Did Timing Matter? Life Cycle Differences in Effects of Exposure to the Great Recession

Kevin Rinz, US Census Bureau

I estimate the effects of exposure to the Great Recession on employment and earnings for groups defined by year of birth over the 10 years following the beginning of the recession. Younger workers experience the largest earnings losses in percentage terms (up to 13%), in part because they remain less likely to work for high-paying employers even as their overall employment recovers more quickly than that of older workers.

I. Introduction

The Great Recession caused enormous disruption in labor markets. Nationally, the hiring rate fell by more than 30%, the number of job openings

Any opinions and conclusions expressed herein are those of the author and do not necessarily reflect the views of the US Census Bureau. All results have been reviewed to ensure that no confidential information is disclosed. The statistical summaries reported in this paper have been cleared by the Census Bureau's Disclosure Review Board (release authorizations CBDRB-FY19-376, CBDRB-FY19-431, CBDRB-FY19-500, CBDRB-FY2019-CES005-005, CBDRB-FY2019-CES005-013, CBDRB-FY2020-CES005-016, CBDRB-FY20-343, CBDRB-FY2020-CES005-046, and CBDRB-FY2021-CES005-010). I thank Leah Clark, John Coglianese, Arin Dube, Harris Eppsteiner, Martha Gimbel, Bryan Stuart, Abigail Wozniak, and participants in the Washington Area Labor Economics Symposium and the Southern Economic Association 2020 conference for helpful comments and conversations, as well as Danny Yagan for sharing the crosswalk used in his analysis. Contact the author at kevin.rinz@census.gov. Information concerning access to the data used in this paper is available as supplemental material online.

Submitted September 1, 2020; Accepted July 8, 2021; Electronically published May 17, 2022.

Journal of Labor Economics, volume 40, number 3, July 2022.

This article is in the public domain. Published by The University of Chicago Press in association with The Society of Labor Economists and The National Opinion Research Center. https://doi.org/10.1086/716346

fell by more than half, and the unemployment rate more than doubled, peaking at 10%. Some groups saw their unemployment rates rise even higher; for workers under 25 years old, it reached 19.5%, and it remains substantially higher than the overall unemployment rate. Locally, more than 90% of commuting zones (CZs) lost establishments on net. The median CZ had 3.8% fewer establishments in 2009 than in 2007, and the hardest-hit quartile lost 5.4%-15.7% of their establishments. Over the same period, all but five CZs saw their unemployment rates increase.¹ Nearly 7.8 million people were laid off in the first quarter of 2009 alone, about 40% more than in the average quarter in 2006. Job loss can have persistent and even permanent negative consequences (Topel 1990; Ruhm 1991; Jacobson, LaLonde, and Sullivan 1993; Neal 1995; Sullivan and von Wachter 2009; Davis and von Wachter 2011; Hershbein 2012; Jarosch 2015; Yagan 2019; Stuart 2019), and recessionrelated changes in hiring and employment dynamics may disadvantage younger workers (Forsythe 2016, 2019). Analyses of the consequences of job loss often focus on prime-age workers with relatively strong labor force attachment and occasionally restrict attention to men, but given changes in the prevalence of work, education, and other activities over the life cycle, the nature of the harm from recession exposure may depend on when in a person's life they are exposed to it. Younger workers considering pursuing additional education might see the opportunity cost of doing so fall (potentially increasing their likelihood of enrolling) and/or liquidity constraints associated with doing so bind more tightly (potentially reducing their likelihood of enrolling). Older workers might be forced to adjust retirement plans or return to the labor force to make ends meet. People not losing jobs themselves can also suffer from environments in which job loss is prevalent if conditions make it harder for them to ask for a raise or change jobs to get one (Beaudry and DiNardo 1991; Topel and Ward 1992). Some prior work has considered the consequences of graduating from college or otherwise entering the labor market during a recession (Kahn 2010; Oreopoulos, von Wachter, and Heisz 2012; Schwandt and von Wachter 2019; Rothstein 2019), but I do not know of any previous work that considers effects of recession exposure on workers of substantially different ages within the same framework and looks for heterogeneity in the effects of exposure across ages.

In this paper, I estimate the effects of exposure to the Great Recession on employment and earnings outcomes for groups of workers who were at different ages (and stages of their working lives) when the recession started in 2007. Fortuitously, the age/life cycle groups I consider correspond to popularly used generational definitions: millennials (born 1981–96) were still in school or entering the workforce at the onset of the recession; members of Generation X (Gen-X, born 1965–80) had largely entered the workforce

¹ The median increase was 3.6 percentage points, and the hardest hit quartile of CZs saw increases of 5.2–11.8 percentage points.

Did Timing Matter?

but were still relatively early in their careers; baby boomers (born 1946–64) were later in their working lives but had generally not yet reached retirement; and the silent generation (born 1928–45) had largely reached retirement age when the recession began. Although generational terminology may call to mind a cohort-based frame for my estimates, I believe the more natural interpretation in this context is that potential heterogeneity in effects across groups arises from differences in age at recession exposure are more likely to be influenced by contemporaneous life cycle considerations that differ across groups and are related to age at the time of the recession. Figure 1,



FIG. 1.—Employment and earnings trends, by generation. Individuals are employed if they have any W-2 earnings in a given year. Earnings include wage and salary income as well as deferred compensation. Source: Form W-2, Census Numident. Release authorization CBDRB-FY19-376. A color version of this figure is available online.

which presents employment and earnings trajectories for these groups over the period considered here, illustrates the life cycle differences well.

I implement a modified version of the approach used in Yagan (2019), estimating the effect of changes in the local unemployment rate between 2007 and 2009 on employment and earnings in each year through 2017. I estimate these effects on an individual-level panel of 4.1 million individuals constructed using data from Form W-2 and a variety of other administrative and survey sources.

I find that increased local unemployment shocks have persistent negative effects on workers' employment and earnings. Exposure to an additional percentage point of increased local unemployment during the recession reduced the average worker's probability of being employed by abut 0.7 percentage points in 2010. This effect declines in magnitude over subsequent years but remains statistically significant through 2015. With regard to earnings, this same exposure cost workers up to about \$600–\$700 annually between 2009 and 2015 and more than \$400 in 2016 and 2017.

The trajectory of the overall employment effect, especially the recovery after 2010, is driven by younger workers, although their earnings do not rebound similarly.² The effect of recession exposure on millennials' employment was most negative in 2010 and began steadily diminishing shortly thereafter. Their earnings, however, have shown little sign of recovery. For millennials, exposure to the average local unemployment shock during the recession reduced earnings in 2017 by nearly \$3,000 on average; from 2007 through 2017, such exposure reduced earnings by more than \$25,000 in total, or about 13% of actual earnings over that period. A similar dynamic emerges for the next youngest workers in Gen-X, although the recovery in their employment effects is less stark because they were less than half as severe as the effects for millennials in absolute terms.

For millennials in particular, changes in educational attainment and employer characteristics due to recession exposure appear to contribute to this divergence. Increased exposure to local unemployment shocks appears to have reduced their likelihood of completing any postsecondary degree by 1.3 percentage points per unit of exposure. These shocks have also persistently reduced the likelihood of working for a high-paying employer.³ This is in line with results from the literature on the "scarring" effects of conditions at labor market entry that suggest that negative long-term effects of these conditions are

² I often use "recovery" in this paper to refer to the reduction in the magnitude of negative effects on employment and earnings. The underlying outcomes have of course changed over time, increasing for some groups and decreasing for others, but these changes are connected to the life cycle.

³ This is true using a variety of measures. To provide one example, exposure to the average shock reduced millennials' probability of working for an employer in the highest-paying quartile by 2.8 percentage points in 2017. About 38% of millennials worked for such employers that year.

due to workers' first jobs being worse (Kahn 2010; Oreopoulos, von Wachter, and Heisz 2012; Schmeider, von Wachter, and Bender 2014).

The adverse effects of exposure to the Great Recession are evident across groups defined by job displacement, gender, and race and ethnicity. They are only modestly mitigated by geographic mobility. They are robust to the inclusion of controls for other changes in local economies during the Great Recession, such as shifts in the industrial composition of employment and changes in local labor market concentration.

This paper makes several contributions relative to both Yagan (2019) and the broader literature on the effects of exposure to recessions. Relative to Yagan, this paper examines heterogeneity in the effects of exposure to the recession across groups defined by age at the onset of the recession, an interesting and important dynamic to consider given significant differences over the life cycle in the prevalence of other opportunities and responsibilities that figure into decisions about work. Finding more persistent and relatively more negative earnings effects among younger workers, it explores potential channels through which these more adverse effects may have been realized; I find that reductions in employer quality and rates of completing postsecondary education likely contribute. Although the negative effects of entering the labor market during a recession are well known (e.g., in the United States: Altonji, Kahn, and Speer 2016; Speer 2016; Abel and Deitz 2016; internationally: Genda, Kondo, and Ohta 2010; Cockx and Ghirelli 2016; Fernández-Kranz and Rodríguez-Planas 2018; and those mentioned above), this paper's consideration of young workers alongside older workers allows it to contribute to that literature as well. Estimates for younger and older workers from the same context (i.e., exposure to the Great Recession) allow me to compare the magnitude and persistence of these effects directly, clarifying the nature of the harm associated with early-career recession exposure. While the adverse earnings effects I find for young workers are the same order of magnitude as those I find for older workers, the higher degree of persistence in these effects for young workers suggests that exposure earlier in life may be more harmful by virtue of affecting a larger share of their working lives. Finally, it was not clear a priori how the effects of exposure to a very severe recession like the Great Recession around labor market entry might differ from the effects of prior, less severe recessions; my estimates suggest that the earnings effects scale up roughly linearly with recession severity.

The rest of this paper proceeds as follows: section II describes the data used in this analysis, section III lays out the empirical strategy, section IV discusses the results, and section V concludes.

II. Data

I draw a 2% sample, stratified by year of birth, from the 2018 vintage of the Census Numident file, which contains one record for each individual who has received a Social Security number and provides information on gender, citizenship, date of birth, and date of death, where applicable, among other things.⁴ I focus on individuals born between 1928 and 1996. For each person in this sample, I obtain information about earnings in each year from 2005 through 2017 from Form W-2. My primary earnings concept combines wage and salary earnings with deferred compensation.⁵ Dollar values are adjusted for inflation using the consumer price index. I also obtain any available location information from Form 1040 (if they filed themselves, I use their own 1040; if they were listed as a dependent, I use information from the form on which they were listed), any 1099 information returns that were issued, the American Community Survey (ACS), and a variety of other administrative sources that have been collected in the Census Bureau's Master Address File–Auxiliary Reference File (MAF-ARF). I also obtain the type(s) of 1099 forms each individual received as well as information on race and ethnicity from the ACS.

From these files, I construct a panel of earnings and geographic information at the individual level. I use location information to identify county and, by extension, CZ of residence.⁶ If geographic information is missing for up to 3 years between observations of an individual in the same CZ, I impute that observed CZ to those years with missing data. Where no earnings information is available for an individual in a given year, I set earnings to zero if location information is available or can be imputed. In my primary analysis sample, I include only individuals who had location information available for each year from 2005 through 2017 and who had not died as of the end of 2017. This gives me a sample of just over 4.1 million individuals, each observed 13 times.⁷ Table 1 reports summary information for individuals in my primary analysis sample. Individuals are assigned to generations using year of birth and the definitions currently used by the Pew Research Center (Dimock 2019).⁸

⁴ Individuals are identified on this file (and on all others used here) via a Protected Identification Key, an anonymized identifier that is assigned by the Census Bureau's Person Identification Validation System and serves as the linkage variable across data sets. See Wagner and Layne (2014) for more details.

⁵ These are the only two types of income provided in the W-2 data available to me. Deferred compensation should be thought of as retirement contributions. I do not observe employee or employer contributions to employer sponsored health insurance premiums.

⁶ More information on how I identify county of residence can be found in app. A (apps. A–F are available online). I use 1990 CZ definitions.

⁷Constructing the analysis sample using more or less stringent data availability requirements has little qualitative effect on my estimates. Key estimates based on alternative samples are presented in app. A.

⁸ It is fairly unlikely that younger millennials or older members of the silent generation were in the labor force for substantial portions of the period considered here. In the case of the youngest millennials, they may not have been legally permitted to work until near the end of the period covered by my data. I do not restrict my analysis to Did Timing Matter?

III. Estimation

I estimate the effects of exposure to the Great Recession using a strategy that is similar to Yagan (2019). Specifically, I estimate

$$y_{it} = \sum_{t} \beta_{t} Shock_{CZ07(i)} \cdot 1[Year = t] + \sum_{a \in ages} \alpha_{a} \cdot 1[Age_{it} = a] + \gamma_{i} + \delta_{t} + \varepsilon_{it}, \qquad (1)$$

where y_{it} is an earnings or employment outcome, α_a are age fixed effects, γ_i are individual fixed effects, δ_i are year fixed effects, and $Shock_{CZO7(i)} = UR_{CZO7(i),2009} - UR_{CZO7(i),2007}$ is the change in the unemployment rate between 2007 and 2009 in individual *i*'s CZ of residence in 2007 (CZO7(*i*)).⁹ The omitted year interaction with the shock variable is 2006, fixing $\beta_{2006} = 0$ mechanically. Standard errors are clustered on 2007 CZ of residence. I estimate separate regressions of this form for each group I analyze.

With 2006 serving as the reference year and the data beginning in 2005, I have essentially no capacity to investigate prerecession trends in outcomes of interest. Fortunately, prerecession effects presented in figure 4 in Yagan (2019) are not suggestive of trends that would generate spurious negative estimates of the effects of recession exposure. Given that my results are very similar to Yagan's in years covered by both analyses, his figure provides support for the strategy employed here as well. One might also be concerned that Yagan's work based on prime-age workers may not carry over to younger workers, and changes in employment or labor force participation among young workers may be correlated with the intensity of the local unemployment shock during the Great Recession. While I do not have enough data to assess prerecession trends in my exact outcomes of interest for young workers, similar measures can be constructed from public-use Current Population Survey data. As figure A1 (figs. A1-A19, B1-B35, C1-C6, D1-D6, E1-E22, F1-F4 are available online) indicates, employment and earnings outcomes for young workers living in places that were exposed to larger and smaller increases in unemployment during the Great Recession followed similar trends for more than a decade prior to the beginning of the recession. This suggests that more and less intensely exposed

ages likely to have been working within these generations. The requirement that individuals survive through 2017 may implicitly exclude some older nonworkers. The youngest millennials turned 16 in 2012, and I want to capture the effects of recession exposure on their early employment and earnings while constructing my sample consistently for all groups. Young individuals who are not working appear in my data as having no earnings and not being employed, just like their older counterparts.

⁹ Because of differences in data availability, I do not use the same specification as Yagan (2019), but our two approaches are conceptually similar. See app. A

			Generation	Baby	Silent
TT 11	All	Millennials	X	Boomers	Generation
Variable	(1)	(2)	(3)	(4)	(5)
Age in 2005 (years)	37.4	16.4	32.6	49.4	66.4
	(17.3)	(4.6)	(4.7)	(5.4)	(4.8)
Male	48.5	50.4	49.0	48.1	43.9
	(49.9)	(49.9)	(50.0)	(50.0)	(49.6)
Female	51.5	49.6	51.0	51.9	56.1
	(49.9)	(49.9)	(50.0)	(50.0)	(49.6)
White, non-Hispanic	68.0	61.3	64.4	73.1	78.7
	(46.7)	(48.7)	(47.9)	(44.4)	(40.9)
Black, non-Hispanic	11.6	14.4	12.0	10.1	7.9
	(32.1)	(35.0)	(32.5)	(30.2)	(26.9)
Asian, non-Hispanic	4.6	4.2	5.6	4.3	3.8
	(20.9)	(20.0)	(23.0)	(20.3)	(19.2)
Hispanic	11.9	15.7	13.3	9.1	6.9
	(32.3)	(36.3)	(34.0)	(28.8)	(25.4)
Other race, non-Hispanic	1.3	1.7	1.4	1.2	.8
	(11.5)	(12.9)	(11.8)	(10.8)	(8.9)
Employed in 2005	66.1	44.5	84.6	78.6	40.8
	(47.3)	(49.7)	(36.1)	(41.0)	(49.2)
Earnings in 2005 (thousands					
of 2017 \$)	30.6	5.8	38.9	49.3	19.5
	(116.5)	(13.5)	(87.8)	(158.0)	(159.7)
Filed/claimed on 1040 in 2005	91.7	93.9	91.8	91.7	85.9
	(27.6)	(23.8)	(27.4)	(27.7)	(34.8)
Employed in 2017	66.5	86.0	80.7	57.0	11.3
	(47.3)	(34.7)	(39.5)	(49.5)	(31.7)
Earnings in 2017 (thousands					
of 2017 \$)	36.2	32.1	54.2	36.5	3.7
	(150.0)	(126.2)	(139.3)	(192.9)	(49.9)
Filed/claimed on 1040 in 2017	85.2	88.7	87.5	84.7	72.8
	(35.5)	(31.7)	(33.1)	(36.0)	(44.5)
Millennial (born 1981–96)	28.7				
	(45.3)				
Generation X (born 1965–80)	27.0				
	(44.5)				
Baby boomer (born 1946–64)	32.8				
	(46.9)				
Silent generation (born 1928-45)	11.5				
,	(31.9)				
Observations	4,121,000	1,184,000	1,113,000	1,350,000	474,000

Table 1 Individual-Level Summary Statistics, Analysis Sample

SOURCE.—Form 1040, Form W-2, Census Numident, ACS (2001–17), decennial census (2000 and 2010). NOTE.—Standard deviations, reported in parentheses, are reconstructed from rounded standard errors and sample sizes. Each individual in the sample is counted once in this table. Sample sizes are rounded for disclosure avoidance purposes. Release authorization CBDRB-FY19-376. places were not diverging significantly in terms of young workers' outcomes of interest prior to the onset of the recession.

Under the assumption that local unemployment shocks are as good as randomly assigned, conditional on age, person, and year fixed effects, the coefficient β_t gives the average change in earnings or employment, relative to 2006, in year *t* due to a 1 percentage point increase in the local unemployment rate between 2007 and 2009.¹⁰ Information on local unemployment rates comes from the Bureau of Labor Statistics' Local Area Unemployment Statistics program. The unemployment rate in the average worker's CZ increased by about 4.6 percentage points between 2007 and 2009. The standard deviation of this unemployment shock is about 2 percentage points.

My primary outcomes of interest are earnings and employment. An individual is employed in year *t* if her earnings are greater than zero. I primarily consider earnings effects in levels. Table B1 (available online) presents the main estimates discussed below as a percentage of the contemporaneous outcome mean for the relevant group for select postrecession years; figures depicting the full set of estimates relative to mean outcomes, as well as estimates using log and inverse hyperbolic sine transformations, are reported in appendix B.

IV. Results

I begin by reporting estimates based on the 1957–76 birth cohorts, as in Yagan (2019), then extend the analysis to consider additional cohorts and examine heterogeneity by groups of cohorts, gender, race and ethnicity, and other factors.

A. Estimates Using Yagan (2019) Cohorts

For the years and cohorts covered by both analyses, the pattern of estimates produced using equation (1) is very similar to that reported in Yagan (2019). Figure 2 presents estimates of the effects of increased local unemployment during the Great Recession on employment and earnings through 2017 for individuals born between 1957 and 1976 in the left column. I replicate the headline result that a Great Recession unemployment shock that was 1 percentage point larger decreased the employment rate in 2015 by more than 0.3 percentage points. I also find that during much of the postrecession period, that same 1 percentage point larger unemployment shock led earnings to be about \$1,000 lower per year on average.

Workers continue to experience adverse employment and earnings effects of unemployment shocks in the two additional years that are included in my analysis, but these effects are less severe than in earlier years. The

¹⁰ Unless otherwise indicated, all effects of unemployment shocks on employment and earnings outcomes will be discussed below on an annual basis, per percentage point of unemployment exposure.



FIG. 2.—Effects of local unemployment shocks on employment and earnings. Points plotted are β coefficients as estimated in equation (1). Shaded regions represent 95% confidence intervals. Source: Form 1040, Form 1099, Form W-2, Census Numident, ACS (2001–17), decennial census (2000 and 2010), MAF-ARF. Release authorization CBDRB-FY19-376. A color version of this figure is available online.

estimates for the 2017 effects on employment and earnings are about 30% and 20% smaller, respectively, than their 2015 counterparts.¹¹ Even as these estimates suggest continued recovery from the worst effects of the recession since 2015, that recovery has been somewhat uneven. For both employment and earnings, 2016 estimates represented statistically significant improvement over 2015, but 2017 was little changed from 2016. This pattern differs somewhat from improvement in the overall prime-age employment rate, which increased by 0.7 percentage points from the end of 2015 to the end

¹¹ Improvement since 2015 is interesting because there was an influential view at the time that the economy was at or near full employment, and individuals who had not yet returned to work were relatively unlikely to do so in the future (Krueger 2015). In its economic projections from March of that year, the central tendency of the Federal Open Market Committee's predicted longer-run unemployment rate was 5.0%–5.2%. Fifteen of the seventeen members polled believed it would be appropriate to start tightening monetary policy within the year (FOMC 2015). The fact that the effects estimated here have diminished since then joins continued improvement in indicators like the unemployment rate as evidence that the labor market recovery had not run its course in 2015.

of 2016 and then by another 1.0 percentage point between the end of 2016 and the end of 2017.

B. Estimates Using Additional Cohorts and by Generation

Estimates of the employment effect follow a strikingly different pattern when estimated using a broader set of cohorts than they do when estimated using only prime-age workers. Because the effects of exposure to unemployment shocks may depend on when in the life cycle one is exposed, I broaden my sample to include individuals born between 1928 and 1996. First, I reestimate full sample effects, as in section IV.A. The right column of figure 2 reports these estimates. Rather than becoming negative during the recession and remaining consistently so for several years after, the employment effect here starts to shrink steadily immediately after reaching its greatest magnitude in 2010. A 1 percentage point larger unemployment shock during the recession reduced an exposed individual's probability of being employed in 2010 by 0.7 percentage points. By 2017, after statistically significant improvement in each year, that effect had declined to almost zero, suggesting full recovery of employment over this period on average for the full sample. Employment for more intensely exposed prime-age workers recovered more gradually over this period and had not returned to 2006 levels by 2017.

Earnings effects in this broader sample are generally smaller in magnitude than they are for prime-age workers, but they follow a similar pattern over time. At its greatest severity, an additional percentage point of unemployment exposure cost the average worker a little over \$700 in 2013, compared with more than \$1,100 for the average prime-age worker that same year. The magnitude of this effect had declined by about 35% by 2017 in both the full sample and the prime-age sample, with most of this decline coming between 2014 and 2016.

Millennials saw substantially larger reductions in their probability of being employed immediately after the recession than did workers in other generations, an important difference revealed by estimating effects separately by generation. Figure 3 reports these estimates. The 1.2 percentage point reduction in millennials' employment rate per 1 percentage point of unemployment shock exposure is more than twice the magnitude of this effect for any other generation. On top of this, millennials were less likely to be working at the time than their prime-age counterparts; per table 1, only 44.5% of millennials were employed in 2005, compared with 84.6% of Gen-X members and 78.6% of baby boomers. Following that initial decline, though, millennial employment recovered consistently and relatively quickly, following a trajectory similar to that seen in the full sample. By 2015, exposure to unemployment shocks had essentially no effect on millennials' probability of being employed, and by 2017 millennials exposed to larger shocks were



FIG. 3.—Effects of local unemployment shocks on employment, by generation. Points plotted are β coefficients as estimated in equation (1). Shaded regions represent 95% confidence intervals. Source: Form 1040, Form 1099, Form W-2, Census Numident, ACS (2001–17), decennial census (2000 and 2010), MAF-ARF. Release authorization CBDRB-FY19-376. A color version of this figure is available online.

actually slightly more likely to be employed than those exposed to smaller shocks.

Older workers also saw steady but slower employment recovery after the recession. In 2017, members of Gen-X exposed to larger shocks were still less likely to be employed, but this effect was relatively small in magnitude (coefficient roughly 0.1 percentage points) and only marginally statistically significant. Baby boomers remained nearly 0.3 percentage points less likely to be employed per unit of unemployment shock exposure in 2017.¹² The oldest workers in my sample, members of the silent generation, had the smallest employment effect in percentage points in 2010) and saw it lose statistical significance in 2012 and remain insignificant in the following years.

¹² For the sake of comparison, the cohorts used in Yagan (2019) correspond to older Gen-X/younger baby boomers here, so the fact that estimates in fig. 2 are basically between those in the panels corresponding to those generations in fig. 3 makes sense.

Did Timing Matter?

Among working-age generations, the earnings trajectories look broadly similar to the full sample estimates, as shown in figure 4. In dollar terms, millennials' earnings effects are smaller in magnitude than those of Gen-X throughout the postrecession period and smaller than the baby boomers' effects for much of it; in combination with positive point estimates for older workers, this explains why the full sample estimates are smaller than the prime-age estimates. However, millennials were in the lower-earning early years of their working lives during and after the recession, and their earnings were most adversely affected by recession exposure relative to their mean.¹³ Exposure to the average local unemployment shock cost millennials about 17% of their earnings on average in each of 2010 and 2011. Cumulatively, from 2007 through 2017, exposure to the average shock cost millennials an amount equivalent to about 13% of total earnings over that period for the average millennial worker. For Gen-X, that figure is about 9.1%, while baby boomers lost about 7.1%. If the silent generation's nonsignificant point estimates are taken at face value, they suggest that exposure to the average shock increased earnings by 7.2% during and after the recession.¹⁴

For millennials, this estimate of cumulative lost earnings is larger than previous estimates of the effects of less severe recessions on young workers. Using Canadian data from 1982 through 1999, Oreopoulos, von Wachter, and Heisz (2012) estimate that graduating from college into a recession cost workers 5% of cumulative income over 10 years. Schwandt and von Wachter (2019) find that entering the labor market during a moderate recession (unemployment increases 3 percentage points) would cost workers 60% of a year of earnings over the first decade of work, an amount comparable in magnitude to the Oreopoulos, von Wachter, and Heisz (2012) estimate. If the Schwandt and

¹³ Estimates of earnings effects as shares of mean earnings are reported in app. B. For these specific estimates, see fig. B3. Log earnings estimates are consistent with the level estimates as shares of mean earnings for this group. In particular, log earnings estimates for millennials, which implicitly exclude individuals too young to work, are more negative than for other generations. This suggests that proportionally larger earnings losses for this group are not just a by-product of lower mean earnings due to a larger share of individuals not working.

¹⁴ Estimates for the oldest workers are also somewhat striking. Although their earnings effects are not statistically significant in any year, point estimates indicate that members of the silent generation who were exposed to larger local unemployment shocks earn more than those exposed to smaller shocks on average. Estimating earnings effects for the silent generation using logged earnings produces negative and significant estimates (see fig. B20). The differences between the log and level estimates suggests an important role for extensive margin adjustments in this group. Although only speculative given the precision of these estimates, this could potentially be consistent with shocks late in the life cycle disrupting retirement plans and leading some workers who might otherwise have left the labor force to increase their labor supply. This highlights the limitations of using earnings as a measure of well-being. In this case, older workers' increased earnings may be a sign of their efforts to hold well-being constant in the face of worsening underlying circumstances.



FIG. 4.—Effects of local unemployment shocks on earnings, by generation. Points plotted are β coefficients as estimated in equation (1). Shaded regions represent 95% confidence intervals. Source: Form 1040, Form 1099, Form W-2, Census Numident, ACS (2001–17), decennial census (2000 and 2010), MAF-ARF. Release authorization CBDRB-FY19-376. A color version of this figure is available online.

von Wachter (2019) estimate scales linearly with recession severity, it implies earnings losses comparable to what I find for millennials, suggesting that my estimates may be large in part because of the greater severity of the 2007–9 recession. For young workers, the earnings losses I find are comparable to some estimates of losing one's job in a mass layoff, even though my analysis does not condition on job loss: Schmeider, von Wachter, and Bender (2014) find annual earnings losses due to layoffs of 15% lasting 15 years in German data. My estimates are smaller than other mass layoff–based estimates like those of von Wachter, Song, and Manchester (2009), who find immediate earnings losses of 30% and long-run losses on 20% for workers who lost jobs during the 1980–82 recession.

My estimates for millennials also differ from estimates of the effects of entering the labor market during a recession, but in ways that make sense given differences in settings considered. Unlike several estimates from that literature, my estimates are based on all young workers, rather than focusing specifically on those completing high school or college, because educational attainment may be responsive to changes in broader economic conditions. As a result, differences in sample composition may explain some differences in estimates. For example, my earnings estimates for millennials are initially somewhat

Did Timing Matter?

more negative than those that Altonji, Kahn, and Speer (2016) find for college graduates, less negative than those that Speer (2016) finds for less educated men, and more persistent than both. Since less educated workers tend to be more adversely affected by recession exposure than more educated workers, finding estimates for a sample containing workers with all levels of education that fall between those corresponding to college graduates and workers who did not finish high school is reasonable. The greater persistence of effects in my setting could be related to the greater persistence of high unemployment rates during the slower labor market recovery that followed the Great Recession.

C. Why Have Employment and Earnings Effects Diverged?

Among millennials and Gen-X members, exposure to the Great Recession no longer depresses employment by the end of my sample period, but effects of exposure on earnings remain near their most negative levels. By contrast, by 2017 baby boomers saw their employment effects recover by 50% relative to 2010, while their earnings effects recovered by 65%. Why have older workers seen their earnings improve more quickly than employment after the recession while younger workers' earnings remain depressed even though the adverse employment effects of recession exposure have worn off entirely for them?

One possibility is that exposure to the recession caused younger workers to accumulate less human capital. If the recession made liquidity constraints more binding, it could have reduced investment in education. Young workers who are unemployed early in their careers may miss out on opportunities to learn skills or receive training on the job, and subsequent cohorts can keep them out of entry-level jobs in future years. Early human capital deficits can be hard to make up, leading to persistently lower productivity and lower earnings over the course of one's career. This channel would be consistent with more persistent adverse earnings effects for younger workers, since older workers would generally have completed their formal education and early career on-the-job training prior to the onset of the recession.

Another possibility is that the recession knocked younger workers down the job ladder to a greater extent than it did older workers.¹⁵ Older workers' more developed professional networks, longer experience in their fields, or previously earned credentials could help them retain or regain their positions on the job ladder. This may have made it more difficult for younger workers to advance or even be hired in the first place, as employers increased skill requirements in areas harder hit by the recession (Hershbein and Kahn 2018).

¹⁵ Still other channels may also contribute to persistently lower earnings for younger workers (e.g., changes in preferences for earnings relative to other job characteristics, weakened bargaining positions due to lower prior earnings), but these will not be testable here.



FIG. 5.—Effects of local unemployment shocks on probability of working for a high-paying employer (top quartile), by generation. Points plotted are β coefficients as estimated in equation (1). Shaded regions represent 95% confidence intervals. Source: Form 1040, Form 1099, Form W-2, Census Numident, ACS (2001–17), decennial census (2000 and 2010). Release authorization CBDRB-FY19-500. A color version of this figure is available online.

I test the job ladder channel by considering the recession's effect on the likelihood of working for high-paying firms by generation. Figure 5 shows, for each generation, the effects of unemployment shocks on the probability of working for an employer that is in the highest-paying quartile, as measured using firm fixed effects estimated via a procedure inspired by Abowd, Kramarz, and Margolis (1999).¹⁶ The empirical specification is the same as that used for the main analysis, shown in equation (1). Millennials, Gen-X members, and baby boomers all see a sustained decline in the probability of working for a high-paying employer during the recession. After 2010, this probability began to recover for Gen-X members and baby boomers, roughly in proportion to the rate at which their overall employment probabilities recovered. Millennials, though, experienced no such recovery. The effects of unemployment shocks on employment at high-paying employers continued to worsen for millennials until 2014, after which year they improved modestly. This is despite the fact that millennials overall employment effects began to recover

¹⁶ This procedure is described in more detail in app. C.

after 2010 and diminished much more quickly than other generations.¹⁷ Together, these estimates suggest that the divergence between the employment and earnings recoveries experienced by millennials may be due in part to the fact that recession exposure has made them persistently less likely to work for higher-paying employers, even as they are no longer less likely to be working.

This result both affirms and extends previous work on the scarring effects of the conditions workers face at labor market entry. Parts of that literature suggest that workers whose initial jobs are worse on various dimensions have persistently lower earnings at the same points in their careers than other workers who faced more favorable initial conditions (Kahn 2010; Oreopoulos, von Wachter, and Heisz 2012; Schmeider, von Wachter, and Bender 2014), consistent with my findings here. However, my analysis also indicates that even though younger workers typically have much more upward job mobility than older workers, it was older workers who were better able to find jobs with higher-paying employers in the aftermath of the Great Recession, aiding their earnings recovery. This further emphasizes the importance of job ladders to understanding employment and earnings dynamics during recessions.

Eligibility rules for programs like unemployment insurance (UI) could contribute to this dynamic. Workers with little to no work history (e.g., new labor market entrants more often found among younger workers) are generally not eligible for UI benefits. As a result, younger workers may have lower reservation wages and be more willing to accept jobs from lower-paying employers and to accept employment more quickly than older workers with access to either UI or self-insurance in the form of savings, contributing to both the faster employment recovery and the more persistent earnings losses among younger workers. Future research should explore this possibility further.

Changes in the likelihood of working for high-paying employers due to the Great Recession are a potentially important contributor to its effects on millennials' earnings. I estimate effects of recession exposure on the average earnings of a person's coworkers. Effects on the earnings of millennials' coworkers follow the same trajectory and have very similar magnitudes as effects on millennials' own earnings.¹⁸ Since about half of changes in mean coworkers' earnings appear to pass through to one's own earnings, this suggests that about half of the reduction in millennials' earnings could be driven by their working for lower-paying employers.¹⁹

Returning to the human capital channel, I consider the effects of exposure to local unemployment shocks on educational attainment among younger

¹⁷ Other formulations of this analysis tell a similar story. Additional estimates are presented in app. C.

¹⁸ See fig. A9.

¹⁹ Regressing log(own earnings) on log(mean coworkers' earnings) with person fixed effects in my sample produces a coefficient of 0.48.

workers. In theory, the sign of the effect of recession exposure on educational attainment is ambiguous. Worsening economic conditions decrease the opportunity cost of higher education and so might increase educational attainment. Alternatively, they could also reduce educational attainment by creating or exacerbating liquidity constraints, leading students to withdraw from or decline to enroll in college. For respondents to the ACS, I observe educational attainment as of survey response. For millennials who responded to the ACS at or after age 25, I can use this information to estimate the effect of recession exposure on educational attainment. Specifically, I can estimate

$$y_i = \beta Shock_{CZ07(i)} + X_i \gamma + \varepsilon_i, \qquad (2)$$

where y_i is an indicator for having completed a given level of educational attainment, *Shock*_{CZ07(i)} is the same local unemployment shock defined above, and X_i is a set of individual characteristics that includes the full set of interactions among year of survey response, age, race/ethnicity, and gender.²⁰

Being exposed to a more severe unemployment shock reduced millennials' likelihood of completing a postsecondary degree. Figure 6 reports the effects of an additional percentage point increase in local unemployment on the probability of completing various amounts of education. Each percentage point of exposure to local unemployment shocks reduced the likelihood of completing any postsecondary degree by 1.3 percentage points.²¹ This is consistent with Stuart (2019), who found that younger children exposed more intensely to the 1980–82 recession were less likely to complete 4-year college degrees. When using the most detailed categories, I also find that recession

²⁰ The data I have contain at most one observation of educational attainment per person, so they do not allow me to observe changes in educational attainment, conduct panel analysis, or examine prerecession trends in educational attainment, and identification in this exercise is entirely cross sectional. One might consider estimating my baseline specification with controls for education included as an alternative way of investigating the contribution of educational attainment to my employment and earnings estimates. Because my baseline specification includes person fixed effects and educational attainment does not vary within person over time in my data, education controls would fall out of my baseline specification, rendering this approach uninformative. I can, however, estimate a model that removes the person fixed effects (instead including CZ fixed effects) and add education controls. Estimates from this model are shown in fig. A10, alongside estimates from the same model without education controls. Estimates of earnings effects for millennials from the version with education controls are about 60% smaller in 2017 than estimates without education controls. Although this is not directly comparable to the main education analysis in this paper, it also suggests a potentially important role for changes in educational attainment.

²¹ Educational attainment is only measurable for ACS respondents as of the response date. Although all individuals in this analysis responded to the ACS at age 25 or older, some could have completed additional education after responding. This estimate is best considered an upper bound on the effect of exposure on final educational attainment, since part of the effect measured here could be delay in educational attainment rather than reduction.



FIG. 6.—Effects of local unemployment shocks on millennials' educational attainment. Points plotted are β coefficients as estimated in equation (2). Black lines represent 95% confidence intervals. Source: Form 1040, Form 1099, Form W-2, Census Numident, ACS (2001–17), decennial census (2000 and 2010), MAF-ARF. Release authorization CBDRB-FY19-376. LTHS = less than high school; HS = high school. A color version of this figure is available online.

exposure increased millennials' likelihood of completing some college but not completing a degree. This is consistent with Foote and Grosz (2019), who find increased postsecondary enrollment in response to mass layoffs. This enrollment is concentrated in associate's degree and certificate programs and increased completions are concentrated in certificate programs, all of which would show up in the "some college" category used here, where I do find an increase associated with recession exposure. Like Foote and Grosz (2019), I do not find evidence of increased associate's degree completion. Kahn (2010) finds that worse conditions at labor market entry increase educational attainment, although that analysis focuses on a sample composed entirely of college graduates, a group not directly comparable to the sample used here, and the educational attainment outcomes considered are entirely postgraduate. The resources available to and considerations faced by this group are not representative of the broader population considered in this paper. Using coefficients from my least aggregated educational attainment specification and a Mincerean earnings regression using ACS-matched people in my analysis sample, changes in educational attainment due to the recession may have reduced earnings by 3.6%-5.7%.²² This represents about a

²² The lower end of this range uses education premiums estimated using all ACSmatched people. The upper end uses education premiums estimated using millennials only.

quarter to a half of the 13% of income millennials lost to recession exposure through 2017.²³

These effects on educational attainment could contribute to the divergence between employment and earnings recovery discussed above. Millennials exposed to more severe shocks may have been forced to pursue work over education, reducing the magnitude of the adverse employment effect over time. This shift away from education could also reduce future earnings growth for these workers, possibly contributing to the persistence of both the reduced likelihood of working for high-paying employers and the negative earnings effects of local unemployment shocks.

D. Estimates by Displacement Status

The harm from losing a job during a recession is well documented, but workers who remain employed may also suffer if prevailing conditions make it more difficult to change jobs or ask for a raise (Beaudry and DiNardo 1991). Annual data are not ideal for identifying employment transitions, but a few broad categories of workers can be identified and compared. Using data from W-2s, I can identify individuals who did not work at all in 2006, prior to the recession (about 1.4 million of the people in my sample); those who were working prior to the recession but then spent at least one full calendar year not working between 2007 and 2009 (about 400,000 people); those who worked at some point in every year during the recession but saw their main employer (i.e., the one that paid the largest share of their earnings) change during that period (about 1.1 million people); and those who worked in all years of the recession and had the same main employer in all years (about 1.3 million people).

In each group, including the two that were employed prior to the recession and did not experience a full year of nonemployment during it, I find significant reductions in earnings following the recession. Figure 7 shows earnings effects for these four groups of workers based on all birth cohorts, estimated using equation (1). At this effect's largest magnitude, workers who remain employed but see their main employer change lose out on about \$1,000 per year per percentage point of local unemployment increase during the recession. This is similar to the dollar amount lost by workers who kept the same main employer throughout the recession. In both cases, this effect amounted to about 2% of annual earnings (or about 10% for a shock of average magnitude) at its largest magnitude.²⁴ Note that the displacement status

²³ Note that this lost income is not necessarily distinct from income lost due to working for lower-paying employers, since one might end up working for a lower-paying employer as a result of having completed less education.

²⁴ Estimates for each generation are reported in app. D. With the exception of the silent generation (for which my main point estimates suggest that recession exposure increased earnings), each generation's estimates are qualitatively similar to these full-sample estimates.



FIG. 7.—Effects of local unemployment shocks on earnings, by displacement status. Points plotted are β coefficients as estimated in equation (1). Shaded regions represent 95% confidence intervals. Source: Form 1040, Form 1099, Form W-2, Census Numident, ACS (2001–17), decennial census (2000 and 2010), MAF-ARF. Release authorization CBDRB-FY19-431. A color version of this figure is available online.

groups used here are constructed by selecting on short-run changes in one outcome of interest. Members of these groups may differ unobservably in important ways. If, for example, those who lost jobs during the recession experienced earnings losses beforehand, a la Ashenfelter (1978), that could affect the interpretation of the estimated earnings effects for that group. The limited prerecession data I have access to do not allow me to thoroughly investigate prerecession earnings trends, so there is little I can do to address potential issues like this. However, the fact that even workers who remain attached to their prerecession primary employer through the duration of the recession (a group that is, if anything, selected not to see an adverse earnings effect) see their earnings reduced by exposure to local unemployment shocks suggests that recession exposure is likely also harmful to nondisplaced workers.

E. Young Workers and Dependency

The youngest millennials turned 11 years old in 2007 and may have experienced different consequences of recession exposure than older millennials,



FIG. 8.—Effects of local unemployment shocks on employment and earnings, subgroups of millennials. Points plotted are β coefficients as estimated in equation (1). Shaded regions represent 95% confidence intervals. Source: Form 1040, Form 1099, Form W-2, Census Numident, ACS (2001–17), decennial census (2000 and 2010), MAF-ARF. Release authorization CBDRB-FY2021-CES005-010. A color version of this figure is available online.

who were old enough to be fully engaged in the labor market at the onset of the Great Recession. The upper panels of figure 8 reports estimates of the effects of recession exposure on employment and earnings for millennials born in 1988 or earlier (i.e., those who turned 19 or older in 2007) and those born in 1989 or later (those who turned no older than 18 in 2007). Given that the unemployment rate remained above its previous peak (6.3% in June 2003) until May 2014 and above its 2006–7 average (4.6%) until February 2017, it is reasonable to think that even millennials who were in middle school in 2007 might suffer adverse labor market consequences due to recession exposure during the period covered by my data. Effects of recession exposure on the parents of these millennials could also subsequently affect their own labor market outcomes (e.g., through loss of resources, disruption of home life/education associated with parental job loss), but detailing intergenerational effects of recession exposure is beyond the scope of this paper.

Employment effects were larger for younger millennials (who were aging into a slack labor market without prior experience or professional networks), while earnings effects were larger in dollar terms for older millennials and proportionally similar for both groups. Both younger and older millennials saw employment recover by the end of the analysis period, while negative



FIG. 9.—Effects of local unemployment shocks on dependent status, millennials. Points plotted are β coefficients as estimated in equation (1). Shaded regions represent 95% confidence intervals. Source: Form 1040, Form 1099, Form W-2, Census Numident, ACS (2001–17), decennial census (2000 and 2010), MAF-ARF. Release authorization CBDRB-FY2021-CES005-010. A color version of this figure is available online.

earnings effects persisted.²⁵ Grouping millennials according to whether they were claimed as dependents on tax forms in 2007 (shown in the lower panels of fig. 8) produces a corresponding pattern of estimates, although the groups appear more similar to each other under this formulation than they do using age. This makes sense because a meaningful share of the older millennials group would have been enrolled in college in 2007, increasing the likelihood that they were claimed as dependents.

Exposure to the recession also made millennials more likely to be claimed as dependents. Figure 9 shows a sharp increase in that likelihood in 2010 in particular, and although the magnitude of the effect declined in subsequent

²⁵ While estimates for narrower cohorts of millennials are too imprecise to support strong conclusions, point estimates suggest that the effects of recession exposure could differ within relatively narrow age ranges depending on when the onset of the recession fell in one's educational trajectory. For example, college graduates who were likely already in the labor market when the recession began appear to have experienced adverse earnings effects immediately, while those likely still in school may have seen them develop more gradually. Non–college graduates who would have been in high school during the recession saw especially large reductions in employment in 2008 and 2009, while effects appear more muted for older workers. See figs. A11 and A12 for estimates for select cohorts.

years, it remained around 70% of its peak magnitude in 2017. Also, as the cohorts in question aged over the course of the analysis period and increasingly lived independently, the size of the effect relative to the underlying average share of millennials being claimed as dependents has increase steadily; the increase in dependent status associated with being exposed to the average local unemployment shock during the Great Recession exceeded 25% of the overall rate of dependent status among millennials in 2017. Although not the only potential driver of this increased frequency of being claimed as a dependent, these estimates are consistent with recent evidence of higher rates of young adults living with their parents since the onset of the Great Recession (e.g., Fry, Passel, and Cohn 2020).

F. Heterogeneity

By a variety of measures, men fared worse than women during and after the Great Recession. Were men and women differentially affected by exposure to local unemployment shocks?

My estimates indicate that these shocks hurt men more in terms of both employment and earnings. Figure 10 shows estimated effects of unemployment shocks on employment and earnings by gender. An additional percentage point of exposure made women about 0.5 percentage points less likely to be employed in each year from 2009 through 2011, while for men this effect reached 0.8 percentage points in 2010. Over the period considered, about 68%–71% of men work in a given year, while the same is true of 64%– 66% of women, so estimated effects for men are also larger relative to the mean employment rate. The absolute difference between employment effects on men and women is not permanent; by 2015, these estimates converge on each other, and neither gender's estimates are statistically significant in 2016 or 2017, nor are they statistically different from each other.

On earnings, men also fare worse than women in absolute terms. On average, men lose between \$900 and \$1,500 per year per percentage point of unemployment shock exposure in each year from 2009 through 2014. The corresponding estimate for women is only about \$400–\$500. Women's estimates, however, show little sign of recovery following the recession, remaining around \$430 in 2017, while the \$630 effect on men that year is less than half the largest magnitude among their annual estimates. Differences between men's and women's estimates are statistically significant in all years. In relative terms, annual effects on men remain larger, reaching 3.3% of annual earnings in 2013; earnings effects for women reach up to 2.2% of annual earnings in 2011.²⁶ Cumulatively, exposure to an unemployment shock of average magnitude cost men an amount equivalent to 9.6% of the average worker's total earnings between 2007 and 2017, while women lost 7.7%.

²⁶ Log earnings estimates tell a similar story.



FIG. 10.—Effects of local unemployment shocks on employment and earnings, 1928–96 birth cohorts, by gender. Points plotted are β coefficients as estimated in equation (1). Shaded regions represent 95% confidence intervals. Source: Form 1040, Form 1099, Form W-2, Census Numident, ACS (2001–17), decennial census (2000 and 2010), MAF-ARF. Release authorization CBDRB-FY19-376. A color version of this figure is available online.

Looking across generations, the difference between overall effects on men and women appear to be driven by Gen-X and the baby boomers. As shown in figure 11, effects on employment for millennial men and women follow very similar trajectories, and for the silent generation, neither gender's employment effects are generally statistically significant. Within Gen-X, men's employment was more adversely affected during the recession, although since 2011 effects on men's employment are not distinguishable from effects on women's employment, and point estimates suggest that men's employment has recovered more quickly and more fully than women's. Among baby boomers, unemployment shocks also had larger employment effects on



FIG. 11.—Effects of local unemployment shocks on employment, by gender and generation. Points plotted are β coefficients as estimated in equation (1). Shaded regions represent 95% confidence intervals. Source: Form 1040, Form 1099, Form W-2, Census Numident, ACS (2001–17), decennial census (2000 and 2010), MAF-ARF. Release authorization CBDRB-FY19-376. A color version of this figure is available online.

men than on women, although that difference largely emerged after the official end of the recession, and it remained statistically significant until 2017. Figure 12 shows a similar pattern for earnings: there is little absolute gender gap in earnings effects among millennials, effects on the silent generation are not statistically significant, and an additional unit of unemployment exposure generally cost Gen-X and baby boomer men up to \$1,800 and \$2,500, respectively, in their worst years postrecession, compared with \$1,000 and \$600 for Gen-X and baby boomer women, respectively.²⁷

Splitting the sample by race and ethnicity, I find that employment and earnings effects for White, Black, Asian, and Hispanic workers followed similar trajectories over the course of the recession and recovery, although earnings effects are generally not statistically distinguishable across groups.²⁸ Taking point estimates literally suggests that recession exposure cost the average Black worker an amount equivalent to 12.1% of the average worker's total earnings between 2007 and 2017. This compares with 8.5% for White workers, 6.6% for Asian workers, and 6.2% for Hispanic workers. These losses

²⁷ Earnings effects for millennial men and women are statistically distinguishable from each other in many years, but the magnitudes of the differences are small.

²⁸ These race/ethnicity groups are constructed such that Hispanic includes individuals of any race, while race groups include only non-Hispanic individuals.

Did Timing Matter?



FIG. 12.—Effects of local unemployment shocks on earnings, by gender and generation. Points plotted are β coefficients as estimated in equation (1). Shaded regions represent 95% confidence intervals. Source: Form 1040, Form 1099, Form W-2, Census Numident, ACS (2001–17), decennial census (2000 and 2010), MAF-ARF. Release authorization CBDRB-FY19-376. A color version of this figure is available online.

also translate into 1.33 years of average earnings for Blacks, 0.94 years for Whites, 0.73 years for Asians, and 0.68 years for Hispanics.²⁹

Note that all local unemployment shocks in this section (and in the paper as a whole) are constructed using the overall unemployment rate rather than group-specific rates (e.g., unemployment rates specifically for millennials or women or Hispanic workers).³⁰ This is important to keep in mind when considering how to interpret estimates in this section. If, for example, men are more likely to work in industries that saw relatively large increases in unemployment, they may be more intensely treated (in terms of either the magnitude or the duration of the shock) by a given increase in the overall unemployment rate than women. Similarly, the unemployment rate among Black workers increased much more than the unemployment rates among White workers during the Great Recession; to the extent that labor markets are racially segregated within localities, Black workers who faced an increase

²⁹ See app. E for figures reporting these underlying point estimates.

³⁰ How a given change in macroeconomic conditions translates into the experience of a particular subgroup is not preordained but endogenous to a wide variety of conditions and policy responses. Using group-specific unemployment shocks in this analysis would therefore have the flavor of controlling for a channel through which that group might be affected by the recession. Moreover, it is not clear that labor markets are sufficiently segmented to make using group-specific unemployment rates appropriate for this exercise.

in the overall unemployment rate of a given magnitude were likely more exposed to the recession than White workers facing the same overall unemployment rate increase. In examples like these, differences in "true" exposure may explain differences in estimated effects of exposure to increases in overall unemployment. As a corollary, the overall unemployment rate may be in some sense a "better" measure of recession exposure for groups whose own unemployment rates more closely track the overall figures (e.g., White workers, men).

G. Effects on Places

One way to mitigate adverse effects of a local unemployment shock is to migrate (Blanchard and Katz 1992; Dao, Furceri, and Loungani 2017); if local unemployment shocks induce substantial migration, estimated effects on places could differ from effects on people. To investigate this possibility, I modify my baseline specification to use CZ fixed effects instead of person fixed effects and to allow recession exposure to vary as individuals move across CZs rather than fixing it based on their 2007 location. Formally, I estimate

$$y_{ict} = \sum_{t} \beta_{t} Shock_{c} \cdot 1[Year = t] + \sum_{a \in ages} \alpha_{a} \cdot 1[Age_{it} = a] + \gamma_{c} + \delta_{t} + \varepsilon_{ict}, \qquad (3)$$

where *c* now indexes CZs, γ_c represents CZ fixed effects, and $Shock_c = UR_{c,2009} - UR_{c,2007}$ is the change in the unemployment rate between 2007 and 2009 in CZ *c*. The interaction term again omits t = 2006. Figure 13 presents the β_t coefficients from this specification, which I call place estimates, alongside the β_t coefficients from my baseline specification, here referred to as person estimates.

Overall, the trajectories of person and place estimates of the effects of recession exposure on employment and earnings are very similar. Where the two diverge, the person estimates are generally smaller in magnitude than the place estimates, although they generally fall within the confidence intervals of the place estimates. This suggests that there may have been some scope for workers to mitigate the effects of recession exposure by migrating, but overall workers fared roughly as well as the places where they lived at the beginning of the recession, even with the ability to move. This is consistent with direct estimates of effects of recession exposure on migration, which do not find evidence that people exposed to larger unemployment shocks were more likely to leave their 2007 CZ.³¹ To the extent that person and place estimates differ, the employment effects differ by more, suggesting that in this context it

³¹ These estimates are presented in app. F.



FIG. 13.—Person versus place estimates, 1928–96 birth cohorts. Points plotted in the solid light gray line are β coefficients as estimated in equation (1). Points plotted in the dashed dark gray line are β coefficients as estimated in equation (3). Shaded regions represent 95% confidence intervals. Source: Form 1040, Form 1099, Form W-2, Census Numident, ACS (2001–17), decennial census (2000 and 2010), MAF-ARF. Release authorization CBDRB-FY19-376. A color version of this figure is available online.

may have been easier to improve one's likelihood of being employed by moving than it was to improve one's earnings.³²

H. Robustness to Controlling for Other Shocks

The Great Recession also affected local labor markets in ways not fully captured by the unemployment rate. For example, average local labor market

 $^{\rm 32}$ Generation-specific estimates, also presented in app. F, suggest that this is especially true of baby boomers.

concentration increased during the Great Recession, breaking with a longrunning downward trend (Rinz 2018). Relatedly, national industry-level changes associated with the recession may have had differential affects on future employment growth across localities based on the geographic distribution of employment by industry. To the extent that these changes are correlated with changes in local unemployment, they could affect the interpretation of my baseline estimates.

As a test of this, I reproduce my main estimates using a specification that includes a similarly constructed local concentration shock and a Bartik shock. This is arguably an overly conservative approach, since the Bartik shock in particular is highly correlated with the unemployment rate and likely reflects the same underlying changes to a significant degree. Despite this, however, estimates from this specification remain significant, have similar magnitudes, and follow the same qualitative patterns described above. Additional details are available in appendix A. Moreover, the coefficients on the concentration shocks suggest that changes in local labor market concentration may have had their own persistent adverse effects on young workers in particular. Future research should investigate this further.

V. Conclusion

Local unemployment shocks associated with the Great Recession had adverse consequences for workers of all ages, but younger workers appear to have suffered the most persistent harms. Despite the fact that the employment rates of millennials and Gen-X members are no longer significantly depressed because of these shocks, their earnings have shown little sign of recovering. For millennials specifically, my estimates indicate that recession exposure reduced educational attainment and made them less likely to work for highpaving employers, two effects that may combine to drive this divergence between the trajectories of their employment and earnings recoveries. The fact that millennials continue to work for lower-paying employers than they otherwise would have indicates that this group in particular may have been knocked off the job ladder during the Great Recession or that it has not fully resumed functioning for them during the recovery (Moscarini and Postel-Vinay 2016). As discussed above, this is consistent with the literature on scarring in the labor market, which suggests that worse initial employment opportunities drive long-term consequences of entering the labor market during a downturn and provides reason to believe these earnings losses could persist in the even-longer run. Here, persistent reductions in working for high-paying employers are primarily observed among younger workers; this dynamic may be specific to conditions at labor market entry.

Shocks to local labor market concentration during the Great Recession also reduced younger workers earnings, even conditioning on unemployment shocks and a Bartik-style labor demand shock. Although smaller in magnitude

Did Timing Matter?

than the effects of unemployment shocks and realized more gradually, the effects of concentration shocks appear to be even more persistent and, if anything, to have grown more negative over time. The story of the labor market recovery from the Great Recession is not just about overcoming a huge demand shock; changes in market structure matter too.

These estimates have implications for how the government might best respond to future recessions. UI is the primary form of assistance provided to workers who lose jobs and is intended to provide some financial security while they look for jobs that best match their skills and preferences. Even though UI benefits were extended substantially during the Great Recession, eligibility requirements that excluded new entrants from receiving them remained in place. Is it a coincidence, then, that younger workers, who were less likely to have the prior work experience that would make them eligible for UI, still saw recession exposure reduce their likelihood of working for high-paying employers in 2017, 10 years after the recession began? And if recession-related changes in market structure are reducing younger workers earnings and making them less likely to be employed, could the structure or value of UI benefits adjust to mitigate these effects? Might some other form of assistance be necessary to improve workers subsequent earnings and employment outcomes in the face of these changes in the labor market? I leave these questions to future research.

References

- Abel, Jaison R., and Richard Deitz. 2016. Underemployment in the early careers of college graduates following the Great Recession. Federal Reserve Bank of New York Staff Report no. 749.
- Abowd, John M., Francis Kramarz, and David N. Margolis. 1999. High wage workers and high wage firms. *Econometrica* 67, no. 2:251–333.
- Altonji, Joseph G., Lisa B. Kahn, and Jamin D. Speer. 2016. Cashier or consultant? Entry labor market conditions, field of study, and career success. *Journal of Labor Economics* 34, no. S1:S361–S401.
- Ashenfelter, Orley. 1978. Estimating the effect of training programs on earnings. *Review of Economics and Statistics* 1:47-57.
- 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, no. 4:665–88.
- Blanchard, Olivier Jean, and Lawrence F. Katz. 1992. Regional evolutions. *Brookings Papers on Economic Activity* 1992, no. 1:1–75.
- Cockx, Bart, and Corinna Ghirelli. 2016. Scars of recessions in a rigid labor market. *Labour Economics* 41:162–76.
- Dao, Mai, Davide Furceri, and Prakash Loungani. 2017. Regional labor market adjustment in the United States: trend and cycle. *Review of Economics and Statistics* 99, no. 2:243–57.

- Davis, Steven J., and Till M. von Wachter. 2011. Recessions and the cost of job loss. *Brookings Papers on Economic Activity* 43, no. 2:1–72.
- Dimock, Michael. 2019. Defining generations: Where Millennials end and Generation Z begins. https://www.pewresearch.org/fact-tank/2019/01 /17/where-millennials-end-and-generation-z-begins/.
- FOMC (Federal Open Market Committee). 2015. Chair's FOMC press conference projections materials, March 18, 2015. https://www.federal reserve.gov/monetarypolicy/files/fomcprojtabl20150318.pdf.
- Fernández-Kranz, Daniel, and Núria Rodríguez-Planas. 2018. The perfect storm: Graduating during a recession in a segmented labor market. *ILR Review* 71, no. 2:492–524.
- Foote, Andrew, and Michel Grosz. 2019. The effect of local labor market downturns on postsecondary enrollment and program choice. *Education Finance and Policy* 15:1–50.
- Forsythe, Eliza. 2016. Why don't firms hire young workers during recessions? Working paper.
- ——. 2019. Careers within firms: Occupational mobility over the lifecycle. *Labour* 33, no. 3:241–77.
- Fry, Richard, Jeffrey S. Passel, and D'Vera Cohn. 2020. A majority of young adults in the U.S. live with their parents for the first time since the Great Depression. https://www.pewresearch.org/fact-tank/2020/09/04 /a-majority-of-young-adults-in-the-u-s-live-with-their-parents-for-the -first-time-since-the-great-depression/.
- Genda, Yuji, Ayako Kondo, and Souichi Ohta. 2010. Long-term effects of a recession at labor market entry in Japan and the United States. *Journal of Human Resources* 45, no. 1:157–96.
- Hershbein, Brad, and Lisa B. Kahn. 2018. Do recessions accelerate routinebiased technological change? Evidence from vacancy postings. *American Economic Review* 108, no. 7:1737–72.
- Hershbein, Brad J. 2012. Graduating high school in a recession: Work, education, and home production. *BE Journal of Economic Analysis and Policy* 12, no. 1:3.
- Jacobson, Louis S., Robert J. LaLonde, and Daniel G. Sullivan. 1993. Earnings losses of displaced workers. American Economic Review 83, no. 4:685–709.
- Jarosch, Gregor. 2015. Searching for job security and the consequences of job loss. Working paper.
- Kahn, Lisa B. 2010. The long-term labor market consequences of graduating from college in a bad economy. *Labour Economics* 17, no. 2:303–16.
- Krueger, Alan B. 2015. How tight is the labor market? The 2015 Martin Feldstein Lecture, National Bureau of Economic Research. https://www.nber.org/reporter/2015number3/.
- Moscarini, Giuseppe, and Fabien Postel-Vinay. 2016. Did the job ladder fail after the Great Recession? *Journal of Labor Economics* 34., no. S1:S55–S93.

734

- Neal, Derek. 1995. Industry-specific human capital: Evidence from displaced workers. *Journal of Labor Economics* 13, no. 4:653–77.
- Oreopoulos, Philip, Till von Wachter, and Andrew Heisz. 2012. The shortand long-term career effects of graduating in a recession. *American Economic Journal: Applied Economics* 4, no. 1:1–29.
- Rinz, Kevin. 2018. Labor market concentration, earnings inequality, and earnings mobility. CARRA Working Paper no. 2018-10.
- Rothstein, Jesse. 2019. The lost generation? Scarring after the Great Recession. Working Paper.
- Ruhm, Christopher J. 1991. Are workers permanently scarred by job displacements? *American Economic Review* 81, no. 1:319–24.
- Schmeider, Johannes F., Till von Wachter, and Stefan Bender. 2014. The costs of job displacement over the business cycle and its sources: Evidence from Germany. Working paper.
- 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, no. S1: S161–S198.
- Speer, Jamin D. 2016. Wages, hours, and the school-to-work transition: The consequences of leaving school in a recession for less-educated men. *BE Journal of Economic Analysis and Policy* 16, no. 1:97–124.
- Stuart, Bryan A. 2019. The long-run effects of recessions on education and income. Working paper.
- Sullivan, Daniel, and Till von Wachter. 2009. Job displacement and mortality: An analysis using administrative data. *Quarterly Journal of Economics* 124, no. 3:1265–306.
- Topel, Robert. 1990. Specific capital and unemployment: Measuring the costs and consequences of job loss. In *Carnegie-Rochester conference series on public policy*, vol. 33, 181–214. Amsterdam: Elsevier.
- Topel, Robert H., and Michael P. Ward. 1992. Job mobility and the careers of young men. *Quarterly Journal of Economics* 107, no. 2:439–79.
- von Wachter, Till, Jae Song, and Joyce Manchester. 2009. Long-term earnings losses due to mass layoffs during the 1982 recession: An analysis using US administrative data from 1974 to 2004. http://www.econ.ucla.edu /tvwachter/papers/mass_layoffs_1982.pdf.
- Wagner, Deborah, and Mary Layne. 2014. The Person Identification Validation System (PVS): Applying the Center for Administrative Records Research and Applications' (CARRA) record linkage software. CARRA Working Paper no. 2014-01.
- Yagan, Danny. 2019. Employment hysteresis from the Great Recession. Journal of Political Economy 127, no. 5:2505–58.