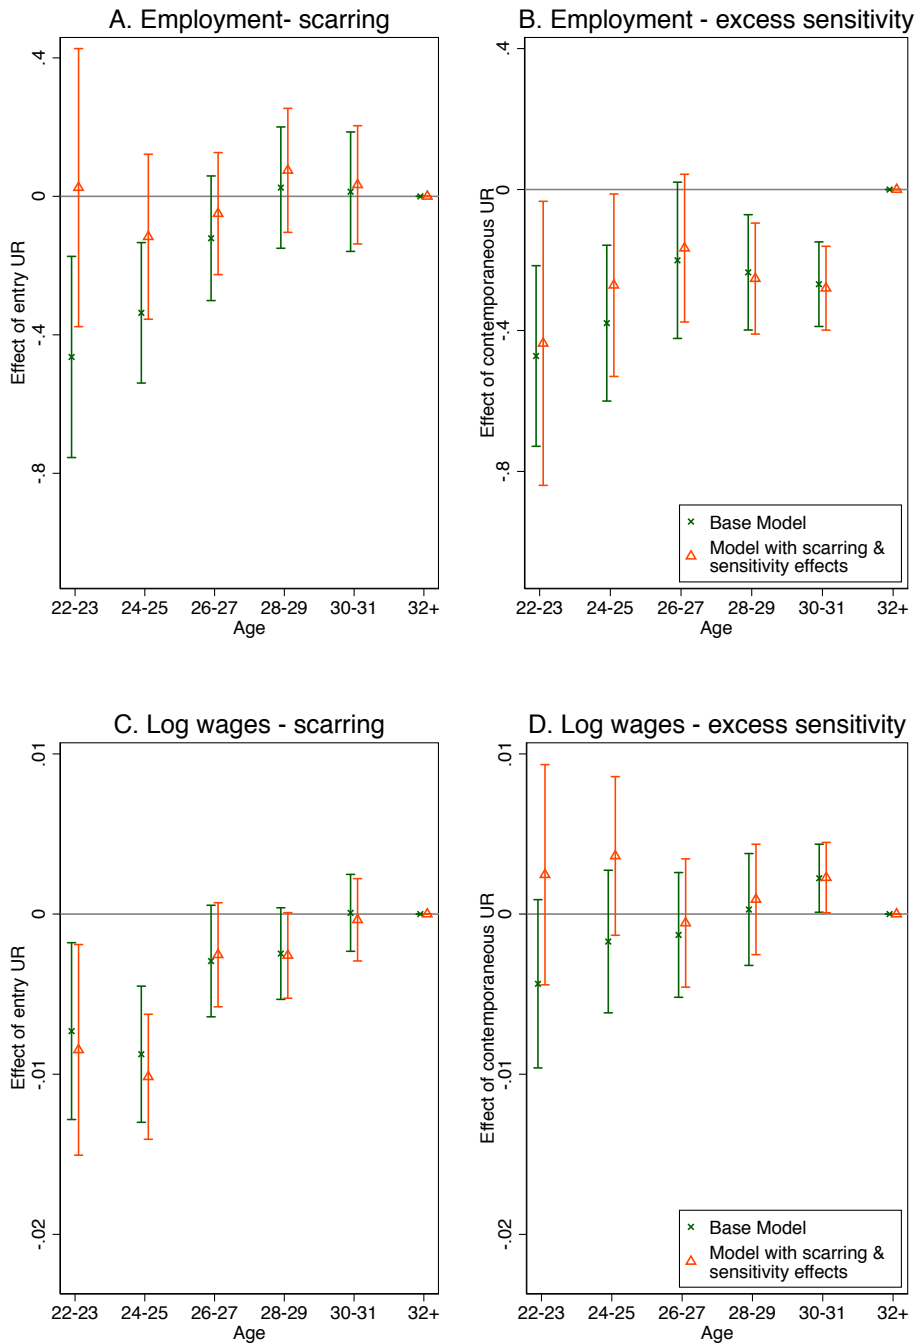
Notes: Figure shows estimated year effects β_t on employment from equation (2), estimated on cohort-age-state-year cells constructed from CPS microdata samples of college graduates aged 22-40. The 2007 coefficient is normalized to zero; the model also normalizes the cohort effects to zero for both the 1984 and 2000 entering cohorts. Recessions are indicated by shaded bars.

Appendix Figure 4. Cohort effects on log wages of college graduates




Notes: Figure shows cohort effects on log real hourly wages from three specifications, computed from CPS Outgoing Rotation Group samples of employed college graduates aged 22-40 and normalized to zero in 1984 and 2000. The “baseline” series repeats the “age-time-cohort decomposition” series from Figure 5. The “pre-recession fit” series estimates equation (2) using just cohorts that turned 22 in 2000 or earlier. Cohort effects for subsequent cohorts are estimated as the mean difference between observed log wages and predictions based on the estimated model. The “age 25+” model includes all cohorts (except the most recent entrants) in the sample, but removes observations when the cohort is less than 25 years old. The 2019 entrants are included in the estimation but their coefficients, which are very noisily estimated, are not plotted. Recessions are indicated by shaded bars.

Appendix Figure 5. Scarring and excess sensitivity estimates based on entrants in 2000 and earlier



Notes: “Base model” derives from estimates of equations (3) (panels A and C) and (4) (panels B and D), and shows coefficients and confidence intervals, clustered at the state level, for ϕ_j and θ_j , respectively. Expanded model shows estimates from a model that includes both terms. Each is estimated on a subsample that excludes cohorts entering the labor market after 2000.

Appendix Figure 6. Instrumental variables estimates of scarring



Notes: OLS estimates are those plotted in panel A of Figure 8. IV estimates are from a specification that instruments for $D_a^j * UR_{sc}$ with $D_a^j * \overline{UR_{sc,22}}$, as discussed in the text.