

The Lost Generation? Scarring after the Great Recession

Jesse Rothstein¹

UC Berkeley, NBER, and IZA

May 2019

Abstract:

I investigate medium- and long-term impacts of the Great Recession on post-recession college graduates. Most so-called “scarring” models emphasize effects of initial conditions that attenuate over the first decade of a worker’s career. But early career recessions may also have permanent effects. I decompose the recent cohorts’ experience into transitory time effects, medium-term scarring, and permanent cohort effects. Cohort effects are strongly cyclical. Medium-term scarring explains only half of this cyclical. The long-run cumulative effect of the recession on graduates’ employment is more than twice as large as the immediate effect.

¹ rothstein@berkeley.edu. I am grateful to Audrey Tiew, Rachel Young, and Nathaniel Ruby for excellent research assistance; to 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

Prolonged unemployment may have long-run consequences for future employability and productivity, perhaps particularly so for young workers for whom weak conditions may mean reduced access to on-the-job experience and training. Consistent with this, college graduates who enter the labor market during downturns are “scarred,” with lower employment and wages throughout their early careers than similar individuals who enter at better times (Kahn 2010; Oreopolous, von Wachter, and Heisz 2012; Schwandt and von Wachter, 2019).

I examine the medium- and long-run effects of the 2007-2009 Great Recession on new entrants’ labor market outcomes. Following the scarring literature, I focus on college graduates, whose human capital is unlikely to have become technologically obsolete during the downturn. The Great Recession provides a natural laboratory due to its depth and duration. The unemployment rate rose from 4.4 percent in May 2007 to 10.0 percent in October 2009, then did not reach 4.4 percent again until March 2017. Employment rates worsened at the same rate as unemployment during the downturn, but have been much slower to recover; to this day, prime-age employment remains below its pre-recession peak.

During the downturn, new labor market entrants had few job opportunities. Job openings were at least 20 percent below their pre-recession peak for nearly five years, and the hiring rate remained significantly depressed until late 2014. The share of individuals aged 20-25 who were neither working nor in school rose by nearly half, from 13.1% in 2007 to a high of 19.0% in 2013.

Comparing younger and older workers in the same labor market, the employment, annual earnings, and wages of those who entered the labor market during and after the recession are

lower than would be expected given the experience of older cohorts at the same time. College graduates whose careers started in 2010 and 2011 have employment rates two percentage points lower than the age- and time-adjusted employment rates of pre-recession cohorts, and those who work have annual earnings that are two percent lower. Employment declines persist for even the most recent labor market entrants, while annual earnings and wages have recovered since the immediate post-recession entrants.

Employment rates of the recession and post-recession entrants have remained low even in more recent data, when older cohorts have recovered. Accordingly, in an age-time-cohort decomposition I find that purely transitory effects account for only about 40% of the employment decline for cohorts entering the labor market in 2009-2011, and none of the decline for cohorts entering later. The remainder is attributed to declines in cohort effects.

I next investigate whether existing scarring specifications capture these patterns. An important feature of the specifications that have typically been estimated is that scarring effects, while long lasting, are not permanent. The effects of the initial employment rate are frequently estimated to dissipate by the individual's 10th year in the labor market (e.g., Borjaschulte and Martorell 2018; Kahn 2010; Oreopoulos et al. 2012). Thus, insofar as these models correctly characterize the scarring effect of the Great Recession, they imply that much of the damage to the post-recession cohorts has already been borne, and that the recession will not cast a large shadow going forward.

I reproduce scarring estimates on employment in Current Population Survey data, using similar specifications to those used by previous authors. Great Recession scarring is closely in line with extrapolations of estimates based on pre-2007 data. Whether I use the full sample or just

the pre-recession subsample, I find meaningful medium-term scarring that fades out over several years. However, this accounts for only about half of the depressed employment of the recent cohorts, relative to other cohorts in the post-recession labor market. The remainder manifests as lower cohort effects for the recent cohorts. These cohort effects, which have generally been treated as nuisance parameters in past scarring work (e.g., Oreopoulos et al., 2012), capture any effect of early career conditions that does not fade out over the early career. I show that these effects are cyclical even when medium-term scarring is controlled separately, and represent a quantitatively important component of the effects of recessions.

I close with a calculation of the contribution of these long-lasting effects to the total cumulative impact of the Great Recession, extrapolating the cohort coefficients through the decades remaining in the affected cohorts' potential careers. My estimates imply that, through a combination of short-, medium-, and long-run effects, the Great Recession will reduce the employment rates of college graduates over the 2007-2060 period by a total of over 15 percentage point-years. Nearly half of this loss reflects medium- and long-term effects of the recession that manifest after 2014, when the unemployment rate was no longer elevated but workers carry the scars of the earlier period. The medium-term effects will gradually fade over the next few years, while the long-run effects will persist as long as the recession entrants remain in the labor market. Cumulatively, the latter are much more important than the former. Effects that will manifest after 2025 account for nearly one-third of the total damage.

This evidence supports the view that business cycles have important hysteresis effects. DeLong and Summers (2012) argue that if post-recession hangovers are large enough counter-cyclical fiscal policies can be self-financing. My estimates of permanent effects on new labor

market entrants capture only a small component of potential hysteresis – and thus should be added to Yagan’s (forthcoming) complementary evidence of effects on older workers – but are on their own large enough to imply a meaningful fiscal offset for countercyclical policies.

II. The Great Recession in the labor market

The unemployment rate rose by 6.5 percentage points between mid-2007 and late 2009, (Appendix Figure A-1). It then declined roughly linearly after mid-2010. It has been below 6% since the third quarter of 2014 and below its pre-recession level since early 2017.

The prime-age (25-54) employment rate fell by over five percentage points by late 2009, and was much slower to recover. Only half of the decline had been erased by the end of 2015; at this writing the employment rate is 79.8%, still half a point below the 2007 peak.

Young people fared particularly poorly. The first panel of Figure 1 shows the employment rate separately for those aged 25-30 and those aged 31-54. Younger workers’ employment is much more cyclically sensitive than that of older workers; their employment rate fell more than five percentage points during the Great Recession, then recovered relatively quickly. The second panel distinguishes college graduates and non-graduates within the 25-30 group. Employment rates are higher and less cyclical for college graduates. However, the graduate employment rate declined substantially during the recession, then did not rebound at all through the first several years of the recovery. By the end of 2014, non-graduates had made up one-third of the ground lost during the recession, while young graduates’ employment was still *below* its end-of-recession level.

Figure A-2 shows average real hourly wages for the same groups broken out in Figure 1. All four mean wage series fell after the recession (though not during it, due to changes in the composition of workers; see Daly, Hobijn, and Wiles 2012), with a larger decline for younger workers and similar declines for young graduates and non-graduates. Older workers' wages have more than recovered in recent years, but young workers' wages, both graduates' and non-graduates', have only approached their 2007 local maxima and remain well below their 2000 peaks. In contrast to the employment series, however, the recovery of graduates' wages was not meaningfully slower than that of non-graduates.

III. Scarring and hysteresis

An influential literature examines the persistent effects of labor market conditions at the time of labor market entry (e.g., Kahn 2010; Oreopoulos et al. 2012; Schwandt and von Wachter, 2019; Borgschulte and Martorell 2018; Oyer 2006, 2008). This research generally finds sizable persistent effects on employment, wages, and earnings that fade out over the decade following graduation.¹ One interpretation is that new graduates are delayed in beginning to climb the job ladder during recessions; another is that labor market weakness at career start translates into a weak bargaining position for many years thereafter (Beaudry and DiNardo 1991). Many of the specifications used in this work (e.g., Oreopoulos et al. 2012) include cohort effects as controls, but treat them as nuisance parameters, focusing on interactions between the entry

¹ Another relevant literature examines the effects of job displacement on subsequent earnings (Jacobson, LaLonde, and Sullivan 1993; von Wachter and Bender 2006; Schmieder, von Wachter, and Heining 2018). These studies also find persistent effects. Davis and von Wachter (2012) find that the displacement effect is larger in recessions.

unemployment rate and early-career potential experience indicators. In contrast, I emphasize the cohort effects as key components of the persistence of recession impacts.

The so-called “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 (forthcoming) 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 the persistent impacts of the Great Recession to those of pre-2007 fluctuations.

IV. Data

I use repeated cross section data from the Current Population Survey (CPS) to examine individuals born between 1948 and 1995 and observed at ages 22 to 40 between 1979 and November 2018.

My primary outcome is the employment rate of young workers, measured in the monthly CPS. I also analyze log real annual earnings, from the Annual Social and Economic Supplement (the March CPS), and log real hourly wages, from the CPS Outgoing Rotation Group (ORG), for those with positive weekly or annual earnings. Data processing and definitions are discussed in the Appendix.

For each outcome and sample, I aggregate to the education-state-year-birth cohort cell. I use two education groups: Bachelor's degrees (BAs) and above, and less than a BA. Following the scarring literature, I focus on the former, but I present results for the latter as well. I merge onto each cell the contemporaneous state unemployment rate and the unemployment rate in the year of the cohort's labor market entry.²

There are three important limitations to the CPS data for my purposes. First, I do not observe the date that a respondent first searched for work. Following the scarring literature, I use the unemployment rate at age 22 (18 for non-graduates). This abstracts from potential endogeneity in the entry date.

Second, I observe only the respondent's current state of residence, not where he/she lived at labor market entry. Schwandt and von Wachter (2019) explore the bias in similar specifications from failing to measure location as of graduation, and conclude that it is "unlikely to be very large." I explore sensitivity to mobility in the Appendix.

Third, some of those observed without degrees at 22 will earn them later. This creates differential selection by age in the populations over which mean outcomes are computed. Figure A-3 shows the graduate share by birth cohort and year. The degree attainment rate for 22-year-olds is only half or one-third 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 1968 birth cohort (turning 22 in 1990; see Bound, Lovenheim, and Turner 2012), but still the great majority of degrees are awarded by 26.

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

Since the mid-1990s, final attainment has grown substantially, with most of this growth is in degrees awarded by age 24.³ There is no indication of a surge in attainment among Great Recession cohorts (though other analyses, not shown here, indicate that there are increases in college going that evidently does not lead to a BA degree). Together, these facts make it plausible that age effects, included in all of my specifications, will absorb most of the bias in cohort mean outcomes averaged across various ages.

I take two approaches to minimizing remaining composition bias. First, I include in all regressions an inverse Mills ratio computed from the cohort-by-age attainment rate (or, for non-graduates, from its complement). 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). Second, I have reestimated my primary specifications excluding observations from ages under 24. Results are largely unchanged.

V. Methods and results

A. Age-time-cohort decomposition

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, or $c=t-a+22$). I estimate age-time-cohort decompositions:

³ 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.

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

Here, β_t are a set of fixed effects for years, γ_a are age effects, and δ_c are birth cohort fixed effects. ζ_s represents state effects, while $\lambda(\cdot)$ is the inverse Mills function applied to the state-cohort-age attainment rate, p_{atc} . The time effects β_t capture both aggregate demand and any supply-side factors that are common across age groups, potentially including the hysteresis effects documented by Yagan (forthcoming). Cohort effects δ_c are permanent, and capture any 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 is not identified due to the linear dependency among them (see, e.g., Schulhofer-Wohl, 2018). I normalize the cohort effects for the 1991 and 2000 entry cohorts (born in 1969 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. My analyses of cohort effects always include or difference out linear time trends to absorb this normalization.

I am interested in the degree to which poor employment rates of young graduates following the 2007-2009 Great Recession can be attributed to transitory demand shocks, captured by β_t in equation (1), versus permanent differences between post-recession cohorts and other cohorts facing the same demand conditions, which would appear in δ_c . To avoid confounding the age and time effects with the experience of post-recession cohorts, who are observed only when young, in my preferred specifications I estimate (1) using data only on cohorts entering the market in 2000 and earlier. I then hold the β_t and γ_a vectors fixed when

estimating the cohort effects for more recent birth cohorts. I obtain nearly identical results, however, when I simply estimate (1) on the entire sample.

The solid line in Figure 2 shows the cohort effects from the basic specification, applied to college graduates and using only pre-recession cohorts to estimate time and age effects.⁴ Cohort effects are largely stable, declining slightly but smoothly, for the entry cohorts between 1975 and 2004. The overall slope is defined by my normalization and not by the data, so the smooth decline could equally well be an increase. However, any change in the slope is identified. There is a sharp downward turn for cohorts entering after about 2005. The cumulative decline from the 2004 to the 2015 cohorts, relative to the slope between 1993 and 2004, is nearly five percentage points, with most of this decline occurring between 2006 and 2010. There are also smaller declines in cohort effects in earlier recessions.

Although the most recent cohorts – those turning 22 after 2014 – have faced a labor market that was much stronger than that faced by 2009-2012 entrants, they have not recovered any of the lost ground, and continue to fare quite poorly relative to past cohorts observed at similar ages. By contrast, the pure time effects (Figure A-5) had fully recovered their 2007 levels by 2014. In other words, the model indicates that the slow recovery of college graduates' employment (shown as the solid line in Figure A-5) since 2011 is not evenly distributed but is specific to younger cohorts. Older cohorts who were already established in the labor market by 2007 had fully recovered by 2013.

⁴ Figures A-4 and A-5 show the age and time effects, respectively, while Figure A-6 shows that cohort effects are nearly identical when the full sample is used to estimate (1).

B. Excess sensitivity and scarring

In specification (1), the decline in cohort effects seen in Figure 2 implies that post-recession cohorts will have lower employment rates in every year of their careers than pre-recession cohorts, after adjusting for any differences in time effects. In other words, the recession has a large, permanent scarring effect, totaling four to five percentage points, on the post-recession cohorts' annual employment rates.

However, equation (1) does not allow for two alternative explanations that would point to rosier futures. First, young people's employment is more cyclically sensitive than that of older people. Because the post-recession cohorts are observed only at young ages, and primarily at times when the unemployment rate is high, excess sensitivity would manifest as downward-biased estimated cohort effects for these cohorts.

To address this, I augment (1) with unemployment rate-age interactions. I create a sequence of age indicators $g_j(a)$, $j=1,\dots,5$, where $g_1(a)$ is an indicator for ages 22-23, $g_2(a)$ for 24-25, $g_3(a)$ for 26-27, $g_4(a)$ for 28-29, and $g_5(a)$ for 30-31. I then control for interactions between these indicators and UR_{st} :

$$Y_{satc} = \alpha + \beta_t + \gamma_a + \delta_c + \zeta_s + \rho\lambda(p_{atc}) + \sum_{j=1}^5 g_j(a) * UR_{st} * \theta_j + (UR_{st} - UR_t)\kappa + \epsilon_{satc}. \quad (2)$$

I also include a main effect for the deviation of the state unemployment rate from the national rate, $(UR_{st} - UR_t)$, to capture effects of local conditions on those aged 32 and over.⁵

⁵ The θ_j coefficients are interpreted as the effect of the state unemployment rate on those from age group j , relative to its effect on those over 32. Because the unemployment rate main effect is specified as the deviation of the state from the national rate, the time effects β_t capture any effect of the national unemployment rate on those aged 32 and older. Thus, the specification

The estimated θ_j s are plotted, with confidence intervals, in Figure 3A. They show the expected excess sensitivity: Employment of the youngest college graduates falls by about 0.5 percentage point for every 1.0 percentage point increase in the unemployment rate, but the effect declines roughly linearly with age.

These estimates are based on the pre-recession cohorts, so generally on lower early-career unemployment rates than prevailed during the Great Recession. The dots in Figure 3A show estimates based on the full sample. These indicate somewhat more sensitivity, suggesting that the youngest graduates did even worse than expected during the Great Recession. But the extrapolation from pre-2007 patterns generally works well.

The short-dashed line in Figure 2 shows the estimated cohort effects from specification (2). The adjustment makes only a modest difference, with slightly better cohort effects for the 2001 to 2010 entry cohorts. It does nothing to improve the estimated cohort effects for the most recent entrants – not surprising, since the unemployment rate has not been elevated during most of their time in the market. Overall, this specification does not indicate that the apparent downturn in cohort effects in the base specification was an artifact of failing to allow for excess sensitivity of younger workers to economic conditions.

A second explanation for the apparent decline in cohort effects is that the recession cohorts may have been scarred by their initial experiences but, as implied by past estimates, the scarring effect will fade away with age. Because the post-recession cohorts are observed only at young ages, the estimated cohort effects from specification (1) will reflect the early-career

allows the national unemployment rate and the state deviation from that to have distinct main effects, but forces their age interactions to be the same.

impacts of elevated unemployment in 2007-2015, and may overstate the permanent damage to the post-recession cohorts.

Again, I address this by including terms in the decomposition that allow effects of entry conditions to fade out gradually. The specification replaces the contemporaneous unemployment rate UR_{st} in (2) with the unemployment rate at labor market entry (i.e., at age 22), $UR22_{sc}$:

$$Y_{satc} = \alpha + \beta_t + \gamma_a + \delta_c + \zeta_s + \rho\lambda(p_{atc}) + \sum_{j=1}^5 g_j(a) * UR22_{sc} * \varphi_j + (UR22_{sc} - UR22_c)\kappa + \epsilon_{satc}. \quad (3)$$

The φ_j coefficients represent age-cohort-state interactions.⁶ Thus, this model has two ways that economic conditions in the first year of the career can have persistent effects: They can have medium-term consequences through the φ_j coefficients, and they can have permanent consequences via the cohort coefficients δ_c .

Medium-term persistence is importantly different in (2) and (3). In (2), a transitory shock in t has a uniform effect via β_t and additional effects on those in the first ten years of their careers through θ_j , but none of these effects persist in $t+1$. In (3), a shock has the same uniform effect via β_t but may have additional effects on the newest entrants; if so, these workers remain damaged for up to ten years. By contrast, those who entered in $t-1$ face no persistent damage in $t+1$ or thereafter. In principle, a specification with both types of persistence is identified, but in

⁶ As before, I omit an interaction with the indicator for those aged 32 and above, and I include a main effect for the unemployment rate at cohort entry, deviated from the national rate in that year. This ensures that any permanent effects of the national unemployment rate at entry is captured by the δ_c cohort effects, while allowing the deviation of the state from the national unemployment rate to have its own effect through κ .

practice the two are difficult to distinguish as each predicts sustained poor outcomes for young workers.

Figure 3B shows the estimated φ_j coefficients from (3). We see strong effects of initial conditions on early career employment, fading out over about six years.⁷ Again, these are slightly larger but basically similar when I include the post-recession cohorts in the sample as when I extrapolate from earlier cohorts, indicating that nonlinearity in the unemployment rate effect is not a major source of bias.

The cohort effects from this specification, plotted as long dashed lines in Figure 2, show a smaller, though still quantitatively important, decline in the period surrounding the Great Recession. This indicates that the earlier specifications may have overstated the degree of permanent damage by failing to allow for persistent but non-permanent effects of entry conditions. In the adjusted series, there is still a sharp change in the series around the 2005 cohort, but the subsequent decline is more gradual and linear. The decline from 2006-2010 is 60% as large as in the baseline specification. Declines in cohort effects during prior recessions are also attenuated.

Appendix Figure A-13 shows cohort effects separately for men and women, using versions of specifications (1) and (3) that interact age, cohort, and scarring variables with gender. Figure A-12 shows the associated scarring coefficients. The post-recession decline in cohort effects is larger for women than for men though apparent for both; semi-transitory scarring is more

⁷ It is possible that the later-age scarring coefficients are attenuated due to geographic mobility. My φ_j estimates are similar when, following von Wachter and Heisz (forthcoming), I instrument for UR_{sc} with a measure that abstracts from mobility. See the appendix.

important for men than for women, but even when it is accounted for we see declines in both male and female cohort effects following the recession.

C. Cyclicalities of estimated cohort effects

To quantify the cyclicalities of the cohort effects in Figure 2, Table 1 shows time series regressions of the three sets of estimated cohort effects on the national unemployment rate when the cohort was 22, controlling for a quadratic time trend (which absorbs the normalization of the cohort effects). Standard errors allow for autocorrelation of the error terms at up to three lags.

Column 1 shows results for the initial decomposition, (1). Each percentage point increase in the unemployment rate at the time of labor market entry is associated with a 0.22 percentage point permanent decline in the cohort employment rate. This is only slightly reduced when the specification allows for excess sensitivity but declines to -0.11 and becomes insignificant when it allows for fadeout of early-career scarring effects. In other words, medium-term scarring effects account for about half of the cyclicalities of scarring effects from (1); the remainder are imprecisely estimated but still appear modestly cyclical. As I discuss below, the remaining response is quantitatively quite important.

Column 4 presents results from a specification that combines (2) and (3), allowing for both excess sensitivity and early-career scarring. In this specification, the θ_j coefficients, capturing excess sensitivity, are similar to those seen in Figure 3A, but the φ_j coefficients are near zero. Cohort effect cyclicalities are as in specification (2).

The lower rows of Table 1, Panel A, compare the estimated impact of the Great Recession to what the cyclical sensitivity model would predict, by comparing the change between the 2006

and 2010 entry cohorts, relative to that between the 2002 and 2006 cohorts, to the predicted change. The model captures the experience of the 2010 cohort reasonably well, though in column 3 the predicted decline is somewhat larger than the actual one.⁸ This is an indication that the parameters of the decompositions (1)-(3) are reasonably stable through the Great Recession period – estimates based on pre-recession cohorts predict accurately the experience of more recent cohorts.

D. Earnings and wages

I have also conducted parallel analyses of log annual earnings and log hourly wages, in each case conditional on employment. Estimates are shown in Figures A-7 and A-8. Cohort effects on annual earnings decline notably for the 2007-2012 entry cohorts, while only the 2009 entry cohort sees a decline in hourly wages (net of transitory time effects). I find stronger evidence of excess sensitivity and scarring effects on hourly wages than on annual earnings, though for each outcome allowing for scarring dramatically improves the estimated cohort effects for post-recession entrants.

The lower panels of Table 1 show the association between each set of cohort effects and the entry unemployment rate. Earnings but not wages covary with unemployment at entry, though the relationship is not large and is only marginally statistically significant. Cyclicalities of earnings is notably reduced when I allow for either early-career scarring or (especially) for excess

⁸ The downturn in cohort effects appears to start before the 2006 cohort, perhaps an indication that conditions throughout the early 20s, not just those at age 22, can have scarring effects. My calculation of the recession effect as the change in trend in 2006 is thus conservative, particularly in specification (3) where the decline in cohort effects is fairly smooth.

sensitivity. Earnings of the post Great Recession cohorts, when employed, are notably better than would be predicted based on past cyclical patterns.

E. Non-college workers

Appendix Table A-1 shows parallel analyses for non-college workers, relating cohort effects on employment, wages, and earnings to the unemployment rate when they were 18 years old. There is little indication of cyclicity of cohort effects on any outcome. There is modest evidence of medium-term scarring on non-graduates' wages and employment, but little for earnings (Appendix Figures A-9, A-10, and A-11). Cohorts that enter the market during recessions have lower employment and wages for the first several years of their careers than do those entering at low-unemployment times (consistent with Schwandt and von Wachter, 2019), but in the long run their wages are if anything higher.

VI. Measuring the shadow cast by the Great Recession

The results in Table 1 and Figure 2 indicate that recessions, and the Great Recession in particular, have hangover effects on college graduates' employment rates. About half of the cyclicity is accounted for by medium-term scarring, but a portion remains even after this is accounted for. This portion, because it is persistent, is quantitatively extremely important. To illustrate this, I use the estimated year, medium-term scarring, and cohort effects from specification (3) to forecast the cumulative short-, medium-, and long-term effects of the Great Recession on college graduates' employment. This is conservative – any of the other specifications would assign larger roles to declines in cohort effects, and thus indicate larger long-term effects.

I simulate a stylized version of the labor market for college graduates from 2000 through 2060. I assume that all cohorts are the same size, and that all workers enter the market at 22 and retire at 65. To forecast the future labor market, I assume that year effects β_t revert to their 2005-2007 average from 2015 through 2060, that the unemployment rate is steady at 4.0 percent throughout that period, and that cohort effects δ_c for post-2014 entrants revert to the average for the 2005-2007 entrants. To measure the effect of the Great Recession, I compare to a counterfactual simulation that shuts off the Great Recession by imposing this stability from 2008 forward rather than just from 2015. The assumed year and cohort patterns in the two simulations are shown in Appendix Figure A-12.

The first panel of Figure 4 shows average employment rates of college graduates, age 22-65, under the two scenarios. They coincide through 2007, declining in the 2001 recession and then recovering. In the observed series, employment plummets in the Great Recession, then recovers about two-thirds of the loss. Beginning in 2015, it begins slowly trending downward. This reflects my assumption that future cohorts will resemble the ones that entered in 2005-2007. These cohorts have lower employment rates than earlier cohorts (Figure 2), so as the older cohorts gradually retire the labor force average trends downward.

The second series shows the counterfactual no-recession scenario. We see the same gradual decline from 2015 onward as in scenario 1, but from a higher level. The two series move approximately in parallel from 2015 through the mid-2050s, when the cohorts that entered following the Great Recession will retire, and then finally converge. Integrating the difference between them, the cumulative effect of the recession is 15.5 percentage point-years, of which

only 7.9 percentage points occurs between 2007 and 2014. The remaining 7.6 percentage points represents the medium- and long-term shadow of the recession.

There are three components to the recession effect indicated in panel A, reflecting the three time-varying terms in equation (3). Panel B shows the pure year effects β_t . These are the short-term direct effects of the recession, common to all participants in the labor market at the same time but evaporating after it ends. By construction, these are stable after 2014. The integrated gap between the two series between 2008 and 2014 is 6.4 percentage point-years.

Panel C shows the medium-term scarring effects, represented by the coefficients ϕ_j in (3). These are specified to last beyond the initial shock for those cohorts that were exposed to it at labor market entry, but to fade out within ten years. To translate these into effects on the overall labor market at a point in time, I average $\sum_j g_j(a) * UR_{sc} * \phi_j$ over all cohorts in the labor market in that year. These averages are graphed in Panel C. The gap between the two series is small because the affected cohorts are a small share of the labor force at any time, and because the ϕ_j coefficients are not large to begin with. The series nearly converge by 2020, as the ϕ_j coefficients are near zero for those six or more years into their careers. Cumulatively, the medium-term scarring effects reduce employment by 1.8 percentage point-years, over one-quarter as large in total as the short-term effect.

Panel D shows long-run effects operating through the cohort coefficients δ_c , averaged over all cohorts in the labor market in each year. They decline steadily in both scenarios over the entire period, due to the composition effect mentioned above. But there is a gap between the two scenarios over the duration of the careers of the 2008-2014 entry cohorts. The gap is not large in any year (though it is larger than the medium-term effects in Panel C, even at their peak),

but it lasts for decades. It integrates to 7.3 percentage-point-years, almost one-third larger than the short-term effect in Panel B and quadruple the cumulative size of the medium-term scarring effects in Panel C.

DeLong and Summers (2012) define a hysteresis parameter η , representing the permanent decline in future output per unit current cyclical reduction. Using a real discount rate of 2.5%, the medium- and long-term effects from Figure 4, 1.8 and 7.3 percentage point-years, respectively, are equivalent in present value terms to a constant, permanent reduction in employment of 0.14 percentage points. Comparing to the short-term effects, 6.4 percentage point-years, this yields $\eta = 0.14/6.4 = 0.022$.⁹ This is at the lower end of DeLong and Summers' (2012) "plausible range," but is nevertheless quantitatively important: They conclude that with $\eta = 0.025$, fiscal stimulus is fully self-financing through its impact on future tax revenues if the multiplier is at least 1.0 and the real government borrowing rate is less than 3.7%, and substantially self-financing for lower multipliers and/or higher borrowing rates. Moreover, my estimates capture only a single, limited source of hysteresis: They should be added to any persistent effects of the recession on any workers who were already in the labor market when the recession began, such as those documented by Yagan (forthcoming), and to any non-labor-side sources of hysteresis. The implication is that avoided long-run consequences of downturns can make a substantial contribution toward the present-value cost of countercyclical policy.

⁹ With higher discount rates of 5% or 10%, the permanent equivalents are 0.19 and 0.21 percentage point reductions per year, and η rises to 0.029 or 0.033, respectively.

VII. Conclusion

The recent U.S. economy, having finally emerged from a decade-long slump, exhibits both historically low unemployment and an unimpressive employment-population ratio, particularly for young college graduates. If the latter is a consequence of the 2007-2009 recession, this has important implications for the economic cost of cyclical volatility.

My analysis of college graduates' employment rates indicates that they are sensitive to the economic conditions that prevailed at the time that the graduates entered the labor market. Cohorts that graduate during recessions have lower employment rates even after the recession is past.

Past research has emphasized persistent effects that fade out over the first ten years after graduation. However, the experience of the cohorts that graduated college during and after the Great Recession has not been consistent with that: Their employment rates have recovered much more slowly than would be predicted based on these so-called scarring effects alone, given the low overall unemployment since 2015. My analysis highlights that cohort effects are a second channel, previously overlooked, by which early career weakness can have long-term effects. This channel, because it is so persistent, is by far the more quantitatively important.

If anything, the Great Recession cohorts are doing better than would be predicted based on the long-run relationship between entry conditions and cohort effects. However, even this better-than-expected performance implies an enormous hangover from the recession that will persist for decades to come. It is concentrated on the employment margin, with some effect as well on log annual earnings, and among college graduates. Although it will be many years before we can measure this effect conclusively, the evidence to date points to long-term effects of the


recession on college graduates' employment that are notably larger than the short-term effect. This type of hysteresis is an important component of the impact of cyclical variation in economic conditions. The imperative to minimize it constitutes a strong argument for counter-cyclical policies aimed at minimizing the initial shock.

References

- Beaudry, Paul, and John DiNardo (1991). "The effect of implicit contracts on the movement of wages over the business cycle: Evidence from micro data." *Journal of Political Economy* 99.4: 665-688.
- 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.
- Clark, Kim B., and Lawrence H. Summers (1982). "Labour force participation: timing and persistence." *The Review of Economic Studies*, 49(5): 825-844.
- 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 Till M. von Wachter (2012). "Recessions and the Cost of Job Loss." *Brookings Papers on Economic Activity* 2012.
- DeLong, J. Bradford, and Lawrence H. Summers (2012). "Fiscal policy in a depressed economy" *Brookings Papers on Economic Activity* 2012: 233-297.
- Gronau, Reuben (1974). "Wage comparisons--A selectivity bias." *Journal of Political Economy* 82.6: 1119-1143.
- Jacobson, Louis S., Robert J. LaLonde, and Daniel G. Sullivan (1993). "Earnings losses of displaced workers." *The American Economic Review* 83.4:685-709.
- Kahn, Lisa B. (2010) "The long-term labor market consequences of graduating from college in a bad economy." *Labour Economics* 17.2: 303-316.
- 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.
- Oyer, Paul (2006). "Initial labor market conditions and long-term outcomes for economists." *Journal of Economic Perspectives* 20.3: 143-160.
- Oyer, Paul (2008). "The making of an investment banker: Stock market shocks, career choices, and lifetime income." *The Journal of Finance*. 63.6:2601-2628.


- Schmieder, Johannes F., Till von Wachter, and Joerg Heining (2018). "The costs of job displacement over the business cycle and its sources: Evidence from Germany." Working paper, March 2018.
- 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.
- von Wachter, Till, & Stefan Bender (2006). "In the right place at the wrong time: The role of firms and luck in young workers' careers." *American Economic Review*, 96(5), 1679-1705.
- Yagan, Danny (forthcoming). "Employment hysteresis from the Great Recession." *Journal of Political Economy*, forthcoming.

Figure 1. Employment rates, by age and education, 1979-2018




Notes: Prior week employment rates, computed from monthly Current Population Survey data, are seasonally adjusted and smoothed using a 3-month triangle smoother.

Figure 2. Cohort effects on employment rates of college graduates aged 22-40



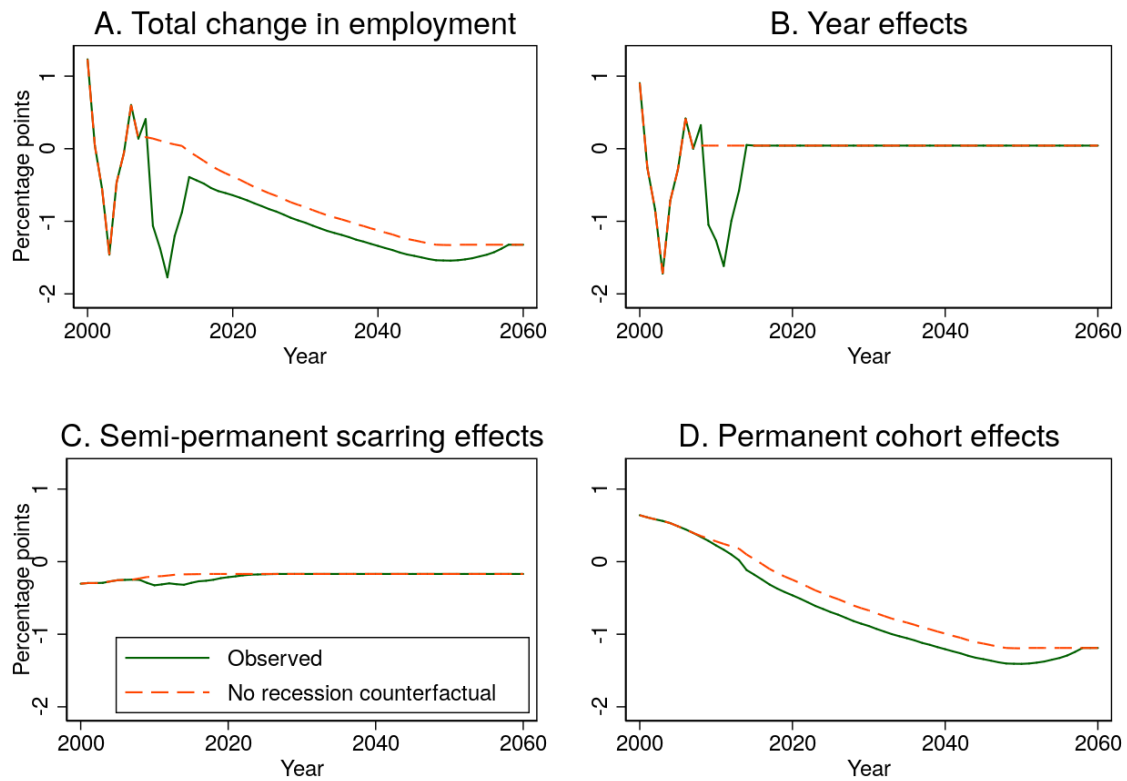
Notes: Prior-week employment rates are estimated from the monthly Current Population Survey and aggregated to the education-cohort-age-year-state cell. Specifications correspond to equations (1), (2), and (3), respectively, and include age, year, and state fixed effects, plus an inverse Mills ratio in the cohort-by-year college attainment rate. Cohorts are indexed by the year they turned 22; cohort effects are normalized to zero for the 1991 and 2000 entry cohorts. Regressions are estimated on a sample of college graduates age 22-40 from the 1970-2000 entry cohorts observed in 1978 to 2018. Cohort effects for post-2000 entrants are estimated assuming that age, time, and other coefficients are as estimated from the earlier cohorts.

Figure 3. Excess sensitivity and scarring coefficients from augmented age-time-cohort decompositions of college graduate employment rates



Notes: See notes to Figure 2. Panel A shows coefficients and confidence intervals for θ_j coefficients, reflecting current unemployment rate-age interactions, from equation (2); panel B shows φ_j coefficients, reflecting entry unemployment rate-age interactions, from equation (3). Connected series and confidence interval correspond to estimates based on 1970-2000 entry cohorts, while dots correspond to estimates using all entry cohorts through 2017.

Figure 4. Decomposition of the long-run effect of the Great Recession on the employment rate of college graduates



Notes: See text for details. Estimates are based on the specification plotted in Figures 2 (“scarring” series), 3B, and Appendix Figures 6 and 7.

Table 1. Time series regressions of estimated cohort effects on the national unemployment rate in the year the cohort entered the labor market

	(1)	(2)	(3)	(4)
Panel A: Employment				
Unemployment rate at age 22	-0.22 (0.05)	-0.17 (0.07)	-0.11 (0.07)	-0.17 (0.06)
Cohort effect model controls for				
Excess sensitivity		Y		Y
Semi-permanent scarring			Y	Y
Change between 2006 and 2010 entry cohorts, vs. 2002 to 2006 (p.p.)				
Predicted	-1.4	-1.1	-0.8	-1.2
Actual	-1.1	-1.2	-0.3	-1.2
Panel B: Mean log annual earnings				
Unemployment rate at age 22 (/100)	-0.74 (0.35)	-0.13 (0.23)	-0.42 (0.23)	-0.56 (0.24)
Change between 2006 and 2010 entry cohorts, vs. 2002 to 2006 (*100)				
Predicted	-4.5	-0.8	-2.6	-3.5
Actual	-1.3	3.9	0.1	0.5
Panel C: Mean log hourly wages				
Unemployment rate at age 22 (/100)	-0.04 (0.24)	0.02 (0.22)	0.26 (0.20)	0.25 (0.20)
Change between 2006 and 2010 entry cohorts, vs. 2002 to 2006 (*100)				
Predicted	-0.4	-0.1	1.4	1.3
Actual	-1.3	-1.3	1.2	1.1

Notes: Each entry corresponds to a separate regression of estimated cohort effects for college graduates on the unemployment rate in the year the cohort was 22, with a quadratic calendar year control and Newey-West standard errors allowing for autocorrelations at lags up to three years. Cohort effects in columns 1-3 are estimated using equations (1), (2), and (3), respectively. In column 4, the specification combines (2) and (3), with controls for both UR_{st} -age and UR_{sc} -age interactions. Panels B and C show estimates for cohort effects on log hourly wages among those employed, from the CPS Outgoing Rotation Groups, and on log annual earnings among those with positive annual employment, from the March CPS. The bottom rows of each panel show the difference between the fitted or actual value for the 2010 and 2006 entry cohorts, less the difference between the 2006 and 2002 cohorts.

Appendices

Appendix A. Data processing and definitions

The primary analysis concerns the cohort-age-year-state employment rate for college graduates. This is estimated from the monthly Current Population Survey (CPS). I construct a synthetic panel from data from each month from January 1979 through November 2018, aggregating to each calendar year.

I divide the population into college graduates and non-graduates. Before 1992, graduates are those with 16 or more years of completed schooling; afterward, it is those with Bachelors or graduate degrees. 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.

My full sample for age-time-cohort decompositions (equations 1-3) consists of people born 1948-1995 (who turned 22 between 1970 and 2017) and observed at ages 22-40. As discussed in the text, my primary estimates identify the age, time, and auxiliary control coefficients only from pre-Recession cohorts. For these specifications, I exclude post-1978 birth cohorts, who turned 22 after 2000. I use the resulting coefficients to form predictions for the subsequent cohorts, and estimate their cohort effects as the difference between observed outcomes and the predictions.

When I examine non-high-school graduates, the sample extends from age 18-40, and begins with the 1952 birth cohort. The full sample extends to the 1999 birth cohort, while the pre-Recession sample again excludes those born after 1978.

In specification (2), I control for the contemporaneous unemployment rate. This is the state average unemployment rate for the calendar year. Specification (3) controls for the unemployment rate when the cohort was 22 years old (18 for non-college-graduates). I use the state where the respondent currently lives for this. State unemployment rates are not available prior to 1976; for the 1948-1954 (or, for non-graduates, 1952-1958) birth cohorts, I use the national unemployment rate.

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

I also present results for log annual earnings and log hourly wages. For annual earnings, I use the 1979-2018 Annual Social and Economic Supplement (ASEC), also known as the March CPS. Each ASEC survey includes information about earnings during the prior calendar year. I measure education and age as of the survey date, but use the prior calendar year's unemployment rate. Earnings are topcoded at the 98th percentile, separately by year; converted to real 2015 dollars using the CPI; and logged. Hourly wages are computed from 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 part time status. Hourly wages are again converted to real 2015 dollars, set to missing if below \$1 or above \$200, and logged.

Appendix B. Additional specifications

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

Educational attainment

I rely on educational attainment as of the CPS survey date. Some people interviewed at age 22 who do not have college degrees at that time will obtain them later. This means that the composition of my college graduate sample changes across age within cohorts. Figure A-3 shows mean attainment by cohort and year. I group entry cohorts into two-year groups. For example, the “1980” series includes those who were 22 in 1980 and 1981. I then show the average attainment of each cohort group in each year until the group is 29/30. Dots indicate the observations when the cohort is 23/24 and when it is 29/30.

For the 1980-1989 entry cohorts, educational attainment was largely stable. Beginning with the 1990 entry cohort, age-30 attainment rose sharply, by about 5 percentage points, with nearly all of this growth occurring among degrees awarded after age 24. Attainment grew slowly from the 1990 to the 2010 entry cohorts, with growth mostly among degrees awarded before 24; between 2002 and 2010, there was even growth in pre-22 awards. Post-2010 entrants are not observed through age 30. There is some indication that post-24 degree awards have risen for these cohorts, but we will not know for sure until their education is complete.

Main decompositions

Figures A-4 and A-5 show estimated age and time effects, respectively, from the main decomposition (1) for employment of college graduates. These are estimated on the pre-Great-Recession cohorts, turning 22 from 1970 through 2000. For comparison, Figure A-4 shows estimated age effects from a specification that includes calendar year but not cohort controls. Age effects are normalized to zero at age 32. The comparison series in Figure A-5 is the unadjusted average employment rate of 22-40-year-old college graduates in each year, without adjustment for cohort differences. Here, year effects are normalized to zero in 2007. In each figure, the cohort effects in the full decomposition are normalized to have zero slope between the 1991 and 2000 entry cohorts.

Figure A-6 shows two sets of estimates of cohort effects from equations (1) and (3). One set, in dashed lines, uses the pre-recession cohorts to identify age, time, and scarring coefficients, then extrapolates these coefficients to subsequent cohorts. The other uses the full sample to estimate the regressions. Results are generally similar.

Figures A-7 and A-8 repeat the decompositions in Figures 2 and 3 of the main text, first for log annual earnings and then for log hourly wages. Samples here exclude those with zero work in the prior year (A-7) or prior week (A-8).

Figures A-9 through A-11 present decompositions for non-college graduates. These are similar to college graduates, except that cohorts enter the sample at age 18, the age-18 unemployment rate is the relevant entry condition, the $g_j(a)$ coefficients refer to years since age 18, and the selection control is an inverse Mills ratio in the share of the cohort that *did not* graduate from college, $\lambda(1 - p_{atc})$. Estimates of the cyclicity of the cohort effect estimates, with respect to the age-18 unemployment rate, are in Table A-1.

Gender differences

Male and female labor force participation patterns are quite different, and it is plausible that recessions could have distinct effects (both permanent and semi-transitory) on them. To explore this, I use employment rates at the age-time-cohort-state-gender level, and augment my basic specifications with gender interactions. Equation (3) becomes:

$$Y_{satcg} = \alpha + \xi_g + \beta_t + \gamma_{ag} + \delta_{cg} + \zeta_s + \rho\lambda(p_{atc}) + \sum_{j=1}^5 g_j(a) * UR22_{sc} * \varphi_{jg} + (UR22_{sc} - UR22_c)\kappa_g + \epsilon_{satcg}. \quad (3')$$

Here, ξ_g represents a gender-specific intercept. The age, cohort, and semi-transitory scarring coefficients are also allowed to vary freely with gender, but time, state, and selection coefficients are held constant across genders.

The non-identification problem is somewhat different here than in the specification used in the main text. As before, I normalize cohort effects to zero for the 1991 and 2000 male entry cohorts. However, I need only one normalization of the female cohort effects: I set the 1991 entry cohort effect to zero, but allow the effects to vary freely before and afterward. This enables me to permit long-run trends in female relative labor force participation, necessary to capture the entry of women into the labor market over time.

Figure A-12 shows the φ_{jg} coefficients from this specification. Semi-transitory scarring is much more severe for men than for women, for whom it is only marginally significant. Figure A-13 shows the cohort effects, both for this specification and for a simpler one that excludes the φ_{jg} and κ_g effects. We see the expected long-run rise in female labor force participation here, though the Great Recession dropoff, which is larger for women than for men, erases essentially all of this progress. For both men and women, unadjusted cohort effects (from the simpler specification without semi-transitory scarring) are about 4 percentage points lower for post-recession entrants than for the 1991 entrants who are the reference group. (Alternative normalizations would change the magnitude of the drop, but would not alter the similarity of the male and female drops.) Allowing for semi-transitory scarring reduces the male decline substantially for the immediate post-recession cohorts, explaining about 2/3 of the falloff relative to the 1991-2000 trend. This has much less of an effect on the recession dropoff for women (though it does have a meaningful impact on the female long-run trend). Overall, there

do appear to be gender differences, but the basic story told by the pooled specifications presented in the main text is not overturned here.

Simulation of the effect of the Great Recession

Figure A-14 presents the assumed time and cohort effects used in the simulation of the long-term effects of the Great Recession, Figure 4. In the series labeled “actual,” I use estimates from specification (3) through 2014. After that year, I assume year effects equal the average of the 2005-7 estimates, and cohort effects equal the average cohort effect estimates for 2005-7 entrants. In the “no recession counterfactual” series, I use these averages for every year and cohort from 2008 onward.

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).¹ 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 birth cohort c would have been exposed. Let $t=c+22$. This average rate is then:

$$\widehat{UR22}_{bc} = \frac{\sum_s UR22_{st} N_{bs,22}}{\sum_s N_{bs,22}}.$$

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 :

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

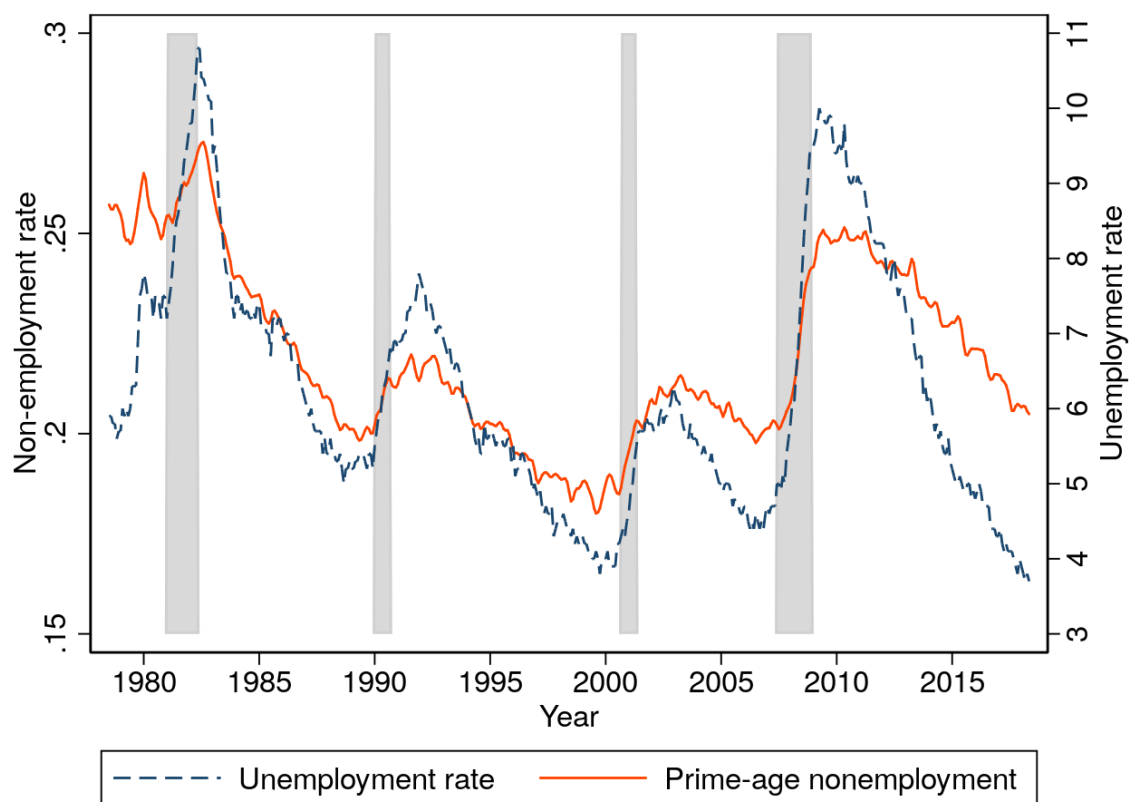
¹ Schwandt and von Wachter (2019) use a double-weighting estimator to abstract from both endogenous mobility and endogenous attainment rates. I treat attainment rates as exogenous, using the selection correction discussed in the text to address changes in cohort composition with age.

Last, I interact this measure with age indicators $g_j(a)$ ($j=1, \dots, 5$) and use the results as instruments for the $g_j(a) * UR22_{sc}$ interactions in equation (3). Appendix Table A-2 shows the estimated scarring coefficients from the base OLS specification and using IV, both using the full sample and limiting to the pre-recession entrants. Coefficients are similar, and if anything smaller in the IV specifications. In other words, accounting for endogenous mobility attributes *more* of the post-Great-Recession decline in young people's outcomes to permanent cohort effects than in my main results.

References


- 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.
- 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.
- 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.

Appendix Figure A-1. Unemployment rate and prime-age non-employment rate




Notes: Unemployment rate is the published, seasonally adjusted rate. Non-employment rate is calculated over 25-54 year olds from the CPS microdata samples, then seasonally adjusted and smoothed using a three-period triangle smoother.

Appendix Figure A-2. Mean log real hourly wages, by age and education, 1979-2018




Notes: Hourly wages are censored at real values of \$1 and \$200 per hour (in 2015 dollars). Monthly values are seasonally adjusted and then smoothed with a 5-month triangle smoother.

Appendix Figure A-3. Share of cohort with a bachelor's degree, by cohort and year




Notes: Each series 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. Triangles and circles indicate observations when respondents were 24 and 30 years old, respectively.

Appendix Figure A-4. Age effects on employment of college graduates




Notes: “Employment rate” series represents estimated cohort-adjusted age effects on employment from a restricted version of equation (1) with no calendar year effects, estimated on college graduates aged 22-40 from the 1970-2000 entry cohorts observed 1979-2018. “Age effects” series is drawn from equation (1), estimated on the same sample and including calendar year adjustments. Age effects are normalized to zero at age 32 (indicated by a dot); the full age-time-cohort decomposition further normalizes the cohort effects to be the same for the 1991 and 2000 entry cohorts.

Appendix Figure A-5. Time effects on employment of college graduates



Notes: “Employment rate” series is the unadjusted average employment rate of 22-40-year-old college graduates in each year, relative to 2007. “Year effects” series is drawn from equation (1), estimated on college graduates aged 22-40 from the 1970-2000 entry cohorts observed 1979-2018 and including calendar year and age adjustments. Year effects are normalized to zero in 2007; the full age-time-cohort decomposition further normalizes the cohort effects to be the same for the 1991 and 2000 entry cohorts.


Appendix Figure A-6. Cohort effects on employment: Extrapolated from pre-recession cohorts and estimated using the full sample




Notes: Figure presents estimated cohort effects from equations (1) and (3). Solid lines show effects estimated using the full sample. Dashed lines show series obtained by estimating the age-time-cohort decomposition using only cohorts aged 22 by 2000, then estimating post-2000 cohort effects as mean residuals by cohort after subtracting time and age effects estimated from the earlier data. In each case, cohort effects are normalized to zero for the 1991 and 2000 entry cohorts.

Appendix Figure A-7. Cohort effects, excess sensitivity, and scarring for college graduates' log annual earnings if employed

A. Cohort effects on log annual earnings of college graduates




B. Excess sensitivity and scarring effects on log annual earnings




Notes: Annual earnings are computed from March CPS data, for those with positive earnings, then logged and averaged to the education-cohort-age-year-state cell. Specifications correspond to equations (1), (2), and (3), respectively, and include age, year, and state fixed effects, plus an inverse Mills ratio in the cohort-by-year college attainment rate. Cohorts are indexed by the year they turned 22; cohort effects are normalized to zero for the 1991 and 2000 entry cohorts. Decompositions are estimated on college graduates aged 22-40 from the 1970-2000 entry cohorts observed 1979-2018, then extrapolated to the post-2000 entry cohorts holding all but cohort effects fixed.

Appendix Figure A-8. Cohort effects, excess sensitivity, and scarring for college graduates' log hourly wages if employed

A. Cohort effects on log hourly wages of college graduates




B. Excess sensitivity and scarring effects on log hourly wages




Notes: Hourly wages are computed from the Merged Outgoing Rotation Groups, for those with positive prior-year earnings, then logged and averaged to the education-cohort-age-year-state cell. Specifications correspond to equations (1), (2), and (3), respectively, and include age, year, and state fixed effects, plus an inverse Mills ratio in the cohort-by-year college attainment rate. Cohorts are indexed by the year they turned 22; cohort effects are normalized to zero for the 1991 and 2000 entry cohorts. Decompositions are estimated on college graduates aged 22-40 from the 1970-2000 entry cohorts observed 1979-2018, then extrapolated to the post-2000 entry cohorts holding all but cohort effects fixed.

Appendix Figure A-9. Cohort effects, excess sensitivity, and scarring for non-graduates' employment

A. Cohort effects on employment rates of non-college graduates




B. Excess sensitivity and scarring effects on employment of non-college graduates




Notes: Prior-week employment rates are estimated from the monthly Current Population Survey and aggregated to the education-cohort-age-year-state cell. Specifications correspond to equations (1), (2), and (3), respectively, and include age, year, and state fixed effects, plus an inverse Mills ratio in the cohort-by-year college non-attainment rate. Cohorts are indexed by the year they turned 18; cohort effects are normalized to zero for the 1991 and 2000 entry cohorts. Regressions are estimated on a sample of non-college graduates age 18-40 from the 1970-1996 entry cohorts observed in 1978 to 2018. Cohort effects for post-1996 entrants are estimated assuming that age, time, and other coefficients are as estimated from the earlier cohorts.

Appendix Figure A-10. Cohort effects, excess sensitivity, and scarring for non-graduates' log annual earnings if employed

A. Cohort effects on log annual earnings of non-college graduates




B. Excess sensitivity and scarring effects on log annual earnings of non-college graduates




Notes: Annual earnings are computed from March CPS data, for those with positive earnings, then logged and averaged to the education-cohort-age-year-state cell. Specifications correspond to equations (1), (2), and (3), respectively, and include age, year, and state fixed effects, plus an inverse Mills ratio in the cohort-by-year college non-attainment rate. Cohorts are indexed by the year they turned 18; cohort effects are normalized to zero for the 1991 and 2000 entry cohorts. Regressions are estimated on a sample of non-college graduates age 18-40 from the 1970-1996 entry cohorts observed in 1978 to 2018. Cohort effects for post-1996 entrants are estimated assuming that age, time, and other coefficients are as estimated from the earlier cohorts.

Appendix Figure A-11. Cohort effects, excess sensitivity, and scarring for non-graduates' log hourly wages if employed

A. Cohort effects on log hourly wages of non-college graduates




B. Excess sensitivity and scarring effects on log hourly wages of non-college graduates



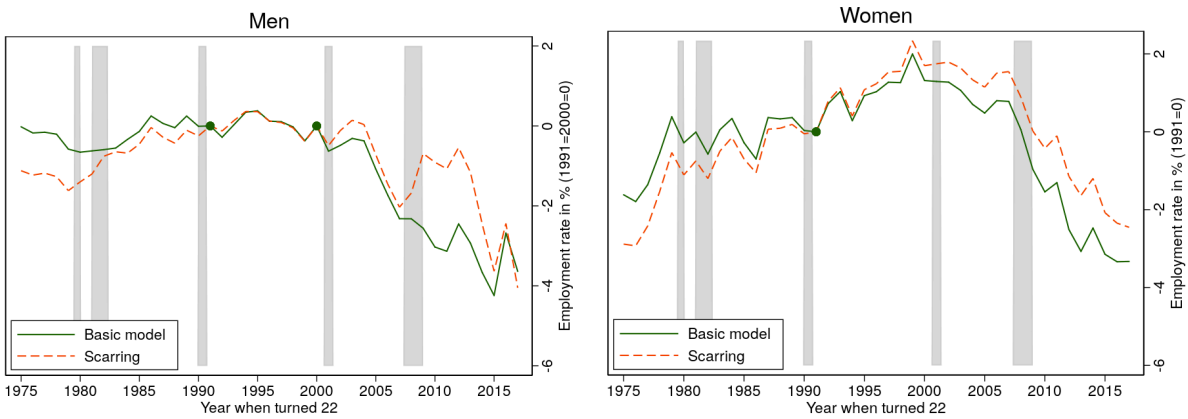
Notes: Hourly wages are computed from the Merged Outgoing Rotation Groups, for those with positive prior-year earnings, then logged and averaged to the education-cohort-age-year-state cell. Specifications correspond to equations (1), (2), and (3), respectively, and include age, year, and state fixed effects, plus an inverse Mills ratio in the cohort-by-year college non-attainment rate. Cohorts are indexed by the year they turned 18; cohort effects are normalized to zero for the 1991 and 2000 entry cohorts. Regressions are estimated on a sample of non-college graduates age 18-40 from the 1970-1996 entry cohorts observed in 1978 to 2018. Cohort effects for post-1996 entrants are estimated assuming that age, time, and other coefficients are as estimated from the earlier cohorts.

Appendix Figure A-12. Scarring effects on employment rates of college graduates, by gender




Notes: See text for description of specification.

Appendix Figure A-13. Cohort effects on employment rates of college graduates, by gender



Notes: See text for description of specification.

Appendix Figure A-14. Assumed time and cohort effects for simulation of the long-term effects of the Great Recession



Notes: “Actual” series is estimated year (panel A) and cohort (panel B) effects from the “scarring” specification plotted in Figure 2 (equation 3), through 2014, and measured in percentage points. The series then revert to their 2005-2007 averages. The “no recession counterfactual” series equals the actual series through 2007, then equals the 2005-2007 average thereafter.

Appendix Table A-1. Time series regressions of estimated cohort effects for non-graduates on the national unemployment rate in the year the cohort was 18

	(1)	(2)	(3)	(4)
Panel A: Employment				
Unemployment rate at age 18	-0.03 (0.04)	-0.05 (0.04)	0.11 (0.07)	-0.09 (0.05)
Cohort effect model controls for				
Excess sensitivity		Y		Y
Semi-permanent scarring			Y	Y
Change between 2006 and 2010 entry cohorts, vs. 2002 to 2006 (p.p.)				
Predicted	-0.5	-0.6	0.4	-0.9
Actual	-1.1	-1.6	-0.2	-2.1
Panel B: Mean log annual earnings				
Unemployment rate at age 22 (/100)	0.17 (0.73)	0.76 (0.50)	0.06 (0.39)	-0.34 (0.46)
Change between 2006 and 2010 entry cohorts, vs. 2002 to 2006 (*100)				
Predicted	0.2	3.8	-0.2	-2.7
Actual	-2.9	1.4	2.6	-2.8
Panel C: Mean log hourly wages				
Unemployment rate at age 22 (/100)	-0.06 (0.25)	0.05 (0.21)	0.22 (0.16)	0.10 (0.18)
Change between 2006 and 2010 entry cohorts, vs. 2002 to 2006 (*100)				
Predicted	-0.2	0.5	1.6	0.9
Actual	0.9	1.1	5.7	4.8

Notes: Each entry corresponds to a separate regression of estimated cohort effects for non-graduates on the unemployment rate in the year the cohort was 18, with a quadratic calendar year control and Newey-West standard errors allowing for autocorrelations at lags up to three years. Cohort effects in columns 1-3 are estimated using equations (1), (2), and (3), respectively. In column 4, the specification combines (2) and (3), with controls for both UR_{st} -age and UR_{sc} -age interactions. Lower rows show estimates for cohort effects on log hourly wages among those employed, from the CPS Outgoing Rotation Groups, and on log annual earnings among those with positive annual employment, from the March CPS. The bottom rows of each panel show the difference between the fitted or actual value for the 2010 and 2006 entry cohorts, less the difference between the 2006 and 2002 cohorts.

Appendix Table A-2. Instrumenting for the age-22 unemployment rate

	Pre-Great Recession cohorts		All cohorts	
	OLS	IV	OLS	IV
	(1)	(2)	(3)	(4)
Coefficient on interaction of age-22 unemployment rate with indicator for age:				
22-23	-0.46 (0.15)	-0.21 (0.15)	-0.67 (0.15)	-0.49 (0.15)
24-25	-0.35 (0.10)	-0.17 (0.10)	-0.47 (0.09)	-0.33 (0.09)
26-27	-0.13 (0.09)	-0.04 (0.08)	-0.20 (0.08)	-0.13 (0.08)
28-29	0.02 (0.09)	0.02 (0.09)	0.05 (0.06)	0.07 (0.06)
30-31	0.01 (0.09)	-0.01 (0.08)	-0.05 (0.07)	-0.07 (0.07)

Notes: Columns 1 and 3 correspond to the scarring estimates from Figure A-6. Columns 2 and 4 repeat the same specifications, instrumenting for the age-22 unemployment rate-age interactions with a simulated unemployment rate (again interacted with age) that abstracts from cohort-specific mobility. See text for details.