The 1980s Recession and U.S. Labor Market Inequality*

Donato Onorato[†] Kamelia Stavreva[‡]
Job Market Paper

October 28, 2026

[Click here for the latest version]

Abstract

Almost 60 percent of the sharp rise in 90–10 wage inequality since 1979 happened between 1980 and 1985, amidst the deep recessions of the early 1980s. This paper studies how this recession contributed to the abrupt transformation of the wage distribution over this period. We construct a new longitudinal employer-employee dataset for 5.3 million workers from Census administrative records and exploit variation in exposure to the recession across local labor markets to estimate the effects on workers' labor market trajectories. We find that the recession persistently reduced earnings and employment, with the largest effects for workers at the bottom of the earnings distribution. Under a less severe recession counterfactual, we estimate that the rise in wage inequality between 1979 and 1988 would have been up to 20 percent smaller. We find no evidence that these effects reflect accelerated technological change or other lasting shifts in relative labor demand. Instead, we find that the recession persistently reduced unionization, and extrapolating our estimates suggests that the recession can account for 42 percent of the national decline in union density over the decade. Decomposing the recession's impacts on aggregate inequality, we find that 35 to 40 percent is accounted for by its effects on unions, through both the reallocation of workers to non-union jobs and weaker union wage spillover effects. Our results suggest that the recession meaningfully contributed to the post-1980s rise in wage inequality and demonstrate how temporary aggregate shocks can have permanent impacts on inequality by reshaping wage-setting institutions.

*We are grateful to Liz Ananat, Sandra Black, Wojciech Kopczuk, Suresh Naidu, Winnie van Dijk, and Ebonya Washington for their constant guidance and advice. Any remaining errors are our own. Any views expressed are those of the authors and not those of the U.S. Census Bureau. The Census Bureau has reviewed this data product to ensure appropriate access, use, and disclosure avoidance protection of the confidential source data used to produce this product. This research was performed at a Federal Statistical Research Data Center under FSRDC Project Number 3420. (CBDRB-FY26-P3420-R13312)

[†]Columbia Economics Department
dao2123@columbia.edu

[‡]Columbia Economics Department
kes2220@columbia.edu

I. Introduction

The 1980s marked the abrupt beginning of a decades-long rise in U.S. labor market inequality, which is one of the defining economic trends of the past half century. Between 1979 and 2020, the ratio of the 90th to the 10th percentile of the hourly wage distribution increased by 38 percent. The causes of this trend have been widely discussed in the economics literature, with prior work emphasizing gradual forces such as technological change (Katz and Murphy, 1992; Autor, Katz and Kearney, 2008; Acemoglu and Autor, 2011), globalization and trade (Feenstra and Hanson, 1996; Greenland, Lake and Lopresti, 2026), deindustrialization (Bluestone and Harrison, 1982), and weakening labor market institutions (Lee, 1999; DiNardo, Fortin and Lemieux, 1996; Farber et al., 2021) as important factors. Yet, around 60 percent of the rise over these four decades occurred in just five years, between 1980 and 1985, amidst the severe twin recessions of the early 1980s.[1]

This sudden rise over a short period is not easily explained by gradual forces that operate over decades, and despite the timing, the recession has been overlooked by the literature examining inequality in this period. However, recessions are temporary shocks, and standard models of local adjustment predict that places hit harder than others should recover as firms adjust and workers move. For the recession to have permanently widened the wage distribution, it must have had effects that did not reverse once the economy recovered, and these effects must have been larger for those at the bottom of the distribution. For instance, it could have accelerated the gradual forces already underway, such as technological change, in more exposed places, or it could have further weakened labor market institutions, like unions, which were already in decline during this period.

This paper asks how the 1980s recession contributed to the sharp rise in wage inequality that followed it. To study the recession, we combine Social Security, IRS, and Census Bureau administrative records to construct a new longitudinal linked employer-employee dataset for a sample of 5.3 million workers that covers the period between 1976 and 1989. Importantly, these data cover a much earlier period than other more commonly used administrative datasets for the United States, which allows us to follow workers through the recession and the decade that followed.

To identify the impacts of the recession, we exploit variation across U.S. labor markets in exposure to the 1980s downturn (Yagan, 2019; Hershbein and Stuart, 2024), which we measure as the change in the local unemployment rate between 1979 and 1982. Leveraging the sharp geographic variation in labor market exposure to the recession enables us to distinguish the

[1] There were two back-to-back recessions, from January to July 1980 and from July 1981 to November 1982, which we refer to together as the 1980s recession.

impacts of the recession from other nationwide secular trends and shocks. We show that a local labor market's sensitivity to the Federal Reserve's monetary policy shocks (Romer and Romer, 2004) during the recession is strongly predictive of its local exposure, and we find no evidence that more exposed places were on different trajectories prior to the recession. The longitudinal data also allow us to address two further potential threats to identification. First, since we follow workers over time, our estimates are not contaminated by compositional changes in who lives in more exposed places. Second, the data allow us to compare similar types of workers by controlling for detailed pre-recession age-earnings-industry fixed effects as well as demographic and human capital characteristics of workers. This comparison mitigates concerns related to aggregate economic trends unrelated to the recession differentially impacting the types of workers in more exposed places.

We find large and persistent negative impacts on the labor market trajectories of workers in more exposed labor markets. Exposure to a one percentage-point larger increase in local unemployment during the 1980s recession reduces the probability of any employment by 0.22 percentage points and annual earnings by 1.9 percent in 1989, a decade after the recession began and well after the aggregate recovery. These effects were largest for workers at the bottom of the pre-recession earnings distribution, where the earnings losses were almost twice as large as for workers at the top of the distribution. The magnitudes are comparable to those documented for the Great Recession (Yagan, 2019), which suggests that these lasting effects on workers may be a common feature of severe recessions.[2] These patterns further demonstrate that the 1980s recession contributed to a long-term increase in employment and earnings inequality between different types of workers.

These worker-level effects point to the recession widening the distribution of wages. However, the recession could have reduced the earnings of workers at the bottom of the pre-recession distribution without affecting aggregate inequality if these effects were concentrated on workers who would have otherwise moved up the distribution. Therefore, to quantify the aggregate distributional consequences of the recession, we construct counterfactual wage distributions as in DiNardo, Fortin and Lemieux (1996) using a distribution regression analogue of our difference-in-differences design (Chernozhukov, Fernández-Val and Melly, 2013; Fernández-Val et al., 2024), which allows us to estimate how the wage distribution in more exposed places would have evolved had it followed a similar path to less exposed ones. We find that had the recession been less severe, the rise in wage inequality between 1979 and 1988 would have been up to 20 percent smaller. The severity of the recession lowered wages throughout the distribution, with the largest effects concentrated between the 20th and 60th percentiles, and a placebo

[2] However, we do find that employment recovers more for workers during the 1980s relative to the Great Recession.

exercise estimating the effects over the pre-recession period produces no comparable effects.

We next explore potential mechanisms for these labor market impacts. These lasting effects could reflect scarring for the cohorts that experienced the recession, in which case more and less exposed places should converge as those cohorts exit the labor force, or they could reflect persistent effects on the places themselves, such as changes to the local wage-setting environment, which would also impact future entrants. Repeating our analysis at the local labor market level using public-use Census microdata, we find that average earnings in more exposed places remain depressed nearly four decades later in 2017, when most of the workers in our panel had left the labor force. These earnings declines are primarily driven by lower wages rather than reduced hours. The dynamics of these effects also indicate that the recession changed the trajectory of these labor markets, rather than just impacting places that were always going to decline or accelerating structural changes that would have occurred regardless. These results demonstrate the remarkable persistence of the recession's impacts and suggest that the earnings effects are driven by persistent changes in the wage structure of these more exposed places.

We find evidence that these local wage-structure effects reflect the recession's effects on labor market institutions rather than on relative labor demand. If the recession had accelerated technological change in more exposed places, we would expect larger effects in routine occupations that were exposed to these changes. However, we find no evidence that the recession had larger effects on employment or earnings in these occupations. We reach a similar conclusion when looking at relative demand for skilled labor, where we find no evidence that the recession reduced the relative wage bill for non-college workers in more exposed places. However, we do find that the recession persistently weakened unions in these places. Each additional percentage point of exposure reduced union membership by 0.7 percentage points by 1989, an effect that operates mainly within industries and is neither explained by the mechanical effects of greater job loss nor by reduced local demand for union representation. The 1980s recession also appears to have been unusual in this respect since its job losses were concentrated in unionized industries to a greater degree than in other postwar recessions, and other recessions did not produce comparable declines in unionization. Extrapolating our estimates, the recession accounts for roughly 42 percent of the national decline in union density between 1979 and 1989.

Finally, we quantify how much of the recession's effect on aggregate wage inequality is accounted for by its effect on unions. We decompose the recession's total effect on inequality into union wage structure effects from spillovers (Fortin, Lemieux and Lloyd, 2021), union composition effects from the reallocation of workers from the union to the non-union sector, and other non-union changes in the wage structure. We estimate that the recession's effects

on unions account for between a third and 40 percent of its total effect on aggregate wage inequality, with both union channels contributing. The remaining share is driven by non-union related changes in the relationship between wages and worker characteristics. These estimates imply that the recession's effect on inequality was roughly 60 percent larger than it would have been absent the effects on unions.

Taken together, our results indicate that the recession contributed to the sharp rise in wage inequality during the 1980s, in part because it reshaped important wage-setting institutions. In particular, the steep decline in unionization that occurred over this period appears to have been partly a consequence of the recession itself. More broadly, our results suggest that the distributional consequences of a recession can depend on the institutional environment in which it occurs and how it impacts these institutions. The 1980s recession scarred workers in ways that resemble the effects of later recessions, but its effects on unions have no counterpart in the recessions that followed, and this union channel accounts for a substantial share of its lasting effect on the wage distribution.

We contribute to a large literature examining the evolution of labor market inequality in the U.S. and its underlying sources (Goldin and Margo, 1992; Katz and Murphy, 1992; Juhn, Murphy and Pierce, 1993; Card and DiNardo, 2002; Autor, Dube and McGrew, 2023). Work on the post-1980s rise has considered the roles of technology (Katz and Murphy, 1992; Autor, Katz and Kearney, 2008; Acemoglu and Autor, 2011), deindustrialization (Bluestone and Harrison, 1982), trade and globalization (Feenstra and Hanson, 1996; Greenland, Lake and Lopresti, 2026), and labor market institutions such as unions and the minimum wage (DiNardo, Fortin and Lemieux, 1996; Lee, 1999). We add to this literature by providing evidence that the recession contributed meaningfully to rising wage inequality in the 1980s, suggesting that this rise in inequality was not purely driven by gradual secular forces. A related literature argues that the compression of the mid-century wage distribution and its subsequent widening track the rise and decline of unions and other institutions that governed wage setting over the postwar period (Levy and Temin, 2007; Collins and Niemesh, 2019; Farber et al., 2021). Our results highlight that a large part of the weakening of unions during the 1980s was due to the impacts of the recession and that a substantial share of the recession's contribution to rising inequality is accounted for through that channel. We believe our paper is one of the first to show that transitory labor market shocks can persistently affect wage inequality by reshaping labor market institutions.

This paper also contributes to a large literature on the impacts of labor market shocks. Earlier work on job displacement (Jacobson, LaLonde and Sullivan, 1993; Davis and von Wachter, 2011) and sector-specific demand shocks (Autor et al., 2014) finds long-run negative impacts on workers. Similar work looking at shocks to local areas, therefore incorporating the effects

on the non-displaced, also finds negative impacts on the labor market outcomes of those local areas (Blanchard and Katz, 1992; Hershbein and Stuart, 2024) and the individuals residing in them (Yagan, 2019; Stuart, 2022). We use new data to study a large recession-induced labor market shock during a period with a different economic and industrial landscape relative to the present. We find similar long-run negative impacts on workers, which suggests that the labor market hysteresis of more recent shocks is not necessarily a modern phenomenon. Moreover, we provide evidence that the incidence of these losses varies across worker types, demonstrating how these transitory shocks can have important implications for employment and earnings inequality, despite the aggregate recovery. We also provide new evidence on how labor market shocks can reshape labor market institutions, highlighting an important channel that can amplify the labor market consequences of these shocks.

Finally, we contribute to a literature on the costs of recessions and disinflation (Lucas, 1987). Prior work on the Volcker disinflation has measured its cost through the sacrifice ratio, the cumulative loss of output per point of inflation removed (Okun, 1978; Gordon, 1982; Ball, 1994), and has considered how much of that cost depends on the credibility of the new policy regime (Sargent, 1982; Goodfriend and King, 2005). Our results suggest that there were substantial costs of the recession induced by the Volcker disinflation that were neither transitory nor evenly distributed, and part of these effects may be overlooked when focusing on aggregate quantities. These costs took the form of increased labor market inequality across places and workers and weakened labor market institutions. A more recent literature considers how the costs of monetary tightening are distributed across households (Coibion et al., 2017; Houtz, Pugsley and Rubinton, 2025). Since our cross-labor market design differences out the common effect of recessions across places and does not account for general equilibrium responses, it cannot measure the aggregate cost of the disinflation and so we view it as complementary to existing structural approaches.

The rest of the paper is organized as follows. Section II provides some background on the recession, its causes, and its consequences, and Section III presents a simple conceptual framework to organize our results. Section IV discusses the data used and the measurement and validation of local labor market exposure. Section V presents our empirical strategy and worker-level results, and Section VI describes our approach to quantifying the distributional consequences of the recession and presents our estimates. Section VII discusses potential mechanisms for these effects, while Section VIII decomposes the recession's effects on aggregate inequality. Finally, Section IX concludes.

II. Background: The 1980s Recession

The 1980s recession, especially its severity, was largely the product of the Federal Reserve's effort to end the inflation of the 1970s. After the 1979 oil shock, in which the Iranian revolution roughly doubled the price of oil, inflation exceeded 13 percent. In October 1979, under newly appointed chairman Paul Volcker, the Federal Reserve responded by adopting new operating procedures that allowed interest rates to rise as far as necessary to slow money growth (Goodfriend and King, 2005). The federal funds rate rose from 11 to over 17 percent within six months and the economy entered a recession in January 1980. The Federal Reserve then cut rates sharply and the economy recovered in July 1980, but inflation did not fall. Having lost credibility by reversing course (Volcker, 1994; Goodfriend and King, 2005), the Federal Reserve tightened again in November 1980 and pushed the funds rate to 19 percent by January 1981, holding it near that level for most of the year. As a result, a second and longer recession began in July 1981, and by the time policy eased in July 1982 the unemployment rate was approaching 10.8, the highest postwar level until the COVID-19 pandemic.

Narrative accounts of the period and the literature both attribute the severity of the downturn primarily to the Federal Reserve's tightening (Romer and Romer, 1989; Mussa, 1994; Goodfriend and King, 2005). Though the oil shock and inherited inflation made a downturn likely, the Federal Reserve's policy deepened and prolonged it (Mussa, 1994).[3] Each recession followed a tightening, and the brief recovery between them followed the easing of 1980. The recessions were deep, and their job losses were concentrated in the industries most sensitive to interest rates and, as high U.S. interest rates drew in capital and the dollar appreciated, to foreign competition (Branson and Love, 1988). Manufacturing accounted for one-fifth of employment but three-quarters of the net jobs lost,[4] and the industrial Midwest was the region most affected (Federal Reserve Bank of Chicago, 1982). We further characterize the incidence of the recession in Section IV.

Unions were central to how the Federal Reserve viewed the inflation it was fighting and the recession's incidence fell disproportionately on them. In 1980, 23 percent of wage and salary workers were union members, and density exceeded 30 percent in manufacturing, mining, construction, and transportation (Adams, 1985). Multi-year contracts in these industries set wage norms, and many contained cost-of-living clauses that raised wages automatically with prices. Past inflation therefore fed into current wages and, through costs, into future prices. Policymakers believed that this wage-price spiral was sustaining inflation (Mitchell and Erick-

[3] In Mussa's account of the Federal Reserve's reasoning, "to establish its credibility, the Federal Reserve had to demonstrate its willingness to spill blood, lots of blood, other people's blood" (Mussa, 1994).

[4] Goods-producing industries accounted for nine of every ten jobs lost in 1982, three-fourths of them in manufacturing (Urquhart and Hewson, 1983).

son, 2005). Volcker argued in December 1980 that inflation had become "institutionalized in three-year wage contracts" (Volcker, 1980), and members of the Federal Open Market Committee watched for concessions in major union contracts as a sign that the policy was working (Federal Open Market Committee, 1981; Mitchell and Erickson, 2005).

The recession had a substantial impact on unions, with some arguing that it was a key factor behind the dismantling of the postwar arrangement between business and labor that characterized the preceding decades (Levy and Temin, 2007).[5] By Volcker's own account, weakening the union position was an important part of how the disinflation worked. He later called Reagan's firing of the striking air traffic controllers in August 1981 "one of the major factors in turning the tide on the inflationary situation" (Volcker, 2000), and by mid-1982 the Board's staff was reporting that labor market slack had "paid handsome dividends" in slowing wage growth (Mitchell and Erickson, 2005). The share of workers under newly negotiated major private settlements accepting a first-year wage freeze or cut rose from zero in 1980 to 44 percent in 1982 (Mitchell, 1985), and between 1980 and 1984 the labor movement lost 2.7 million members, 1.9 million of them in goods-producing industries where non-union employment instead rose by 1.1 million. The recession hit hardest where unions were strongest, and the jobs returning in the recovery were disproportionately non-union (Adams, 1985).[6]

By late 1987 the unemployment rate had returned to its 1979 level. What remained was a permanent decline in union density, which would have taken twenty-one years at the pace of the 1970s,[7] and a sudden rise in wage inequality that did not reverse. This paper asks whether and how the recession contributed to both.

III. Conceptual Framework

To organize our results and think about how the recession may have impacted inequality, we use a simple model of a local labor market with monopsonistic employers and a union sector that threatens non-union employers. Appendix F formally discusses the model and results in detail.

Firms differ in their productivity p, and since workers do not move costlessly to the highest paying employer, each firm faces an upward-sloping labor supply curve with elasticity ε and

[5] Levy and Temin (2007) refer to this arrangement, in which unions traded labor peace for wages indexed to prices and productivity, as the "Treaty of Detroit" after the 1950 agreement between General Motors and the United Auto Workers, and date its dismantling to the combination of Volcker's monetary policy and the Reagan administration's decisions of 1981.

[6] Adams (1985) documents this pattern industry by industry. He attributes the shift to employers responding to the downturn by moving work to non-union facilities, contracting out, and buying inputs from non-union or foreign suppliers, alongside import competition and the deregulation of trucking and airlines.

[7] Union density fell from 25.4 percent in 1979 to 17.9 percent in 1987. A linear trend fitted to 1969–1979 declines at 0.35 points a year, at which rate the 1979–87 decline would have taken 21 years.

pays a share $\varepsilon/(1+\varepsilon)$ of its marginal product (Manning, 2003; Card et al., 2018). A fraction ψ of firms are unionized and must pay at least a wage floor $\bar{w}$ set outside the local market by pattern bargaining. This floor binds at low-productivity union firms, so union workers at low- and middle-productivity firms all earn $\bar{w}$ (Freeman, 1980). Non-union firms face a threat of unionization that rises with local union density and falls with their own wage, so they pay a premium to keep their workers from organizing (Rosen, 1969; Freeman and Medoff, 1981). Thus, for any two non-union firms with productivities $p_L < p_H$, the wage gap is

$$\ln w_L - \ln w_H = \frac{\ln(p_L + \Delta) - \ln(p_H + \Delta)}{1+\varepsilon}, \qquad \Delta = \frac{\psi T \kappa}{\varepsilon} \tag{1}$$

where T is the firm's loss in profit if its workers organize and κ measures how easily they can organize. The firm-level labor supply elasticity ε determines how much of the productivity gap between firms passes through to wages (Autor, Dube and McGrew, 2023), while the threat of unionization Δ acts like a common increase in productivity that raises wages the most at low-productivity firms.

We model the recession as a temporary fall in every firm's productivity, $p \to \lambda p$ with $\lambda < 1$, which causes the least productive firms to exit. We assume that exit is irreversible, since re-entry requires sunk investments (Dixit, 1989), and that competition for workers depends on the number of employers, so ε falls when firms exit.[8] Union firms are more likely to exit, since the floor makes them less profitable than non-union firms with the same productivity. This means that after productivity recovers, the local labor market has fewer employers and fewer unions, which affects wages through three channels that are illustrated in equation (1). First, with less competition, surviving firms pay a smaller share of their marginal product, so wages fall and fewer workers participate. From equation (1), more of the productivity differences between firms also passes through to wages, so the gap between low-wage and high-wage workers increases. Employment losses are also concentrated among low-wage workers, since the least productive firms are the ones that exit. Second, lower union density ψ lowers threat effects Δ, which lowers wages by relatively more at low-productivity non-union firms. Finally, displaced union workers leave the compressed union sector and some are re-employed in the non-union distribution.

The model also shows how unions interact with other features of the wage-setting environment, such as employer market power, to mediate the effects of the recession. A union sector that survives the recession dampens inequality since, holding union density fixed, the

[8] This can occur because fewer employers reduce the rate at which workers receive outside offers (Burdett and Mortensen, 1998) or raise each remaining firm's share of the local labor market (Berger, Herkenhoff and Mongey, 2022).

threat premium rises when competition weakens, because $\Delta = \psi T\kappa/\varepsilon$ rises. The incentive to raise wages to attract workers falls with ε, but the incentive to prevent organizing does not, so the threat becomes a larger part of why non-union firms pay above their monopsony wage. Since the threat premium is largest at low-productivity firms, this offsets part of the increasing dispersion in wages. When the recession instead weakens the union sector, these effects are amplified. Workers move from the relatively compressed union distribution into a non-union distribution with rising dispersion from weaker competition. At the same time, the threat premium that was reducing wage markdowns falls with lower union density. Therefore, the recession's impacts on the union sector can increase inequality directly and by exacerbating non-union channels like employer wage-setting power.

IV. Data and Recession Exposure

We compile several public-use and administrative datasets to study the effects of the 1980s recession on local labor markets and workers. This section focuses on the new longitudinal employer-employee dataset we compiled to study the recession as well as our measure of local exposure to the recession.[9]

IV.A WORKER PANEL

Main Sample

To analyze the worker-level effects of the recession, we construct a new worker-level panel spanning 1976–1989 by combining Social Security Administration (SSA), IRS, and other administrative records housed at the Census Bureau.

The main sample consists of individuals that appear either in the Current Population Survey (CPS) between 1973 and 2024, or the Survey of Income and Program Participation (SIPP) between 1984 and 2024, or the American Community Survey (ACS) between 2015 and 2019. For each of these individuals we observe their full earnings histories that are reported to the Social Security Administration (SSA). Specifically, we observe information from the SSA's Summary Earnings Record from 1951 to the present and the Detailed Earnings Record from 1978 to the present. The Summary Earnings Record contains FICA-eligible[10] employment and earnings information aggregated across employers, while the Detailed Earnings Record contains information for each W-2 form filed for that worker (including employer EINs).

[9] Further details on other supplemental data sources are provided when they are used and in the Data Appendix (Appendix A).

[10] The U.S. payroll tax that funds programs like Social Security and Medicare.

We further restrict this sample to individuals that are working-age during our sample period (i.e., 23–54 in 1979) and that did not die prior to 1989.[11] In order to assign individuals to their recession exposure, we must observe their location of residence prior to the recession. We only observe residential information for the bulk of this sample through income tax return (Form 1040) filings. Therefore, we further restrict the sample to individuals that filed an income tax return in 1979 and we use the residential information on that tax return to assign workers to their pre-recession local labor market. Lastly, to focus on workers that had some pre-recession labor force attachment, we restrict to individuals that reported positive wage earnings to the SSA in at least one year between 1976 and 1979. The final sample consists of around 5.3 million individuals.

Main Outcomes

Our main labor market outcomes of interest are employment and earnings. Following Davis and von Wachter (2011), Autor et al. (2014), and Yagan (2019) we define *employment* in a given year as an indicator for having positive FICA-eligible wage earnings for that year. This annual employment definition is an extreme extensive margin measure of having been employed at any time during the year. To capture short-term intensive margin employment adjustments, we construct a secondary measure of whether the individual has *high labor force attachment*, which we define as having annual earnings that exceed the annual earnings of a worker working 1600 hours (forty weeks at forty hours per week) at the prevailing federal minimum wage.

We also examine annual wage *earnings*, which we measure in two ways. Our primary measure is the annual earnings information reported in each individual's W-2 tax forms. This earnings type is only observable starting in 1978, which means we only have two pre-recession years (1978 and 1979) to observe earnings pre-trends. Therefore, we supplement this with an alternative measure of annual earnings reported in the Summary Earnings Record. These only correspond to earnings that are subject to social security taxes and are topcoded at the annual maximum taxable amount.[12] All dollar values are inflated to 2015 dollars. Other variables used in the main analysis are defined when introduced and discussed in Appendix A.1.

Summary Statistics

Table 1 compares individuals in our main analysis sample to those in the same birth cohorts (1925 to 1956) that appear in the Current Population Survey March (ASEC) Supplement in 1979 and 1989. Our main sample is broadly representative of the overall population compared

[11] Date of death is obtained from the SSA Numident.

[12] In 1979 this was $22,900 (~$80,000 in 2015$) or roughly the 90th percentile of the annual earnings distribution.

to the CPS. However, the sample is selected along a few dimensions for several reasons. First, our starting sample is individuals that ever appear in one of the CPS (1973–2024), SIPP (1984–2024), or ACS (2015–2019) for whom we can observe SSA earnings histories. The probability of appearing in one of these surveys is decreasing in age in 1979. A 24-year-old in 1979 is much more likely to appear in one of these surveys over the next 40 years, and therefore enter our sample, compared to a 54-year-old in 1979. Thus, our sample is younger than the CPS population with an average age of 33.3 years versus 36.8. As expected, the disparity is greatest for those older than 45 in 1979 (9.9 percent of our sample compared to 23.3 for the CPS). Second, we restrict to workers that had positive earnings in at least one year between 1976 and 1979 and that filed a tax return in 1979. Both of these conditions will tend to select workers that are more attached to the labor force and have higher earnings. The share with positive wage earnings in 1979 in our sample is 91.2 percent relative to 74.7 in CPS and the average earnings in 1979 in our sample is slightly higher (35,530 versus 32,277). As a result of this positive labor market selection, we also overrepresent the demographic characteristics associated with those outcomes. Our sample is whiter (85.8 percent versus 81.9 percent) and has fewer individuals with less than a college education (71.3 percent versus 79.6 percent).

To address this selection, we reweight our sample to match the CPS March Supplement in 1979 based on the demographic and labor market characteristics listed in Table 1. Specifically, we reweight our sample to match the age, race, education, and earnings (including percent with no earnings) distributions of the corresponding birth cohorts in the CPS March Supplement. We also provide unweighted estimates for our main results in Appendix C.

IV.B LOCAL RECESSION EXPOSURE

Our local labor markets are commuting zones (CZs), which are clusters of counties defined based on commuting patterns.[13] Our analysis focuses on 722 commuting zones covering the contiguous United States. The majority of the data used in our analysis are available at the county level, which we then aggregate to the CZ level using the crosswalk from Autor, Dorn and Hanson (2013).

Following Yagan (2019), we measure exposure to the 1980s recession using the change in the CZ unemployment rate between 1979 and 1982, which we view as a reduced-form measure of the local labor demand shock associated with the recession. This also allows us to directly compare our estimates to the Great Recession, whose aggregate unemployment shock was comparable in magnitude. We measure CZ unemployment using the Bureau of Labor Statistics Local Area Unemployment Statistics (LAUS).[14]

[13] We use 1990 commuting zone definitions.

[14] Official county-level LAUS estimates only begin in 1990, so we use unofficial BLS county estimates available

Exposure to the 1980s recession varied substantially across both industries and local labor markets in the United States. The average CZ experienced a 4.18 percentage point increase in unemployment with an interquartile range of 3.17 percentage points, ranging from negligible increases to roughly 11 percentage points in the most exposed CZs. Figure 1A shows that exposure was concentrated in the Midwest and Mid-Atlantic, the Pacific Northwest, and parts of the South, but also varied substantially within regions. This partly reflects differences in industrial composition and in sensitivity to the Federal Reserve's monetary policy shocks.[15] Figure 1B displays the change in unemployment across industries. The most exposed industries are in manufacturing, mining, and construction, although there is also substantial variation within these sectors. For instance, unemployment rose by 20 percentage points in the blast furnace and steel industry compared to around 2 percentage points in the ordnance industry, which was insulated by steady military demand.

Although we attribute these changes in unemployment to the recession, they may reflect other contemporaneous factors. Table B.1 shows which pre-recession characteristics of CZs predict exposure. The 1979 manufacturing share is the largest determinant, and broad industrial composition alone explains around one-third of the variation across labor markets, while manufacturing employment growth in the preceding decade is not predictive. Therefore, our exposure measure could pick up pre-existing trends in these industries, such as deindustrialization. However, Table 2 shows that once we condition on the 1979 manufacturing share, exposure does not predict pre-recession changes in the manufacturing share, population, employment, earnings, or union density. This suggests that our exposure measure is not picking up places that were already declining or mean reversion from a pre-recession local boom.

V. Worker Level Effects of the 1980s Recession

For the recession to have contributed to a persistent rise in inequality during the 1980s, this temporary shock must have had persistent effects on some workers despite the aggregate recovery, and those effects must have been relatively larger for workers at the bottom of the wage distribution. This section examines this by estimating the impacts of local exposure to the 1980s recession on individuals' labor market outcomes.

starting in 1976. Appendix A.2 describes these data and validates them against other sources.

[15] As further validation, we show in Appendix B that each CZ's sensitivity to monetary policy shocks based on industry composition is strongly predictive of our exposure measure, consistent with the fact that the recession was exacerbated by the Federal Reserve's efforts to combat inflation.

V.A EMPIRICAL STRATEGY

We adopt an empirical strategy similar to earlier work that uses administrative data to examine the impacts of labor market shocks (Jacobson, LaLonde and Sullivan, 1993; Davis and von Wachter, 2011; Autor et al., 2014; Yagan, 2019; Stuart, 2022). We estimate variations of the following regression:

$$\Delta y_{it} = \beta_t \Delta \text{Unemp}_{c(i,79)} + \theta_{g(i,79)t} + \mathbf{X}_i' \Gamma_t + \mathbf{W}_{c(i,79)}' \Pi_t + \varepsilon_{it} \tag{2}$$

In our main analysis, the outcomes are demeaned relative to the individual's pre-recession mean: $\Delta y_{it} = y_{it} - \bar{y}_{i,\text{pre}}$ where y_{it} denotes the outcome of interest for individual i in year t.[16] $\Delta \text{Unemp}_{c(i,79)}$ is our 1980s recession exposure measure, and $\mathbf{X}_i$ are time-invariant individual wage determinants such as race, sex, and education. The $\theta_{g(i,79)t}$ terms denote fixed effects for the individual's group g, which are interactions between age, industry (4-digit NAICS), and earnings decile all measured in 1979. In our robustness exercises, we also control for $\mathbf{W}_{c(i,79)}$ which are pre-recession characteristics of the individual's 1979 commuting zone such as the manufacturing share of employment in 1979. We run this specification separately in each year t, so the coefficients on all terms are allowed to vary over time. Each regression therefore compares changes in y relative to pre-recession baselines across workers facing different local recession exposure. The coefficients of interest are the β_t coefficients, which together trace out the effect of recession exposure over time.

The β_t identify the average causal effect of exposure under a parallel trends assumption that conditional on the detailed fixed effects and individual characteristics, the outcomes of workers facing different levels of exposure would have followed parallel paths in the absence of the shock. Since our outcomes are relative to the worker's own baseline, any fixed worker or location determinants of y that are correlated with exposure, such as ability or location premia, are differenced out and do not bias our estimates of β_t. The main threat to identification is pre-existing differential trends for workers in exposed areas, for instance due to the secular decline of manufacturing. The group fixed effects help mitigate this concern by comparing similar types of workers that face similar nationwide shocks. We further assess this threat by examining the pre-recession coefficients for evidence of any confounding pre-trends.

[16] The pre-recession mean is typically measured between 1976–1979. However, some outcomes are only available starting in 1978, in which case the pre-recession mean is measured between 1978–1979. See Section IV.A for details on which outcomes are available over which years.

V.B EFFECTS ON EMPLOYMENT AND EARNINGS

We find that the recession persistently reduced employment for workers in more exposed labor markets, well after the aggregate recovery. Figure 2 displays the estimated effects of local exposure to the 1980s recession on our measure of employment between 1976 and 1989. The coefficients prior to the recession are centered around zero and show no evidence of any pre-recession negative employment trajectories, consistent with our identification assumption. The estimated effects become more negative during the recession period (1979–1982), reaching a nadir of −0.46 in 1983, and then attenuate to −0.22 by 1989, so that a worker exposed to a 1 percentage point larger increase in local unemployment remains 0.22 percentage points less likely to be employed a decade after the recession. Put differently, the long-run decline in employment is about a fifth of the size of the recession's initial unemployment shock. Scaled by the interquartile range of exposure, this implies a decline in the probability of employment of 0.7 percentage points, which is substantial relative to the 3 percentage point increase in the aggregate employment-to-population ratio over the same period.

Unemployment spells that last less than a year, even long ones, leave our annual employment measure unchanged, so we also estimate effects on our high labor force attachment measure. Figure 2 displays these estimates alongside the coefficients on employment. The pattern is similar but the effects are larger, with each additional percentage point of local unemployment reducing the probability of high labor force attachment by 0.37 percentage points in 1989. That we find negative impacts on high labor force attachment indicates that our employment effects are not just driven by the displacement of workers with weak labor force attachment to begin with. The gap between the two estimates is consistent with some workers moving from high to low attachment status, perhaps due to longer or more frequent unemployment spells. However, since we do not observe labor supply directly, this comparison should be interpreted with caution.

When looking at earnings instead, we also find that the recession reduced long-run earnings for those that remain employed with effects that are even more persistent than what we find for employment. Figure 3 displays the estimated effects of exposure on our log annual wage earnings outcome. Similar to our employment estimates, the coefficients display no pre-trends prior to the recession period and decline during the recession years. Unlike employment, the estimated effect on earnings shows no signs of recovery and remains persistently depressed following the recession. Our estimates indicate that exposure to each additional point of unemployment reduced earnings by 1.9 percent. Scaling by the interquartile range of exposure, cumulative 1983–1989 earnings fall by around $14,000, or roughly half a year of earnings.

Since log earnings are only observed for employed workers, these estimates must be interpreted alongside the disemployment effects documented above. If disemployment falls disproportionately on low-wage workers, a possibility we examine in Section V.C, the earnings estimates understate the true effect, while the reverse is true if high-wage workers are displaced. The trajectory of the two sets of estimates is informative about the importance of this selection effect. Under the first case, as employment recovers and low-earnings workers return to work, the measured earnings effect should become more negative. We observe no such movement, as the earnings estimates are stable while the employment estimates attenuate by half, which we interpret as evidence that our earnings effects are not primarily driven by selective disemployment.

These results demonstrate that greater exposure to the recession had lasting effects on workers. Two features of the results are worth highlighting. First, the negative impacts extend well beyond the aggregate recovery. National unemployment reached its 1979 levels by 1987, yet workers in more exposed places remained less likely to be employed and earned less. Second, the employment effects partially recover by 1989, with the effect on any employment halving from its trough, whereas the earnings effects do not, suggesting a persistent change in how workers are compensated, not only whether they work. These long-lasting effects are surprising relative to what standard models of local adjustment predict (Blanchard and Katz, 1992) and contemporaneous accounts of the period, which reserved the hysteresis diagnosis for Europe relative to the resilient U.S. labor market (Blanchard and Summers, 1986). However, our results are consistent with more recent evidence of the persistent impacts of recession shocks on individuals (Kahn, 2010; Yagan, 2019; Stuart, 2022) and suggest that labor market hysteresis may be a common feature of recessions. This persistence provides the first piece of evidence that this temporary shock could have contributed to a persistent rise in labor market inequality.

V.C EFFECTS ACROSS THE EARNINGS DISTRIBUTION

The previous section established that the recession had lasting impacts on average for workers. This section asks whether these lasting effects were relatively larger for workers at the bottom of the earnings distribution. We estimate our main specification (equation (2)) separately for workers with earnings in each decile of the pre-recession earnings distribution, pooling the years 1987 to 1989.[17] Within each decile we continue to control for 1979 age-industry cells, so each comparison holds these other dimensions fixed. The differences we document across groups therefore net out how those groups are sorted across industries and the earnings dis-

[17] Earnings are measured in 1979 and we include a decile 0 category for workers with zero earnings.

tribution. As a result, these estimates are likely conservative, and any measure of the total incidence should also account for any concentration in more exposed industries.

We find that the labor market impacts of greater local recession exposure are substantially larger for workers at the bottom of the pre-recession distribution. Figure 4A displays the estimated effect of exposure on the long-run (1987–1989) probability of being employed separately for workers in each decile of the pre-recession earnings distribution. The long-run estimated effect on employment is negative across all wage deciles, but the effects are largest at the bottom deciles. The estimates indicate that the impact on the probability of being employed is roughly three and a half times as large for a bottom decile worker (−0.46 p.p.) compared to a top decile worker (−0.13 p.p.). Running a joint regression with all workers and interacting our exposure measure with indicators for the worker's pre-recession wage quintile, we can reject that the effects are the same in the top and bottom quintiles.

We find a similar pattern when looking at earnings instead. Figure 4B displays the same estimates for log earnings and the estimates are 1.7 times as large for a bottom decile worker (−0.025) compared to a top decile worker (−0.015) and a joint regression interacting exposure with pre-recession wage quintile indicators rejects equality of the top and bottom quintile effects. The earnings gradient is flatter than the employment gradient, where bottom decile workers are affected roughly three and a half times as much as those at the top. The pattern of employment selection works against finding an earnings gradient. Since employment losses are concentrated at the bottom, the measured earnings effect is biased most where we estimate the largest declines, suggesting the true earnings gradient is likely steeper than we estimate.[18]

These results indicate that local recession exposure increased both employment and earnings inequality across the pre-recession earnings distribution. Our results are consistent with prior work (Hoynes, Miller and Schaller, 2012) demonstrating that the unemployment rates of low-wage workers are more sensitive to business cycle fluctuations, but they further highlight that these unequal impacts are persistent, extending beyond the temporary fluctuation. We discuss additional heterogeneity across non-wage margins in Appendix C. Together, our results on the persistence of the recession's impact and its larger incidence for low-wage workers establish the conditions under which the recession could have contributed to rising aggregate wage inequality, which we investigate further in Section VI.

[18] This assumes that workers are also positively selected on employment within earnings deciles.

V.D COMPARISON TO THE GREAT RECESSION AND ROBUSTNESS

Comparison to the Great Recession

The Great Recession offers a natural benchmark to contextualize these effects, since it was comparable in severity. Using a similar empirical setup and definition of local exposure, Yagan (2019) finds that each additional percentage point of local unemployment during the Great Recession reduced the probability of employment by 0.39 percentage points in 2015, eight years after the start of the recession. Our employment estimates are of comparable magnitude at their trough, 0.46 percentage points in 1983, but attenuate to 0.22 by 1989, roughly half the persistent effect documented for the Great Recession at a similar horizon. The earnings effects are more similar. Each additional percentage point of exposure reduces cumulative earnings over the seven post-recession years (1983–1989) by $4,419, against $6,212 for the Great Recession (2009–2015).[19] The major difference is that employment recovered partially in the 1980s relative to the Great Recession. Yagan (2019) also finds similar patterns of heterogeneity in employment and earnings for workers with lower pre-recession earnings. Despite the vastly different contexts between the two recessions, they produced remarkably similar impacts on workers, suggesting that the hysteresis following the Great Recession was not unique to that episode.

Robustness

Our estimates are not sensitive to how we control for worker characteristics or to a wide range of pre-recession commuting zone characteristics, and they are not driven by certain regions or types of workers. Table 3 reports the 1989 earnings estimate across various alternative specifications, with coefficients ranging from −0.012 to −0.029, roughly centered around our main estimate of −0.019.[20] Columns 1–4 vary the worker-level controls, addressing the concern that exposed places contained different types of workers, for instance less-educated workers who also faced falling demand nationally. Clustering at the state level (column 3) leaves inference unchanged. Columns 5, 6, and 8 address the concern that more exposed places were declining for reasons unrelated to the recession, such as deindustrialization, technological change, import competition, or the lingering effects of earlier recessions, by controlling for various commuting zone characteristics. Column 7 adds state fixed effects, addressing concerns that changes in state policy unrelated to the recession, such as deregulation, drive our estimates. Including all of these together (column 9) gives a coefficient of −0.012. Column 10 drops workers

[19] Both dollar amounts are reported in 2015$ and reflect total earnings in dollars, which includes individuals with zero earnings.

[20] Results for employment are similar and reported in Appendix C.

in sectors that are sensitive to the business cycle (manufacturing, mining, and construction), and the estimate is unchanged, indicating that the results are not just driven by workers in cyclical sectors. Finally, column 11 interacts exposure with regional indicators. The effects are negative in every region, largest on the West Coast and smallest in the South, indicating that no single region drives the result.

VI. The Recession's Impacts on Aggregate Wage Inequality

The previous section established that the recession had lasting effects on workers' labor market outcomes that were largest for those with lower earnings. These results point to the recession widening the earnings distribution, but they cannot tell us how much it contributed to aggregate wage inequality for several reasons. First, they are average effects for workers ranked by their pre-recession earnings. If workers with low pre-recession earnings would otherwise have moved up the distribution, a loss concentrated among them lowers their average earnings without affecting the bottom of the post-recession distribution. Second, they are average effects within pre-recession age-earnings-industry cells, whereas the effect on the overall distribution depends on how the entire distribution of wages shifts within each cell.[21] Finally, they are effects per unit of exposure and are not informative about the size of the recession's contribution relative to the overall rise in inequality. Therefore, in this section we estimate the recession's impact on the aggregate hourly wage distribution directly by constructing counterfactual distributions as in DiNardo, Fortin and Lemieux (1996). Specifically, we estimate the hourly wage distribution in 1988 had the distribution in states with above-median increases in unemployment evolved similarly to the distribution in states with below-median increases.

VI.A Estimating the Counterfactual Distribution

Let $S \in \{0, 1\}$ indicate whether a worker lives in a state with an above-median increase in unemployment between 1979 and 1982 (hereafter, severely exposed), let $T \in \{0, 1\}$ index the pre- and post-recession years 1979 and 1988, and let Y denote the worker's log hourly wage.[22] The object of interest is the post-recession distribution of wages that would have prevailed had the distribution in high-exposure states followed a similar trajectory to the low-exposure one. Letting $Y(s)$ denote the worker's potential outcomes under each scenario, we can write this

[21] The regression also weights cells by their variance in exposure rather than by their population shares.

[22] Here, we discretize exposure to make the interpretation and exposition simpler, though this exercise can be extended to use a continuous measure of recession exposure.

counterfactual as:[23]

$$F^{c}_{Y|T}(y \mid 1) = F_{Y(0)|S,T}(y \mid 1,1)\Pr(S=1 \mid T=1) + F_{Y|S,T}(y \mid 0,1)\Pr(S=0 \mid T=1) \quad (3)$$

where we replaced the observed distribution of wages for high-exposure workers $F_{Y(1)|S,T}(y \mid 1,1)$ with the low-exposure counterfactual $F_{Y(0)|S,T}(y \mid 1,1)$, which is not observed. We interpret this as a less severe recession counterfactual, in which the severity of exposure across places matched that of less exposed labor markets. This is not a no-recession counterfactual since less exposed labor markets still experienced a downturn and our cross-sectional comparisons net out the common impact of the recession, likely understating its effect. Moreover, this exercise does not directly account for any general equilibrium responses.

To identify the distribution of wages high-exposure workers would have earned under the low-exposure counterfactual, we use a distribution regression analogue of our main difference-in-differences design (Chernozhukov, Fernández-Val and Melly, 2013; Fernández-Val et al., 2024). This DR-DID approach generalizes the difference-in-differences strategy to estimate effects on conditional distributions. We can write the conditional distribution of potential wages as:[24]

$$F_{Y(0)|S,T,X}(y \mid s,t,x) = \Lambda\big(\alpha(y,x) + \beta(y,x)t + \gamma(y,x)s + \delta(y,x)st\big)$$

where Λ is an invertible link function and $\alpha(y,x), \beta(y,x), \gamma(y,x)$, and $\delta(y,x)$ are parameters governing its evolution.

The main identifying assumption is that:

Assumption A1 (Distributional parallel trends). *For each (y,x) in the relevant support, we assume that*

$$\begin{aligned}\Lambda^{-1}\Big[F_{Y(0)|S,T,X}\big(y \mid 1,1,x\big)\Big] - \Lambda^{-1}\Big[F_{Y(0)|S,T,X}\big(y \mid 1,0,x\big)\Big] \\ = \Lambda^{-1}\Big[F_{Y(0)|S,T,X}\big(y \mid 0,1,x\big)\Big] - \Lambda^{-1}\Big[F_{Y(0)|S,T,X}\big(y \mid 0,0,x\big)\Big]\end{aligned}$$

i.e. $\delta(y,x) = 0$.

Assumption A1 is the distributional analogue of the parallel trends assumption maintained throughout the paper, but it is stated on a transformation of the distribution of potential outcomes and therefore depends on the link function. Under the logistic link we use, it requires that absent severe exposure the log-odds of earning less than y among workers with charac-

[23] Here $F_{A|B}(a \mid b)$ denotes the conditional distribution of A given $B = b$.

[24] Note that this formulation does not impose any parametric assumptions on the underlying distribution. See Fernández-Val et al. (2024) for details.

teristics x would have changed by the same amount in high- and low-exposure states. The absence of pre-trends in Section V supports this assumption at the mean, and we assess it further below with a placebo exercise over the pre-recession period. Under A1 and a common support assumption, the counterfactual conditional distribution is identified by a difference-in-differences estimator (Fernández-Val et al., 2024):

$$F_{Y(0)|S,T,X}(y \mid 1,1,x) = \Lambda\Big[\Lambda^{-1}\big[F_{Y(0)|S,T,X}(y \mid 1,0,x)\big] + \Lambda^{-1}\big[F_{Y(0)|S,T,X}(y \mid 0,1,x)\big] - \Lambda^{-1}\big[F_{Y(0)|S,T,X}(y \mid 0,0,x)\big]\Big]$$

We follow the literature (Chernozhukov, Fernández-Val and Melly, 2013; Fernández-Val et al., 2024) and estimate this counterfactual by fitting the following distribution regression under a logit link, separately at 981 thresholds spanning the support of log wages y

$$F_{Y|S,T,X}(y \mid s,t,x) = \Lambda\big(p_\alpha(x)'\alpha(y) + p_\beta(x)'\beta(y)t + \gamma(y)s + \theta(y)st\big) \tag{4}$$

where $p_\alpha(x)$ and $p_\beta(x)$ are vectors of covariates and their transformations and $\alpha(y)$, $\beta(y)$, $\gamma(y)$, and $\theta(y)$ are the corresponding coefficient vectors that vary at each threshold. The counterfactual conditional distribution $\widehat{F}_{Y(0)|S,T,X}(y \mid 1,1,x)$ is estimated from these fitted values, setting $\theta(y)$ to zero. The counterfactual unconditional distribution $\widehat{F}_{Y(0)|S,T}(y \mid 1,1)$ is the average of these fitted values over the empirical distribution of X among high-exposure workers in 1988. We then recover the unconditional distribution $\widehat{F}^c_{Y|T}(y \mid 1)$ using equation (3) with the observed exposure shares $\Pr(S = s|T = 1)$.

We estimate these counterfactual distributions on the CPS data assembled by Fortin, Lemieux and Lloyd (2021), so that our estimates are comparable to similar decompositions in the literature. The data combine the Merged Outgoing Rotation Group files for 1979–1988 with the May supplements, which provide union status before 1983, and are described further in Appendix A.6. We restrict our attention to men aged 16–64 and we measure hourly wages directly for hourly workers and as usual weekly earnings divided by hours for others. We also drop workers with allocated wages to avoid known biases from these imputations and we weight by the CPS sampling weights. The distribution regressions control for schooling, experience, and their interaction, marital status, race, industry, occupation, part-time status and metropolitan residence, together with state, year and quarter fixed effects and 1979 state industry shares and 1979 state-by-industry union coverage, both interacted with year. The full specification is listed in Table 5 and discussed further in Appendix E.1. Since the CPS identifies state but not CZ of residence, exposure in this section is measured using state-level changes in unemployment.

VI.B ESTIMATES OF THE DISTRIBUTIONAL EFFECTS OF THE RECESSION

Table 4 describes the change in the wage distribution for men between 1979 and 1988. As is well documented, measures of overall wage inequality for men rose over this period. The 90–10 log wage gap increased by 0.221 log points, with 0.136 of the rise coming from the top half of the distribution and the remaining 0.085 from the bottom. This rise was larger in high-exposure states on every measure except the 50–10 gap, which rose by the same amount in both groups. The 90–10 gap rose by 0.24 log points in high-exposure states against 0.21 in low-exposure states, despite starting from a lower level. We next investigate the extent to which these descriptive patterns can be attributed to the impacts of the recession.

Greater exposure to the recession lowered wages throughout the distribution, with the largest losses in the bottom half. Figure 5A displays the difference between the quantiles of the observed and the counterfactual wage distribution for men in 1988. The largest negative effects are at the bottom, where wages fall by around 8 percent at the 20th percentile against a 3 percent decline at the 90th, with the effect declining across the distribution. The shape of the estimates is similar to our worker-level heterogeneity in Section V.C, which is consistent with wage declines falling on workers who remained near the bottom of the distribution rather than workers who would otherwise have moved up it, although the two are estimated on different data and populations.

Table 5 converts these quantile effects into effects on different measures of wage inequality. Relative to the less severe counterfactual, the rise in the 90–10 gap between 1979 and 1988 would have been 13.5 percent smaller, the rise in the wage Gini coefficient 19 percent smaller, and the rise in the 80–20 gap 20 percent smaller. The effect is concentrated in the upper-half ratios, where the rise in the 90–50 gap would have been 18 percent smaller against 6 percent for the 50–10 gap, with roughly 80 percent of the recession's contribution to the 90–10 gap operating through the 90–50 gap. These are estimated effects on the national distribution, which dilutes the effect on high-exposure states. For high-exposure states themselves, the rise in the 90–10 gap would have been 25 percent smaller and the rise in the 90–50 gap 29 percent smaller.

To assess our distributional parallel trends assumption, we perform a placebo exercise estimating the same model over the pre-recession period from 1977 to 1979 instead. Figure 6A and Table E.1 display the quantile effects and estimates on measures of wage inequality. The placebo quantile effects are centered around zero and the estimated effect on the 90–10 gap is −0.009 (0.016) and indistinguishable from zero. Overall, we find no evidence of any effects on the pre-recession distribution, consistent with our identifying assumption.

These results indicate that the recession reduced wages throughout the distribution, with

the largest losses at the bottom of the distribution, mirroring the worker-level heterogeneity documented in Section V.C. Together, they suggest that the recession was a meaningful contributor to rising wage inequality during the 1980s. We next explore several mechanisms that can explain why the recession had such a persistent impact on inequality.

VII. Mechanisms

The previous sections established that the recession had persistent effects on workers' earnings and on the aggregate wage distribution that were still present years after the recession ended. This persistence is surprising in light of standard models of local adjustment. In these models, a temporary shock to local labor demand should dissipate through entry and migration, so that more and less exposed places eventually converge (Blanchard and Katz, 1992). The documented lack of convergence implies that something about more exposed places changed persistently. We consider two broad explanations in this section. The first is that the recession hit places that were already in relative decline due to longer-run secular shifts in labor demand, so that the gaps we document would have emerged regardless. The second is that the recession itself changed these places in ways that did not reverse once aggregate demand recovered, either through persistent changes in local labor demand or the wage-setting environment.

VII.A THE RECESSION'S IMPACT: PEOPLE VERSUS PLACES

The persistent effects documented in Section V are consistent with two different explanations that make different predictions at longer time horizons. The first is worker-level scarring, where displaced workers lose job-specific rents and firm-specific human capital, and general human capital decays over long spells of unemployment (Jacobson, LaLonde and Sullivan, 1993; Davis and von Wachter, 2011; Schmieder, von Wachter and Heining, 2023). Under this explanation, the effects of the recession are concentrated among the cohorts that experienced it, and we should expect more and less exposed places to converge as these cohorts exit the labor force and as workers migrate (Blanchard and Katz, 1992). The second is that the recession persistently changed these local labor markets themselves by altering the local structure of wages, with consequences both for workers who were never displaced and for those who entered these labor markets long after the recession. Under this explanation, the effects should persist even as exposed cohorts exit the labor market.

To distinguish between these explanations, we repeat our worker-level analysis at the commuting zone level using public-use Census microdata, which allow us to extend the analysis to 2017 and to separate earnings into wages and hours worked. We estimate the following

long-difference regressions that mirror our worker-level specification:

$$\Delta y_{ct} = \alpha + \beta \Delta \text{Unemp}_c + \mathbf{X}'_{ct}\Gamma + \delta_{d(c)} + \varepsilon_{ct} \tag{5}$$

where Δy_{ct} is the change in a commuting zone outcome over a decade or longer, $\mathbf{X}_{ct}$ are controls for pre-recession trends and commuting zone characteristics, and $\delta_{d(c)}$ are Census division fixed effects.

Our commuting zone estimates reproduce the worker-level estimates. Columns 3 and 4 of Table 6 indicate that a 1 percentage point larger increase in unemployment reduced average earnings by 1.6 percent and the share of adults with positive earnings by 0.33 percentage points between 1980 and 1990, which are very similar to our worker-level estimates. Columns 1 and 2 estimate the same specification for changes between 1970 and 1980 and we find no evidence of negative pre-trends. If anything, earnings in more exposed places were growing faster prior to the recession.

Extending this analysis into the present, we find that earnings in more exposed commuting zones remain depressed nearly four decades later, and this gap is driven by wages. Columns 6 and 8 indicate that a 1 percentage point larger increase in unemployment reduced average earnings by 1.5 percent and hourly wages by 1.6 percent between 1980 and 2017, while the effect on the share with positive earnings is smaller and statistically insignificant (column 7). By 2017, most of the workers in our 1979 panel had exited the labor force, so these wage effects are largely borne by workers who were not directly affected by the recession. These results are consistent with the recession having persistent effects on the local wage structures in more exposed places.

The dynamics of these effects are also informative about whether the recession changed these places or merely hit places that were going to decline regardless. If the recession merely marked places that were already losing out to longer-run shifts in labor demand, we would expect these places to continue falling behind as those shifts proceeded. If instead it merely accelerated adjustments that would eventually have occurred everywhere, we would expect less exposed places to catch up and the gap to close. Figure 7 displays event study estimates of the effect of exposure on log average earnings in each commuting zone between 1975 and 2019 using county-level QCEW data. The estimates show no pre-trend and decline during and after the recession, reaching 1.4 percent per percentage point of exposure by 1985 and 1.7 percent by 1990. They then remain stable at around 2.2 percent from 2000 through 2019, with no sign of convergence.[25] The estimates indicate that the gap did not continue to widen, as we would

[25] Employment in more exposed commuting zones fell by around 9 percent per interquartile range of exposure by 1985 and never recovered (Figure D.1). Appendix D reports the corresponding event studies for unemployment and labor force participation.

expect if the recession had marked places in decline, and it did not close, as we would expect if the recession had accelerated adjustments that were going to occur regardless.

Together, these results indicate that the recession persistently changed more exposed labor markets. Since the persistent effects primarily operate through wages, we next ask how the recession changed the wage structure in these places. We assess whether these changes reflect persistent changes in relative labor demand as well as the role of labor market institutions.

VII.B PERSISTENT CHANGES IN RELATIVE LABOR DEMAND

The persistent decline in earnings and wages documented above is consistent with a lasting change in the fundamentals of more exposed places, such as their productivity or attractiveness to workers. However, this would lower wages for all workers in a place, so for the recession to have widened the distribution of wages through labor demand it must have persistently shifted relative demand against lower-wage workers in these places. Technological change is the natural candidate, both because it is a leading explanation for the rise in wage inequality over this period (Katz and Murphy, 1992; Autor, Katz and Kearney, 2008; Acemoglu and Autor, 2011) and because recessions can accelerate technological change (Hershbein and Kahn, 2018; Jaimovich and Siu, 2020). If the recession accelerated technological change in more exposed places, our distributional results would reflect its timing rather than a separate channel.

If the recession was operating through this channel, we would expect employment and earnings in routine occupations, which were more exposed to these changes, to decline by more in more exposed places. However, we find no differential effects. Table 7 reports the estimated effects of exposure on routine employment and on earnings by occupation between 1980 and 1990 using our long-differences specification (equation (5)).[26] Column 1 indicates that exposure had no effect on the change in the routine share of employment, and we can rule out declines larger than 0.05 percentage points per unit of exposure, relative to an average decline of 1 percentage point across commuting zones over the decade. Columns 4 and 5 indicate that the earnings losses from exposure were similar for workers in routine and non-routine occupations.[27] One limitation of this analysis is that these measures capture the reallocation of employment across occupations rather than technology adoption directly, so we cannot rule out adoption that left the occupational distribution unchanged.

The same conclusion also holds when we consider shifts in relative demand by skill rather than by occupation type. A persistent shift in relative demand against less-educated workers would reduce both their relative employment and their relative wages, and would draw

[26] Routine occupations are defined using the task measures of Autor and Dorn (2013).

[27] Columns 2 and 3 split routine employment into routine cognitive and routine manual occupations and find no effect on either.

relatively more educated workers into more exposed places. We do not find this pattern in Table D.1, which estimates our long-differences specification separately by education. The employment rate of non-college workers fell by 0.31 percentage points per unit of exposure compared to a statistically insignificant 0.06 percentage points for college workers, but hourly wages fell by similar amounts for both groups, and the college population of more exposed places fell by twice as much as the non-college population. Relative employment fell while relative wages did not, and college-educated workers left these places rather than moving to them, which is inconsistent with a shift in relative demand toward more-skilled workers.[28]

The results in this section cannot rule out every form of a shift in relative labor demand, but we find no evidence of the patterns that this type of shift would produce. We therefore turn to the wage-setting environment in more exposed places, focusing on a prominent labor market institution that changed substantially over this period, unions.

VII.C IMPACTS ON UNIONS

The recession coincided with a turning point for unionization in the United States. Although the union share of employment had been declining since the mid-1950s, the number of union members did not begin to fall until the recession, and almost 30 percent of the decline in union density between 1954 and the present occurred in the five years that followed it. Union jobs were unusually exposed to this recession because multi-year contracts with cost-of-living clauses kept union wages high even as the economy contracted, and because deregulation and rising foreign competition increased pressure on union employers. Since unions compress wages, both directly and through threat effects on non-union employers (DiNardo, Fortin and Lemieux, 1996; Western and Rosenfeld, 2011; Fortin, Lemieux and Lloyd, 2021), a persistent decline in union jobs in more exposed places could alter the wage structure in these places and explain part of the recession's impact on overall inequality. We next investigate the extent to which the recession contributed to this decline.

To measure the impact of the recession on unionization, we use the May supplement and Merged Outgoing Rotation Group files of the CPS between 1977 and 1989, which provide union membership for employed workers.[29] Since the union question is only asked of the employed, we code the nonemployed as nonmembers, so our outcome is union membership as a share of the working-age population rather than the conventional density measure defined over employed workers. Since the CPS only identifies state of residence, we measure recession expo-

[28] Table D.2 reports migration rates by education between 1985 and 1990, which indicate that more exposed places attracted fewer college-educated in-migrants, with no significant effect on net migration for either group.

[29] Farber et al. (2021) discuss these data in more detail.

sure at the state level in this analysis. We estimate event study specifications of the form

$$y_{it} = \alpha + \sum_k \beta_k \mathbb{1}\{t = k\} \cdot \Delta\text{Unemp}_{s(i)} + \mathbf{X}_i'\Gamma + \mathbf{W}_{s(i),t}'\Pi + \delta_{s(i)} + \gamma_t + \varepsilon_{it}$$

where y_{it} is an indicator for union membership, $\mathbf{X}_i$ are individual controls, $\mathbf{W}_{s(i),t}$ are pre-recession state characteristics interacted with year, and $\delta_{s(i)}$ and γ_t are state and year fixed effects.

We find that the recession persistently reduced unionization in more exposed states. The estimates in Figure 8A show no pre-trend, decline over the recession period, and stabilize by 1983 at a level that persists through the end of the decade. This timing also helps distinguish the recession's effect from that of the firing of striking air traffic controllers in August 1981, which is widely thought to have emboldened employer opposition to unions. The decline in more exposed states begins in 1980, a year before the firings, and we find no effect on public sector membership, which the firings might have affected more directly (Appendix D), though we cannot rule out that they amplified the effects of local recession shocks. By 1989, individuals living in a state with a 1 percentage point larger increase in unemployment were 0.7 percentage points less likely to be union members. Since we code the nonemployed as nonmembers, part of this effect reflects the recession's effect on employment itself, which in the same specification is a reduction of 0.65 percentage points and accounts for 23 percent of the membership decline.[30] The remainder reflects the destruction of union jobs that were not replaced, the movement of workers from union to non-union employment, and new hires flowing disproportionately to non-union jobs.

This decline occurred mainly within industries, and it does not reflect a decline in the local demand for union representation. Adding industry-by-year fixed effects reduces the estimated effect from 0.68 to 0.51 percentage points (Table D.3, column 6), so roughly a quarter of the decline reflects reallocation toward less unionized industries and the remainder reflects declining membership within them, a split consistent with the broader literature (Farber and Western, 2001). When looking at worker support for unions, we find no evidence of a decline in demand for union representation in more exposed places. We find no effect of exposure on the average union vote share in NLRB elections, where we can rule out effects larger than 1.5 percentage points relative to a pre-recession average of 53 percent (Figure 8B).[31] The decline in membership therefore reflects a decline in union jobs rather than in worker demand, consistent with contemporary observations that employment in the affected industries recovered

[30] Since union membership is the product of membership among the employed (d) and the employment rate (e), our union membership effect is $\beta_m = e \cdot \beta_d + d \cdot \beta_e$. The mechanical disemployment component is $d \cdot \beta_e = 0.24 \times 0.65 = 0.16$.

[31] We also find no persistent effect on the number of NLRB elections per worker (Figure D.3).

while union employment did not (Adams, 1985).

The 1980s recession also stands out among postwar recessions in its impact on unions. In Figure D.4, we compare two measures of how each postwar recession affected unions. We find that employment losses during the 1980s recession were concentrated in highly unionized cells to a much greater degree than in any other recession, and that local recession shocks translated into persistent declines in union membership in 1980–82 and, to a lesser extent, in 1973–75, but not in any recession before or after this period in our data. This suggests that the late 1970s and early 1980s were a period in which recession shocks translated into union losses, perhaps reflecting the vulnerability of unions at the time, and that the 1980s recession was the largest such episode.

Applying our estimates to the 4.8 percentage point national increase in unemployment during the recession implies a recession-induced decline in union density of 3.7 percentage points, which accounts for roughly 42 percent of the decline between 1979 and 1989.[32] However, this extrapolation assumes that the local and aggregate effects coincide, since our design differences out the common national effect of the recession. It also ignores general equilibrium considerations such as the erosion of national union resources, which would lead it to understate the aggregate effect.

Together, the results in this section indicate that the recession persistently reduced unionization in more exposed places, and that this decline mainly reflects the relative decline in union jobs rather than the mechanical effects of job loss, reallocation across industries, or reduced demand for union representation.[33] Given the evidence that unions compress wages, these results establish unions as a channel through which the recession could have reshaped the wage structure of more exposed places. We quantify how much of the recession's effect on wage inequality is accounted for through this channel in the next section.

VIII. Decomposing the Recession's Distributional Effects

The previous two sections established that the recession widened the distribution of wages in more exposed places and that it persistently reduced unionization in these places. This section asks how consequential the recession's effects on unions were for its impacts on wage inequality. We decompose the total effect of severe exposure estimated in Section VI into the part operating through unions and the part operating through other changes to the wage structure.

[32] The implied effect on union density (membership among the employed) is $\beta_d = (\beta_m - d \cdot \beta_e)/e = 0.78$ percentage points, which gives $4.8 \times 0.78 = 3.74$. Union density declines by around 9 percentage points between 1979 and 1989 in our data.

[33] These estimates are also robust to controlling for pre-recession state characteristics and region-by-year fixed effects (Appendix D).

The decomposition is an accounting exercise in the tradition of DiNardo, Fortin and Lemieux (1996), and we discuss the assumptions under which its components have a causal interpretation below.

To illustrate the channels, we write the conditional distribution of wages for severely exposed workers as a mixture of the union F^u and non-union F^n distributions

$$F_{Y|S,T,X}(y \mid 1,1,x) = \pi_u(x)F^u_{Y|S,T,X}(y \mid 1,1,x) + (1-\pi_u(x))F^n_{Y|S,T,X}(y \mid 1,1,x)$$

where $\pi_u(x) = \Pr(U = 1|S = 1, T = 1, X = x)$ is the conditional probability of being covered by a union and $F^u_{Y|S,T,X}(y \mid 1,1,x)$ is shorthand for $F_{Y|S,T,X,U}(y \mid 1,1,x,U = 1)$.

The recession can affect this distribution through its effect on unions in two ways. First, by reducing the probability of union coverage $\pi_u(x)$, it moves workers from the more compressed union distribution F^u to the less compressed non-union distribution F^n. We refer to this as the union composition effect. Second, by reducing the level of union density in a local market, it weakens the threat effects that raise wages at non-union employers, which changes the shape of both F^u and F^n. We refer to this as the union spillover effect (Western and Rosenfeld, 2011; Fortin, Lemieux and Lloyd, 2021), and we measure it by additionally conditioning the wage distributions on the market-level (state-by-industry) union density D. All remaining changes to F^u and F^n in more exposed places are the non-union wage structure effects, which capture changes in how wages are set that do not operate through unions.

We estimate the three components by extending the distribution regression approach of Section VI. We first construct a counterfactual local union density for each market as the average counterfactual probability of coverage, using a difference-in-differences logit for coverage analogous to the distribution regressions in Section VI. We then estimate the conditional wage distributions separately for union and non-union workers, conditioning on local union density, in a stacked distribution regression following Fortin, Lemieux and Lloyd (2021), and construct counterfactual union and non-union distributions as in Section VI. Combining these objects yields three intermediate distributions, Ψ_1, which holds only local union density at its counterfactual value, Ψ_2, which additionally holds the union and non-union wage structures at their counterfactual values, and $\Psi_3 = F^c$, which additionally holds the coverage probabilities at their counterfactual values and recovers the total counterfactual distribution from Section VI.

The decomposition of the total effect is then

$$F_{Y|S,T,X}(y \mid 1,1,x) - F^{c}_{Y|S,T,X}(y \mid 1,1,x) = \underbrace{\left[F_{Y|S,T,X}(y \mid 1,1,x) - \Psi_1\right]}_{\text{union spillover wage structure}} + \underbrace{\left[\Psi_1 - \Psi_2\right]}_{\substack{\text{non-union} \\ \text{wage structure}}} + \underbrace{\left[\Psi_2 - F^{c}_{Y|S,T,X}(y \mid 1,1,x)\right]}_{\text{union composition}}$$

Appendix E.3 defines each intermediate distribution and describes how each difference isolates the corresponding channel.

Each component is a well-defined difference of estimated distributions under our parallel trends assumptions, and the decomposition retains the standard accounting interpretation (DiNardo, Fortin and Lemieux, 1996; Chernozhukov, Fernández-Val and Melly, 2013). Interpreting the channels as causal requires additional assumptions, which we discuss in Appendix E.4. Absent these, the union composition component measures the change in the wage distribution accounted for by the movement of workers between the union and non-union sectors, valued at the imputed wage difference between those groups inclusive of any selection it embodies. The spillover component measures the change associated with cross-market variation in local union density among observably similar workers, which combines the causal effect of union density with the selection of workers into markets with different levels of unionization.

We find that the recession's effects on unions account for a large share of its effect on wage inequality. Table 8 presents the decomposition for different measures of wage dispersion. Across measures, the union spillover and composition channels together account for between 33 and 40 percent of the total effect of severe exposure, with both channels playing an important role. The remaining 60 to 67 percent reflects residual non-union related changes to the wage structure in more exposed places.

These estimates suggest that the recession's effect on wage inequality was roughly 60 percent larger than it would have been absent its effects on unions. This offers a candidate explanation for why the 1980s recession had an outsized effect on the wage distribution. While its effects on workers in Section V are comparable to those of later downturns, a substantial share of its effect on wage inequality operated through unions, and as Section VII shows, the decline in unionization that accompanied this recession was the largest of any postwar recession. Further, as discussed in Section III, unions can counter employer wage-setting power, which the recession may also have impacted. Thus, the decline of unions can also amplify the recession's effect on the non-union sources of wage dispersion. Together, these results indicate that the interaction between a transitory demand shock and labor market institutions can have lasting

consequences for wage inequality.

IX. Conclusion

This paper studies how the 1980s recession contributed to the rise in U.S. wage inequality that followed it. Exploiting variation across local labor markets in exposure to the recession, we find large and persistent negative effects on workers in more exposed places that extend well beyond the aggregate economic recovery. The labor market impacts are comparable in magnitude to those of the Great Recession, though the recovery in employment was larger during the 1980s. Importantly, these effects are concentrated among workers with pre-recession earnings near the bottom of the distribution. Together, these results suggest that the recession increased labor market inequality, and we estimate that had the wage distribution in more exposed states evolved similarly to that in less exposed ones, the rise in wage inequality between 1979 and 1988 would have been up to 20 percent smaller.

These reductions in wages and earnings remain present almost 40 years after the recession, suggesting that it persistently changed these labor markets. The dynamics of these long-run effects also indicate that the recession did not just hit places that would have declined anyway, nor did it speed up structural changes that were already underway, as the gap between more and less exposed places opens up during the recession and remains relatively stable into the present. We find limited evidence that these changes reflect a lasting shift in relative labor demand. Instead, we find that the recession persistently weakened unions in these places, and our estimates suggest that the recession's effects on wage inequality were roughly 60 percent larger than they would have been absent its effects on unions.

We believe these results speak to a long debate about the sources of rising wage inequality in the U.S. since 1980, which has emphasized secular forces such as technological change, trade, deindustrialization, and weakening labor market institutions. We provide evidence that the sharp rise of the early 1980s was not purely the product of these forces, and that a transitory demand shock, deepened by the Federal Reserve's efforts to end inflation, contributed meaningfully to it. We also provide a new explanation for the decline in unionization over the decade, which prior work has identified as an important contributor to rising inequality (DiNardo, Fortin and Lemieux, 1996; Farber et al., 2021; Fortin, Lemieux and Lloyd, 2021). The recession accounts for a substantial share of this decline, and a large share of its contribution to rising inequality is accounted for through this channel. Institutional decline was therefore not only an independent secular trend but partly a consequence of this aggregate shock.

At least two limitations are worth considering when interpreting our results. First, our distributional analysis operates on the wage distribution among employed workers, as is stan-

dard in the literature. However, as we document, the recession also pushed workers, disproportionately low-wage ones, out of employment, and this is itself an important dimension of labor market inequality. Our counterfactual analysis does not incorporate this directly, and bringing both margins into a single distributional framework is a promising extension. Second, our research design compares different labor markets with different levels of exposure and nets out the common component of the recession. Therefore, our estimates identify the consequences of the relative severity, rather than the recession's total effect, and they similarly do not account for general equilibrium responses. We view this evidence as complementary to existing structural approaches that incorporate general equilibrium effects.

Our findings also inform how the costs of economic contractions are assessed. Accounts that focus on the transitory loss of output and employment can miss persistent redistribution across places and workers. Our results also suggest that the distributional consequences for the labor market depend on the institutional landscape. The same contraction arriving in an economy with a different wage-setting environment may carry different distributional consequences, which makes labor market institutions an important determinant of who bears the burden of economic adjustments.

Lastly, our results point, more broadly, toward further questions about the inequality consequences of recessions. This recession's effects on workers resemble those of the Great Recession, yet its effects on unions have no counterpart in recessions since then. If institutions mediate how transitory shocks impact labor market inequality, then the distributional consequences may vary across episodes depending on the institutional landscape while still scarring workers in similar ways. Comparing the distributional effects of recessions across episodes and institutional settings would be informative and could also shed light on the non-union channels that account for a majority of the effects we document. Finally, future work should incorporate the employment margin into distributional analyses, since disemployment is a large part of how recessions affect labor market inequality and current approaches typically focus on those that remain employed.

References

Acemoglu, Daron, and David Autor. 2011. "Skills, Tasks and Technologies: Implications for Employment and Earnings." In *Handbook of Labor Economics, Volume 4B*, ed. Orley Ashenfelter and David Card, 1043–1171. Elsevier.

Adams, Larry T. 1985. "Changing Employment Patterns of Organized Workers." *Monthly Labor Review*, 108(2): 25–31.

Alexander, J. Trent, David Bleckley, Jonathan Fisher, Katie Genadek, Susan Hautaniemi Leonard, and Aristotle Magganas. 2024. "Newly Available Individual-Level U.S. Tax Data from 1969–1994." U.S. Census Bureau, Associate Directorate for Economic Programs ADEP Working Paper 2024-02.

Autor, David, Arindrajit Dube, and Annie McGrew. 2023. "The Unexpected Compression: Competition at Work in the Low Wage Labor Market." National Bureau of Economic Research Working Paper 31010.

Autor, David H., and David Dorn. 2013. "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market." *American Economic Review*, 103(5): 1553–1597.

Autor, David H., David Dorn, and Gordon H. Hanson. 2013. "The China Syndrome: Local Labor Market Effects of Import Competition in the United States." *American Economic Review*, 103(6): 2121–2168.

Autor, David H., David Dorn, Gordon H. Hanson, and Jae Song. 2014. "Trade Adjustment: Worker-Level Evidence." *Quarterly Journal of Economics*, 129(4): 1799–1860.

Autor, David H., Lawrence F. Katz, and Melissa S. Kearney. 2008. "Trends in U.S. Wage Inequality: Revising the Revisionists." *Review of Economics and Statistics*, 90(2): 300–323.

Ball, Laurence. 1994. "What Determines the Sacrifice Ratio?" In *Monetary Policy*, ed. N. Gregory Mankiw, 155–193. University of Chicago Press.

Berger, David, Kyle Herkenhoff, and Simon Mongey. 2022. "Labor Market Power." *American Economic Review*, 112(4): 1147–1193.

Blanchard, Olivier J., and Lawrence H. Summers. 1986. "Hysteresis and the European Unemployment Problem." In *NBER Macroeconomics Annual 1986, Volume 1*, ed. Stanley Fischer, 15–90. MIT Press.

Blanchard, Olivier Jean, and Lawrence F. Katz. 1992. "Regional Evolutions." *Brookings Papers on Economic Activity*, 1992(1): 1–75.

Bluestone, Barry, and Bennett Harrison. 1982. *The Deindustrialization of America: Plant Closings, Community Abandonment, and the Dismantling of Basic Industry.* New York:Basic Books.

Branson, William H., and James P. Love. 1988. "U.S. Manufacturing and the Real Exchange Rate." In *Misalignment of Exchange Rates: Effects on Trade and Industry*, ed. Richard C. Marston, 241–276. Chicago:University of Chicago Press.

Burdett, Kenneth, and Dale T. Mortensen. 1998. "Wage Differentials, Employer Size, and Unemployment." *International Economic Review*, 39(2): 257–273.

Card, David. 1996. "The Effect of Unions on the Structure of Wages: A Longitudinal Analysis." *Econometrica*, 64(4): 957–979.

Card, David, Ana Rute Cardoso, Jörg Heining, and Patrick Kline. 2018. "Firms and Labor Market Inequality: Evidence and Some Theory." *Journal of Labor Economics*, 36(S1): S13–S70.

Card, David, and John E. DiNardo. 2002. "Skill-Biased Technological Change and Rising Wage Inequality: Some Problems and Puzzles." *Journal of Labor Economics*, 20(4): 733–783.

Chernozhukov, Victor, Iván Fernández-Val, and Alfred Galichon. 2010. "Quantile and Probability Curves Without Crossing." *Econometrica*, 78(3): 1093–1125.

Chernozhukov, Victor, Iván Fernández-Val, and Blaise Melly. 2013. "Inference on Counterfactual Distributions." *Econometrica*, 81(6): 2205–2268.

Chow, Melissa, Teresa C. Fort, Christopher Goetz, Nathan Goldschlag, James Lawrence, Elisabeth Ruth Perlman, Martha Stinson, and T. Kirk White. 2021. "Redesigning the Longitudinal Business Database." U.S. Census Bureau, Center for Economic Studies CES Working Paper 21-08.

Coibion, Olivier, Yuriy Gorodnichenko, Lorenz Kueng, and John Silvia. 2017. "Innocent Bystanders? Monetary Policy and Inequality." *Journal of Monetary Economics*, 88: 70–89.

Collins, William J., and Gregory T. Niemesh. 2019. "Unions and the Great Compression of Wage Inequality in the US at Mid-Century: Evidence from Local Labour Markets." *Economic History Review*, 72(2): 691–715.

Davis, Steven J., and Till von Wachter. 2011. "Recessions and the Costs of Job Loss." *Brookings Papers on Economic Activity*, 2011(2): 1–72.

DiNardo, John, Nicole M. Fortin, and Thomas Lemieux. 1996. "Labor Market Institutions and the Distribution of Wages, 1973–1992: A Semiparametric Approach." *Econometrica*, 64(5): 1001–1044.

Dixit, Avinash. 1989. "Entry and Exit Decisions Under Uncertainty." *Journal of Political Economy*, 97(3): 620–638.

Dorn, David. 2009. "Essays on Inequality, Spatial Interaction, and the Demand for Skills." PhD diss. University of St. Gallen.

Farber, Henry S. 2005. "Nonunion Wage Rates and the Threat of Unionization." *Industrial and Labor Relations Review*, 58(3): 335–352.

Farber, Henry S., and Bruce Western. 2001. "Accounting for the Decline of Unions in the Private Sector, 1973–1998." *Journal of Labor Research*, 22(3): 459–485.

Farber, Henry S., Daniel Herbst, Ilyana Kuziemko, and Suresh Naidu. 2021. "Unions and Inequality over the Twentieth Century: New Evidence from Survey Data." *Quarterly Journal of Economics*, 136(3): 1325–1385.

Federal Open Market Committee. 1981. "Transcript of the Meeting of October 5–6, 1981." Board of Governors of the Federal Reserve System.

Federal Reserve Bank of Chicago. 1982. "Review and Outlook: 1981–82. Hard Times—The Midwest in Trauma." *Economic Perspectives*, 6(Midyear): 3–9.

Feenstra, Robert C., and Gordon H. Hanson. 1996. "Globalization, Outsourcing, and Wage Inequality." *American Economic Review*, 86(2): 240–245.

Fernández-Val, Iván, Jonas Meier, Aico van Vuuren, and Francis Vella. 2024. "A Simple Distributional Difference-in-Differences Estimator for Univariate and Bivariate Outcomes." arXiv preprint 2409.02311.

Flood, Sarah, Miriam King, Renae Rodgers, Steven Ruggles, J. Robert Warren, Daniel Backman, Annie Chen, Grace Cooper, Stephanie Richards, Megan Schouweiler, and Michael Westberry. 2024. "IPUMS CPS: Version 12.0 [dataset]." *Minneapolis, MN: IPUMS*.

Fortin, Nicole M., Thomas Lemieux, and Neil Lloyd. 2021. "Labor Market Institutions and the Distribution of Wages: The Role of Spillover Effects." *Journal of Labor Economics*, 39(S2): S369–S412.

Fortin, Nicole M., Thomas Lemieux, and Neil Lloyd. 2022. "Right-to-Work Laws, Unionization, and Wage Setting." National Bureau of Economic Research Working Paper 30098.

Fort, Teresa C., and Shawn D. Klimek. 2018. "The Effects of Industry Classification Changes on US Employment Composition." U.S. Census Bureau, Center for Economic Studies CES Working Paper 18-28.

Freeman, Richard B. 1980. "Unionism and the Dispersion of Wages." *Industrial and Labor Relations Review*, 34(1): 3–23.

Freeman, Richard B., and James L. Medoff. 1981. "The Impact of the Percentage Organized on Union and Nonunion Wages." *Review of Economics and Statistics*, 63(4): 561–572.

Genadek, Katie R., Charles Hokayem, and Philip Pendergast. 2021. "The Summary Earnings Record and Detailed Earnings Record Extracts." U.S. Census Bureau, Associate Directorate for Economic Programs ADEP Working Paper 2021-05.

Goldin, Claudia, and Robert A. Margo. 1992. "The Great Compression: The Wage Structure in the United States at Mid-Century." *Quarterly Journal of Economics*, 107(1): 1–34.

Goodfriend, Marvin, and Robert G. King. 2005. "The Incredible Volcker Disinflation." *Journal of Monetary Economics*, 52(5): 981–1015.

Gordon, Robert J. 1982. "Why Stopping Inflation May Be Costly: Evidence from Fourteen Historical Episodes." In *Inflation: Causes and Effects*, ed. Robert E. Hall, 11–40. University of Chicago Press.

Greenland, Andrew, James Lake, and John Lopresti. 2026. "Trade and US Inequality in the Tokyo Round." National Bureau of Economic Research Working Paper 34785.

Hershbein, Brad, and Bryan A. Stuart. 2024. "The Evolution of Local Labor Markets After Recessions." *American Economic Journal: Applied Economics*, 16(3): 399–435.

Hershbein, Brad, and Lisa B. Kahn. 2018. "Do Recessions Accelerate Routine-Biased Technological Change? Evidence from Vacancy Postings." *American Economic Review*, 108(7): 1737–1772.

Hirsch, Barry T., and David A. Macpherson. 2003. "Union Membership and Coverage Database from the Current Population Survey: Note." *Industrial and Labor Relations Review*, 56(2): 349–354.

Houtz, Alexander, Benjamin Pugsley, and Hannah Rubinton. 2025. "The Uneven Welfare Costs of the Volcker Disinflation." Working paper.

Hoynes, Hilary, Douglas L. Miller, and Jessamyn Schaller. 2012. "Who Suffers During Recessions?" *Journal of Economic Perspectives*, 26(3): 27–48.

Jacobson, Louis S., Robert J. LaLonde, and Daniel G. Sullivan. 1993. "Earnings Losses of Displaced Workers." *American Economic Review*, 83(4): 685–709.

Jaimovich, Nir, and Henry E. Siu. 2020. "Job Polarization and Jobless Recoveries." *Review of Economics and Statistics*, 102(1): 129–147.

Jordà, Òscar. 2005. "Estimation and Inference of Impulse Responses by Local Projections." *American Economic Review*, 95(1): 161–182.

Juhn, Chinhui, Kevin M. Murphy, and Brooks Pierce. 1993. "Wage Inequality and the Rise in Returns to Skill." *Journal of Political Economy*, 101(3): 410–442.

Kahn, Lisa B. 2010. "The Long-Term Labor Market Consequences of Graduating from College in a Bad Economy." *Labour Economics*, 17(2): 303–316.

Kamphorst, Jonne, Zachary Schaller, and Samuel Young. 2025. "Union Election Stats." *https://unionelectionstats.com*, Data downloaded August 2025.

Katz, Lawrence F., and Kevin M. Murphy. 1992. "Changes in Relative Wages, 1963–1987: Supply and Demand Factors." *Quarterly Journal of Economics*, 107(1): 35–78.

Kilss, Beth, and Frederick J. Scheuren. 1978. "The 1973 CPS-IRS-SSA Exact Match Study." *Social Security Bulletin*, 41(10): 14–22.

Lee, David S. 1999. "Wage Inequality in the United States During the 1980s: Rising Dispersion or Falling Minimum Wage?" *Quarterly Journal of Economics*, 114(3): 977–1023.

Levy, Frank, and Peter Temin. 2007. "Inequality and Institutions in 20th Century America." National Bureau of Economic Research Working Paper 13106.

Lucas, Robert E., Jr. 1987. *Models of Business Cycles.* Oxford:Basil Blackwell.

Manning, Alan. 2003. *Monopsony in Motion: Imperfect Competition in Labor Markets.* Princeton, NJ:Princeton University Press.

Mitchell, Daniel J. B. 1985. "Shifting Norms in Wage Determination." *Brookings Papers on Economic Activity*, 1985(2): 575–608.

Mitchell, Daniel J. B., and Christopher L. Erickson. 2005. "Not Yet Dead at the Fed: Unions, Worker Bargaining, and Economy-Wide Wage Determination." *Industrial Relations*, 44(4): 565–606.

Mussa, Michael. 1994. "U.S. Monetary Policy in the 1980s." In *American Economic Policy in the 1980s*, ed. Martin Feldstein, 81–145. Chicago:University of Chicago Press.

Neumark, David, and Michael L. Wachter. 1995. "Union Effects on Nonunion Wages: Evidence from Panel Data on Industries and Cities." *Industrial and Labor Relations Review*, 49(1): 20–38.

Okun, Arthur M. 1978. "Efficient Disinflationary Policies." *American Economic Review*, 68(2): 348–352.

Piketty, Thomas, and Emmanuel Saez. 2003. "Income Inequality in the United States, 1913–1998." *Quarterly Journal of Economics*, 118(1): 1–41.

Romer, Christina D., and David H. Romer. 1989. "Does Monetary Policy Matter? A New Test in the Spirit of Friedman and Schwartz." In *NBER Macroeconomics Annual 1989, Volume 4*, ed. Olivier Jean Blanchard and Stanley Fischer, 121–184. Cambridge, MA:MIT Press.

Romer, Christina D., and David H. Romer. 2004. "A New Measure of Monetary Shocks: Derivation and Implications." *American Economic Review*, 94(4): 1055–1084.

Rosen, Sherwin. 1969. "Trade Union Power, Threat Effects and the Extent of Organization." *Review of Economic Studies*, 36(2): 185–196.

Ruggles, Steven, Sarah Flood, Matthew Sobek, Daniel Backman, Annie Chen, Grace Cooper, Stephanie Richards, Renae Rogers, and Megan Schouweiler. 2024. "IPUMS USA: Version 15.0 [dataset]." *Minneapolis, MN: IPUMS*.

Sargent, Thomas J. 1982. "The Ends of Four Big Inflations." In *Inflation: Causes and Effects*, ed. Robert E. Hall, 41–98. University of Chicago Press.

Schmieder, Johannes F., Till von Wachter, and Jörg Heining. 2023. "The Costs of Job Displacement over the Business Cycle and Its Sources: Evidence from Germany." *American Economic Review*, 113(5): 1208–1254.

Schott, Peter K. 2008. "The Relative Sophistication of Chinese Exports." *Economic Policy*, 23(53): 5–49.

Stuart, Bryan A. 2022. "The Long-Run Effects of Recessions on Education and Income." *American Economic Journal: Applied Economics*, 14(1): 42–74.

Taschereau-Dumouchel, Mathieu. 2020. "The Union Threat." *Review of Economic Studies*, 87(6): 2859–2892.

Tolbert, Charles M., and Molly Sizer. 1996. "U.S. Commuting Zones and Labor Market Areas: A 1990 Update." Economic Research Service, U.S. Department of Agriculture ERS Staff Paper AGES-9614.

Urquhart, Michael A., and Marillyn A. Hewson. 1983. "Unemployment Continued to Rise in 1982 as Recession Deepened." *Monthly Labor Review*, 106(2): 3–12.

U.S. Bureau of Labor Statistics. 1998. "Historical Comparability of Local Area Unemployment Statistics (LAUS) Data." *Documentation provided with the archived 1976–1989 substate LAUS files.*

U.S. Bureau of Labor Statistics. 2026. "Local Area Unemployment Statistics." *BLS Handbook of Methods*, https://www.bls.gov/opub/hom/lau/. Accessed October 10, 2026.

U.S. Census Bureau. 1992. "Measuring the Effect of Benefits and Taxes on Income and Poverty: 1979 to 1991." U.S. Government Printing Office Current Population Reports P60-182RD, Washington, DC.

Volcker, Paul A. 1980. "A Rare Opportunity." *Remarks before the Tax Foundation, New York, December 3, 1980*, Available from FRASER, Federal Reserve Bank of St. Louis.

Volcker, Paul A. 1994. "Comment on "U.S. Monetary Policy in the 1980s"." In *American Economic Policy in the 1980s*, ed. Martin Feldstein, 145–150. Chicago:University of Chicago Press.

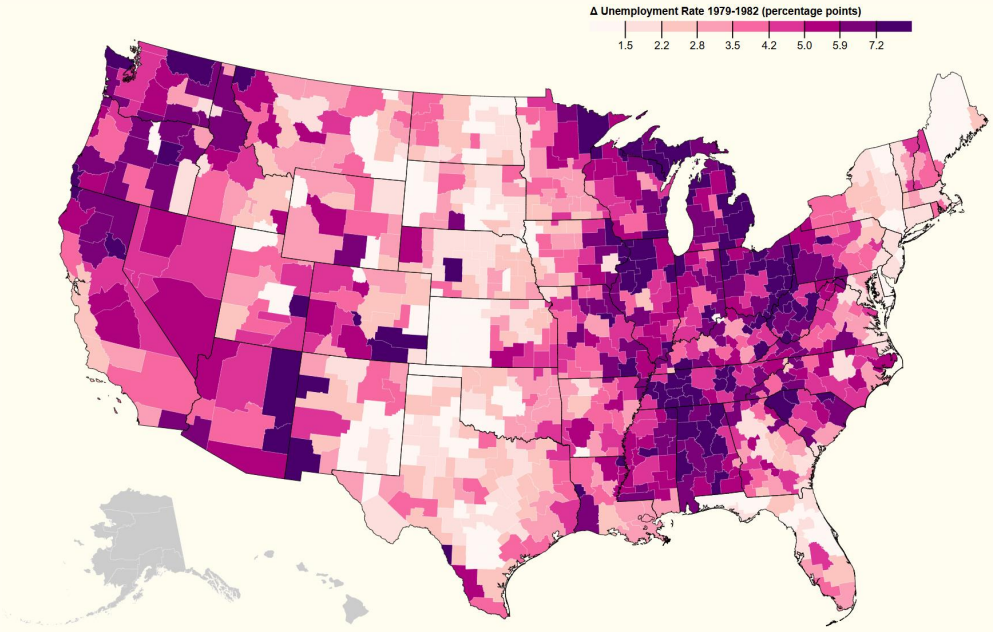
Volcker, Paul A. 2000. "Interview for *Commanding Heights*." *Public Broadcasting Service, September 26, 2000*, Transcript available at https://www.pbs.org/wgbh/commandingheights.

Western, Bruce, and Jake Rosenfeld. 2011. "Unions, Norms, and the Rise in U.S. Wage Inequality." *American Sociological Review*, 76(4): 513–537.

Yagan, Danny. 2019. "Employment Hysteresis from the Great Recession." *Journal of Political Economy*, 127(5): 2505–2558.

FIGURES

A. Variation in Unemployment Change Across Commuting Zones



B. Distribution of Industry Changes in Unemployment

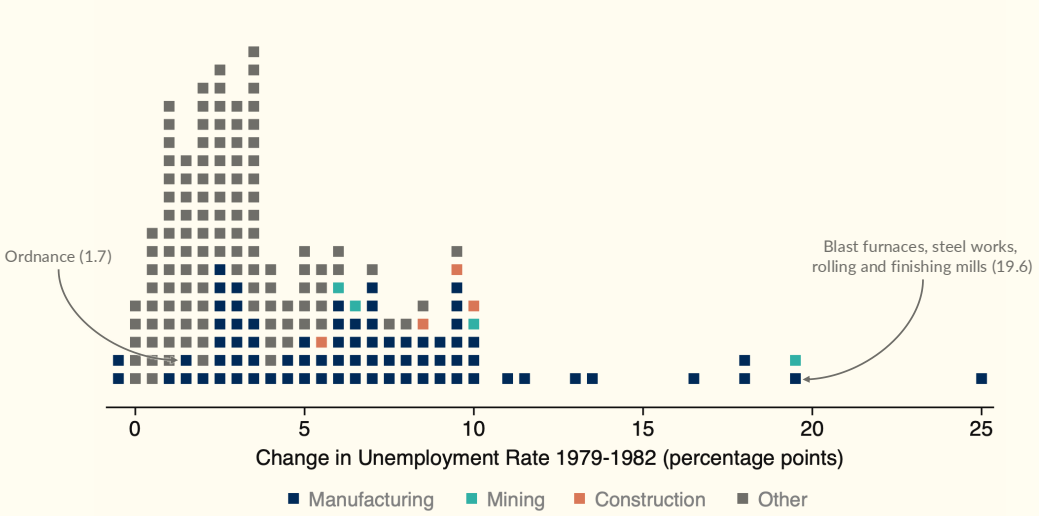


FIGURE 1
Variation in Recession Exposure Across CZs and Industries

Notes: This figure displays the distribution of our measure of recession exposure across places and industries. Panel A displays the change in unemployment between 1979 and 1982 across commuting zones, measured using the BLS Local Area Unemployment Statistics. Panel B displays the change in unemployment between 1979 and 1982 across 203 industries, measured using the Current Population Survey.

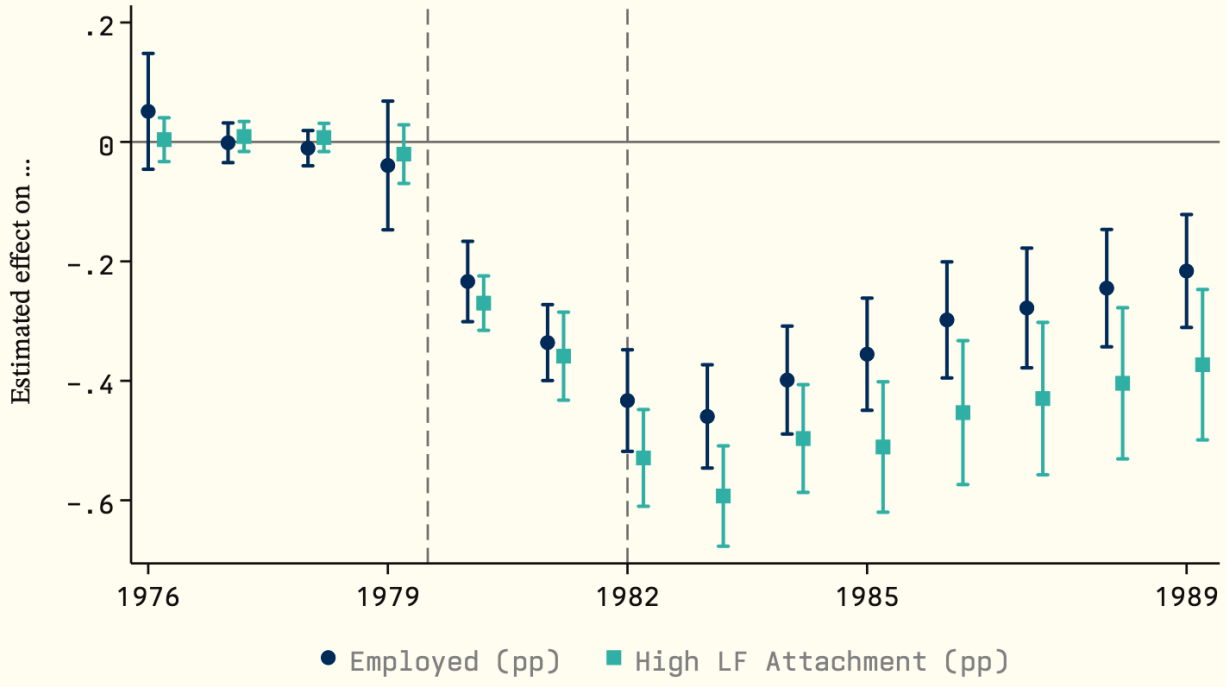


FIGURE 2
Impacts of the 1980s Recession on Employment

Notes: This figure displays regression estimates of the effect of local exposure to the 1980s recession on employment and high labor force attachment. Each point is the coefficient on our recession exposure measure from a separate regression for year t, controlling for worker demographic characteristics and educational attainment as well as pre-recession age×industry×earnings decile fixed effects. The employment and high labor force attachment outcomes are relative to each worker's pre-recession baseline (i.e., subtracting their 1976–1979 mean). Standard errors are clustered at the commuting zone level. Regressions are weighted so that our main sample matches the CPS (see Section IV.A for details). Estimates used in this figure were disclosed under FSRDC Project Number 3420 (CBDRB-FY26-P3420-R13312).

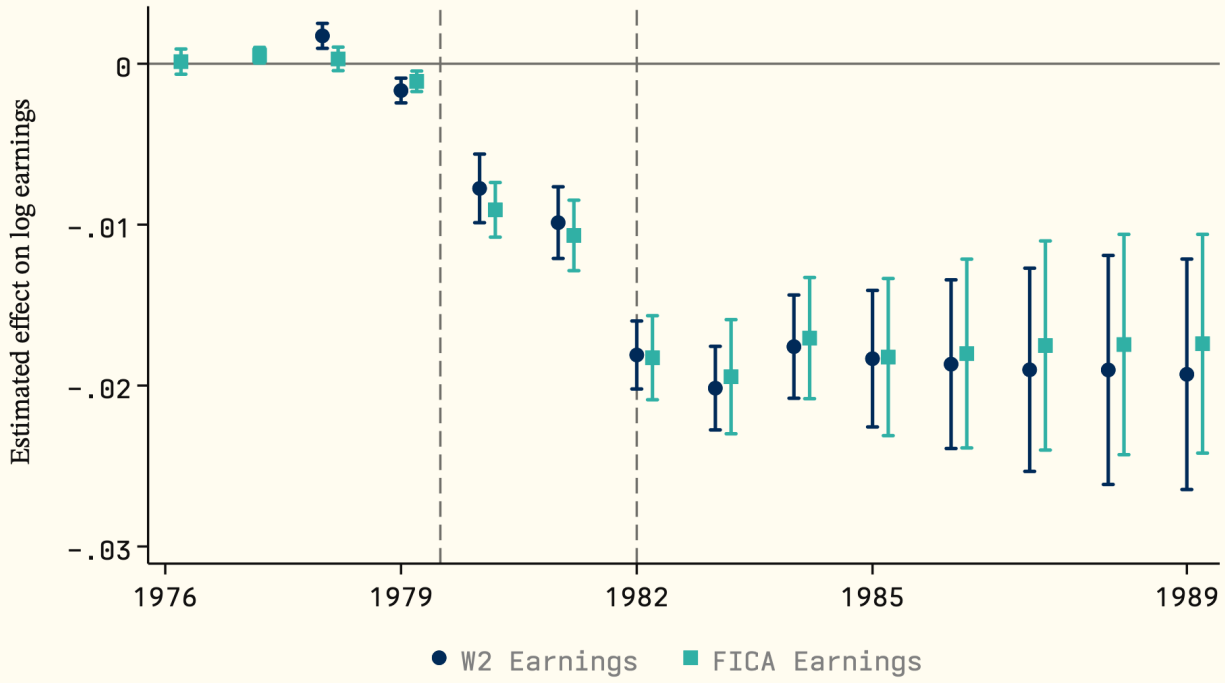
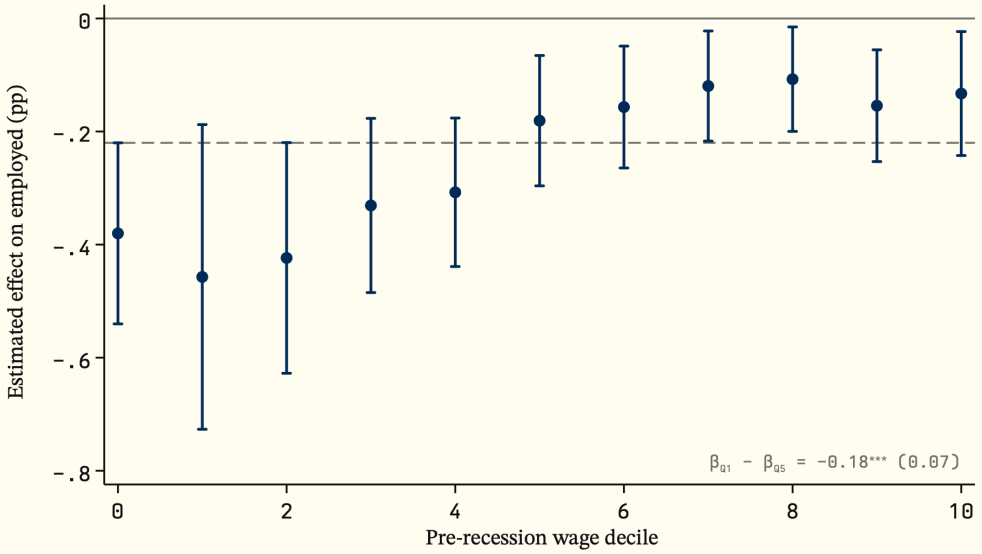


FIGURE 3
Impacts of the 1980s Recession on Earnings

Notes: This figure displays regression estimates of the effect of local exposure to the 1980s recession on log annual W-2 earnings and log FICA earnings. Each point is the coefficient on our recession exposure measure from a separate regression for year t, controlling for worker demographic characteristics and educational attainment as well as pre-recession age×industry×earnings decile fixed effects. The outcomes are relative to each worker's pre-recession baseline (i.e., subtracting their 1978–1979 mean for W-2 earnings, which are only available from 1978, and their 1976–1979 mean for FICA earnings). Since the outcomes are measured in logs, individuals with zero earnings are dropped from the regression sample. Standard errors are clustered at the commuting zone level. Regressions are weighted so that our main sample matches the CPS (see Section IV.A for details). Estimates used in this figure were disclosed under FSRDC Project Number 3420 (CBDRB-FY26-P3420-R13312).

A. Employment



B. Earnings

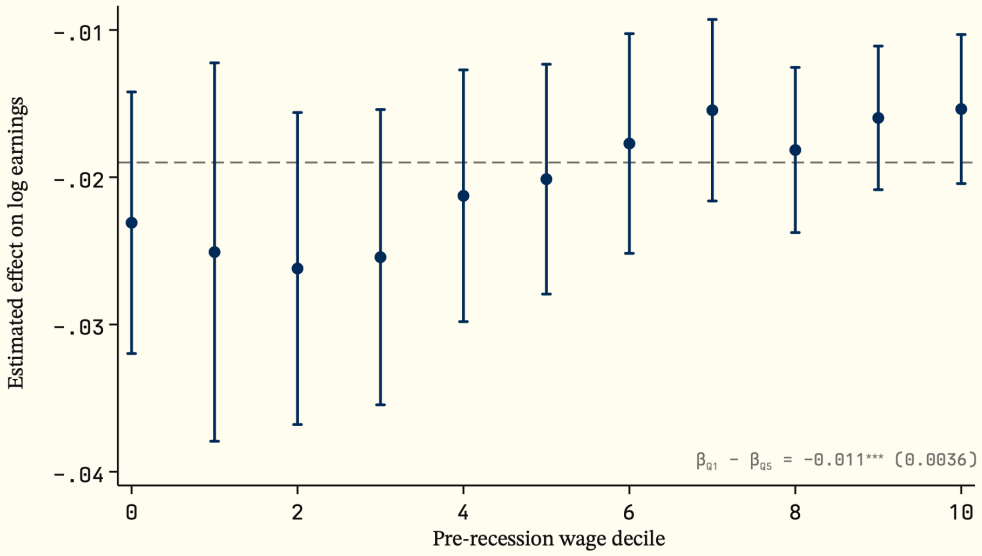
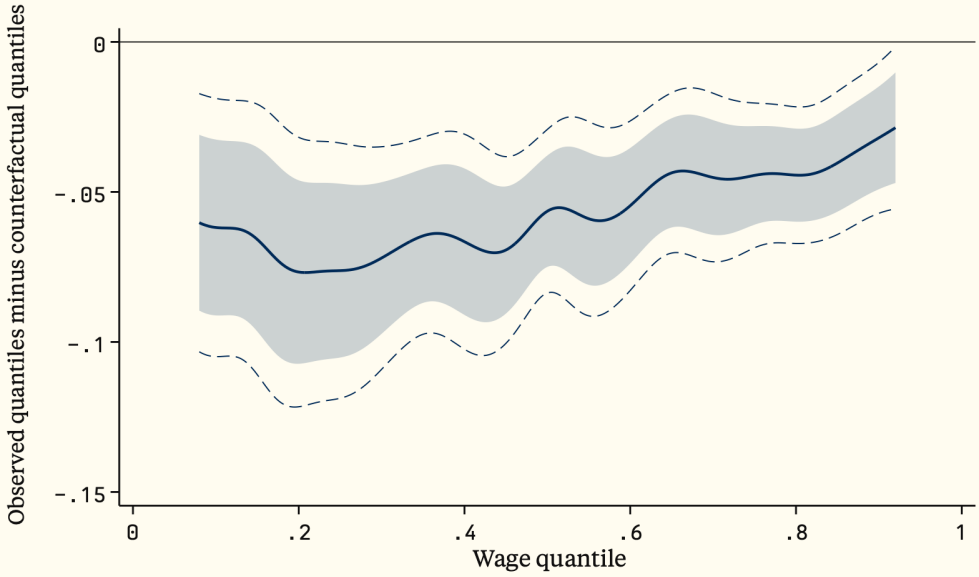


FIGURE 4

Long-Run Impacts of the 1980s Recession on Employment and Earnings — Heterogeneity by Pre-recession Earnings Decile

Notes: This figure displays regression estimates of the effect of local exposure to the 1980s recession on employment and log annual wage earnings pooling 1987–1989. Each point is the coefficient on our recession exposure measure from a separate regression pooling the years 1987–1989 and restricting to workers in each pre-recession earnings decile. Each regression controls for worker demographic characteristics and educational attainment as well as pre-recession age×industry fixed effects. Deciles are based on each worker's pre-recession earnings in 1979. Workers with zero earnings are grouped into decile 0. The outcomes are relative to each worker's pre-recession baseline (1976–1979 for employment and 1978–1979 for earnings). Standard errors are clustered at the commuting zone level. Regressions are weighted so that our main sample matches the CPS (see Section IV.A for details). Estimates used in this figure were disclosed under FSRDC Project Number 3420 (CBDRB-FY26-P3420-R13312).

A. Quantile Effects



B. CDF Effects

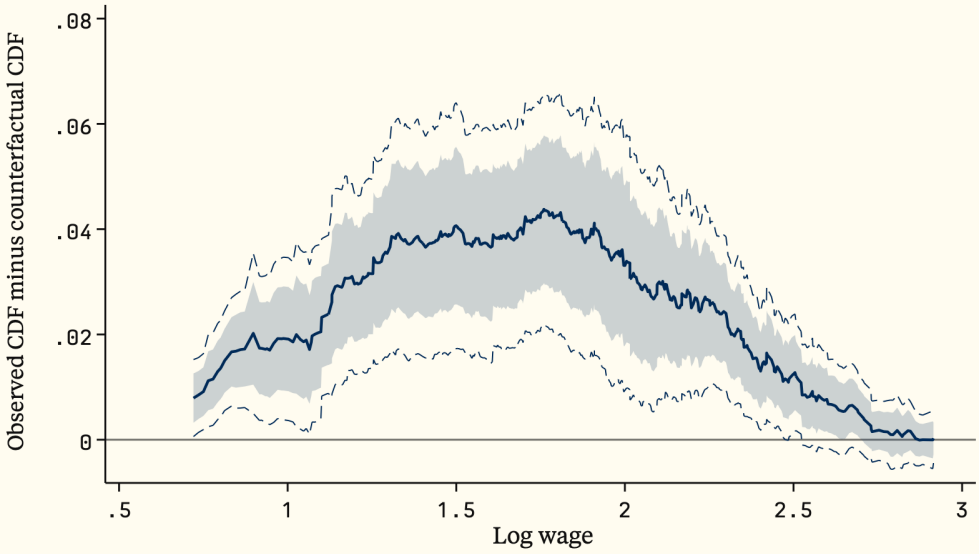
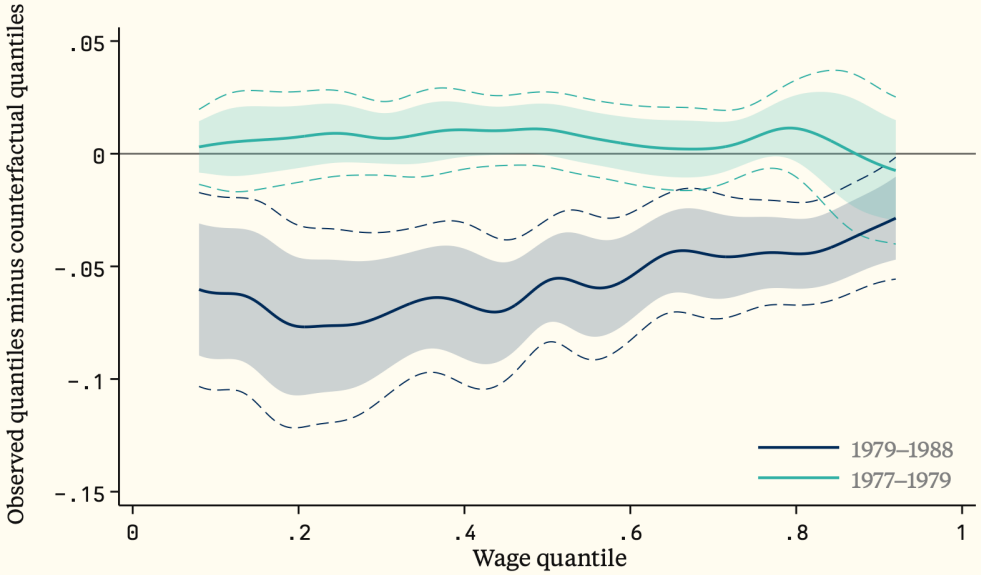


FIGURE 5
Impacts of the 1980s Recession on the Distribution of Wages for Men in 1988 Relative to the Less Severe Counterfactual

Notes: This figure displays estimates of the quantile and CDF effects of severe recession exposure relative to the less severe exposure counterfactual. Each panel displays the difference between the observed distribution of wages for men in 1988 and the counterfactual distribution in 1988. Due to bunching in the wage distribution that creates noise when inverting the CDF to compute quantiles, we smooth the estimates in Panel A using a Gaussian kernel with bandwidth 0.035. The estimates in Panel B are not smoothed and represent the raw difference. The shaded regions correspond to the pointwise 95 percent confidence intervals, while the dashed lines correspond to the uniform 95 percent confidence bands. Standard errors are computed using a weighted bootstrap clustered at the state level with 100 replications (see Appendix E.1).

A. Quantile Effects



B. CDF Effects

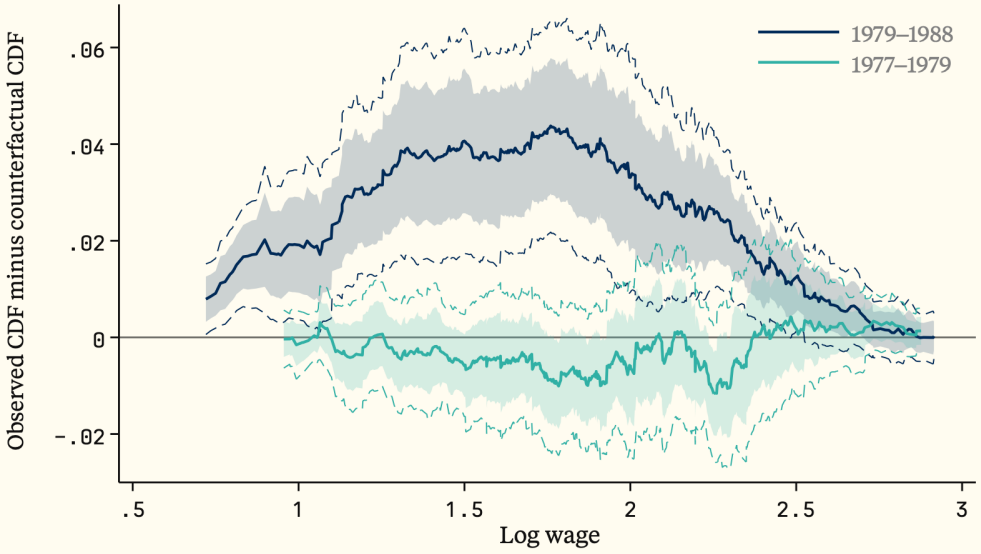


FIGURE 6
Impacts of the 1980s Recession on the Distribution of Wages for Men — Pre-Recession (1977–1979) Placebo

Notes: This figure displays estimates of the quantile and CDF effects of severe recession exposure relative to the less severe exposure counterfactual. Each panel displays the difference between the observed distribution of wages for men in 1988 and the counterfactual distribution in 1988, replicating our main results. Each panel also plots a placebo difference between the observed distribution in 1979 and the counterfactual distribution in 1979, using 1977 as the pre-period. Due to bunching in the wage distribution that creates noise when inverting the CDF to compute quantiles, we smooth the estimates in Panel A using a Gaussian kernel with bandwidth 0.035. The estimates in Panel B are not smoothed and represent the raw difference. The shaded regions correspond to the pointwise 95 percent confidence intervals, while the dashed lines correspond to the uniform 95 percent confidence bands. Standard errors are computed using a weighted bootstrap clustered at the state level with 100 replications (see Appendix E.1).

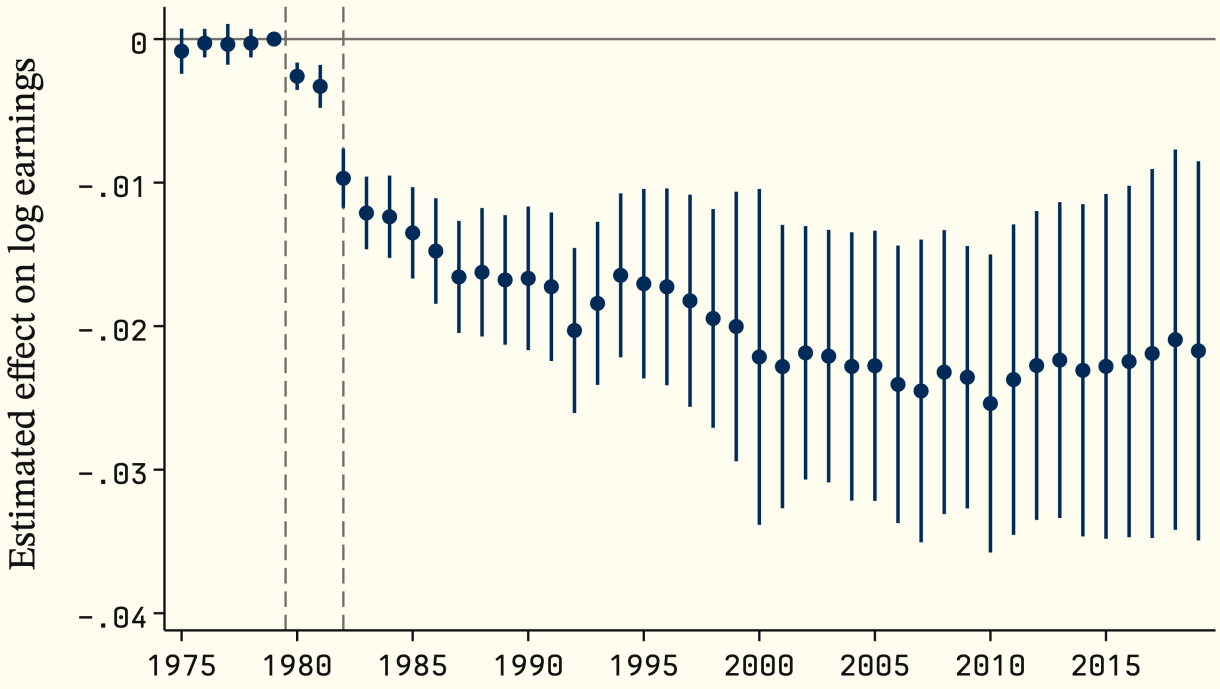
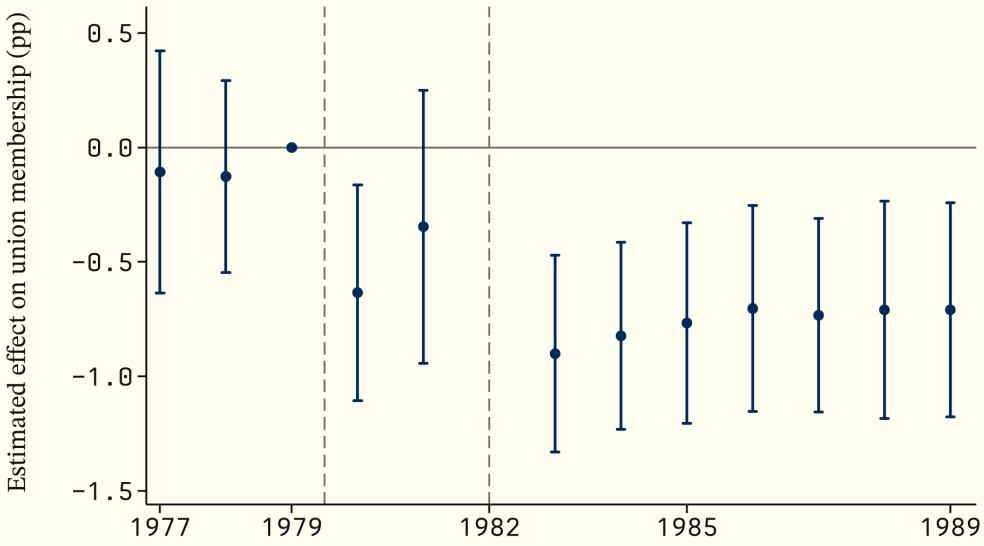


FIGURE 7
Long-Run Effects of Recession Exposure on Commuting Zone Earnings, 1975–2019

Notes: This figure displays event-study estimates of the effect of local recession exposure on the log average annual earnings in a commuting zone between 1975 and 2019, using county-level data from the BLS Quarterly Census of Employment and Wages. Each coefficient is on the interaction between the change in the commuting zone unemployment rate between 1979 and 1982 and a year indicator, relative to 1979, from a specification with commuting zone and division-by-year fixed effects and pre-recession commuting zone characteristics interacted with year (exposure to the 1973–1975 recession, pre-recession earnings and population growth, and the 1979 manufacturing share, routine share, relative supply of college workers, foreign-born share of the non-college population, share aged 65 and over, and female employment share). Regressions are weighted by 1979 employment and standard errors are clustered by commuting zone.

A. Union Membership



B. Average Union Vote Share in NLRB Elections

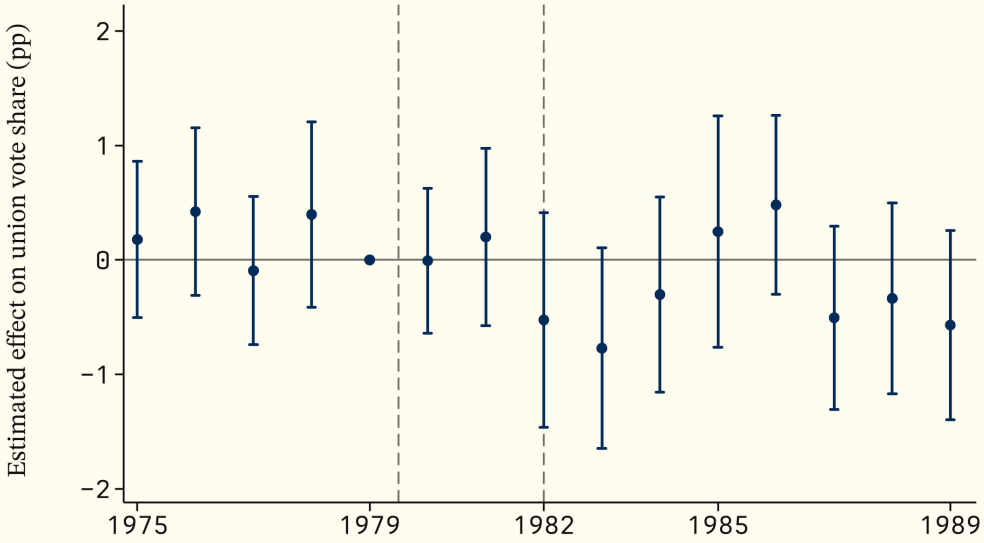


FIGURE 8
Impacts of the 1980s Recession on Unionization

Notes: This figure displays regression estimates of the effect of local exposure to the 1980s recession on union membership and the average union vote share in NLRB elections. Panel A uses the CPS and measures exposure as the change in the state unemployment rate between 1979 and 1982. The outcome is an indicator for being a union member among all adults aged 21–65, so that non-workers are counted as non-members, and is measured in percentage points. The regressions control for a quadratic in age, sex, race, five education group fixed effects and their interaction with age, and state and year fixed effects. The CPS does not include union membership in 1982. Standard errors are clustered at the state level and regressions are weighted using CPS weights. Panel B measures exposure as the change in the CZ unemployment rate between 1979 and 1982. The outcome is the average share of votes in favor of the union in NLRB elections in each CZ and year, in percentage points. The regressions control for CZ and division-by-year fixed effects. Standard errors are clustered at the CZ level and regressions are weighted by employment in 1979. Both panels exclude Alaska, Hawaii, and the District of Columbia. Vertical bars show 95 percent confidence intervals.

TABLES

TABLE 1
Summary Statistics — Main Sample

	Main Sample		CPS ASEC	
	Mean (1)	Std. Dev. (2)	Mean (3)	Std. Dev. (4)
Outcomes (1989):				
Employed (%)	84.9	35.8	71.1	45.3
High Labor Force Attachment (%)	75.0	43.3	62.5	48.4
Earnings (2015$)	41,370	42,410	34,284	38,670
Pre-recession Characteristics (1979):				
Employed (%)	91.2	28.3	74.7	43.5
High Labor Force Attachment (%)	72.3	44.8	60.6	48.9
Earnings (2015$)	35,530	33,720	32,277	34,895
Female (%)	48.8	50.0	51.7	50.0
Age	33.3	7.94	36.8	9.37
Less than 25 (%)	12.1	32.6	8.5	28.0
25–45 (%)	78.0	41.4	68.2	46.6
Greater than 45 (%)	9.9	30.0	23.3	42.3
White (%)	85.8	35.0	81.9	38.5
Less than College (%)	71.3	45.2	79.6	40.3
Manufacturing (%)	16.2	36.8	19.2	39.4
Constr./Mining (%)	6.7	25.0	5.9	23.6
Obs.	5,309,000			

Notes: This table displays summary statistics describing the main worker-level estimation sample relative to the Current Population Survey ASEC/March Supplement for the same birth cohorts. See Section IV.A for more details about the construction of the main sample. Employed is an indicator for having positive wage earnings during the year. High Labor Force Attachment is an indicator for having wage earnings above what a worker would earn working 1,600 hours at the federal minimum wage. Earnings are annual wage earnings inflated to 2015$. Manufacturing and Constr./Mining are indicators for working at an employer in a manufacturing or construction and mining industry. Estimates used in this table were disclosed under FSRDC Project Number 3420 (CBDRB-FY26-P3420-R13312).

TABLE 2
Recession Exposure and Pre-recession Labor Market Trends

	Δ Manuf. share	Δ Pop.	Δ Emp.	Δ Earnings	Δ Union Density
	(1)	(2)	(3)	(4)	(5)
ΔUnemp.	-0.009	0.003	-0.002	-0.001	0.023
	(0.011)	(0.003)	(0.003)	(0.001)	(0.016)
Obs.	717	722	722	721	713
Adj. R^2	0.12	0.53	0.52	0.19	0.17

Notes: This table displays coefficients from regressing pre-recession changes in various labor market characteristics on the change in CZ unemployment between 1979 and 1982. Column 1 uses the change in the log manufacturing share of employment between 1969 and 1978, column 2 uses the change in log population between 1969 and 1979, column 3 uses the change in log employment between 1969 and 1978, column 4 uses the change in log average annual earnings between 1975 and 1978, and column 5 uses the change in log state union density between 1969 and 1978. Each regression controls for the 1979 share of employment in agriculture, mining, construction, manufacturing, transportation and utilities, wholesale trade, retail trade, finance, services, and public administration, as well as Census division fixed effects. Standard errors are clustered at the state level.

TABLE 3
Impacts of the 1980s Recession on Earnings — Robustness

	Log Earnings (t = 1989)										
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
ΔUnemp.	-0.020***	-0.019***	-0.019***	-0.019***	-0.029***	-0.020***	-0.018***	-0.015***	-0.012***	-0.020***	-0.031***
	(0.002)	(0.004)	(0.004)	(0.004)	(0.003)	(0.002)	(0.002)	(0.002)	(0.002)	(0.004)	(0.008)
ΔUnemp. × Rust Belt											0.006
											(0.009)
ΔUnemp. × South											0.023***
											(0.009)
ΔUnemp. × West Coast											-0.020**
											(0.010)
Demographic controls	X	X	X	X	X	X	X	X	X	X	X
Age-earnings-industry FE		X	X	X	X	X	X	X	X	X	X
State clustered SEs			X								
Additional individual controls				X					X		
Growth in pre-recession CZ characteristics					X				X		
Exposure to 73–75 recession						X			X		
State FE							X		X		
Other CZ controls								X	X		
No manuf/mining/constr. workers										X	
R^2	0.03	0.22	0.22	0.22	0.22	0.22	0.22	0.22	0.22	0.22	0.22

Notes: This table displays regression estimates of the effect of local exposure to the 1980s recession on log earnings in 1989 using equation (2) estimated on our main worker panel. All regressions control for sex, race, age, and education fixed effects. Column 2 adds 1979 age-earnings-industry fixed effects, which replicates our main specification. Column 3 clusters the standard errors at the state level instead of the CZ level. Column 4 controls for additional worker characteristics: an indicator for being married in 1979 and an indicator for having self-employment income in 1979. Column 5 controls for employment and population growth in the worker's 1979 CZ. Column 6 controls for exposure to the 1973–1975 recession in the worker's 1979 CZ. Column 7 controls for the worker's state fixed effects. Column 8 controls for other characteristics of the worker's 1979 CZ: earnings and population growth in the preceding decade, the share of employment in manufacturing, the routine share of employment, female labor force participation rate, share of population 65+, the ratio of college to non-college educated workers, and the share of the non-college population that is foreign born all measured from the 1980 Census. Column 9 controls for all of these simultaneously. Column 10 drops workers whose 1979 industry of employment was in manufacturing, mining, or construction. Column 11 interacts the CZ change in unemployment with Rust Belt, South, and West Coast indicators based on the worker's 1979 state of residence. Estimates used in this table were disclosed under FSRDC Project Number 3420 (CBDRB-FY26-P3420-R13312).

TABLE 4
Changes in Wage Dispersion for Men 1979–1988

	Aggregate Distribution			High-exposure ($S = 1$)			Low-exposure ($S = 0$)		
Statistic	1979	1988	Δ	1979	1988	Δ	1979	1988	Δ
90–10	1.306	1.527	0.221	1.272	1.511	0.239	1.340	1.547	0.207
90–50	0.586	0.722	0.136	0.548	0.695	0.147	0.644	0.757	0.113
50–10	0.720	0.805	0.085	0.725	0.816	0.092	0.696	0.790	0.094
80–20	0.862	1.022	0.160	0.825	1.014	0.189	0.899	1.035	0.136
Std. Dev.	0.442	0.523	0.081	0.429	0.518	0.089	0.456	0.530	0.074
Gini	0.243	0.286	0.044	0.233	0.282	0.048	0.255	0.292	0.038
Obs.	20,706	71,391		11,297	37,472		9,409	33,919	

Notes: This table displays measures of wage dispersion in 1979 and 1988 for men aged 16–64 in the CPS, for the aggregate distribution and separately for high- and low-exposure states. Hourly wages are measured directly for hourly workers and as weekly earnings divided by usual hours worked for non-hourly workers. Individuals with allocated wages or earnings are dropped. 90–10 refers to the difference between the 90th and 10th percentiles of the log hourly wage distribution, and similarly for 90–50, 50–10, and 80–20. Std. Dev. refers to the standard deviation of log hourly wages, and Gini to the Gini coefficient of hourly wages. Exposure is measured as the change in the state unemployment rate between 1979 and 1982, and high exposure is defined as above-median exposure. All estimates are weighted by CPS earnings weights.

TABLE 5

Impacts of the 1980s Recession on Wage Inequality 1979–1988 Relative to the Less Severe Counterfactual

	Aggregate Distribution			High-exposure ($S = 1$) Distribution		
Statistic	Total change	Recession effect	Share (%)	Total change	Recession effect	Share (%)
90–10	0.221	0.030	13.5	0.239	0.060	25.3
		(0.014)			(0.024)	
90–50	0.136	0.024	18.0	0.147	0.042	28.6
		(0.010)			(0.018)	
50–10	0.085	0.005	6.4	0.092	0.018	20.0
		(0.011)			(0.021)	
80–20	0.160	0.032	20.1	0.189	0.068	35.9
		(0.011)			(0.019)	
Std. Dev.	0.081	0.012	14.4	0.089	0.024	27.3
		(0.004)			(0.007)	
Gini	0.044	0.008	18.6	0.048	0.015	31.6
		(0.002)			(0.004)	

Notes: This table displays estimates of the effect of severe exposure on changes in each measure of wage dispersion between 1979–1988 relative to the less severe counterfactual. The sample is restricted to men aged 16–64 in the CPS. Column 1 displays the change in each measure between 1979–1988. Column 2 displays the recession's contribution, measured as the difference between the statistic computed on the observed distribution and on the counterfactual distribution. Column 3 displays the share of the total change explained by the recession. Columns 4–6 repeat the same for workers in high-exposure states. All estimates are weighted by CPS earnings weights. Standard errors are reported in parentheses below the effects and are computed using a weighted bootstrap clustered at the state level with 100 replications (see Appendix E.1).

TABLE 6

Impacts of the 1980s Recession on Long-Run Commuting Zone Labor Market Outcomes

	1970–1980		1980–1990			1980–2017		
	(1) Log Earnings	(2) Positive Earnings (pp)	(3) Log Earnings	(4) Positive Earnings (pp)	(5) Log Hourly Wage	(6) Log Earnings	(7) Positive Earnings (pp)	(8) Log Hourly Wage
ΔUnemp.	0.0065***	-0.0333	-0.0157***	-0.3282***	-0.0111***	-0.0146***	-0.1307	-0.0161***
	(0.0021)	(0.1000)	(0.0033)	(0.0590)	(0.0022)	(0.0032)	(0.1705)	(0.0030)
Obs.	722	722	722	722	722	722	722	722
R^2	0.56	0.48	0.80	0.65	0.82	0.59	0.57	0.65

Notes: This table displays coefficients from regressing changes in average CZ labor market outcomes in the Census on the change in CZ unemployment between 1979 and 1982. Log earnings refers to the change in average log annual earnings in the CZ during the corresponding decade. Positive earnings refers to the change in the share of individuals that reported positive wage earnings, in percentage points. Log hourly wage refers to the change in the average log hourly earnings in the CZ. Hourly wages are measured as the annual wage earnings divided by annual weeks worked times usual weekly hours worked. Individuals with allocated wage earnings are dropped. Columns 1 and 2 use changes between the 1970 and 1980 Census, columns 3–5 use changes between the 1980 and 1990 Census, and columns 6–8 use changes between the 1980 Census and the 2015–2019 5-year ACS. Each regression controls for division fixed effects and earnings and population growth in the preceding decade, the share of employment in manufacturing, the routine share of employment, female employment share, share of the population aged 65 and over, ratio of college to non-college educated workers, and share of the non-college population that is foreign born, all measured in the base year. We also control for exposure to the 1969–1970 recession in columns 1 and 2 and exposure to the 1973–1975 recession in columns 3 to 8. Regressions are weighted by employment in 1979 and standard errors are clustered at the state level.

TABLE 7

Impacts of the 1980s Recession on Routine Occupations

Dependent Variable: CZ change between 1980–1990

	Employment (pp)			Log Earnings	
	(1)	(2)	(3)	(4)	(5)
	Routine	Routine Cognitive	Routine Manual	Routine	Non-Routine
ΔUnemp.	0.0139	0.0424	-0.0800	-0.0147***	-0.0170***
	(0.0340)	(0.0460)	(0.0636)	(0.0034)	(0.0034)
Obs.	722	722	722	722	722
R^2	0.69	0.56	0.48	0.76	0.76

Notes: This table displays coefficients from regressing changes in CZ labor market outcomes in the Census between 1980 and 1990 on the change in CZ unemployment between 1979 and 1982. Columns 1–3 use the share of workers in routine, routine cognitive, and routine manual occupations, based on each occupation's routine task measure, in percentage points. Columns 4 and 5 use average log annual earnings of workers in routine and non-routine occupations. Individuals with allocated wage earnings are dropped. Each regression controls for division fixed effects and earnings and population growth in the preceding decade, the share of employment in manufacturing, the routine share of employment, female employment share, share of the population aged 65 and over, ratio of college to non-college educated workers, and share of the non-college population that is foreign born, all measured in the base year. We also control for exposure to the 1973–1975 recession. Regressions are weighted by employment in 1979 and standard errors are clustered at the state level.

TABLE 8
Decomposition of the Recession's Impacts on Wage Inequality

Statistic	Recession effect	Decomposition		
		Union spillover	Non-union wage structure	Union composition
90–10	3.0	0.6	1.9	0.6
		(19.8)	(61.6)	(18.6)
90–50	2.4	0.3	1.5	0.6
		(13.7)	(62.1)	(24.2)
50–10	0.6	0.3	0.3	-0.0
		(45.5)	(59.6)	(-5.1)
80–20	3.3	0.6	2.0	0.7
		(17.0)	(61.5)	(21.4)
Std. Dev.	1.2	0.2	0.8	0.2
		(17.8)	(64.3)	(18.0)
Gini	0.8	0.1	0.5	0.2
		(14.5)	(67.0)	(18.5)

Notes: This table displays estimates of the decomposition of the recession's total effect into union spillover, non-union wage structure, and union composition effects for each measure of wage dispersion. Effects are multiplied by 100, so column 1 equals 100 times column 2 of Table 5. The sample is restricted to men aged 16–64 in the CPS. Column 1 repeats the effect of severe exposure on each measure. Columns 2, 3, and 4 display the effects attributed to the union spillover, non-union wage structure, and union composition channels, with each channel's share of the recession's effect (in percent) displayed below in parentheses. All estimates are weighted by CPS earnings weights.

Online Appendix

"The 1980s Recession and U.S. Labor Market Inequality"

Donato Onorato Kamelia Stavreva

A Data Appendix **A–1**
A.1 Administrative Worker Panel . . . A–1
A.2 Commuting Zones and Recession Exposure . . . A–8
A.3 Commuting Zone Characteristics . . . A–18
A.4 Decennial Census and American Community Survey . . . A–22
A.5 CPS Union Membership . . . A–24
A.6 CPS Wage Data for the Distributional Analysis . . . A–26
A.7 CPS Annual Social and Economic Supplement . . . A–28
A.8 NLRB Representation Elections . . . A–30
A.9 Aggregate Time Series . . . A–31

B Recession Exposure: Measurement and Validation **A–32**
B.1 What Predicts Recession Exposure? . . . A–32
B.2 Monetary Policy and the Industry Incidence of the Recession . . . A–32
B.3 Monetary Policy and the Geographic Incidence of the Recession . . . A–37

C Worker-Level Effects: Additional Results and Robustness **A–42**
C.1 Unweighted Estimates . . . A–42
C.2 Robustness . . . A–42
C.3 Heterogeneity . . . A–47
C.4 Comparison to the Great Recession . . . A–51

D Mechanisms: Additional Results **A–52**
D.1 The Recession's Impact: People versus Places . . . A–52
D.2 Persistent Changes in Relative Labor Demand . . . A–53
D.3 Impacts on Unions . . . A–56

E Distributional Analysis **A–62**

E.1 Estimating the Counterfactual Distribution A–62

E.2 Pre-recession Placebo A–69

E.3 Estimating the Decomposition A–71

E.4 Causal Interpretation of the Decomposition A–79

F Model Appendix **A–84**

F.1 Environment A–84

F.2 Wage Setting A–85

F.3 Union Firms and Exit A–87

F.4 The Recession A–88

F.5 Heterogeneity and Inequality A–89

F.6 Discussion of Assumptions A–90

List of Figures

A.1 Change in State Unemployment 1979–1982: Unofficial County and Official State LAUS Estimates . . . A–12
A.2 CZ Unemployment Rates in the LAUS and the Decennial Census . . . A–13
A.3 Unemployment Rates Across the 1989/1990 Seam . . . A–14
A.4 CZ Recession Exposure and Private Employment Growth 1979–1982 . . . A–15
A.5 Distribution of State Recession Exposure . . . A–16
A.6 Distribution of Industry Recession Exposure . . . A–17
B.1 What Predicts Recession Exposure? . . . A–33
B.2 Industry Employment Responses to Monetary Policy Shocks, Selected Industries . A–36
B.3 Industry Employment Changes and Monetary Policy Predicted Changes, 1979–1982 A–37
B.4 CZ Employment Changes and Shift-Share Predicted Changes, 1979–1982 . . . A–39
B.5 CZ Recession Exposure and Shift-Share Predicted Changes . . . A–40
B.6 Monetary Exposure and CZ Employment Changes Across Periods . . . A–41
C.1 Impacts of the 1980s Recession on Employment and Earnings — Unweighted . . . A–43
C.2 Impacts of the 1980s Recession on High Labor Force Attachment — Heterogeneity by Pre-recession Earnings Decile . . . A–47
C.3 Impacts of the 1980s Recession on Employment — Heterogeneity by Worker Characteristics . . . A–48
C.4 Impacts of the 1980s Recession on High Labor Force Attachment — Heterogeneity by Worker Characteristics . . . A–49
C.5 Impacts of the 1980s Recession on Earnings — Heterogeneity by Worker Characteristics . . . A–50
D.1 Long-Run Effects of Recession Exposure on Commuting Zone Employment, 1975–2019 . . . A–52
D.2 Impacts of the 1980s Recession on CZ Unemployment and Labor Force Participation A–54
D.3 Impacts of the 1980s Recession on CZ NLRB Elections per 10,000 Workers . . . A–59
D.4 Recessions and Unionization Across Postwar Recessions . . . A–60
D.5 Evolution of Union Density Around Recessions . . . A–61
E.1 Impacts of the 1980s Recession on the Distribution of Wages for Men in Severely Exposed States in 1988 Relative to the Less Severe Counterfactual . . . A–68
E.2 Effect of Severe Exposure on Union Coverage by Predicted Wage . . . A–75
E.3 Effect of Union Density on the Distribution of Wages . . . A–76
E.4 Effect of Union Density on the Distribution of Wages — Alternative Specifications A–77

List of Tables

B.1 What Predicts Recession Exposure? A–34
C.1 Impacts of the 1980s Recession on Employment — Robustness A–44
C.2 Impacts of the 1980s Recession on High Labor Force Attachment — Robustness . A–45
C.3 Impacts of the 1980s Recession on FICA Earnings — Robustness A–46
C.4 Comparison to the Great Recession A–51
D.1 Recession Exposure and Labor Market Outcomes by Education A–55
D.2 Recession Exposure and Migration by Education A–56
D.3 Impacts of the 1980s Recession on Unions — Robustness A–58
E.1 Impacts of the 1980s Recession on Wage Inequality — Pre-recession (1977–1979) Placebo A–70

A. Data Appendix

This appendix describes the data sources used in the paper, the construction of the main variables, and the samples used in each analysis.

A.1 ADMINISTRATIVE WORKER PANEL

Our main worker-level analysis uses a new longitudinal panel covering 1976–1989 that we construct by linking together various administrative records at the Census Bureau. This section describes the source data, the construction of the analysis sample, and the main variables. The underlying microdata are confidential and were accessed through a Federal Statistical Research Data Center under FSRDC Project Number 3420. All estimates from these data reported in the paper were reviewed and approved for release by the Census Bureau's Disclosure Review Board (CBDRB-FY26-P3420-R13312).

A.1.1 *Data Sources*

Sample frame. The Census Bureau assigns survey respondents an anonymized person identifier, the Protected Identification Key (PIK), by matching their personally identifiable information (name, address, date of birth, and sex) to a reference file built from administrative records (Genadek, Hokayem and Pendergast, 2021). The PIK allows us to link the same individual across each of the different administrative records. Our starting population is every individual with a PIK who appears in the following surveys in the years that are available to us in the FSRDC: Current Population Survey Annual Social and Economic Supplement (CPS ASEC) between 1973–2024, the Survey of Income and Program Participation (SIPP) between 1984–2024, or the American Community Survey (ACS) between 2015–2019.[34]

Since individuals enter our frame by appearing in a survey at some point over a 50 year window, rather than through a sample drawn in 1979, the frame is not a random sample of the 1979 working-age population. Most importantly, the probability of entering the frame is decreasing in age in 1979, since older cohorts are less likely to survive into the later survey years where most of the frame comes from. We discuss the implications of this for the composition of our sample and how we address it in Appendix A.1.4.

Summary Earnings Record (SER). The SER is an extract from the Social Security Administration's (SSA) Master Earnings File and contains each individual's total annual earnings

[34] Not every survey year is linked to the Social Security Administration's earnings record extracts that are sent to the Census Bureau.

subject to Social Security (FICA) taxes from 1951 onwards, aggregated across all of their employers (Genadek, Hokayem and Pendergast, 2021). The version of the SER we use is topcoded at the annual Social Security taxable maximum.[35] In 1979 the taxable maximum was roughly the 90th percentile of the annual earnings distribution, so SER earnings contain little information about earnings at the top of the distribution. Since the SER covers the full sample period, we use it to construct our employment measures and to extend our pre-recession period for earnings to 1976–1977, which the Detailed Earnings Record does not cover.

Detailed Earnings Record (DER). The DER contains the information reported on every W-2 form filed for an individual from 1978 onwards, so that an individual appears once for each employer they worked for in a given year (Genadek, Hokayem and Pendergast, 2021). Unlike the SER, DER earnings are not topcoded. Each record also includes the Employer Identification Number (EIN) of the employer that filed the W-2, which allows us to link workers to their employers in the Longitudinal Business Database. We measure annual wage earnings as the sum of wage and salary earnings across all W-2s filed for the individual in that year.

Social Security coverage of public sector workers. Social Security coverage was nearly universal for private sector wage and salary employment during our sample period, but several groups of workers were not covered. In the early 1970s, around 9 out of 10 workers were in covered employment, with the main exclusions being federal civil service employees, railroad workers, and some employees of state and local governments and nonprofit institutions (Kilss and Scheuren, 1978). Federal civilian employees were only brought into the system in 1983 and 1984, and many state and local government employees remained uncovered until 1991 (Genadek, Hokayem and Pendergast, 2021). Since our employment measures are based on FICA-covered earnings in the SER, workers in uncovered jobs appear as non-employed. In addition, the Longitudinal Business Database excludes government employers (see below), so these workers, along with any other public sector workers, are assigned to a missing industry category.

Longitudinal Business Database (LBD). The LBD is a longitudinal database of the universe of private, non-farm business establishments with paid employees in the United States from 1976 onwards (Chow et al., 2021).[36] It is constructed by linking annual snapshots of the

[35] The taxable maximum was $15,300 in 1976, $16,500 in 1977, $17,700 in 1978, $22,900 in 1979, $25,900 in 1980, $29,700 in 1981, $32,400 in 1982, $35,700 in 1983, $37,800 in 1984, $39,600 in 1985, $42,000 in 1986, $43,800 in 1987, $45,000 in 1988, and $48,000 in 1989 (nominal dollars).

[36] The underlying Business Register also contains government-owned or operated establishments, but these are classified as out of scope following the County Business Patterns, along with establishments in out of scope

Census Bureau's Business Register, which is built from business tax filings and Census Bureau surveys, and contains each establishment's location, employment, payroll, industry, and firm identifiers. Importantly for our purposes, the LBD provides vintage-consistent NAICS industry codes for every establishment, including those observed before the SIC-to-NAICS transition in 1997 (Fort and Klimek, 2018; Chow et al., 2021). This allows us to assign 1979 employers an industry code that is consistent with the more recent industry classification. We link workers to the LBD using the EIN reported on their W-2s. We assign each EIN a single industry in each year based on the industry with the most employment across its establishments.

Numident. The SSA Numident file contains information collected when an individual applies for a Social Security Number, including their sex, date of birth, and date of death. We use it to measure age, to drop individuals who died before the end of our sample period, and to fill in any missing demographic information (sex and race) that couldn't be assigned from other sources.

Decennial Census. We also link individuals to their responses in the 2000 and 2010 Decennial Census, which we use to measure race and sex. Education is only available for individuals that responded to the 2000 Census long form, since the 2010 Census did not ask about education.

IRS Form 1040. We observe the residential location for the bulk of our sample through the address reported on their individual income tax return (Form 1040). These data come from the historical IRS Form 1040 files held at the Census Bureau, which contain the near universe of individual income tax returns filed for tax years 1969, 1974, 1979, 1984, 1989, and 1994 (Alexander et al., 2024).[37] Each record contains the filer's mailing address, filing status, several income amounts, and the identifiers of the filer and, for joint returns, the spouse. The Census Bureau has replaced these identifiers with PIKs, which were assigned to 99.5 percent of filers and 96.9 percent of spouses in 1979 (Alexander et al., 2024). This means that we observe the location of both spouses on a joint return. We assign each individual to the county, and therefore the commuting zone, of the address on their 1979 tax return.[38] Since we measure location prior to the recession, our exposure measure is not affected by any migration that

industries such as crop production (Chow et al., 2021).

[37]The 1979 file contains 90.8 million returns, which is 97.9 percent of the number of returns reported in IRS publications for that year (Alexander et al., 2024).

[38]The files contain several county identifiers, including those reported on the original return, those added by the Census Bureau for population estimates, and those from a more recent linkage of the tax addresses to the Census Bureau's Master Address File (Alexander et al., 2024). Mailing addresses may also differ from residential addresses, for instance for filers using a PO Box.

occurs in response to the recession.

A.1.2 *Sample Construction*

Starting from the sample frame, we impose the following restrictions. First, we restrict to individuals between the ages of 23 and 54 in 1979 (birth cohorts 1925–1956), so that individuals in our sample remain of working age through 1989. Second, we drop individuals that died before 1989 according to the Numident. Third, we require that the individual appears as a filer or spouse on a Form 1040 for tax year 1979, since we cannot otherwise observe their pre-recession location. Lastly, to focus on workers that had some pre-recession labor force attachment, we restrict to individuals that reported positive wage earnings to the SSA in at least one year between 1976 and 1979. The final sample consists of around 5.3 million individuals.

The tax filing restriction is the most consequential of these for the composition of our sample. For tax year 1979, an individual was required to file a return if their gross income exceeded $3,300 if single, or if their combined gross income exceeded $5,400 if married and filing jointly, with higher thresholds for filers aged 65 or older.[39] The single threshold corresponds to annual earnings from working around 1,140 hours at the 1979 federal minimum wage of $2.90, which is below our high labor force attachment threshold and places a worker at roughly the 5th percentile of the 1979 annual earnings distribution among workers in our cohorts with positive earnings in the CPS ASEC. Using the tax variables that the Census Bureau releases with the March 1980 CPS, which assign each respondent a filing status for tax year 1979 based on their reported income and family characteristics, we estimate that 93.7 percent of individuals in our birth cohorts filed a return or were the spouse on a joint return, including 98.8 percent of those with positive wage earnings and 77.5 percent of those without.[40]

Filing a tax return was nonetheless the norm during this period. In 1979, the 92.7 million individual income tax returns filed were equal to around 95 percent of the estimated 97.5 million potential tax units in the United States (Piketty and Saez, 2003; Alexander et al., 2024), and contemporary assessments considered compliance with the filing requirements to be high (Kilss and Scheuren, 1978).[41] Three features of the filing requirements limit how much this

[39] The thresholds were $4,300 for single filers aged 65 or older, $6,400 for married couples with one spouse aged 65 or older, $7,400 for married couples with both spouses aged 65 or older, and $1,000 for married individuals filing separately. Individuals with net self-employment earnings of $400 or more were required to file regardless of their gross income.

[40] We approximate our birth cohorts using individuals aged 24–55 in the March 1980 CPS, since the survey records age in March of the year following the tax year. Shares are weighted using CPS person weights. The March CPS does not ask respondents about taxes. Instead, the Census Bureau simulates each respondent's income taxes, including their filing status, by applying the tax rules to their reported income and family characteristics, and releases these variables in files that accompany the public-use March CPS (U.S. Census Bureau, 1992). These filing statuses therefore reflect who was required or likely to file rather than actual filing behavior.

[41] Piketty and Saez (2003) define potential tax units as married couples and single adults aged 20 or older. Since

restriction selects on earnings. First, because the threshold for married couples applies to combined income, a married worker with low earnings is required to file whenever the couple's combined income exceeds $5,400, and we observe both spouses on a joint return. Second, many workers below the threshold filed anyway to receive a refund of income taxes withheld from their wages. Third, low-income workers with children had an additional incentive to file to claim the Earned Income Credit, which was worth up to $500 in 1979.[42] Nonetheless, the restriction excludes individuals that had little or no income in 1979 and were not married to a filer, so it selects individuals that are more attached to the labor force and have higher earnings. Combined with the age gradient in the survey frame, this means that our sample is younger, more likely to be employed, and has higher earnings than the corresponding cohorts in the 1979 CPS ASEC (see Table 1).

A.1.3 *Variable Definitions*

Employment. Following Davis and von Wachter (2011), Autor et al. (2014), and Yagan (2019), an individual is *employed* in a given year if they have positive wage earnings in the SER in that year. This captures employment at any point during the year. Since the SER covers the full sample period, this measure is available from 1976 onwards.

High labor force attachment. An individual has *high labor force attachment* in a given year if their annual SER wage earnings exceed what a worker would earn working 1600 hours (forty weeks at forty hours per week) at the prevailing federal minimum wage.[43] Since these thresholds are well below the taxable maximum, the SER topcoding does not affect this measure. This measure captures intensive margin employment adjustments, such as longer or more frequent unemployment spells lasting less than a year, that our employment measure misses.

Earnings. Our primary earnings measure is annual wage earnings from the DER, which is available from 1978 onwards. We also use annual earnings from the SER, which are available for the full 1976–1989 period but topcoded at the taxable maximum. Our main earnings outcome is log annual earnings, which is only defined for individuals with positive earnings in

some returns are filed by individuals that are claimed as dependents on another return, this ratio somewhat overstates the share of tax units that filed.

[42] In 1979 the Earned Income Credit was equal to 10 percent of the first $5,000 of earned income and was phased out between $6,000 and $10,000 of income. It was only available to filers with a qualifying child.

[43] The federal minimum wage was $2.30 in 1976 and 1977, $2.65 in 1978, $2.90 in 1979, $3.10 in 1980, and $3.35 between 1981 and 1989. The corresponding thresholds are $3,680 in 1976, $4,640 in 1979, and $5,360 between 1981 and 1989 (nominal dollars).

that year, and we also report effects on earnings in levels which include zeros. All dollar values are converted to 2015 dollars using the CPI-U.

Pre-recession earnings decile. We assign each individual to a decile of the 1979 annual earnings distribution among individuals in our sample with positive 1979 earnings. Individuals with zero earnings in 1979 are assigned to a separate decile 0 category. These deciles enter the group fixed effects in our main specification (equation (2)) and define the groups in our heterogeneity analysis (Figure 4).

Pre-recession industry. We assign each individual the industry of their 1979 employer, which we obtain by linking the EIN on their 1979 W-2 to the LBD. For individuals with multiple employers in 1979, we use the employer from which they received the highest earnings. Industry is measured using the 4-digit vintage-consistent NAICS code that the LBD assigns to the EIN in 1979. Individuals without a 1979 W-2, and those whose 1979 employer does not appear in the LBD, such as government employers, are assigned to a separate missing industry category. Together with age and the pre-recession earnings decile, industry defines the 1979 group cells $g(i, 79)$ in our main specification (equation (2)), so that each regression compares workers of the same age, in the same 4-digit industry, and in the same part of the earnings distribution in 1979.

Demographics. We measure age using the date of birth in the Numident. We measure race and sex using the 2000 and 2010 Decennial Census and the individual's survey responses, and fill in any remaining missing values using the Numident. Education is measured using the individual's survey and Census responses. When multiple responses are available for a characteristic, we use the modal response, using the most recent response as the tiebreaker. Marital status is measured using the filing status on the individual's 1979 tax return. An individual is self-employed if they have positive self-employment income recorded in the SER or DER in 1979.

A.1.4 *Reweighting*

As discussed in Section IV.A, our sample is selected relative to the overall population along several dimensions. To address this, we reweight our sample to match the 1979 CPS ASEC for the same birth cohorts (1925–1956), which we describe in more detail in Appendix A.7. Specifically, we use raking to match the marginal distributions of age (shares in age bins 23–29, 30–34, 35–39, 40–44, 45–49, and 50–54), sex, race (white, Black, and other), education (less than high school, high school, some college, and college or more), and 1979 annual wage

earnings (the share with zero earnings and the share in each decile of the positive earnings distribution in the CPS ASEC). All of our main estimates use these weights, and Appendix C reports unweighted versions of our main results.

A.2 COMMUTING ZONES AND RECESSION EXPOSURE

This section describes our definition of local labor markets and the construction of our measures of local exposure to the 1980s recession. Following Yagan (2019), our main measure of exposure is the change in the commuting zone unemployment rate between 1979 and 1982. For our analyses using the CPS, which only identifies an individual's state of residence, we instead measure exposure at the state level.

A.2.1 *Commuting Zones*

Our local labor markets are commuting zones (CZs), which are clusters of counties with strong commuting ties between them (Tolbert and Sizer, 1996). We use the 1990 CZ definitions and map counties to CZs using the crosswalk from Autor, Dorn and Hanson (2013). We restrict our attention to the 722 CZs that cover the contiguous United States, dropping Alaska and Hawaii. In the worker panel, we assign each individual to the CZ of the county on their 1979 tax return (see Appendix A.1). All of our other CZ-level measures are constructed by aggregating county-level data to CZs using the same crosswalk.

A.2.2 *Recession Exposure*

We measure local unemployment using data from the Bureau of Labor Statistics (BLS) Local Area Unemployment Statistics (LAUS) program, which produces monthly and annual estimates of employment, unemployment, and the labor force for states and substate areas. Since the CPS sample is too small to produce reliable estimates for most areas, LAUS estimates are constructed by combining survey data with administrative records (U.S. Bureau of Labor Statistics, 2026).

During our sample period, statewide estimates for 10 large states came directly from the monthly CPS starting in 1978, while estimates for the remaining states were produced using the Handbook method and adjusted to CPS controls.[44] County estimates are produced using the Handbook method, a building-block approach in which each component of employment and unemployment is estimated separately and then summed. Employment combines payroll employment from employer reports, adjusted from place of work to place of residence, with estimates of agricultural, self-employed, and other workers not on payrolls. Unemployment combines counts of unemployment insurance claimants with estimates of those who have exhausted their benefits and of new entrants and reentrants to the labor force, who are

[44] The 10 states were California, Florida, Illinois, Massachusetts, Michigan, New Jersey, New York, Ohio, Pennsylvania, and Texas. North Carolina was added in 1985. Starting in 1989, estimates for the other states were instead produced using time-series models (U.S. Bureau of Labor Statistics, 2026).

not covered by unemployment insurance. The resulting county estimates are then scaled so that they sum to the statewide totals (U.S. Bureau of Labor Statistics, 2026).

Official LAUS estimates for counties are only available starting in 1990. For the period before 1990, we use unofficial county-level tabulations covering 1976–1989 that we obtained by request from the BLS. These data contain annual average employment, unemployment, and labor force for every county in the United States, and they appear to be the same estimates that underlie contemporaneous BLS publications on regional unemployment during this period.

We make a small number of adjustments to these data so that county definitions are consistent with the CZ crosswalk. We combine Columbus city with Muscogee County in Georgia, which consolidated in the early 1970s, and Washabaugh County with Jackson County in South Dakota, which merged in 1983. We then sum employment, unemployment, and the labor force across the counties in each CZ, and compute the CZ unemployment rate as total unemployment divided by the total labor force. Our measure of exposure to the recession is the change in this unemployment rate between 1979 and 1982

$$\Delta\text{Unemp}_c = u_{c,1982} - u_{c,1979}$$

measured in percentage points. Both years are annual averages, so the measure spans the full period covering the two recessions of January–July 1980 and July 1981–November 1982. As discussed in Section IV, the average CZ experienced a 4.18 percentage point increase in unemployment, and the interquartile range across CZs is 3.17 percentage points.

A.2.3 *Validation of the Unofficial County Estimates*

Since our main exposure measure relies on unofficial data, we assess their quality in several ways. It is worth first noting what distinguishes these data from the official LAUS estimates. According to documentation provided by the BLS with the archived files, county estimates for 1976–1989 were stored as they were last published for each reference year and have not been revised since, so they are not consistent with the official estimates from 1990 onwards (U.S. Bureau of Labor Statistics, 1998). Within the pre-1990 period, the main methodological change for county estimates was the factors used to adjust employment from place of work to place of residence, which were based on the 1970 Census for 1976–1982 and on the 1980 Census for 1983–1989. A new procedure for allocating employment to counties within labor market areas was also introduced in 1985 (U.S. Bureau of Labor Statistics, 1998). Our exposure measure is constructed from the 1979 and 1982 estimates, so both years are produced under the same methodology. In contrast, the official statewide estimates were re-estimated back to 1978 using time-series models in 1989, 1994, and 1996, so we should not expect the unofficial

county estimates to sum exactly to the current official state estimates.

Comparison to official state estimates. We first sum the unofficial county estimates to the state level and compare them to the official LAUS estimates for each state. Across the 48 contiguous states, the two unemployment rates differ by between 0.12 and 0.25 percentage points on average in each year between 1976 and 1989, and their correlation is between 0.98 and 0.996 in every year. The largest average difference is in 1983, the first year in which the 1980 Census-based residency adjustment was used. Figure A.1 compares the change in the unemployment rate between 1979 and 1982 implied by the two sources. The correlation between them is 0.98 (0.99 weighting by the 1979 labor force), and a regression of one on the other has a slope of 1.04. Only Delaware differs substantially, which may reflect subsequent revisions to its official statewide estimates. The two sources also agree on which states were severely exposed. Of the 48 states, 44 are on the same side of the median under both measures, and the four that are not (Louisiana, Minnesota, Rhode Island, and Utah) are within 0.1 percentage points of the median.

Summed to the national level, the two sources are nearly identical. The national unemployment rate implied by the unofficial county estimates is within 0.01 percentage points of the sum of the official state estimates in every year between 1976 and 1989 (see Figure A.3A).

Comparison to the Decennial Census. Since the unofficial estimates are controlled to statewide totals, the comparison to official state estimates is not informative about how well they allocate unemployment across areas within a state, which is the variation our design relies on. To assess this, we compare the 1980 CZ unemployment rate implied by the unofficial estimates to the CZ unemployment rate in the 1980 Decennial Census, which is measured independently of the LAUS. We construct the Census unemployment rate for the civilian labor force aged 16 and older outside of institutional group quarters, using the same geographic crosswalks as in Appendix A.4. The Census measures labor force status in a reference week in April rather than over the full year, so we do not expect the two to agree exactly. As a benchmark, we repeat the same comparison using the official county estimates and the 1990 Census.

Figure A.2 displays the results. The unofficial estimates agree with the Census at least as well as the official estimates do. The labor force weighted correlation between the unofficial LAUS and the Census unemployment rate in 1980 is 0.88, compared to 0.78 between the official LAUS and the Census in 1990. Removing state means, so that the comparison only uses variation across CZs within the same state, the correlations are 0.68 in 1980 and 0.65 in 1990.

Continuity at the 1989/1990 seam. We next examine whether the unofficial estimates line up with the official county estimates that begin in 1990. If the unofficial estimates were systematically biased, we would expect a jump in CZ unemployment rates between 1989 and 1990 relative to the corresponding change in the official statewide estimates, which are consistent across the two years. We do not find evidence of a systematic shift. Weighting by the labor force, the average difference between the change in a CZ's unemployment rate and the change in its state's official unemployment rate between 1989 and 1990 is −0.01 percentage points, which is similar to every other year between 1985 and 1995.[45] Changes across the seam are somewhat noisier than in the years that follow. The standard deviation of the CZ change relative to the state change is 1.01 percentage points between 1989 and 1990, compared to between 0.81 and 0.87 percentage points between 1990 and 1992. However, they are not noisier than some changes within the unofficial data, such as those between 1984 and 1985 and between 1985 and 1986 (1.07 and 1.19 percentage points), when the procedure for allocating employment to counties changed. Figure A.3 displays the national series and the series for the largest CZ in four states, none of which show a large discontinuity at the seam.

Comparison to establishment employment. Finally, we verify that our exposure measure reflects reductions in local labor demand by comparing it to the change in private sector employment between 1979 and 1982 in the QCEW, which is based on establishment records rather than the LAUS. Figure A.4 shows that the two are strongly related. The labor force weighted correlation is −0.69, and each additional percentage point increase in a CZ's unemployment rate is associated with 2.7 log points lower private sector employment growth over the recession. The results are similar using total employment from the BEA (a correlation of −0.66).

Together, these results suggest that the unofficial county estimates are a reliable measure of local unemployment during this period. They are consistent with the official statewide estimates, agree with an independent measure of local unemployment at least as well as the official county estimates do, show no systematic break when the official series begins, and track establishment-based measures of the local employment shock.

[45] The unweighted average difference is −0.10 percentage points, suggesting that the unemployment rate in smaller CZs falls slightly across the seam.

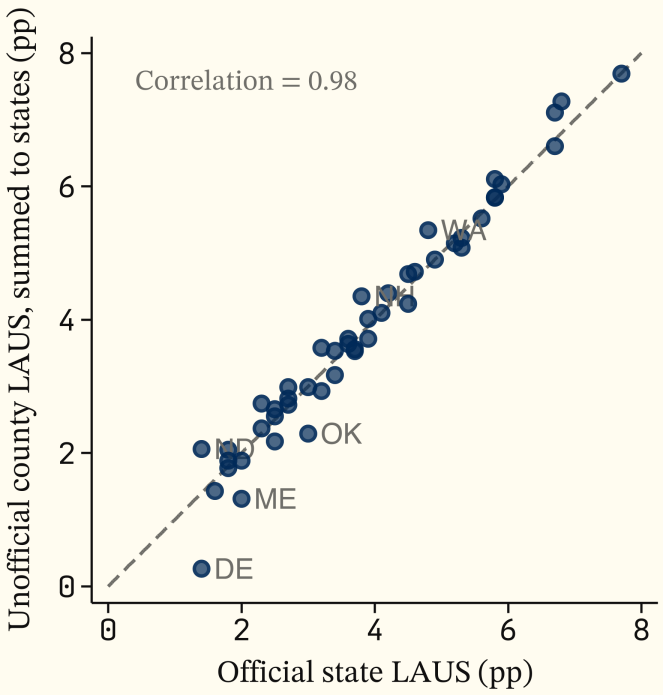


FIGURE A.1
Change in State Unemployment 1979–1982: Unofficial County and Official State LAUS Estimates

Notes: This figure compares the change in the state unemployment rate between 1979 and 1982 implied by the unofficial county LAUS estimates summed to the state level (vertical axis) to the change in the official state LAUS estimates (horizontal axis) for the 48 contiguous states. The dashed line is the 45 degree line. States whose two measures differ by more than 0.5 percentage points are labeled.

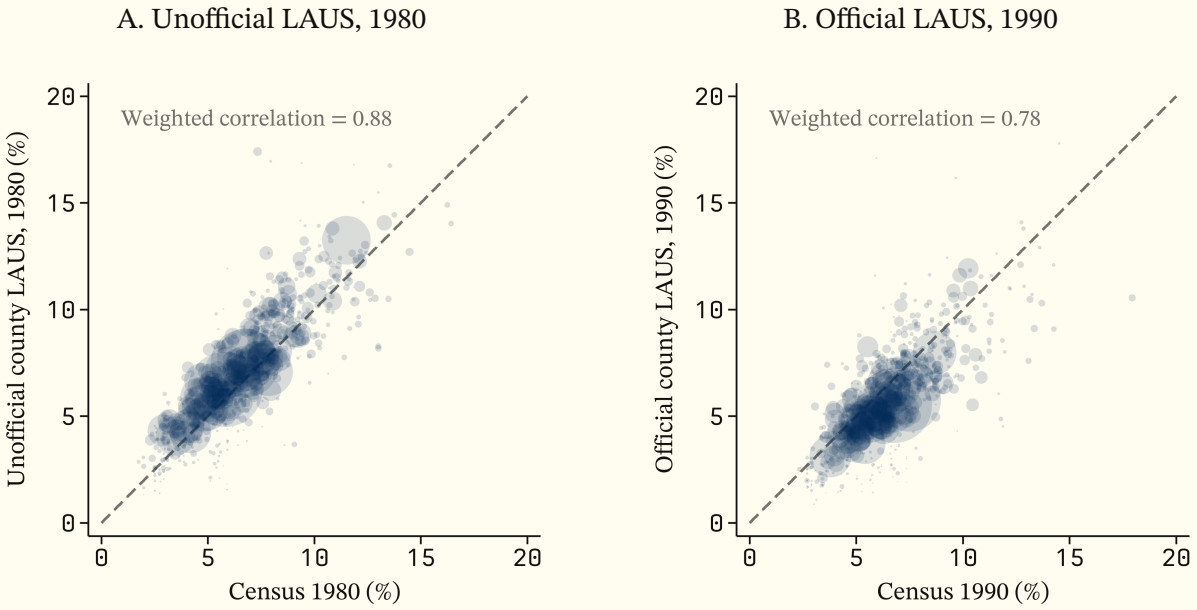


FIGURE A.2
CZ Unemployment Rates in the LAUS and the Decennial Census

Notes: This figure compares the CZ unemployment rate implied by the county LAUS estimates (vertical axis) to the CZ unemployment rate in the Decennial Census (horizontal axis). Panel A uses the unofficial county estimates and the 1980 Census, and Panel B uses the official county estimates and the 1990 Census. The Census unemployment rate is computed for the civilian labor force aged 16 and older outside of institutional group quarters. Markers are proportional to the CZ labor force. The dashed line is the 45 degree line. For legibility, the figure only displays CZs with unemployment rates between 0 and 20 percent in both sources, which excludes 2 CZs in Panel A and 3 CZs in Panel B. Correlations are weighted by the labor force and computed using all CZs.

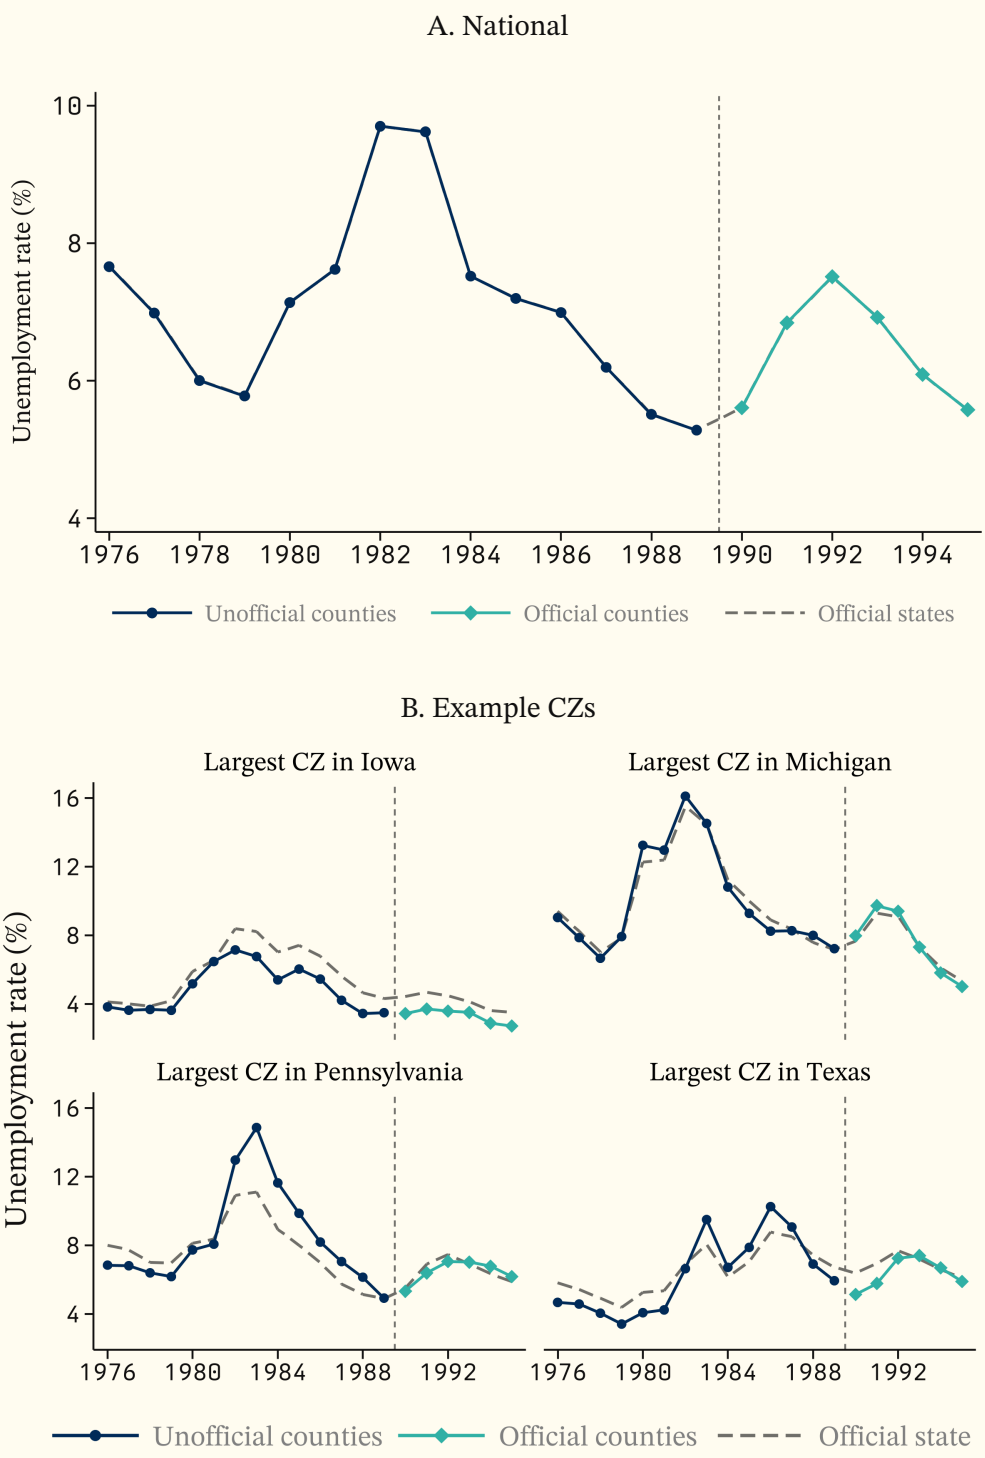


FIGURE A.3
Unemployment Rates Across the 1989/1990 Seam

Notes: This figure displays annual average unemployment rates implied by the unofficial county LAUS estimates (1976–1989), the official county LAUS estimates (1990–1995), and the official state LAUS estimates (dashed). Panel A sums each source across the 48 contiguous states. Panel B displays the largest CZ (by 1979 labor force) in Iowa, Michigan, Pennsylvania, and Texas alongside the official estimates for its state. The vertical line marks the transition from the unofficial to the official county estimates.

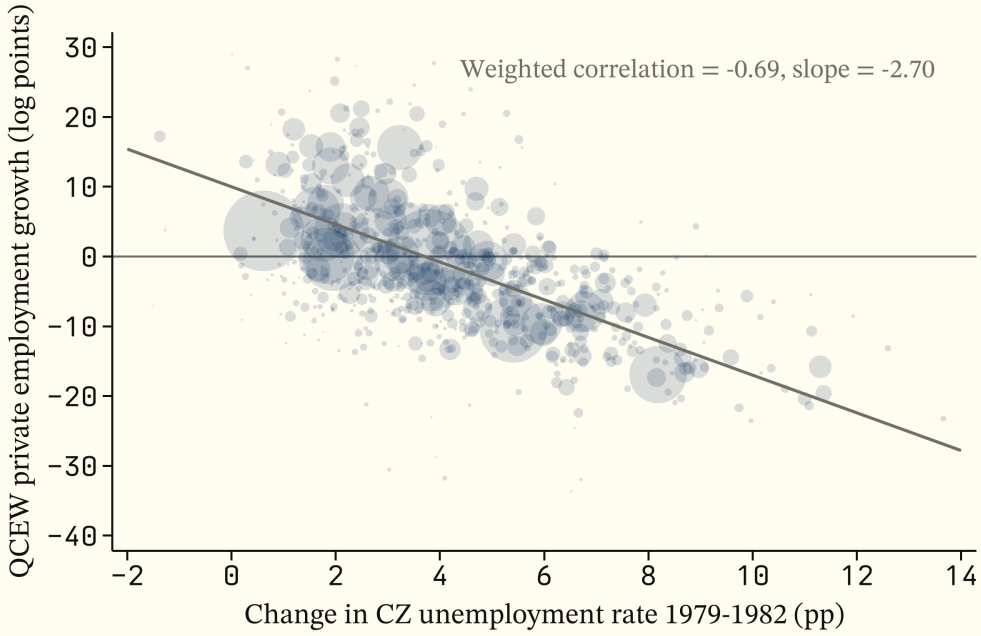


FIGURE A.4
CZ Recession Exposure and Private Employment Growth 1979–1982

Notes: This figure displays the relationship between the change in the CZ unemployment rate between 1979 and 1982, constructed from the unofficial county LAUS estimates, and the log change in private sector employment between 1979 and 1982 in the QCEW. Markers are proportional to the CZ labor force in 1979, and the line is a labor force weighted linear fit. For legibility, the figure only displays CZs with changes in unemployment between −2 and 14 percentage points and employment growth between −40 and 30 log points, which excludes 18 small CZs. The linear fit and the correlation are computed using all CZs.

A.2.4 *State-Level Exposure*

The CPS only identifies an individual's state of residence, so in our analyses of union membership (Section VII) and of the wage distribution (Sections VI and VIII) we measure exposure at the state level. We use official LAUS annual average unemployment rates for states, which are available from 1976 onwards, and define state exposure analogously as $\Delta\text{Unemp}_s = u_{s,1982} - u_{s,1979}$. We restrict to the 48 contiguous states, dropping Alaska, Hawaii, and the District of Columbia. In the distributional analysis, a state is *severely exposed* ($S = 1$) if its increase in unemployment was above the median across these 48 states.

Figure A.5 displays the distribution of state exposure along with the median that separates the severely and less severely exposed states.

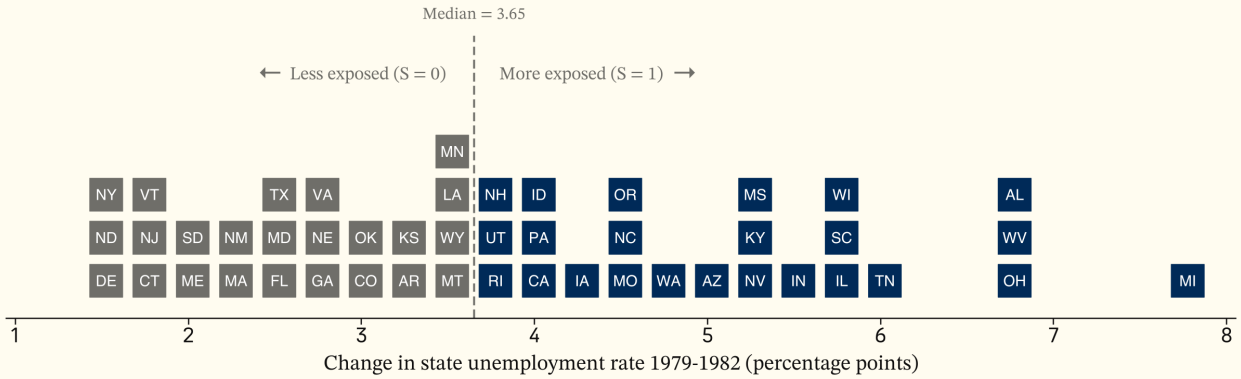


FIGURE A.5
Distribution of State Recession Exposure

Notes: This figure displays the distribution of the change in the state unemployment rate between 1979 and 1982 across the 48 contiguous states, measured using official LAUS annual averages. Each square is a state, labeled by its abbreviation and stacked in bins of 0.25 percentage points. The dashed line marks the median of 3.65 percentage points, which separates the severely exposed ($S = 1$, blue) and less severely exposed ($S = 0$, gray) states used in the distributional analysis.

A.2.5 *Industry Exposure*

Figure 1B displays the change in unemployment between 1979 and 1982 across industries. We measure industry unemployment rates using the monthly CPS for 1979 and 1982, pooling all months in each year. Industries are defined using the 3-digit 1970 Census industry classification, and unemployed individuals are assigned to the industry of their most recent job. We restrict to individuals aged 16–64 with a reported industry and compute the unemployment rate in each industry and year as the weighted share of labor force participants that are unemployed. We use the CPS final weights, except for respondents to the March supplement

for whom we use the supplement weights. To limit noise from small cells, we restrict to industries with at least 200 observations in each year. For the figure, we group industries into manufacturing, mining, construction, and all other industries.

Figure A.6 reproduces the distribution, with each square representing an industry.

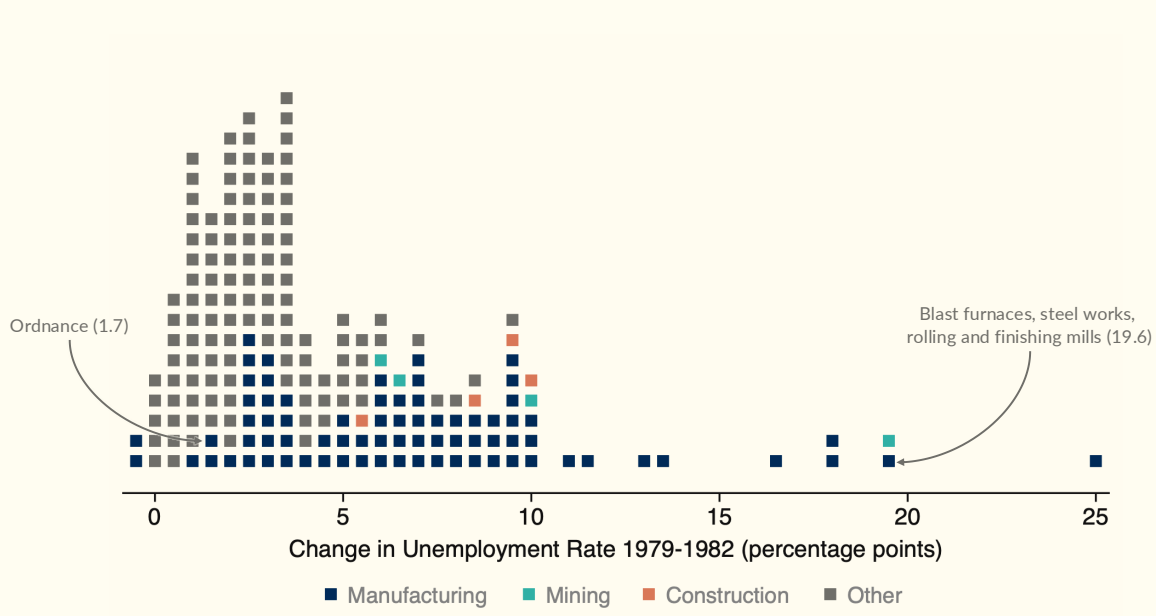


FIGURE A.6
Distribution of Industry Recession Exposure

Notes: This figure displays the distribution of the change in the unemployment rate between 1979 and 1982 across 3-digit 1970 Census industries, measured using the monthly CPS. Each square is an industry, colored by its sector. Industries with fewer than 200 observations in either year are excluded. Unemployed individuals are assigned to the industry of their most recent job.

A.3 COMMUTING ZONE CHARACTERISTICS

Throughout the paper, we use a number of pre-recession characteristics of commuting zones, both to characterize where the recession hit hardest (Tables 2 and B.1) and as controls in our CZ-level and worker-level analyses. This section describes the underlying data sources and then lists each variable. Unless otherwise noted, all county-level data are aggregated to CZs using the same county to CZ crosswalk described in Appendix A.2, and we restrict to the 722 CZs in the contiguous United States.

A.3.1 *Data Sources*

BEA Regional Economic Accounts. The Bureau of Economic Analysis (BEA) publishes annual county-level estimates of employment and earnings by industry as part of its Regional Economic Accounts. We use county employment by industry (CAEMP25) and earnings by industry (CAINC5) between 1969 and 2000, which classify industries using the SIC system during this period. BEA employment counts the number of jobs by place of work, including both wage and salary jobs and self-employment. Some county-industry cells are suppressed to avoid disclosing information about individual employers. We follow the cleaning procedures in Hershbein and Stuart (2024), which, among other things, combine independent cities in Virginia with their surrounding counties so that county definitions are consistent over time.

Quarterly Census of Employment and Wages. The Quarterly Census of Employment and Wages (QCEW), produced by the BLS, contains employment and total wages for the near universe of establishments covered by state unemployment insurance programs. We use annual county-level data on private sector employment and wages, which are available starting in 1975. We restrict to private sector employment since coverage of government employment changes over the period. Wages are converted to 2015 dollars using the CPI-U.

SEER population estimates. We measure county population using the population estimates produced by the National Cancer Institute's Surveillance, Epidemiology, and End Results (SEER) program, which adjust the Census Bureau's county population estimates to be consistent over time and are available by age starting in 1969.

Commuting zone characteristics from Autor, Dorn and Hanson (2013). We use several CZ characteristics constructed by Autor, Dorn and Hanson (2013) from the Decennial Census, which are available for 1970 and 1980. These include the share of employment in manufacturing, the share of employment in routine occupations, the ratio of college to non-college

educated adults, the share of the non-college population that is foreign born, the female share of employment, and the share of the population aged 65 and older. Unless otherwise noted, we use the 1980 values, which measure characteristics as of the 1980 Census, just prior to the recession.

Decennial Census. For our CZ long-difference regressions, we also control for the change in average log earnings and the change in log working-age population over the preceding decade, which we construct from the Decennial Census microdata as described in Appendix A.4.

Import penetration. We measure local exposure to import competition using data on import penetration by 4-digit SIC industry from Schott (2008), defined as imports divided by imports plus domestic production net of exports. We assign import penetration to CZs by weighting each industry's import penetration in 1979 by its share of CZ manufacturing employment in 1979, measured using the County Business Patterns (CBP). Since the CBP suppresses employment counts for many small county-industry cells, we use county-industry employment imputed following the procedure in Autor and Dorn (2013) and Autor, Dorn and Hanson (2013), which bounds each suppressed cell using the reported number of establishments by size class. We convert industry codes between SIC vintages using the NBER-CES Manufacturing Industry Database concordance.

State union density. We measure state-level union density using the CPS-based state union membership series from the Farber et al. (2021) replication files. For pre-recession changes in union density, which require years before the CPS consistently identifies union membership, we use the series in the same files that combines the CPS with earlier survey data.

Census regions and divisions. We assign states to the nine Census divisions and four Census regions. In some robustness exercises, we also define the Rust Belt as Illinois, Indiana, Michigan, New York, Ohio, Pennsylvania, and Wisconsin; the South as it is defined by the Census Region; and the West Coast as California, Oregon, and Washington.

A.3.2 *Exposure to Earlier Recessions*

Since LAUS data are not available before 1976, we cannot measure exposure to earlier recessions using changes in unemployment. Instead, we measure CZ exposure to the 1969–1970 and 1973–1975 recessions using the log change in total BEA employment over each recession, $\ln E_{c,1970} - \ln E_{c,1969}$ and $\ln E_{c,1975} - \ln E_{c,1973}$. Note that, unlike our main exposure measure, more negative values of these measures indicate greater exposure. We use these measures as

controls for the lingering effects of earlier downturns, for instance in our CZ long-difference regressions (Table 6) and in our worker-level robustness checks (Table 3).

A.3.3 *Variable Definitions*

The following lists each CZ-level variable used in the paper, along with its source and the years over which it is measured.

- *Industry employment shares (1979):* share of total CZ employment in each of ten broad industries (agriculture; mining; construction; manufacturing; transportation and utilities; wholesale trade; retail trade; finance, insurance, and real estate; services; and government). Source: BEA.
- *Manufacturing employment growth (1969–1978):* log change in CZ manufacturing employment. Source: BEA.
- *Change in manufacturing share (1969–1978):* log change in the manufacturing share of CZ employment. Source: BEA.
- *Employment growth (1969–1978):* log change in total CZ employment. Source: BEA.
- *Population growth (1969–1978):* log change in total CZ population. Source: SEER.
- *Earnings growth (1975–1978):* log change in average annual private sector earnings per worker, measured in 2015 dollars. Source: QCEW.
- *Exposure to the 1973–1975 recession:* $\ln E_{c,1975} - \ln E_{c,1973}$, where $E_{c,t}$ is total CZ employment. Source: BEA.
- *Exposure to the 1969–1970 recession:* $\ln E_{c,1970} - \ln E_{c,1969}$. Source: BEA.
- *Import penetration (1979):* CZ manufacturing employment-weighted average of industry import penetration. Source: Schott (2008) and CBP.
- *Manufacturing share of employment (1980):* Source: Autor, Dorn and Hanson (2013).
- *Routine share of employment (1980):* share of CZ employment in routine-task-intensive occupations. Source: Autor, Dorn and Hanson (2013).
- *Relative supply of college educated workers (1980):* ratio of college to non-college educated adults. Source: Autor, Dorn and Hanson (2013).

- *Foreign-born share of the non-college population (1980):* Source: Autor, Dorn and Hanson (2013).
- *Female share of employment (1980):* Source: Autor, Dorn and Hanson (2013).
- *Population share aged 65 and older (1980):* Source: Autor, Dorn and Hanson (2013).
- *Lagged earnings growth (1970–1980):* change in average log annual wage earnings among wage and salary workers aged 16–64 over the preceding decade, used in the CZ long-difference regressions. For the 1970–1980 placebo specifications, we instead use the change between 1960 and 1970. Source: Decennial Census.
- *Lagged population growth (1970–1980):* log change in the population aged 16–64 over the preceding decade (1960–1970 for the placebo specifications). Source: Decennial Census.
- *State union density (1979):* union membership as a share of employed wage and salary workers in the CZ's state. Source: Farber et al. (2021).
- *Change in state union density (1969–1978):* log change in state union density. Source: Farber et al. (2021).

A.4 DECENNIAL CENSUS AND AMERICAN COMMUNITY SURVEY

To study the long-run impacts of the recession on CZ labor markets (Tables 6 and 7), we use public-use microdata from the Decennial Census and the American Community Survey (ACS) obtained from IPUMS USA (Ruggles et al., 2024). This section describes the samples, the construction of our earnings and wage measures, and how we aggregate them to CZs.

A.4.1 *Samples*

We use the 1970 1% Form 1 Metro sample, the 1980 and 1990 5% samples, and the 2015–2019 5-year ACS sample. We also use the 1960 5% sample to construct the pre-period controls for our 1970–1980 placebo regressions. In each sample, we restrict to individuals between the ages of 16 and 64 and drop individuals living in institutional and other group quarters. Since the Census asks about earnings and weeks worked during the previous calendar year, the 1970, 1980, and 1990 samples measure labor market outcomes in 1969, 1979, and 1989, and the 2015–2019 ACS measures outcomes over the preceding twelve months.

A.4.2 *Geography*

The public-use Census samples do not identify counties for most individuals. Instead, the smallest geographic units are county groups in 1970 and 1980 and Public Use Microdata Areas (PUMAs) from 1990 onwards, which can span more than one CZ. We assign these units to CZs using the crosswalks developed by Dorn (2009) and Autor, Dorn and Hanson (2013), which give the share of each county group's or PUMA's population that lives in each CZ. When a unit spans several CZs, we split each individual across the corresponding CZs, multiplying their person weight by the share of the unit's population in each CZ.

A.4.3 *Earnings and Wages*

Earnings. We measure annual earnings using wage and salary income. We set earnings to missing for individuals whose wage and salary income was allocated by the Census Bureau, to avoid the known biases associated with these imputations. Earnings above the topcode are multiplied by 1.4.[46] All dollar values are converted to 2015 dollars using the CPI-U for the year in which the earnings were earned. Our *positive earnings* outcome is an indicator for having any wage and salary income during the year. Our *log earnings* outcome is the log of annual earnings, which we construct for wage and salary workers with positive earnings.

[46]The topcodes are \$50,000 in 1970, \$75,000 in 1980, and \$140,000 in 1990. In the ACS, the top 0.5 percent of earnings in each state are replaced with the state mean of earnings above that threshold, so we multiply earnings above the 99th percentile of the earnings distribution in each state by 1.4.

Hourly wages. We construct hourly wages by dividing annual earnings by annual hours worked, which is the product of weeks worked and usual weekly hours during the year. We restrict hourly wages to full-time full-year wage and salary workers, defined as those who usually worked at least 35 hours per week and at least 40 weeks during the year. To limit the influence of measurement error, we drop hourly wages below $1 and above $100 in 1979 dollars. Since the 1970 Census only reports weeks worked in intervals, we do not construct hourly wages for 1970, which is why Table 6 does not report effects on hourly wages for the 1970–1980 placebo period.

A.4.4 *Occupations*

To measure employment in routine occupations (Table 7), we map each individual's occupation to the consistent occupation classification (*occ1990dd*) developed by Dorn (2009) and Autor and Dorn (2013). We classify an occupation as *routine* using the routine task-intensity classification from Autor, Dorn and Hanson (2013), in which routine occupations are those in the top third of the 1980 employment-weighted distribution of routine task intensity. We also group occupations into four broad categories based on their major occupation group. Routine cognitive occupations are sales and administrative support occupations, routine manual occupations are precision production, craft, and repair occupations and operators, fabricators, and laborers, nonroutine cognitive occupations are managerial, professional, and technical occupations, and nonroutine manual occupations are service occupations. Individuals without a reported occupation are excluded from these measures.

A.4.5 *Aggregation to Commuting Zones*

We aggregate each outcome to the CZ level by taking the weighted mean across individuals in each CZ, using person weights multiplied by the CZ allocation factors described above. Shares of employment in routine occupations are computed among individuals with a reported occupation, and occupation-specific earnings are computed among wage and salary workers in each occupation group. We also construct the CZ working-age population as the sum of these weights among individuals aged 16–64. Our long-difference outcomes are the changes in these CZ means between Census years, between 1970 and 1980 for our placebo, between 1980 and 1990, and between 1980 and the 2015–2019 ACS.

A.5 CPS UNION MEMBERSHIP

Our analysis of the recession's effects on union membership (Figure 8A and Table D.3) uses individual-level data from the Current Population Survey (CPS) compiled by Farber et al. (2021). This section describes the data, the construction of our union membership outcome, and our controls.

A.5.1 *Data*

The CPS first asked about union membership on a regular basis in the May supplement between 1973 and 1981, and has asked it every month of the outgoing rotation groups (ORG) since 1983. The union question was not asked in 1982. We therefore combine the May supplements for 1977–1981 with the monthly ORG samples for 1983–1989, using the harmonized individual-level file from the Farber et al. (2021) replication package. The ORG samples are roughly twelve times larger than the May supplements, so our post-recession estimates are considerably more precise than our pre-recession estimates. The file covers individuals aged 21 to 65, and we drop Alaska, Hawaii, and the District of Columbia. All estimates are weighted by the CPS individual weights, which are normalized to average one within each year. As a result, the larger ORG samples receive more weight than the May supplements in our pooled regressions.

A.5.2 *Union Membership*

The union membership question is only asked of employed wage and salary workers. Since we are interested in the recession's effect on union membership in the population, including through its effects on employment, we code individuals who are not employed as non-members. Our outcome is therefore an indicator for being a union member as a share of the working-age population, rather than the conventional measure of union density, which is defined over employed wage and salary workers. In 1979, 16.7 percent of our sample were union members and 66.6 percent were employed, which implies a union density of around 25 percent among the employed. In Section VII, we use these two components to separate the mechanical effect of the recession's disemployment effects on union membership from its effect on union density.

We also separately measure membership in private and public sector unions, based on whether the individual is employed in the private or public sector. Individuals who are employed but neither in the private nor the public sector, such as the self-employed, are coded as non-members in both.

A.5.3 *Controls*

All of our specifications control for a quadratic in age, sex, race (Black and other race indicators), education (indicators for high school dropouts, high school graduates, some college, college graduates, and graduate degrees), and education interacted with age, along with state and year fixed effects. Column 4 of Table D.3 additionally controls for the pre-recession state characteristics described in Appendix A.3, each interacted with year fixed effects. Since these characteristics are measured at the CZ level, we aggregate them to the state level by taking the average across CZs in each state, weighting by 1979 population. Column 5 includes Census region-by-year fixed effects, and column 6 includes industry-by-year fixed effects, where industries are defined using the industry classification in the Farber et al. (2021) file.

A.5.4 *Data Considerations*

Two features of the data are worth noting when interpreting our estimates. First, the switch from the May supplements to the ORG samples in 1983 coincides with our post-recession period. Since we include year fixed effects, any change in the level of measured union membership due to the change in survey instruments is absorbed, and only differential changes across more and less exposed states would affect our estimates.

A.6 CPS WAGE DATA FOR THE DISTRIBUTIONAL ANALYSIS

Our analysis of the recession's impact on the wage distribution (Sections VI and VIII) uses the CPS wage data constructed by Fortin, Lemieux and Lloyd (2021), which follow the sample selection criteria and variable definitions in DiNardo, Fortin and Lemieux (1996). Using the same data allows us to compare our estimates directly to the existing decompositions of the role of unions in the rise in wage inequality during this period. This section describes the data and the construction of our union coverage measures.

A.6.1 *Samples*

Our main analysis compares the wage distribution in 1979 and 1988. For 1988 (and the intermediate years 1983–1987), we use the Merged Outgoing Rotation Group (MORG) files of the CPS. Since the CPS did not ask about union status in the outgoing rotation groups until 1983, our 1979 sample instead comes from the May 1979 CPS supplement, which asked about union coverage, matched to the May–August 1979 MORG as in DiNardo, Fortin and Lemieux (1996) and Fortin, Lemieux and Lloyd (2021). For our pre-recession placebo exercise (Appendix E.2), we use 1977 as the pre-period, which also comes from the May CPS. Since the 1977 and 1979 samples both come from a single May supplement, they are roughly a third to a quarter of the size of the samples from 1983 onwards.

We restrict to men between the ages of 16 and 64 and drop Alaska, Hawaii, and the District of Columbia. We also drop observations with missing or allocated wages. Since the CPS imputation procedure for wages does not condition on union status, including allocated wages would attenuate the estimated union wage gap (Fortin, Lemieux and Lloyd, 2021). Our final sample consists of 20,706 men in 1979 and 71,391 men in 1988. All estimates are weighted using the CPS earnings weights.

A.6.2 *Wages*

Our wage measure is the log hourly wage constructed by Fortin, Lemieux and Lloyd (2021). For workers paid by the hour, the hourly wage is the hourly wage that the worker directly reports. For other workers, the hourly wage is usual weekly earnings divided by usual weekly hours. Topcoded earnings are imputed using a stochastic Pareto imputation (Fortin, Lemieux and Lloyd, 2021). We estimate all of our distribution regressions using wages measured in 1979 dollars, and we construct the grid of thresholds for the distribution regressions on this scale.

A.6.3 *Union Coverage*

We measure unionization using union coverage, which includes both union members and workers who are not members but whose job is covered by a union contract, following Fortin, Lemieux and Lloyd (2021). In addition to each worker's own coverage status, our decomposition requires a measure of union density in each worker's labor market. Following Fortin, Lemieux and Lloyd (2021), we define a market as a state-by-industry cell, using the ten industries listed below, and measure market union density as the share of workers in the cell who are covered by a union contract in each year.

Many state-by-industry cells contain few observations, particularly in the smaller 1979 sample. For cells with at least 25 observations in a year, we use the observed coverage rate. For smaller cells, we instead use the predicted coverage rate from a logit model of individual coverage status on state, year, and industry fixed effects, state-by-year and industry-by-year fixed effects, and industry-specific linear time trends, estimated on the pooled 1979 and 1983–1988 samples. The baseline market coverage rate that we include as a control in our distribution regressions is the market's coverage rate in 1979. For the handful of state-by-industry cells with no observations in 1979, we assign the average 1979 coverage rate across the state's other industries.

A.6.4 *Covariates*

The covariates in our distribution regressions follow Fortin, Lemieux and Lloyd (2021) and consist of:

- *Education and experience:* years of schooling, a quartic in potential experience, years of schooling interacted with experience, and 16 education-by-experience cells (four education groups by four ten-year experience groups).
- *Industry:* ten industry groups (agriculture, forestry, fishing, and mining; construction; durable manufacturing; nondurable manufacturing; wholesale and retail trade; transportation, communication, and utilities; finance, insurance, and real estate; other services; professional, health, education, and social services; and public administration).
- *Occupation:* four occupation groups (management, technical, and professional; clerical and sales; service; and production, transportation, and repair).
- *Other worker characteristics:* indicators for being married, non-white, working part-time, and living in a metropolitan area.

- *State characteristics:* the share of state employment in each of the ten industries in 1979 from the BEA (see Appendix A.3), and the 1979 market coverage rate described above.

To avoid cells that perfectly predict wages below a threshold, we pool the industry-by-occupation cells with fewer than 100 men in 1979 into a single residual category. All covariates are interacted with year indicators, as described in Section VI.A.

A.6.5 *Data Considerations*

Two features of the data are worth noting. First, as discussed above, the 1979 sample comes from the May supplement while the post-recession samples come from the MORG. Since all of our estimates are differences between severely and less severely exposed states, any common change in measured wages due to the change in survey instrument is differenced out. Our placebo exercise, which compares 1977 to 1979, uses the same instrument in both years.

A.7 CPS ANNUAL SOCIAL AND ECONOMIC SUPPLEMENT

We use the CPS Annual Social and Economic Supplement (ASEC), also known as the March supplement, as a benchmark for our worker panel. We use it both to assess how our sample compares to the overall population (Table 1) and as the target population when reweighting the worker panel (Appendix A.1.4). We obtain the ASEC from IPUMS CPS (Flood et al., 2024).

A.7.1 *Sample*

Since the ASEC asks about earnings during the previous calendar year, we use the 1980 and 1990 ASEC to measure outcomes in 1979 and 1989. To match the worker panel, we restrict to the same birth cohorts (1925–1956), which we assign using the respondent's age during the earnings year, so that the 1979 sample consists of individuals aged 23 to 54. We exclude members of the armed forces. All estimates are weighted by the ASEC person weights.

A.7.2 *Variable Definitions*

We construct each variable in Table 1 to match its definition in the worker panel as closely as possible.

- *Earnings:* annual wage and salary income during the previous calendar year, in 2015 dollars. We set earnings to missing for respondents whose wage and salary income was allocated by the Census Bureau, and multiply earnings above the topcode by 1.4. In the 1980 ASEC, wage and salary income is topcoded at $50,000. In the 1990 ASEC, earnings

from the longest job and from other jobs are reported and topcoded separately at $99,999, so we adjust each component and sum them. Average earnings include zeros.

- *Employed:* an indicator for positive wage and salary income during the year, matching our employment definition in the worker panel.
- *High labor force attachment:* an indicator for annual wage and salary income exceeding 1,600 times the prevailing federal minimum wage, as in the worker panel.
- *Demographics:* sex, age, race, and education. White refers to non-Hispanic white respondents, and less than college refers to respondents without a four-year college degree.
- *Industry:* indicators for manufacturing and for construction and mining, based on the industry of the respondent's longest job during the previous year. These use the 1970 Census industry classification in the 1980 ASEC and the 1980 classification in the 1990 ASEC. In the worker panel, industry is instead measured using the NAICS code of the worker's highest-paying 1979 employer (see Appendix A.1), so the industry shares are not exactly comparable across the two sources.

A.7.3 *Reweighting Targets*

The targets for our reweighting of the worker panel are the 1979 marginal distributions in this sample: the shares in six age groups (23–29, 30–34, 35–39, 40–44, 45–49, and 50–54), the female share, the shares in each race group, the shares in four education groups (less than high school, high school, some college, and college or more), and the shares with zero earnings and in each decile of the positive earnings distribution in 1979.

A.8 NLRB REPRESENTATION ELECTIONS

To study whether the recession changed workers' demand for unionization or unions' organizing activity (Figures 8B and D.3), we use administrative records of union representation elections conducted by the National Labor Relations Board (NLRB). In the private sector, workers who want a union to represent them can petition the NLRB to hold a secret ballot election, and the union is certified if a majority of voters vote in favor. We obtain these records from the Union Election Stats database (Kamphorst, Schaller and Young, 2025), which compiles NLRB election records from 1962 onwards and builds on earlier data collection efforts by Farber and Western (2001) and others.

A.8.1 *Sample*

The data contain one record for each bargaining unit in each election, including the number of eligible employees, the votes cast for and against union representation, the dates of the petition and the election, and the location of the employer. We restrict to elections held between 1975 and 1989. We further restrict to representation (RC) petitions, which are filed by employees or a union seeking to be certified as the bargaining representative, and exclude decertification (RD) petitions and petitions filed by employers (RM). This leaves 75,234 elections nationwide. We assign each election to the CZ of the employer's county and to the year in which the election was held, and drop elections in Alaska, Hawaii, and the District of Columbia.

A.8.2 *Outcomes*

We construct two outcomes for each CZ and year.

- *Average union vote share:* the share of votes cast in favor of union representation, averaged across all elections held in the CZ that year. This outcome is only defined for CZ-years with at least one election.
- *Elections per 10,000 workers:* the number of elections held in the CZ that year per 10,000 private sector workers in the CZ in 1979, measured using the QCEW (see Appendix A.3). We hold the denominator fixed at its 1979 value so that changes in the outcome reflect changes in organizing activity rather than changes in local employment. CZ-years with no elections are assigned zero elections.

The vote share captures workers' support for unions conditional on an election being held, while the number of elections captures unions' organizing activity and workers' willingness to petition for an election. Before the recession, the average union vote share was 53 percent.

A.9 AGGREGATE TIME SERIES

We use two national time series, union density and the manufacturing share of employment, in our discussion of the 1980s recession in Section II.

Union density. We measure national union density as union membership as a share of wage and salary workers. From 1973 onwards, the series is based on the CPS as compiled by Hirsch and Macpherson (2003) and updated at unionstats.com, and earlier years come from historical estimates of union membership. In this series, union density fell from 25.4 percent in 1979 to 17.9 percent in 1987.

Manufacturing employment. We measure the manufacturing share of employment as total employment in durable and nondurable goods manufacturing divided by total nonfarm payroll employment, using the BLS Current Employment Statistics (CES) series obtained from FRED. We use end-of-year values for each year.

B. Recession Exposure: Measurement and Validation

This appendix provides additional evidence on our measure of local exposure to the recession. We first characterize which pre-recession characteristics of commuting zones predict exposure. We then show that the industry and geographic incidence of the recession closely matches what we would predict based on each industry's sensitivity to monetary policy before the recession, consistent with the recession being largely driven by the Federal Reserve's efforts to reduce inflation.

B.1 What predicts recession exposure?

Figure B.1 displays the coefficients from regressing the change in the CZ unemployment rate between 1979 and 1982 on the full set of pre-recession CZ characteristics, corresponding to column 4 of Table B.1. All characteristics are standardized, so each coefficient is the change in exposure associated with a one standard deviation increase in the characteristic, holding the others fixed. The manufacturing share of employment is the strongest predictor of exposure. Conditional on the manufacturing share, CZs with larger mining shares were also more exposed, while CZs with higher female employment shares, more college educated workers, and larger wholesale trade shares were less exposed. Pre-recession manufacturing employment growth, exposure to the 1973–1975 recession, and import exposure do not predict exposure once we account for industry composition.

B.2 Monetary policy and the industry incidence of the recession

The recession was concentrated in particular industries, which in turn explains much of its geographic incidence. If the recession was primarily driven by monetary policy, then the industries that contracted the most should be those that were most sensitive to monetary policy before the recession. We test this by estimating each industry's employment response to monetary policy shocks using only data from before the recession and using these responses to predict each industry's change in employment between 1979 and 1982.

Data. We measure industry employment using the BLS Current Employment Statistics (CES) monthly, seasonally adjusted, SIC-based employment series by 2-digit industry. We restrict to the 43 industries whose series are available continuously from 1968 to 1987. We measure monetary policy shocks using the narrative monetary policy shocks of Romer and Romer (2004), which are available monthly from 1969 to 1996.

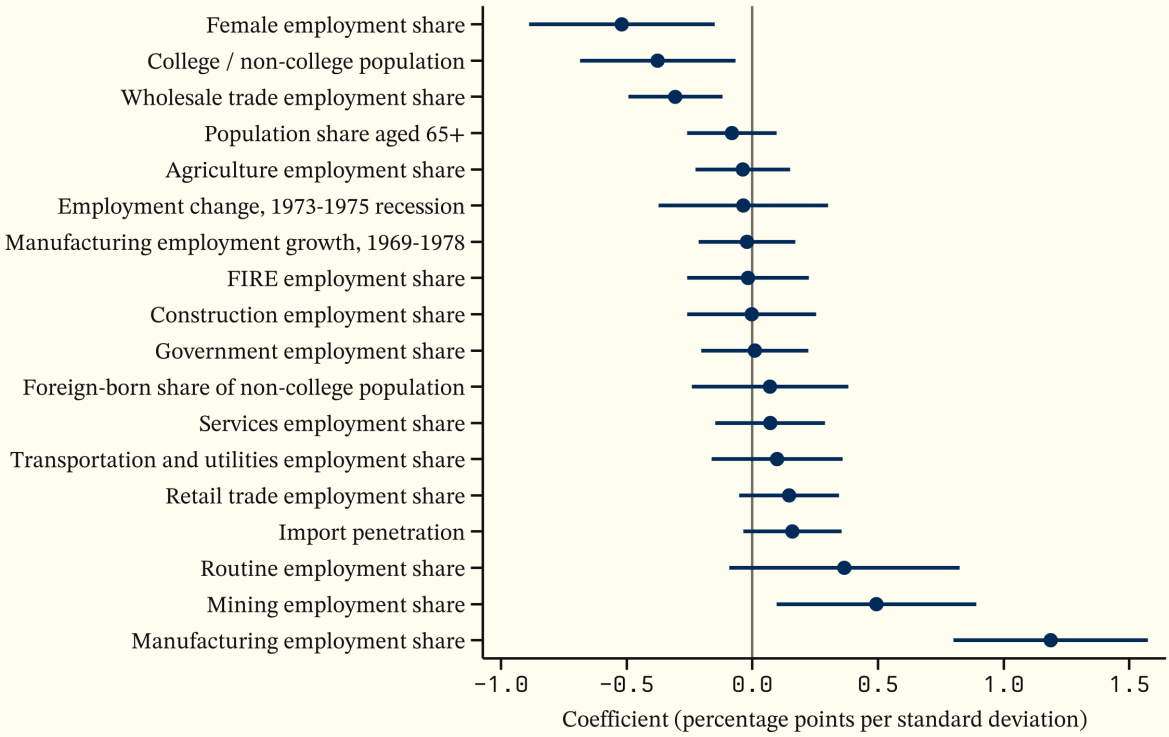


FIGURE B.1
What Predicts Recession Exposure?

Notes: This figure displays coefficients and 95 percent confidence intervals from a regression of the change in the CZ unemployment rate between 1979 and 1982 on pre-recession CZ characteristics, corresponding to column 4 of Table B.1. All characteristics are standardized to have mean zero and unit standard deviation across CZs. The regression includes Census division fixed effects and is unweighted. Standard errors are clustered at the state level. See Appendix A.3 for variable definitions.

TABLE B.1

What Predicts Recession Exposure?

Dependent Variable: Change in Unemployment Rate 1979–1982

	ΔU			
	(1)	(2)	(3)	(4)
Manuf. emp. share (1979)	1.399***	1.435***	1.356***	1.188***
	(0.185)	(0.152)	(0.158)	(0.192)
Δ Manuf. log emp. 69-78	-0.114	-0.086	-0.080	-0.021
	(0.105)	(0.102)	(0.101)	(0.096)
Emp. shock, 73-75 recession			-0.089	-0.036
			(0.157)	(0.168)
Import exposure (1979)			0.184*	0.160
			(0.103)	(0.097)
Routine emp. share (1979)			-0.014	0.366
			(0.155)	(0.227)
College/non-college pop. (1979)				-0.376**
				(0.154)
Foreign born/non-college pop. (1979)				0.071
				(0.155)
Female emp. share (1979)				-0.519***
				(0.184)
Share population 65+ (1979)				-0.081
				(0.089)
Obs.	722	722	720	720
Adj. R^2	0.32	0.47	0.48	0.51
Division FE		✓	✓	✓

Notes: This table displays how different CZ characteristics predict changes in unemployment between 1979 and 1982. All independent variables are centered and standardized so that the coefficients should be interpreted as the effect associated with a 1 standard deviation change in each variable. Column 1 includes the 1979 shares of employment in agriculture, mining, construction, manufacturing, transportation and utilities, wholesale trade, retail trade, finance, services, and public administration, along with the change in log manufacturing employment between 1969 and 1978. Only the coefficient on the manufacturing share is displayed. Column 2 adds Census division fixed effects. Column 3 adds the employment shock in the 1973–1975 recession, import exposure, and the routine share of employment. Column 4 adds the ratio of college to non-college population, the foreign-born share of the non-college population, the female share of employment, and the share of the population aged 65 and over. All covariates added in columns 3 and 4 are measured in 1979, and these columns lose two CZs to missing covariates. The table reports the adjusted R^2. Standard errors are clustered at the state level.

* $p < 0.10$,** $p < 0.05$,*** $p < 0.01$

Industry sensitivity to monetary policy. For each industry j, we estimate local projections (Jordà, 2005) of the form

$$\ln E_{j,t+h} - \ln E_{j,t-1} = \beta_j(h)\varepsilon_t + \sum_{p=1}^{6} a_{j,p}(h)\varepsilon_{t-p} + \sum_{p=1}^{6} c_{j,p}(h)\Delta \ln E_{j,t-p} + g_j(h)t + \alpha_j(h) + u_{j,t+h}$$

for horizons $h = 0, \ldots, 36$ months, where $E_{j,t}$ is employment in industry j in month t and ε_t is the Romer and Romer (2004) shock. The coefficients $\beta_j(h)$ trace out industry j's cumulative employment response h months after a monetary policy shock. Importantly, we estimate these responses using only data between January 1969 and September 1976, so that even the 36-month horizon ends before the Federal Reserve's change in operating procedures in October 1979. Nothing about the estimated sensitivities is therefore informed by the recession itself. We estimate the responses for all horizons jointly so that we obtain their full covariance matrix, using Newey-West standard errors with a bandwidth of 37 months and a small-sample degrees of freedom correction. Figure B.2 displays the estimated responses for four example industries. Employment in construction and transportation equipment falls by around 6 to 10 percent in the two to three years following a 1 percentage point monetary policy shock, while employment in food manufacturing and health services is largely unaffected.

We combine each industry's estimated responses with the shocks that occurred around the recession to predict its change in log employment between September 1979 and September 1982,

$$\hat{d}_j = \sum_{h=0}^{36} \hat{\beta}_j(h)\left(\varepsilon_{T_1-h} - \varepsilon_{T_0-h}\right),$$

where T_0 is September 1979 and T_1 is September 1982. The predicted change $\hat{d}_j$ is the change in log employment that industry j's pre-recession sensitivity to monetary policy implies given the path of monetary policy shocks before and during the recession.

Figure B.3 compares each industry's actual change in log employment between September 1979 and September 1982 to its predicted change. Weighting by industry employment, the slope is 1.37 (0.25) and the predicted changes explain 45 percent of the variation in actual employment changes across industries. Since the predicted changes are estimated with sampling error, which attenuates the slope, we also report the slope corrected for errors-in-variables using the estimated sampling variance of each prediction, which is 1.81 (0.32). Correcting for the sampling error, the predicted changes explain 59 percent of the variation. Industries that were more sensitive to monetary policy before the recession, such as construction and durable goods manufacturing, experienced the largest employment declines. The largest deviations from the predicted changes are in industries affected by other contemporaneous shocks, such as pri-

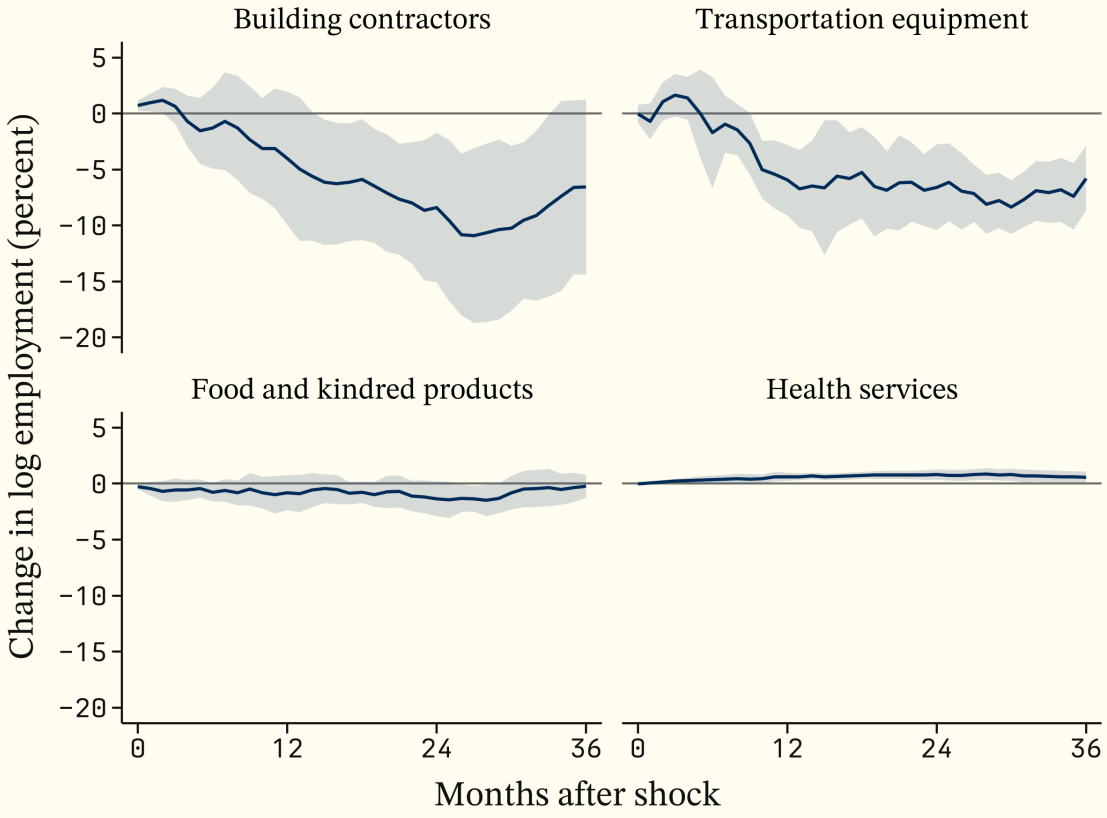


FIGURE B.2
Industry Employment Responses to Monetary Policy Shocks, Selected Industries

Notes: This figure displays the estimated response of log industry employment to a 1 percentage point Romer and Romer (2004) monetary policy shock at horizons of 0 to 36 months, estimated using local projections on CES data from January 1969 to September 1976 (see Appendix B.2). The shaded regions are 95 percent confidence intervals based on Newey-West standard errors.

mary metals and railroads, which declined by more than predicted, and oil and gas extraction, which grew during the oil price boom.

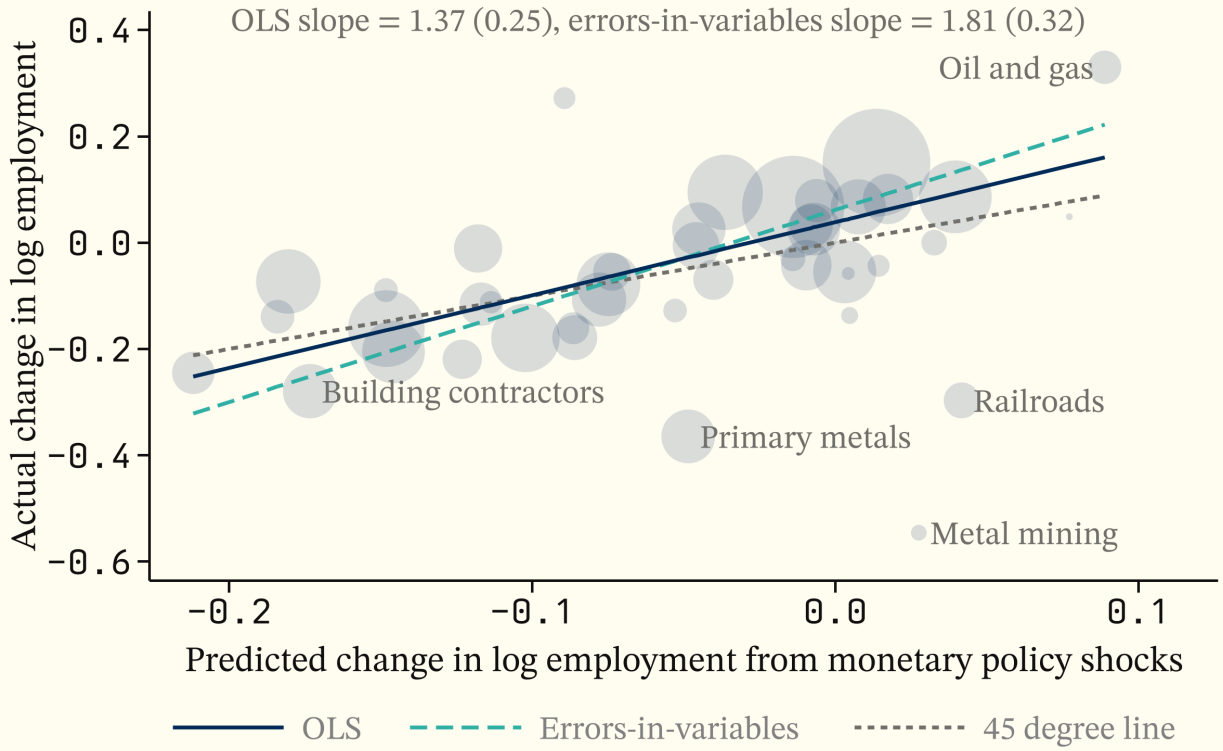


FIGURE B.3
Industry Employment Changes and Monetary Policy Predicted Changes, 1979–1982

Notes: This figure compares the actual change in log employment between September 1979 and September 1982 for 43 2-digit SIC industries in the CES to the change predicted by each industry's estimated response to Romer and Romer (2004) monetary policy shocks. Responses are estimated using local projections on data from January 1969 to September 1976 (see Appendix B.2). Markers are proportional to industry employment in September 1979. The solid line is the employment-weighted OLS fit, and the dashed line is the employment-weighted fit corrected for errors-in-variables using the estimated sampling variance of each predicted change. The dotted line is the 45 degree line. Standard errors are robust to heteroskedasticity.

B.3 MONETARY POLICY AND THE GEOGRAPHIC INCIDENCE OF THE RECESSION

We next ask whether these industry predictions also account for the geographic incidence of the recession. We construct two shift-share measures for each CZ c,

$$\text{MonExp}_c = \sum_j s_{jc,1979}\,\hat{d}_j \qquad \text{and} \qquad \text{Bartik}_c = \sum_j s_{jc,1979}\,\Delta_j,$$

where $s_{jc,1979}$ is industry j's share of CZ employment in 1979 from the County Business Patterns (see Appendix A.3), $\hat{d}_j$ is the monetary policy predicted change in industry employment, and Δ_j is the realized national change in industry employment between March 1979 and March 1982, also measured using the County Business Patterns. The two measures use identical shares and differ only in the industry employment changes. The realized shift-share captures all industry-level shocks during the recession, while the monetary exposure measure only captures the variation in industry employment that is predicted by pre-recession sensitivity to monetary policy. Both are predicted changes in log employment, so more exposed CZs have more negative values.

Figure B.4 displays the relationship between each measure and the change in log CZ employment between 1979 and 1982. Monetary exposure is strongly predictive of local employment changes, with a slope of 1.81 (0.31) and an R^2 of 0.22. This is 86 percent of the R^2 of the realized shift-share (0.25), even though the monetary exposure measure only uses each industry's pre-recession sensitivity to monetary policy and the path of monetary policy shocks. Both slopes are larger than one, consistent with local multiplier effects, as declines in an area's traded sectors can reduce local demand for non-traded goods and services.

Figure B.5 displays the same relationships for our measure of recession exposure, the change in the CZ unemployment rate between 1979 and 1982. Monetary exposure is also strongly predictive of the change in local unemployment. A one log point lower predicted employment change is associated with a 0.53 percentage point larger increase in the unemployment rate, and monetary exposure explains 27 percent of the variation in the change in unemployment across CZs (compared to 39 percent for the realized shift-share).

One concern is that monetary exposure simply captures persistent differences across places, such as the decline of durable goods manufacturing in the Rust Belt, rather than the recession itself. If so, monetary exposure should predict local employment changes in any period. To test this, we re-estimate the relationship between local employment changes and each shift-share measure for every three-year window ending between 1977 and 1987. We update the realized shift-share in each window using that window's national industry employment changes, but hold the industry changes in the monetary exposure measure fixed at their 1979–1982 predicted values. Both measures use industry shares at the start of each window.

If monetary exposure captured persistent differences across places, it should predict local employment changes in every window. Instead, Figure B.6 shows that monetary exposure only predicts local employment changes in windows that overlap with the recession. In the 1976–1979 window, monetary exposure has no predictive power (a slope of 0.31 with a standard error of 0.37), while the realized shift-share does (a slope of 1.49 with a standard error of 0.27). The predictive power of monetary exposure rises as more of the recession enters the window,

peaks in the 1979–1982 window, and disappears by the 1981–1984 window. The realized shift-share remains predictive in every window after 1979. Before 1979, both measures have little predictive power, since there was relatively little variation in national industry employment growth.

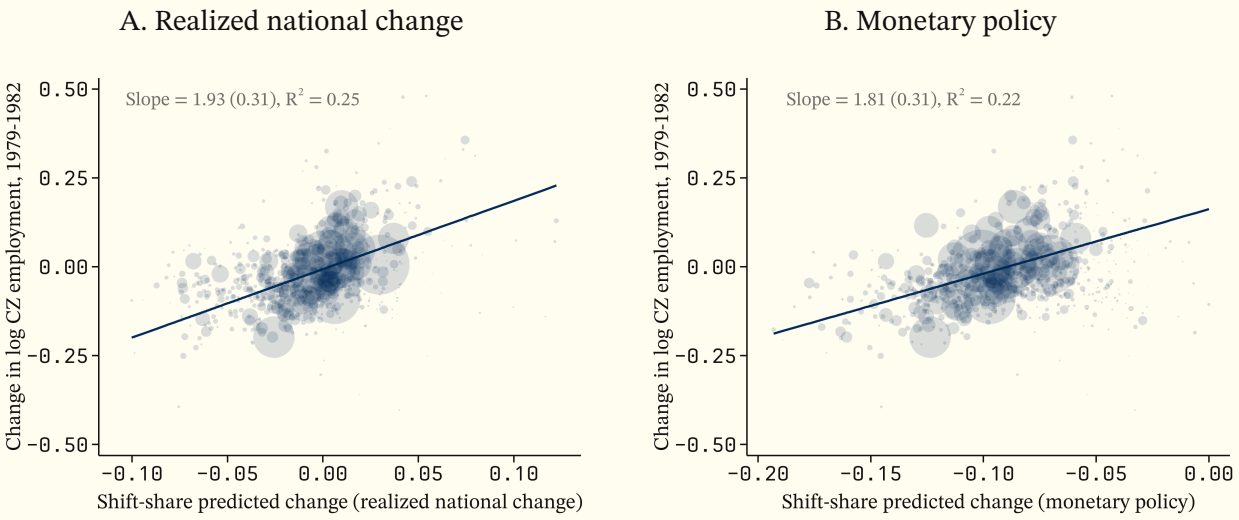


FIGURE B.4
CZ Employment Changes and Shift-Share Predicted Changes, 1979–1982

Notes: This figure displays the relationship between the change in log CZ employment between March 1979 and March 1982 from the County Business Patterns and two shift-share predicted changes in log employment. Panel A uses the realized national change in industry employment, and Panel B uses the change predicted by each industry's pre-recession sensitivity to monetary policy (see Appendix B.2). Both use 1979 CZ industry employment shares. Markers are proportional to CZ employment in 1979, and the lines are employment-weighted linear fits. Standard errors are robust to heteroskedasticity.

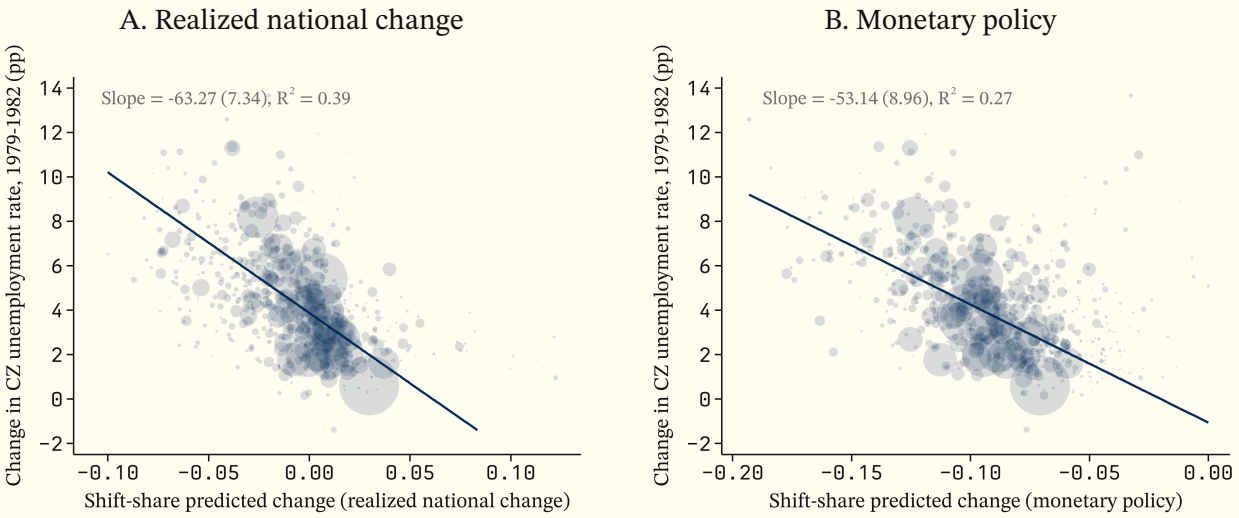


FIGURE B.5
CZ Recession Exposure and Shift-Share Predicted Changes

Notes: This figure displays the relationship between the change in the CZ unemployment rate between 1979 and 1982 (our measure of recession exposure, see Appendix A.2) and the two shift-share predicted changes in log employment described in Figure B.4. Markers are proportional to CZ employment in 1979, and the lines are employment-weighted linear fits. Standard errors are robust to heteroskedasticity.

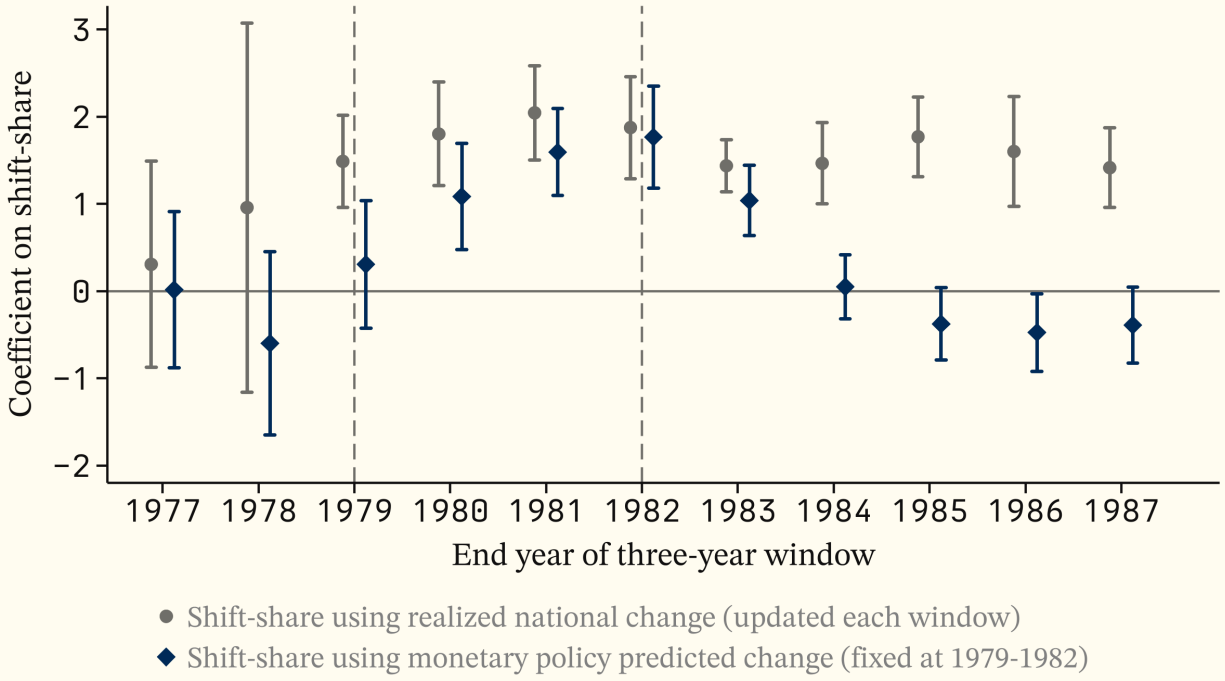


FIGURE B.6
Monetary Exposure and CZ Employment Changes Across Periods

Notes: This figure displays coefficients and 95 percent confidence intervals from regressions of the change in log CZ employment over each three-year window on the two shift-share measures. The realized shift-share (gray) is reconstructed in each window using that window's national change in industry employment and industry shares at the start of the window. The monetary exposure measure (blue) holds each industry's predicted change fixed at its 1979–1982 value and uses industry shares at the start of each window. The dashed lines mark the 1976–1979 and 1979–1982 windows. Regressions are weighted by CZ employment at the start of each window, and standard errors are robust to heteroskedasticity.

C. Worker-Level Effects: Additional Results and Robustness

This appendix presents additional worker-level results using our administrative worker panel (see Appendix A.1). All estimates in this appendix were disclosed under FSRDC Project Number 3420 (CBDRB-FY26-P3420-R13312).

C.1 Unweighted Estimates

Our main estimates are weighted so that our sample matches the 1979 CPS ASEC for the same birth cohorts (see Appendix A.1.4). Figure C.1 displays our main event study estimates without these weights. The unweighted estimates follow a similar pattern during the recession, although the employment effects are smaller and less persistent. Each additional percentage point of exposure reduces the probability of employment by 0.30 percentage points in 1983 (compared to 0.46 in our main estimates), but by only 0.03 percentage points in 1989 (standard error of 0.04). The effect on high labor force attachment remains negative in 1989 (-0.17 percentage points with a standard error of 0.05). The unweighted earnings estimates are similar to our main estimates, with each additional percentage point of exposure reducing W-2 earnings by 1.8 percent in 1989. Our weights increase the importance of older and less educated workers, who are underrepresented in our sample and experience larger and more persistent employment losses (Appendix C.3), which can explain the difference.

C.2 Robustness

Tables C.1 to C.3 repeat the robustness exercises in Table 3 for employment, high labor force attachment, and log FICA earnings. The effects on high labor force attachment and FICA earnings are negative and statistically significant in every specification, and the FICA earnings estimates closely track our main W-2 earnings estimates. The employment effects are less robust. They range from -0.12 to -0.25 percentage points across columns 2 to 10, are smallest when we control for other CZ characteristics (columns 8 and 9), and are not statistically significant without the age-earnings-industry fixed effects (column 1). Column 11 shows that the persistent employment effects are concentrated outside the South and the West Coast.

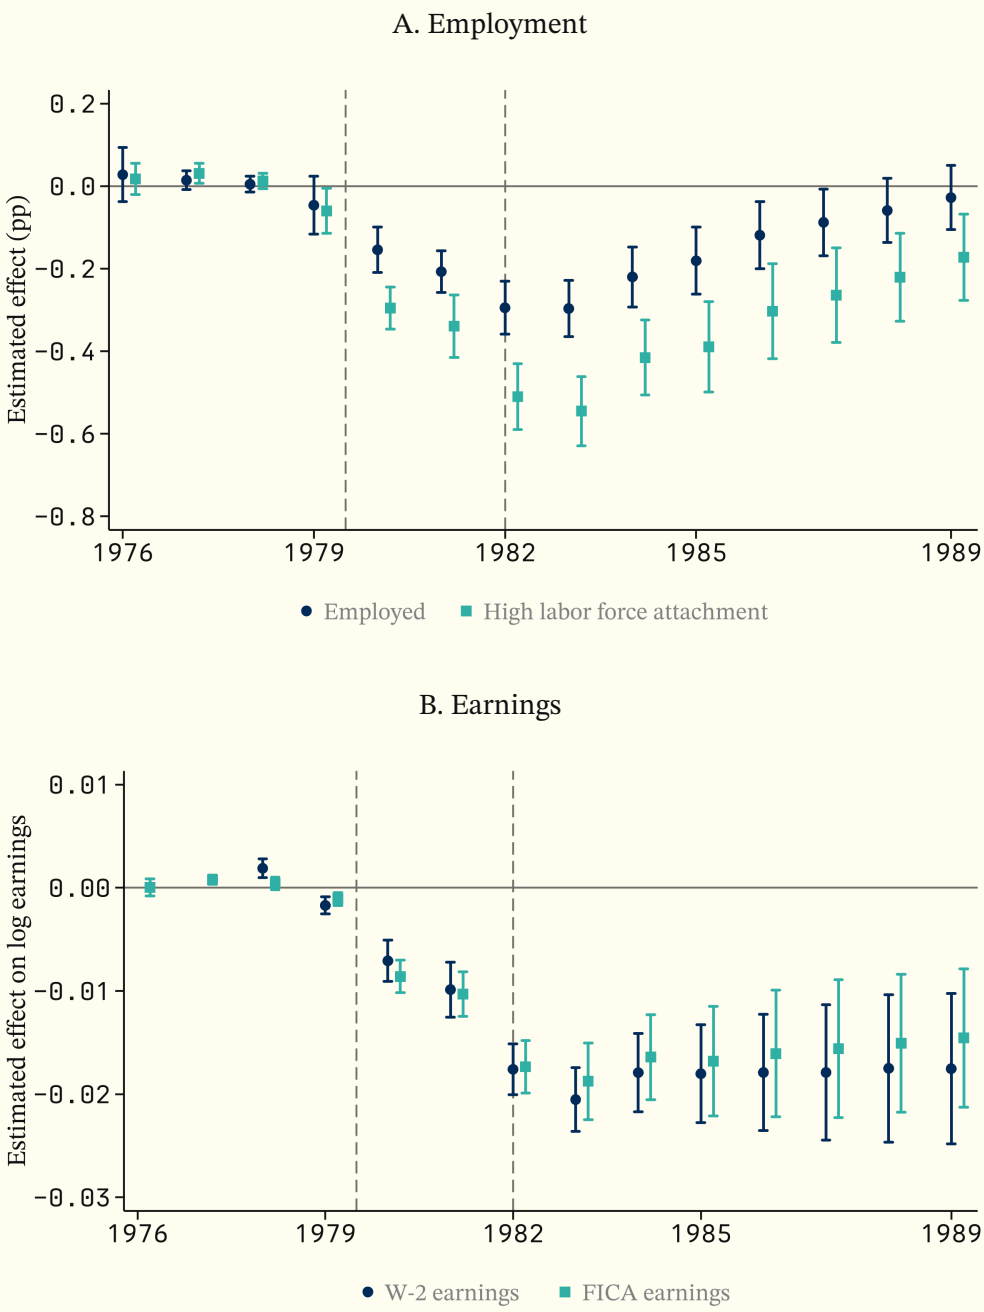


FIGURE C.1
Impacts of the 1980s Recession on Employment and Earnings — Unweighted

Notes: This figure replicates Figures 2 and 3 without reweighting our sample to match the CPS. Panel A displays estimates of the effect of local exposure to the 1980s recession on employment and high labor force attachment, and Panel B displays estimates for log W-2 earnings and log FICA earnings from the SER, from equation (2). The outcomes are relative to each worker's pre-recession baseline. Standard errors are clustered at the CZ level. Estimates used in this figure were disclosed under FSRDC Project Number 3420 (CBDRB-FY26-P3420-R13312).

TABLE C.1
Impacts of the 1980s Recession on Employment — Robustness

	Employed (pp, $t = 1989$)										
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
ΔUnemp.	-0.09	-0.22***	-0.22**	-0.22***	-0.24***	-0.24***	-0.25***	-0.15*	-0.12**	-0.22***	-0.31***
	(0.06)	(0.05)	(0.09)	(0.05)	(0.05)	(0.06)	(0.05)	(0.08)	(0.05)	(0.05)	(0.10)
ΔUnemp. × Rust Belt											-0.03
											(0.12)
ΔUnemp. × South											0.35***
											(0.13)
ΔUnemp. × West Coast											0.51**
											(0.22)
Demographic controls	X	X	X	X	X	X	X	X	X	X	X
Age-earnings-industry FE		X	X	X	X	X	X	X	X	X	X
State clustered SEs			X								
Additional individual controls				X					X		
Growth in pre-recession CZ characteristics					X				X		
Exposure to 73–75 recession						X			X		
State FE							X		X		
Other CZ controls								X	X		
No manuf/mining/constr. workers										X	
R^2	0.01	0.09	0.09	0.10	0.10	0.09	0.10	0.10	0.10	0.09	0.09

Notes: This table displays regression estimates of the effect of local exposure to the 1980s recession on employment in 1989 using equation (2), in percentage points. The columns follow Table 3. Standard errors are clustered at the CZ level except in column 3. Estimates used in this table were disclosed under FSRDC Project Number 3420 (CBDRB-FY26-P3420-R13312).
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

TABLE C.2
Impacts of the 1980s Recession on High Labor Force Attachment — Robustness

	High Labor Force Attachment (pp, $t = 1989$)										
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
ΔUnemp.	-0.31***	-0.37***	-0.37***	-0.38***	-0.53***	-0.41***	-0.37***	-0.29***	-0.21***	-0.38***	-0.65***
	(0.10)	(0.06)	(0.11)	(0.06)	(0.07)	(0.07)	(0.06)	(0.07)	(0.06)	(0.06)	(0.13)
ΔUnemp. × Rust Belt											0.12
											(0.15)
ΔUnemp. × South											0.69***
											(0.18)
ΔUnemp. × West Coast											0.42**
											(0.21)
Demographic controls	X	X	X	X	X	X	X	X	X	X	X
Age-earnings-industry FE		X	X	X	X	X	X	X	X	X	X
State clustered SEs			X								
Additional individual controls				X					X		
Growth in pre-recession CZ characteristics					X				X		
Exposure to 73–75 recession						X			X		
State FE							X		X		
Other CZ controls								X	X		
No manuf/mining/constr. workers										X	
R^2	0.03	0.17	0.17	0.17	0.17	0.17	0.17	0.17	0.17	0.16	0.17

Notes: This table displays regression estimates of the effect of local exposure to the 1980s recession on high labor force attachment in 1989 using equation (2), in percentage points. The columns follow Table 3. Standard errors are clustered at the CZ level except in column 3. Estimates used in this table were disclosed under FSRDC Project Number 3420 (CBDRB-FY26-P3420-R13312).
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

TABLE C.3
Impacts of the 1980s Recession on FICA Earnings — Robustness

	Log FICA Earnings ($t = 1989$)										
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
ΔUnemp.	-0.018***	-0.017***	-0.017***	-0.017***	-0.027***	-0.018***	-0.016***	-0.014***	-0.011***	-0.018***	-0.031***
	(0.002)	(0.003)	(0.004)	(0.003)	(0.003)	(0.002)	(0.002)	(0.003)	(0.002)	(0.003)	(0.008)
ΔUnemp. × Rust Belt											0.009
											(0.009)
ΔUnemp. × South											0.022**
											(0.009)
ΔUnemp. × West Coast											-0.012
											(0.009)
Demographic controls	X	X	X	X	X	X	X	X	X	X	X
Age-earnings-industry FE		X	X	X	X	X	X	X	X	X	X
State clustered SEs			X								
Additional individual controls				X					X		
Growth in pre-recession CZ characteristics					X				X		
Exposure to 73–75 recession						X			X		
State FE							X		X		
Other CZ controls								X	X		
No manuf/mining/constr. workers										X	
R^2	0.04	0.14	0.14	0.14	0.14	0.14	0.14	0.14	0.15	0.13	0.14

Notes: This table displays regression estimates of the effect of local exposure to the 1980s recession on log FICA earnings from the SER in 1989 using equation (2). FICA earnings are topcoded at the Social Security taxable maximum (see Appendix A.1). The columns follow Table 3. Standard errors are clustered at the CZ level except in column 3. Estimates used in this table were disclosed under FSRDC Project Number 3420 (CBDRB-FY26-P3420-R13312).
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

C.3 HETEROGENEITY

Figure C.2 displays the effect of exposure on high labor force attachment between 1987 and 1989 separately by each worker's 1979 earnings decile. Similar to employment and earnings (Figure 4), the effects are largest for workers at the bottom of the pre-recession earnings distribution. Each additional percentage point of exposure reduces the probability of high labor force attachment by around 0.75 percentage points for workers in the first and second deciles, compared to 0.14 percentage points for workers in the top decile.

Figures C.3 to C.5 display the effects separately by worker characteristics. The effects on employment and high labor force attachment are larger for older workers, Black workers, and workers without a college degree. For instance, each additional percentage point of exposure reduces high labor force attachment by 0.78 percentage points for workers aged 45 to 54 in 1979, compared to 0.18 percentage points for workers aged 23 to 29. The earnings effects are more similar across groups, although they are also larger for older and Black workers.

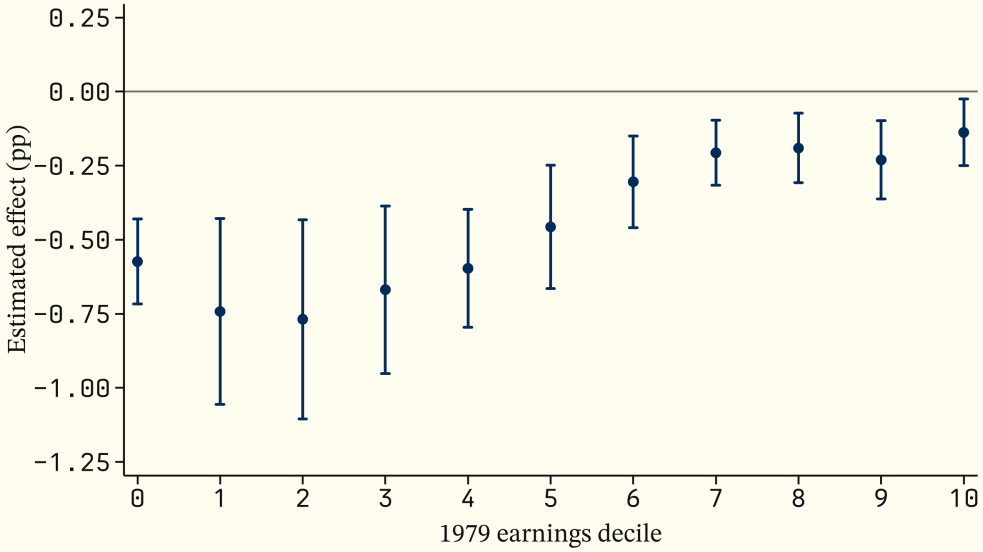


FIGURE C.2
Impacts of the 1980s Recession on High Labor Force Attachment — Heterogeneity by Pre-recession Earnings Decile

Notes: This figure displays regression estimates of the effect of local exposure to the 1980s recession on high labor force attachment between 1987 and 1989 from equation (2). Each coefficient is from a separate regression restricting to workers in each decile of the 1979 earnings distribution. Workers with zero earnings in 1979 are grouped into decile 0. Standard errors are clustered at the CZ level. Regressions are weighted so that our main sample matches the CPS. Estimates used in this figure were disclosed under FSRDC Project Number 3420 (CBDRB-FY26-P3420-R13312).

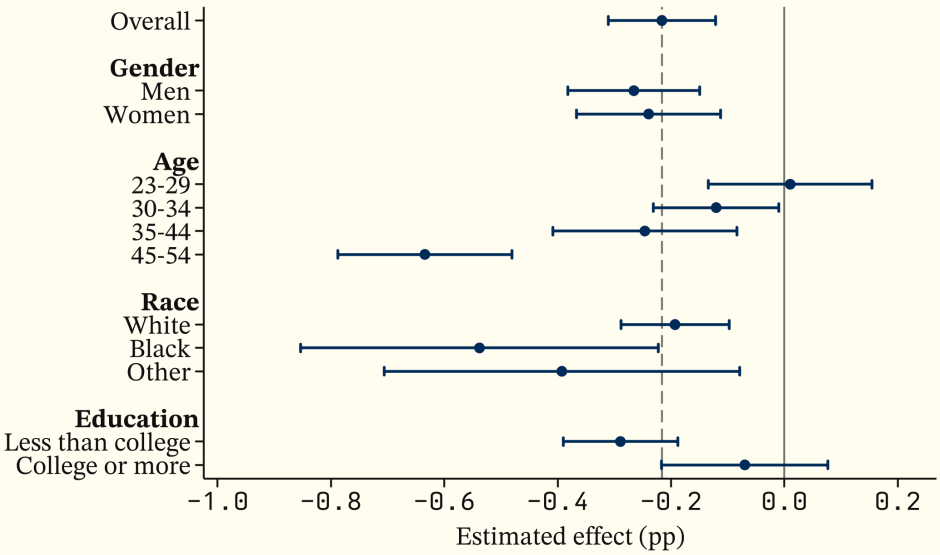


FIGURE C.3

Impacts of the 1980s Recession on Employment — Heterogeneity by Worker Characteristics

Notes: This figure displays regression estimates of the effect of local exposure to the 1980s recession on employment between 1987 and 1989 from equation (2). Each coefficient is from a separate regression restricting to workers in each group. Age is measured in 1979. The dashed line marks the full-sample estimate in 1989. Standard errors are clustered at the CZ level. Regressions are weighted so that our main sample matches the CPS. Estimates used in this figure were disclosed under FSRDC Project Number 3420 (CBDRB-FY26-P3420-R13312).

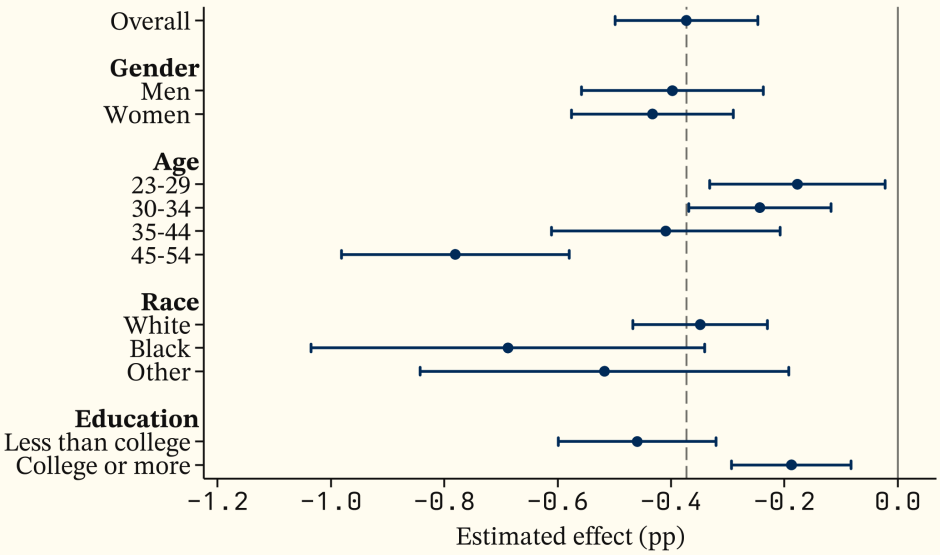


FIGURE C.4
Impacts of the 1980s Recession on High Labor Force Attachment — Heterogeneity by Worker Characteristics

Notes: This figure displays regression estimates of the effect of local exposure to the 1980s recession on high labor force attachment between 1987 and 1989 from equation (2). Each coefficient is from a separate regression restricting to workers in each group. Age is measured in 1979. The dashed line marks the full-sample estimate in 1989. Standard errors are clustered at the CZ level. Regressions are weighted so that our main sample matches the CPS. Estimates used in this figure were disclosed under FSRDC Project Number 3420 (CBDRB-FY26-P3420-R13312).

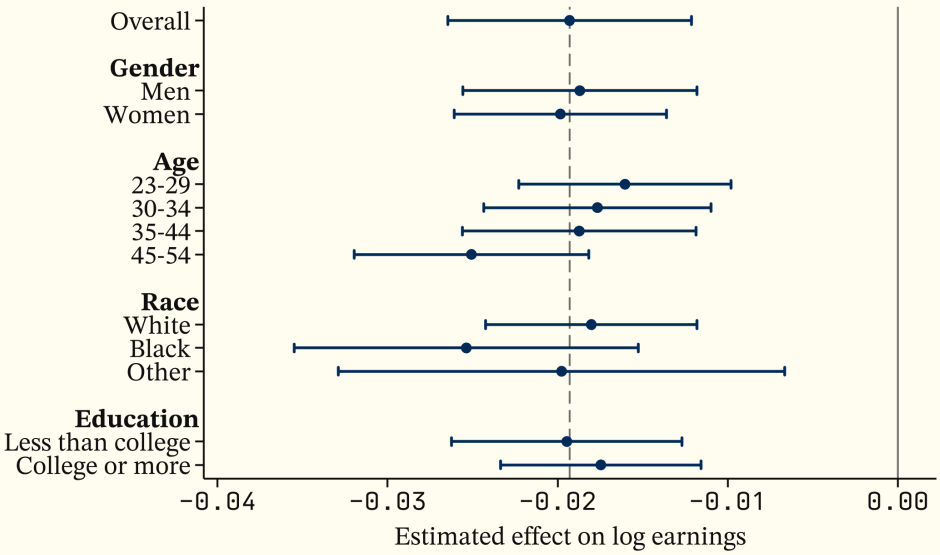


FIGURE C.5

Impacts of the 1980s Recession on Earnings — Heterogeneity by Worker Characteristics

Notes: This figure displays regression estimates of the effect of local exposure to the 1980s recession on log W-2 earnings between 1987 and 1989 from equation (2). Each coefficient is from a separate regression restricting to workers in each group. Age is measured in 1979. The dashed line marks the full-sample estimate in 1989. Standard errors are clustered at the CZ level. Regressions are weighted so that our main sample matches the CPS. Estimates used in this figure were disclosed under FSRDC Project Number 3420 (CBDRB-FY26-P3420-R13312).

C.4 COMPARISON TO THE GREAT RECESSION

Table C.4 collects the comparison between our estimates and those in Yagan (2019) for the Great Recession discussed in Section V. Both papers measure exposure as the change in the CZ unemployment rate during the recession and define employment as having positive wage earnings in a year. Cumulative earnings include years with zero earnings and are measured over the seven years after each recession.

TABLE C.4
Comparison to the Great Recession

	1980s recession (this paper)	Great Recession (Yagan, 2019)
Employment (pp), trough	−0.46 (1983)	
Employment (pp), 8 years after start	−0.22 (1989)	−0.39 (2015)
Cumulative earnings (2015$), 7 years	−4,419	−6,212
	(693)	

Notes: This table compares the effects of each additional percentage point of local exposure in the 1980s recession to the effects in the Great Recession estimated by Yagan (2019). Cumulative earnings are total wage earnings over 1983–1989 (this paper) and 2009–2015 (Yagan, 2019), including zeros. Standard errors in parentheses. Estimates for the 1980s recession were disclosed under FSRDC Project Number 3420 (CBDRB-FY26-P3420-R13312).

D. Mechanisms: Additional Results

This appendix reports additional results for each of the mechanisms in Section VII.

D.1 The Recession's Impact: People versus Places

Figure D.1 displays the same long-run event study as Figure 7 for log private sector employment. Employment in more exposed CZs fell during the recession and did not recover. Each additional percentage point of exposure reduced private sector employment by 2.8 percent in 1985 and by 3.3 percent in 2019.

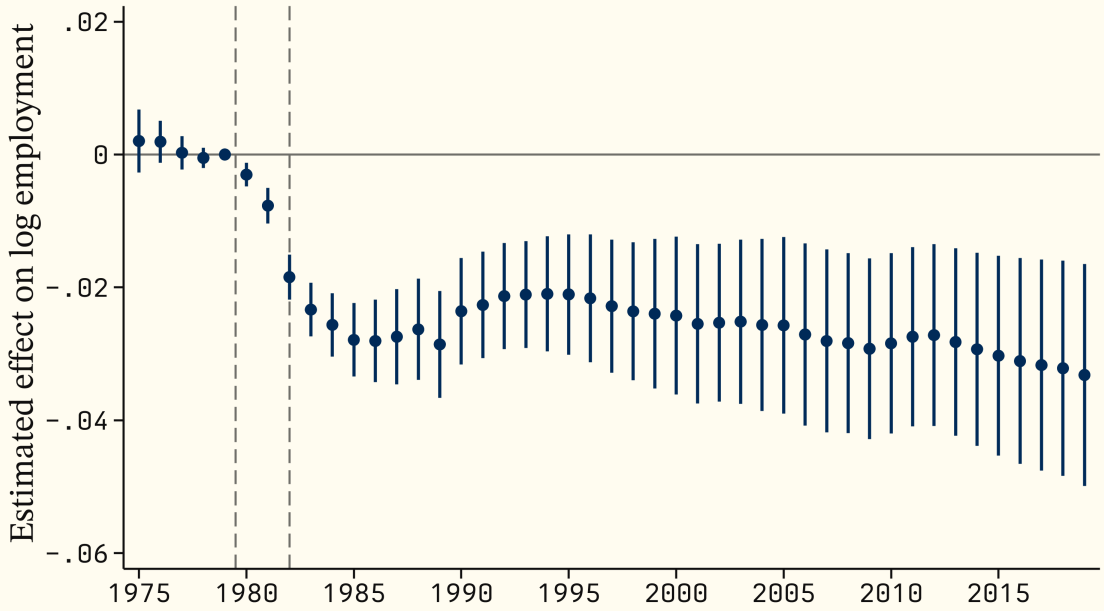


FIGURE D.1
Long-Run Effects of Recession Exposure on Commuting Zone Employment, 1975–2019

Notes: This figure displays event-study estimates of the effect of local recession exposure on log private sector employment between 1975 and 2019, using county-level data from the BLS Quarterly Census of Employment and Wages aggregated to CZs. The specification, controls, and weights are the same as in Figure 7. Standard errors are clustered at the CZ level. Vertical bars show 95 percent confidence intervals.

To understand how the lower employment in more exposed places was absorbed, we estimate the effects of exposure on the CZ unemployment rate and labor force participation using event study specifications of the form

$$y_{ct} = \alpha + \sum_k \beta_k \mathbb{1}\{t = k\} \cdot \Delta\text{Unemp.}_c + \mathbf{X}'_{ct}\Gamma + \delta_{d(c),t} + \varepsilon_{ct} \tag{D.1}$$

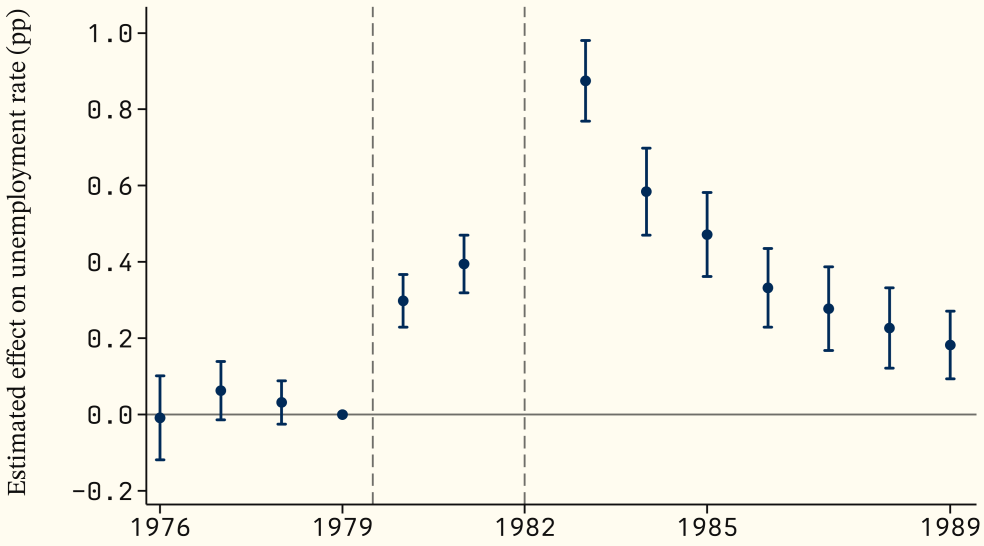
where y_{ct} is the CZ outcome of interest, $\mathbf{X}_{ct}$ are CZ pre-recession characteristics interacted with year fixed effects, and $\delta_{d(c),t}$ are division-by-year fixed effects.

Figure D.2A displays the estimates for the CZ unemployment rate. The effect falls after the recession but remains positive through the end of the decade. CZs that experienced an additional percentage point increase in unemployment during the recession have unemployment rates that are 0.18 percentage points higher in 1989. Figure D.2B displays the estimates for the share of the population not in the labor force. The rise in nonparticipation lags the rise in unemployment by several years, consistent with discouraged workers leaving the labor force, and shows no sign of recovering. Each additional percentage point of exposure increases the share not in the labor force by 0.42 percentage points in 1989. Therefore, most of the persistent employment response to the recession operated through lower labor force participation rather than higher unemployment.

D.2 PERSISTENT CHANGES IN RELATIVE LABOR DEMAND

Table D.1 estimates our long-differences specification separately for workers with and without a college degree. Exposure reduced hourly wages and annual hours for both groups by similar amounts. However, the employment rate of non-college workers fell by 0.31 percentage points per unit of exposure, while the employment rate of college workers did not change, since the college population fell along with college employment. This is consistent with the migration patterns in Table D.2. More exposed places attracted fewer college-educated in-migrants between 1985 and 1990, with no significant effect on net migration for either group.

A. Unemployment Rate



B. Share Not in the Labor Force (NILF)

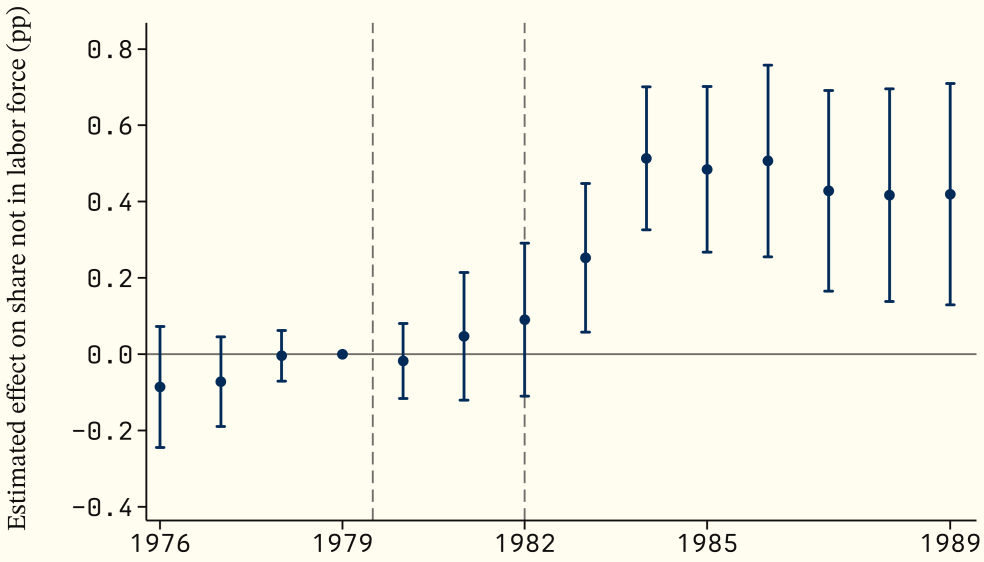


FIGURE D.2
Impacts of the 1980s Recession on CZ Unemployment and Labor Force Participation

Notes: This figure displays regression estimates of the effect of local exposure to the 1980s recession on the CZ unemployment rate and the share of the population not in the labor force from equation (D.1), in percentage points. The unemployment rate is measured using the BLS Local Area Unemployment Statistics. The share not in the labor force is one minus the LAUS labor force divided by the population aged 16 and older from the SEER population estimates. Since exposure is the change in the unemployment rate between 1979 and 1982, the 1982 coefficient for the unemployment rate is mechanically one and is omitted. Each regression controls for CZ and division-by-year fixed effects, as well as exposure to the 1973–1975 recession, pre-recession earnings and population growth, and the manufacturing share of employment, routine share of employment, female employment share, share of the population aged 65 and over, ratio of college to non-college educated workers, and share of the non-college population that is foreign born measured in the 1980 Census, all interacted with year fixed effects. Regressions are weighted by employment in 1979 and standard errors are clustered at the CZ level. Vertical bars show 95 percent confidence intervals.

TABLE D.1
Recession Exposure and Labor Market Outcomes by Education

	1980–1990					
	(1) Log Hourly Wage	(2) Log Annual Hours	(3) Log Wage Bill	(4) Log Employment	(5) Log Population	(6) Employment Rate (pp)
Less than College	-0.012***	-0.015***	-0.027***	-0.013***	-0.008***	-0.308***
	(0.002)	(0.003)	(0.005)	(0.003)	(0.002)	(0.061)
College Plus	-0.009***	-0.019***	-0.027***	-0.017***	-0.017***	-0.057
	(0.002)	(0.004)	(0.005)	(0.004)	(0.004)	(0.050)
Obs.	722	722	722	722	722	722

Notes: This table displays coefficients from regressing changes in CZ labor market outcomes for college and non-college workers in the Census between 1980 and 1990 on the change in CZ unemployment between 1979 and 1982. Each coefficient is from a separate regression. College refers to individuals with 16 or more years of schooling. Log hourly wage refers to the change in the average log hourly wage of full-time, full-year wage and salary workers after subtracting the national average log hourly wage of each worker's sex, education, and potential experience cell in that year. Hourly wages are measured as annual wage earnings divided by annual weeks worked times usual weekly hours worked. Individuals with allocated wage earnings are dropped from the wage measure. Log annual hours refers to the change in the log of total annual hours worked by wage and salary workers aged 16–64, measured as annual weeks worked times usual weekly hours worked. Log employment refers to the change in the log number of wage and salary workers aged 16–64 who worked in the previous year. Log population refers to the change in the log number of individuals aged 16–64 not living in institutional or other group quarters. Employment rate refers to the change in the ratio of employment to population, in percentage points. Log wage bill refers to the change in the log of total annual wage and salary income of these workers, including allocated earnings. Each regression controls for division fixed effects, exposure to the 1973–1975 recession, earnings and population growth between 1969 and 1979, the share of employment in manufacturing, the routine share of employment, the female employment share, the share of the population aged 65 and over, the ratio of college to non-college educated workers, and the share of the non-college population that is foreign born all measured in 1980. Earnings growth is measured as the change in log BEA wages per wage and salary job and population growth is measured using SEER population estimates. Regressions are weighted by employment in 1979 and standard errors are clustered at the state level.

TABLE D.2
Recession Exposure and Migration by Education

	1985–1990			
	(1) Out-Migration	(2) In-Migration	(3) Arrivals from Abroad	(4) Net Migration
Less than College	-0.384***	-0.284*	-0.059*	0.109
	(0.131)	(0.158)	(0.031)	(0.166)
College Plus	-0.345*	-0.488***	-0.059	-0.196
	(0.193)	(0.145)	(0.039)	(0.218)
Mean, Less than College	11.84	11.23	1.92	1.73
Mean, College Plus	22.01	20.86	2.48	2.03
Obs.	722	722	722	722

Notes: This table displays coefficients from regressing migration rates between 1985 and 1990 for college and non-college individuals on the change in CZ unemployment between 1979 and 1982. The sample is restricted to individuals aged 25–64 in the 1990 Census not living in institutional or other group quarters. College refers to individuals with 16 or more years of schooling. Out-migration refers to the share of individuals living in the CZ in 1985 who lived in a different CZ in 1990. In-migration refers to individuals who lived in a different CZ in 1985 as a share of the CZ population in 1990. Arrivals from abroad refers to individuals who lived outside the United States in 1985 as a share of the CZ population in 1990. Net migration refers to in-migrants plus arrivals from abroad minus out-migrants as a share of the CZ population in 1985. All rates are measured in percent and the bottom panel reports the mean of each outcome. Residence in 1985 is measured using the Census migration PUMA, and individuals whose 1985 migration PUMA spans multiple CZs are allocated across CZs using population shares. Each regression controls for division fixed effects, exposure to the 1973–1975 recession, earnings and population growth between 1969 and 1979, the share of employment in manufacturing, the routine share of employment, the female employment share, the share of the population aged 65 and over, the ratio of college to non-college educated workers, and the share of the non-college population that is foreign born all measured in 1980. Earnings growth is measured as the change in log BEA wages per wage and salary job and population growth is measured using SEER population estimates. Regressions are weighted by employment in 1979 and standard errors are clustered at the state level.

D.3 IMPACTS ON UNIONS

Table D.3 assesses threats to the union membership estimates from contemporaneous shocks correlated with recession exposure. Columns 4 and 5 control for pre-recession state characteristics interacted with year fixed effects and for region-by-year fixed effects, and the estimated effects are similar. A more specific concern is Ronald Reagan's firing of striking air traffic controllers in August 1981, which is widely considered to have emboldened employer opposition to unions. Any national effect of this event is absorbed by the year fixed effects, so it would only affect our estimates if its effect were correlated with local exposure to the recession. Two

pieces of evidence suggest it is not driving our results. First, the effect on union membership is already negative in 1980, a year before the firings (Figure 8A). Second, columns 2 and 3 estimate the effects on private and public sector membership separately, and we find no effect on public sector membership, which might have been more affected since the air traffic controllers' union (PATCO) was a public sector union. We cannot, however, rule out that the firings amplified the effects of the local recession shocks on unionization.

Figure D.3 displays the effect of exposure on the number of NLRB representation elections per worker in each CZ. We find no persistent effect on organizing activity. Figure D.4 compares the 1980s recession to other postwar recessions. Panel A shows that employment losses were concentrated in highly unionized state-group-by-industry cells to a greater degree in the 1980s recession than in other recessions. Panel B shows that local recession shocks reduced household union membership in the 1980s recession and, to a lesser extent, in the 1973–1975 recession, but not in other recessions. Figure D.5 makes the same comparison for national union density, relative to its trend in the years before each recession. Union density fell by around 17 log points relative to trend in the five years after the start of the 1980s recession, compared to around 5 log points after the 1969–1970 and 1973–1975 recessions, and it did not fall relative to trend after the more recent recessions.

TABLE D.3
Impacts of the 1980s Recession on Unions — Robustness

	Union Membership					
	(1)	(2)	(3)	(4)	(5)	(6)
$\Delta U \times \mathbb{1}\{t > 1979\}$	-0.680***	-0.676***	-0.004	-0.872***	-0.770***	-0.510***
	(0.143)	(0.142)	(0.061)	(0.276)	(0.239)	(0.136)
Obs.	1,806,332	1,806,332	1,806,332	1,806,332	1,806,332	1,559,006
R^2	0.06	0.06	0.04	0.06	0.06	0.19
Indiv. Controls	✓	✓	✓	✓	✓	✓
State Chars. × Year				✓		
Region × Year					✓	
Industry × Year						✓

Notes: This table displays estimates of the effect of local exposure to the 1980s recession on union membership in the CPS, in percentage points. Each coefficient is on the interaction between state recession exposure and an indicator for years after 1979. The outcome is an indicator for being a union member among all adults aged 21–65, so that non-workers are counted as non-members. All regressions control for a quadratic in age, sex, race, five education group fixed effects and their interaction with age, and state and year fixed effects. Columns 2 and 3 use private and public sector union membership as the outcome. Column 4 adds controls for exposure to the 1973–1975 recession, pre-recession earnings and population growth, the manufacturing share of employment in 1979, and the routine share of employment, female employment share, share of the population aged 65 and over, ratio of college to non-college educated workers, and share of the non-college population that is foreign born measured in the 1980 Census, all interacted with year fixed effects. Column 5 adds region-by-year fixed effects and column 6 adds industry-by-year fixed effects. Standard errors are clustered at the state level and regressions are weighted using CPS weights.

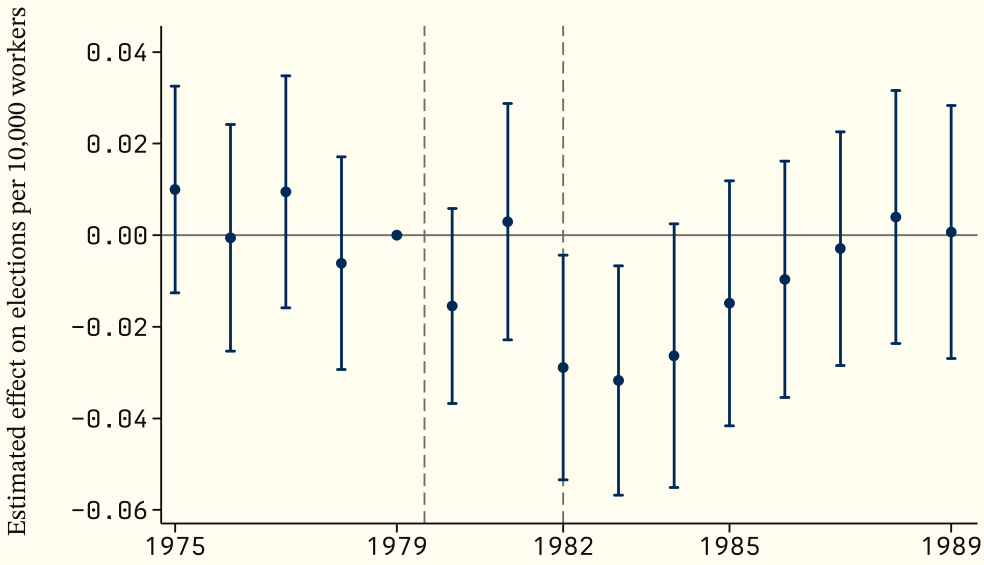
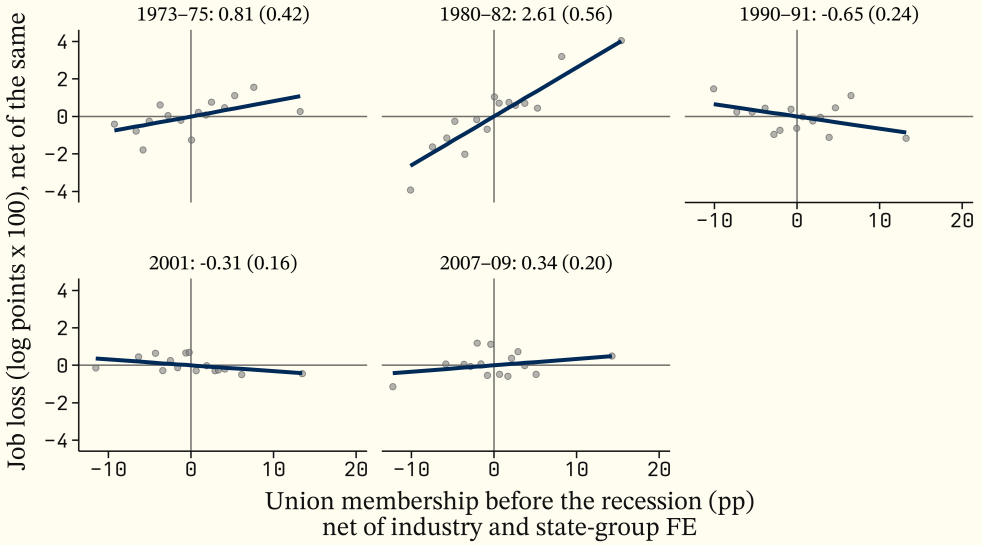


FIGURE D.3
Impacts of the 1980s Recession on CZ NLRB Elections per 10,000 Workers

Notes: This figure displays regression estimates of the effect of local exposure to the 1980s recession on the number of NLRB representation elections per 10,000 workers in each CZ between 1975 and 1989. Exposure is the change in the CZ unemployment rate between 1979 and 1982. The regressions control for CZ and division-by-year fixed effects, exclude Alaska, Hawaii, and the District of Columbia, and are weighted by employment in 1979. Standard errors are clustered at the CZ level. Vertical bars show 95 percent confidence intervals.

A. Employment losses and pre-recession union density, by recession



B. Effect of each recession's local shock on household union membership

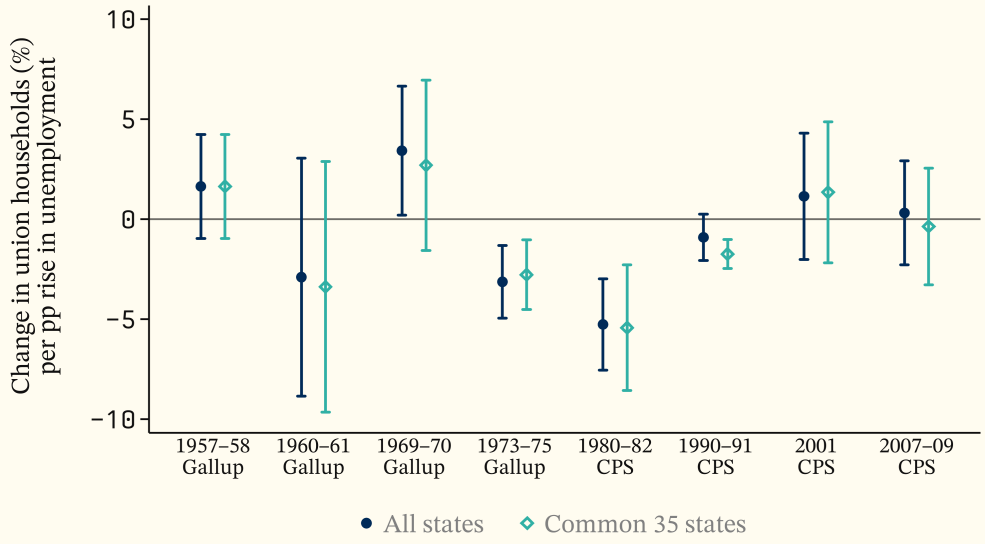


FIGURE D.4
Recessions and Unionization Across Postwar Recessions

Notes: Panel A displays, for each recession since 1973–1975, a binned scatter plot of job losses during the recession (log points × 100) against pre-recession union membership (percentage points) across state-group-by-industry cells, both net of state-group and industry fixed effects. Each point is one of 15 employment-weighted bins and the line is the fitted slope. The number above each plot is that slope, the additional job loss per 10 percentage points of pre-recession union membership, with its standard error clustered by state group in parentheses. Panel A begins with the 1973–1975 recession since union membership by state and industry is first available in the May 1973 CPS. Panel B displays, for each recession since 1957–1958, the effect of the rise in the state unemployment rate during the recession on the share of respondents aged 21–65 in a union household, as a percent of the pre-recession share. Each estimate is from a difference-in-differences regression comparing the five years after the recession's base year to the three years before, with state and year fixed effects, a quadratic in age, sex, race, and education, and the base-year state manufacturing share interacted with year, weighted using survey weights. Estimates use Gallup data before 1980 and the CPS from 1980 onwards. Filled markers use all states with a measured shock and hollow markers use the 35 states with a shock in every recession. Vertical bars are 95 percent confidence intervals with standard errors clustered at the state level.

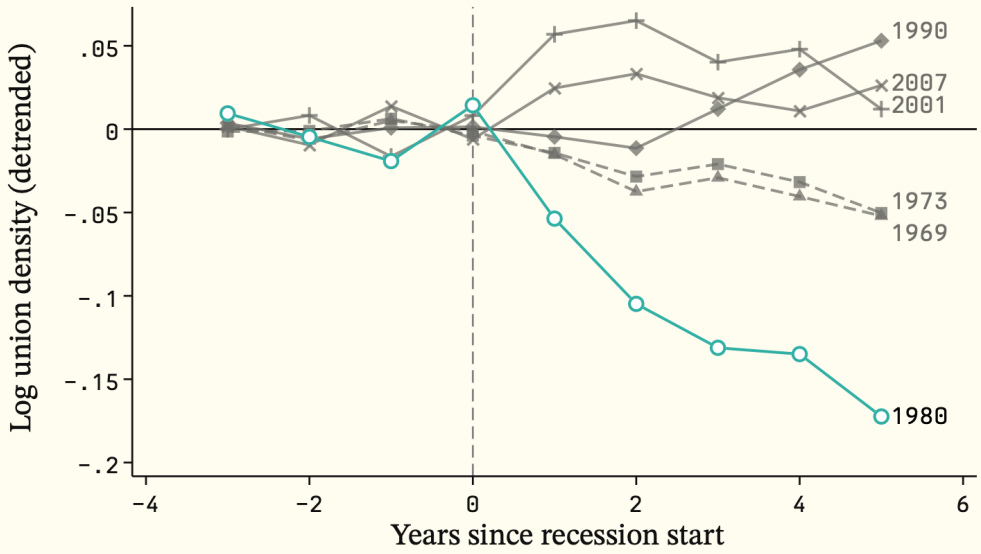


FIGURE D.5
Evolution of Union Density Around Recessions

Notes: This figure plots the evolution of the log union density around each postwar recession, measured as union membership as a share of wage and salary workers (see Appendix A.9). Each series is detrended against a linear trend estimated on the four years prior to the start of each recession ($t = -3$ to $t = 0$).

E. Distributional Analysis

This appendix provides additional details on the distributional analysis in Section VI and the decomposition in Section VIII.

E.1 Estimating the Counterfactual Distribution

This section describes the target parameters of the analysis in Section VI, the identification assumptions, and estimation procedure. Throughout, we follow the distribution regression difference-in-differences framework of Fernández-Val et al. (2024). Our assumptions and estimation procedure follow theirs closely.

Setup and target parameters

Let $S \in \{0,1\}$ indicate whether a worker lives in a severely exposed state, which plays the role of the treatment group, let $T \in \{0,1\}$ index the years 1979 and 1988, let X denote the worker's characteristics, and let Y denote the worker's log hourly wage. Severe exposure applies only to workers in severely exposed states in 1988, $ST = 1$. Let $Y(0)$ and $Y(1)$ denote potential wages under the less severe exposure of less exposed states and under severe exposure. We observe

$$Y = Y(0)(1 - ST) + Y(1)ST$$

so we observe $Y(0)$ for both groups in 1979, $Y(0)$ for less exposed states in 1988, and $Y(1)$ for severely exposed states in 1988. This implicitly imposes that wages in 1979 are unaffected by the recession.

The distribution of $Y(0)$ conditional on S, T, and X can be written as

$$F_{Y(0)|S,T,X}(y \mid s,t,x) = \Lambda\big(\alpha(y,x) + \beta(y,x)t + \gamma(y,x)s + \delta(y,x)st\big) \tag{E.1}$$

where Λ is a strictly increasing distribution function and $(y,x) \mapsto (\alpha(y,x), \beta(y,x), \gamma(y,x), \delta(y,x))$ are unrestricted functions. This representation is fully saturated in S and T, so it imposes no restrictions on the distribution of $Y(0)$.

We are interested in the distributions of the potential wages of severely exposed workers in 1988. The distribution of $Y(1)$ is identified by the observed distribution $F_{Y(1)|S,T}(y \mid 1,1) = F_{Y|S,T}(y \mid 1,1)$. The counterfactual distribution of $Y(0)$ is

$$F_{Y(0)|S,T}(y \mid 1,1) = \int_{\mathcal{X}_{11}} F_{Y(0)|S,T,X}(y \mid 1,1,x)\, dF_{X|S,T}(x \mid 1,1) \tag{E.2}$$

where $\mathcal{X}_{11}$ is the support of X among severely exposed workers in 1988. The distributional treatment effect and the quantile treatment effect are

$$\tau(y) = F_{Y(1)|S,T}(y \mid 1,1) - F_{Y(0)|S,T}(y \mid 1,1), \qquad \tau_q^* = F^{-1}_{Y(1)|S,T}(q \mid 1,1) - F^{-1}_{Y(0)|S,T}(q \mid 1,1) \quad \text{(E.3)}$$

for $y \in \mathbb{R}$ and $q \in (0,1)$. $\tau(y)$ measures the change in the probability that a severely exposed worker earns less than y and τ_q^* measures the change in the q-th quantile of their wage distribution. For any measure of inequality $\nu(\cdot)$, such as the 90–10 gap or the Gini coefficient, the effect of severe exposure is $\nu(F_{Y(1)|S,T}(\cdot \mid 1,1)) - \nu(F_{Y(0)|S,T}(\cdot \mid 1,1))$.

Equation (E.2) integrates over the observed 1988 characteristics of severely exposed workers. If severe exposure did not affect X, this is the distribution of $Y(0)$ among severely exposed workers in 1988. Since X includes characteristics such as industry and occupation that the recession may have affected, we interpret equation (E.2) more generally as a counterfactual in the tradition of DiNardo, Fortin and Lemieux (1996), in which the conditional distribution of wages changes while the composition of the workforce does not change. The effects of the recession that operate through changes in X, including changes in the industry and occupation mix, are therefore not part of $\tau(y)$. Since wages are only observed for workers with positive earnings, the analysis also conditions on having positive earnings during the year.

We also consider the aggregate distribution of wages. The aggregate counterfactual distribution combines equation (E.2) with the distribution of wages in less exposed states

$$F^c_{Y|T}(y \mid 1) = p\, F_{Y(0)|S,T}(y \mid 1,1) + (1-p)\, F_{Y|S,T}(y \mid 0,1), \qquad p = \Pr(S=1 \mid T=1) \quad \text{(E.4)}$$

which is equation (3) in the main text, so the effect on the aggregate distribution is $F_{Y|T}(y \mid 1) - F^c_{Y|T}(y \mid 1) = p\,\tau(y)$.

Identification

Let $\mathcal{Y}_0(s,t;x)$ denote the support of $Y(0)$ conditional on $S=s$, $T=t$, and $X=x$. We make the following assumptions.

Assumption A1 (No interaction). $\delta(y,X) = 0$ *almost surely for all* $y \in \mathbb{R}$ *in equation* (E.1).

Assumption E1 (Support). $\mathcal{Y}_0(1,1;X) \subseteq \mathcal{Y}_0(0,1;X) \cup \mathcal{Y}_0(1,0;X) \cup \mathcal{Y}_0(0,0;X)$ *almost surely.*

Assumption A1 is the distributional parallel trends assumption in Section VI.A. Substitut-

ing $\delta(y,x)=0$ into equation (E.1) shows that it is equivalent to

$$\Lambda^{-1}\big[F_{Y(0)|S,T,X}(y \mid 1,1,x)\big]-\Lambda^{-1}\big[F_{Y(0)|S,T,X}(y \mid 1,0,x)\big]=\Lambda^{-1}\big[F_{Y(0)|S,T,X}(y \mid 0,1,x)\big] \\ -\Lambda^{-1}\big[F_{Y(0)|S,T,X}(y \mid 0,0,x)\big]$$

which is a parallel trends assumption on the transformation Λ^{-1} of each conditional distribution rather than on the mean. Since this transformation is nonlinear, the choice of link function is part of the identifying assumption. Under the logit link we use, Assumption A1 requires that absent severe exposure the log odds of earning less than y would have changed by the same amount in both groups of states for workers with the same characteristics. Assumption E1 ensures that the expression in equation (E.5) below is well defined for severely exposed workers in 1988.

Lemma E1. *Under Assumptions A1 and E1,* $(y,x)\mapsto F_{Y(0)|S,T,X}(y \mid 1,1,x)$ *is identified on* $\mathbb{R}\times\mathcal{X}_{11}$ *by*

$$F_{Y(0)|S,T,X}(y \mid 1,1,x)=\Lambda\Big(\Lambda^{-1}\big[F_{Y|S,T,X}(y \mid 1,0,x)\big]+\Lambda^{-1}\big[F_{Y|S,T,X}(y \mid 0,1,x)\big]-\Lambda^{-1}\big[F_{Y|S,T,X}(y \mid 0,0,x)\big]\Big) \tag{E.5}$$

Proof. Under Assumption A1, $F_{Y(0)|S,T,X}(y \mid 1,1,x)=\Lambda(\alpha(y,x)+\beta(y,x)+\gamma(y,x))$, and evaluating equation (E.1) at the other three combinations of (s,t) gives $\alpha(y,x)=\Lambda^{-1}[F_{Y(0)|S,T,X}(y \mid 0,0,x)]$, $\beta(y,x)=\Lambda^{-1}[F_{Y(0)|S,T,X}(y \mid 0,1,x)]-\alpha(y,x)$, and $\gamma(y,x)=\Lambda^{-1}[F_{Y(0)|S,T,X}(y \mid 1,0,x)]-\alpha(y,x)$. Substituting, and replacing the distributions of $Y(0)$ with the observed distributions in the three groups where $Y=Y(0)$, gives equation (E.5). Assumption E1 ensures that each term is well defined. □

This is Lemma 2 in Fernández-Val et al. (2024). By Lemma E1 and equation (E.2), $F_{Y(0)|S,T}(y \mid 1,1)$ is identified, and so are $\tau(y)$, τ_q^*, and the effects on any measure of inequality, since these are functionals of identified distributions. Identifying the aggregate counterfactual in equation (E.4) requires one additional assumption.

Assumption E2 (No effects on less exposed states). *Severe exposure in other states does not affect the distribution of wages in less exposed states or the share of workers in severely exposed states* p.

Under Assumption E2, the observed distribution of wages in less exposed states and the observed share p are the same under the less severe counterfactual, so $F^c_{Y|T}(y \mid 1)$ is identified by equation (E.4). The first part of Assumption E2 also underlies the interpretation of less exposed states as a control group. If the recession lowered wages in less exposed states, for instance through ripple effects or migration, our estimates would understate the true effects.

Estimation

We use the CPS wage data assembled by Fortin, Lemieux and Lloyd (2021) (see Appendix A.6). Our pre-recession sample is the May 1979 CPS supplement, which records union coverage, and our post-recession sample is the 1988 Merged Outgoing Rotation Groups. We restrict the sample to men aged 16–64 with a non-allocated hourly wage and drop workers in Alaska, Hawaii, and the District of Columbia. Wages are measured in 1979 dollars, directly for hourly workers and as usual weekly earnings divided by usual weekly hours for others. All estimates are weighted by the CPS earnings weights. A worker is severely exposed if their state's unemployment rate increased by more than the median among the 48 states between 1979 and 1982, using annual averages from the BLS Local Area Unemployment Statistics.

We replace the functions in equation (E.1) with linear indexes in the covariates and estimate the distribution regression model in equation (4) for the observed wages of all workers, including severely exposed workers in 1988

$$F_{Y|S,T,X}(y \mid s,t,x) = \Lambda\big(p_\alpha(x)'\alpha(y) + p_\beta(x)'\beta(y)t + \gamma(y)s + \theta(y)st\big)$$

using a separate logit of $\mathbb{1}\{Y \le y\}$ at each threshold y in a grid $\mathcal{Y}$. We set $p_\theta(x) = 1$, so at each threshold the model restricts the effect of severe exposure on the log odds to be the same for all workers, but allows it to vary freely across thresholds. The grid $\mathcal{Y}$ consists of the 981 weighted quantiles of the pooled 1979 and 1988 wage distribution at the percentiles $0.010, 0.011, \ldots, 0.990$, following Chernozhukov, Fernández-Val and Melly (2013). Because reported wages are heaped at round dollar amounts, many of these quantiles coincide, and the grid contains 465 distinct thresholds.

The covariates $p_\beta(x)$, which are interacted with the year, include years of schooling, a quartic in potential experience, their interaction, 16 education-by-experience cells, marital status, race, 10 industries, and 4 occupation groups, as well as the 1979 share of state employment in each of the 10 industries and the 1979 union coverage rate in the worker's state-by-industry market. The covariates $p_\alpha(x)$ additionally include part-time status, metropolitan area of residence, quarter, and state fixed effects. Since exposure is assigned at the state level, the state fixed effects absorb the main effect of S, and $\theta(y)$ is the coefficient on $S \times T$.

We then construct the plug-in estimators

$$\widehat{F}_{Y(0)|S,T}(y \mid 1,1) = \frac{1}{N_{11}} \sum_{i:\, S_i = T_i = 1} \Lambda\big(p_\alpha(X_i)'\hat{\alpha}(y) + p_\beta(X_i)'\hat{\beta}(y) + \hat{\gamma}(y)\big)$$

and $\widehat{F}_{Y(1)|S,T}(y \mid 1,1)$, which includes $\hat{\theta}(y)$ inside Λ, where the averages use the CPS weights

and N_{11} is their sum among severely exposed workers in 1988. Since we use the logit link, $\widehat{F}_{Y(1)|S,T}(y \mid 1,1)$ equals the empirical distribution of wages among severely exposed workers in 1988. We estimate the aggregate counterfactual by combining $\widehat{F}_{Y(0)|S,T}(y \mid 1,1)$ with the empirical distribution in less exposed states in 1988 using equation (E.4), with $\hat{p} = 0.56$. We rearrange each estimated distribution to be monotonic (Chernozhukov, Fernández-Val and Galichon, 2010) and estimate quantiles as $\widehat{F}^{-1}(q) = \inf\{y : \widehat{F}(y) \geq q\}$ at $q = 0.020, 0.021, \ldots, 0.980$. For the aggregate distribution, we combine the two distributions before rearranging and inverting.

Because of the heaping in reported wages, the estimated quantile functions are step functions, and differences between them are noisy. We therefore smooth each quantile function using a Gaussian kernel in q with bandwidth 0.035 before taking differences, and we report quantile effects between the 8th and 92nd percentiles. The percentile gaps in Table 5 are read from the smoothed quantile functions. The standard deviation of log wages and the Gini coefficient of wage levels are computed from the unsmoothed quantile functions between the 2nd and 98th percentiles. This trimming lowers the levels of both measures but largely cancels in the differences we report.

The share of the change in each measure explained by the recession divides the effect of severe exposure by the observed change in the same measure between 1979 and 1988, computed in the same way. We use the aggregate change for the aggregate effects and the change within severely exposed states for the effects within those states. Since this share compares a counterfactual in 1988 to a change over the decade, it measures the share of the rise in inequality between 1979 and 1988 attributable to the difference in exposure across states.

Inference

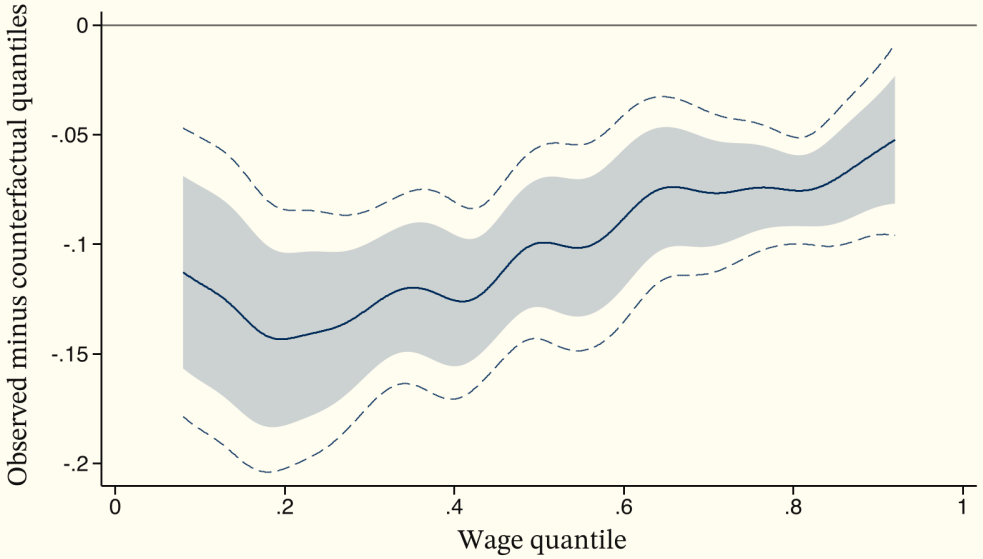
We use a weighted bootstrap with weights drawn at the state level to allow for correlation within states. In each of 100 replications, we draw a weight from a standard exponential distribution for each of the 48 states, multiply each worker's CPS weight by the weight of their state, and repeat every step of the estimation, including the distribution regressions, $\hat{p}$, the rearrangement, the inversion, and the smoothing. The standard error of each estimate is its standard deviation across replications, and pointwise 95 percent confidence intervals are calculated using this standard error. We also compute uniform confidence bands for the effects $\tau(y)$ and τ^*_q. Their critical value is the 95th percentile across replications of the largest deviation from the estimate over the reported range, standardized by the bootstrap standard error rather than the rescaled interquartile range in Fernández-Val et al. (2024). This gives critical values of 2.88 for the aggregate quantile effects and 3.58 for the aggregate CDF effects. The

threshold grid is held fixed across replications.

Effects within severely exposed states

Figure E.1 displays the estimates of τ_q^* and $\tau(y)$ within severely exposed states. The effects follow the same pattern as the aggregate effects in Figure 5 and are about twice as large. Severe exposure lowered wages by around 14 percent at the 20th percentile and around 6 percent at the 90th percentile.

A. Quantile Effects



B. CDF Effects



FIGURE E.1
Impacts of the 1980s Recession on the Distribution of Wages for Men in Severely Exposed States in 1988 Relative to the Less Severe Counterfactual

Notes: This figure displays estimates of the quantile and CDF effects of severe recession exposure within severely exposed states. Each panel displays the difference between the observed 1988 distribution of wages for men in severely exposed states and the counterfactual distribution in which their wage distribution evolved like that of less exposed states, as described in Appendix E.1. The quantile effects in Panel A are smoothed using a Gaussian kernel with bandwidth 0.035. The shaded regions correspond to pointwise 95 percent confidence intervals and the dashed lines to uniform 95 percent confidence bands, from a weighted bootstrap clustered at the state level with 100 replications.

E.2 PRE-RECESSION PLACEBO

Assumption A1 cannot be tested directly, since it restricts the distribution of $Y(0)$ for severely exposed workers in 1988, which we do not observe. However, with more than one pre-recession year we can examine whether the same restriction holds before the recession. We therefore re-estimate the counterfactual in Appendix E.1 using 1977 as the pre-period ($T = 0$) and 1979 as the post-period ($T = 1$). Since the recession had not yet occurred in 1979, the distributional treatment effect $\tau(y)$ in this placebo should be zero for all y if severely and less exposed states were on parallel trends in the sense of Assumption A1 before the recession.

The placebo uses the same estimator as our main estimates. Severely exposed states are the same states as in our main analysis, defined by the change in their unemployment rate between 1979 and 1982, and the covariates are the same, including the 1979 state industry shares and 1979 coverage rates interacted with year. We keep the 1979 controls because the placebo tests the same conditional assumption as our main design. Both years come from the May CPS supplements, so the placebo compares the same survey instrument in each year. We construct the threshold grid from the pooled 1977 and 1979 wage distribution and compute standard errors and uniform confidence bands using the same weighted bootstrap as our main estimates. We do not report the share of the change explained by severe exposure, since the observed changes in inequality between 1977 and 1979 are small and negative, which makes these shares uninterpretable.

Figure 6 displays the placebo quantile and CDF effects for the aggregate distribution, and Table E.1 reports the effects on each measure of inequality. The placebo quantile effects are centered around zero, and the uniform confidence bands contain zero at every percentile. None of the effects on the measures of inequality are statistically different from zero, either in the aggregate distribution or within severely exposed states. For instance, the placebo effect on the aggregate 90–10 gap is −0.009 (0.016), compared to 0.030 (0.014) in our main estimates. The largest placebo estimate is for the 90–50 gap, −0.016 (0.011), which has the opposite sign to our main estimate of 0.024 (0.010). If this difference in trends had continued after 1979, our main estimate of the effect on the 90–50 gap would be understated.

The placebo examines Assumption A1 over the two years before the recession. It cannot rule out differences in trends that emerged only after 1979, such as an acceleration of deindustrialization in more exposed states. Our controls for the 1979 industry composition of each state interacted with year account for differences in trends of this kind that are related to industry composition.

TABLE E.1

Impacts of the 1980s Recession on Wage Inequality — Pre-recession (1977–1979) Placebo

	Aggregate Distribution		High-exposure ($S = 1$) Distribution	
Statistic	Total change	Recession effect	Total change	Recession effect
90–10	-0.045	-0.009	-0.051	-0.018
		(0.016)		(0.029)
90–50	-0.026	-0.016	-0.031	-0.023
		(0.011)		(0.018)
50–10	-0.018	0.006	-0.020	0.006
		(0.008)		(0.016)
80–20	-0.024	0.004	-0.028	0.003
		(0.011)		(0.020)
Std. Dev.	-0.017	-0.003	-0.019	-0.005
		(0.004)		(0.008)
Gini	-0.008	-0.002	-0.010	-0.004
		(0.003)		(0.004)

Notes: This table displays estimates of the effect of severe exposure on each measure of wage dispersion in 1979 relative to the less severe counterfactual, using 1977 as the pre-period. The sample is restricted to men aged 16–64 in the CPS May supplements. Column 1 displays the observed change in each measure between 1977 and 1979. Column 2 displays the placebo effect of severe exposure, measured as the difference between the statistic computed on the observed and counterfactual distributions. Columns 3 and 4 repeat the same for workers in severely exposed states. All estimates are weighted by CPS weights. Standard errors are reported in parentheses below the effects and are computed using a weighted bootstrap clustered at the state level with 100 replications (see Appendix E.1).

E.3 ESTIMATING THE DECOMPOSITION

This section describes the intermediate distributions used in the decomposition in Section VIII, the identification assumptions, and our estimation procedure. We discuss the additional assumptions under which each component of the decomposition has a causal interpretation in Appendix E.4.

Setup and intermediate distributions

We extend the notation in Appendix E.1 to include each worker's own union coverage $U \in \{0, 1\}$ and their market M, which is their state-by-industry cell, using the 10 broad industries in Fortin, Lemieux and Lloyd (2021). The union density of market m in year t is $d_t(m) = \mathbb{E}[U \mid M = m, T = t]$, and $D = d_T(M)$ denotes the density of the worker's market. In addition to the potential wages $Y(0)$ and $Y(1)$, let $U(0)$ and $U(1)$ denote a worker's potential coverage under the less severe and severe exposure scenarios. For severely exposed workers in 1988, we define:

- $G_1(y \mid u, x, d) = F_{Y|S,T,U,X,D}(y \mid 1, 1, u, x, d)$, the observed distribution of wages among workers with coverage u and characteristics x in markets with union density d;
- $p_1(x) = \Pr(U = 1 \mid S = 1, T = 1, X = x)$, the observed probability of coverage;
- $\pi^c(x) = \Pr(U(0) = 1 \mid S = 1, T = 1, X = x)$, the probability of coverage under the less severe counterfactual;
- $d^c(m) = \mathbb{E}[U(0) \mid M = m, T = 1]$, the union density of market m under the less severe counterfactual;
- $G_0(y \mid u, x) = F_{Y(0)|S,T,U(0),X}(y \mid 1, 1, u, x)$, the counterfactual distribution of wages among workers who would have coverage u under the less severe counterfactual.

In the notation of Section VIII, the union and non-union distributions F^u and F^n correspond to G_1 in the observed distribution and to $G_0(\cdot \mid 1, x)$ and $G_0(\cdot \mid 0, x)$ in the counterfactual, and the probability of coverage $\pi_u(x)$ corresponds to $p_1(x)$ and $\pi^c(x)$.

The decomposition compares four distributions of wages for severely exposed workers in

1988:

$$\Psi_0(y) = \int G_1\big(y \mid u, x, d_1(m)\big)\, dF_{U,X,M|S,T}(u, x, m \mid 1, 1) = F_{Y|S,T}(y \mid 1, 1)$$
$$\Psi_1(y) = \int G_1\big(y \mid u, x, d^c(m)\big)\, dF_{U,X,M|S,T}(u, x, m \mid 1, 1)$$
$$\Psi_2(y) = \int \Big[p_1(x)\, G_0(y \mid 1, x) + \big(1 - p_1(x)\big)\, G_0(y \mid 0, x)\Big]\, dF_{X|S,T}(x \mid 1, 1)$$
$$\Psi_3(y) = F_{Y(0)|S,T}(y \mid 1, 1)$$

Ψ_0 is the observed distribution. Ψ_1 evaluates the observed wage distribution of each worker's coverage group and characteristics at the counterfactual density of their market, holding everything else at its observed value. Ψ_2 replaces the wage distributions of covered and uncovered workers with their counterfactual distributions, holding the probability of coverage at its observed value. Ψ_3 is the counterfactual distribution in Appendix E.1. The union spillover, nonunion wage structure, and union composition components of the decomposition are $\Psi_0 - \Psi_1$, $\Psi_1 - \Psi_2$, and $\Psi_2 - \Psi_3$, which sum to the total effect $\tau(y) = \Psi_0(y) - \Psi_3(y)$. As in Appendix E.1, the aggregate version of each distribution combines it with the observed distribution in less exposed states using equation (E.4), so each component of the aggregate decomposition is p times the corresponding component in severely exposed states.

Identification

Ψ_0 is observed, and Ψ_3 is identified by Lemma E1. Identifying Ψ_1 and Ψ_2 requires the following assumptions, where $p_{st}(x) = \Pr(U = 1 \mid S = s, T = t, X = x)$.

Assumption E3 (Coverage parallel trends). *The probability of coverage under the less severe counterfactual satisfies the no-interaction restriction on the logit scale*

$$\Lambda^{-1}\big[\pi^c(x)\big] - \Lambda^{-1}\big[p_{10}(x)\big] = \Lambda^{-1}\big[p_{01}(x)\big] - \Lambda^{-1}\big[p_{00}(x)\big]$$

with $0 < p_{st}(x) < 1$ *for* $(s, t) \neq (1, 1)$*, for almost every* x *among severely exposed workers in 1988.*

Assumption E4 (Coverage and markets). $U(0) \perp M \mid S, T, X$.

Assumption E5 (Within-coverage parallel trends). *Assumptions A1 and E1 hold for the distribution of* $Y(0)$ *conditional on* $U(0)$ *and* X.

Assumption E6 (Density support). *For almost every severely exposed worker in 1988,* $(U, X, d^c(M))$ *lies in the support of* (U, X, D) *among severely exposed workers in 1988.*

Assumption E3 applies the same difference-in-differences restriction as Assumption A1 to the probability of coverage, so that the difference in the log odds of coverage between severely and less exposed states would have remained constant absent severe exposure among workers with the same characteristics. Assumption E4 states that, among workers with the same exposure and characteristics, the market a worker is in carries no additional information about their counterfactual coverage. Assumption E5 applies the identification argument in Appendix E.1 separately to workers who would and would not be covered under the less severe counterfactual. Assumption E6 ensures that the counterfactual density in each market is within the range of densities observed among severely exposed workers with the same coverage and characteristics.

Lemma E2 (Counterfactual coverage and density). *Under Assumption E3,* $\pi^c(x) = \Lambda\big(\Lambda^{-1}[p_{10}(x)] + \Lambda^{-1}[p_{01}(x)] - \Lambda^{-1}[p_{00}(x)]\big)$. *Under Assumptions E3 and E4,* $d^c(m) = \mathbb{E}[\pi^c(X) \mid M = m, T = 1]$ *for every severely exposed market* m.

Proof. In the three groups with $ST = 0$, observed coverage equals counterfactual coverage, $U = U(0)$, so $p_{st}(x) = \Pr(U(0) = 1 \mid S = s, T = t, X = x)$. Substituting into Assumption E3 and applying Λ gives the first result. By the law of iterated expectations, $d^c(m) = \mathbb{E}\big[\mathbb{E}[U(0) \mid X, M, T = 1] \mid M = m, T = 1\big]$. Since exposure is determined by the market, Assumption E4 implies $\mathbb{E}[U(0) \mid X, M = m, T = 1] = \mathbb{E}[U(0) \mid X, S = 1, T = 1] = \pi^c(X)$ in every severely exposed market, which gives the second result. □

Lemma E3 (Intermediate distributions). *Under Assumptions E3, E4, and E6,* Ψ_1 *is identified. Under Assumption E5 and* $0 < p_{st}(x) < 1$, $G_0(y \mid u, x)$ *is identified by equation (E.5) with each distribution also conditioned on* $U = u$, *and* Ψ_2 *is identified.*

Proof. G_1 is a conditional distribution of observed variables and $d^c(m)$ is identified by Lemma E2, so under Assumption E6 every evaluation of G_1 in Ψ_1 is identified. In the three groups with $ST = 0$, $U = U(0)$ and $Y = Y(0)$, so the observed distributions of wages conditional on (U, X) in these groups are the distributions of $Y(0)$ conditional on $(U(0), X)$. Applying Lemma E1 conditional on $U(0) = u$ under Assumption E5 identifies $G_0(y \mid u, x)$. $p_1(x)$ and $F_{X|S,T}(x \mid 1, 1)$ are observed, so Ψ_2 is identified. □

Estimation

All of the intermediate distributions use the sample, threshold grid, covariates, and weights described in Appendix E.1.

Counterfactual coverage and density. We estimate the probability of coverage using a logit of U on the same covariates as in the distribution regressions, pooling both groups of states

and both years. To allow the effect of severe exposure on coverage to differ across workers, we interact $S \times T$ with a quartic in a predicted wage index, constructed as the fitted value from a regression of 1979 log wages on schooling, a quartic in experience, their interaction, education-by-experience cells, marital status, race, metropolitan area, and region. We estimate $p_1(x)$ from the fitted values and $\pi^c(x)$ by setting every $S \times T$ term to zero. We estimate $d_1(m)$ and $d^c(m)$ by averaging these fitted values over severely exposed workers in each of the 240 severely exposed markets in 1988.

Severe exposure lowered the probability of coverage of severely exposed workers by 1.9 percentage points on average, from 27.1 to 25.2 percent, and lowered union density by between 0.2 and 4.7 percentage points across markets. Figure E.2 shows how this effect varies with the predicted wage index. The decline in coverage was largest for workers with lower predicted wages, at up to 4 percentage points, and close to zero for workers with the highest predicted wages. This pattern is not driven by the quartic functional form. Replacing it with a separate effect for each ventile of the predicted wage index gives a similar pattern and the same average effect of 1.9 percentage points, as does a model with a single effect for all workers.

The wage distribution conditional on density. Following Fortin, Lemieux and Lloyd (2021), we estimate G_1 using a stacked distribution regression, separately for covered and uncovered workers in severely exposed states. To increase the variation in union density, we pool the years 1984–1988, and we measure union density as the contemporaneous coverage rate of the worker's market from the CPS (see Appendix A.6). We divide wages into 58 bins of 0.05 log points and estimate a single logit for $\mathbb{1}\{Y \leq y\}$ at the upper edge of every bin, with each worker contributing one observation per bin. Union density enters linearly and interacted with a quartic in the position of the bin, so the effect of density on the log odds of earning less than y varies smoothly across the distribution. The regressions also include the covariates in Appendix E.1, state-by-year and industry-by-year fixed effects, bin-by-year fixed effects, the 1979 coverage rate of the worker's market interacted with year, and interactions of the bin position with education-by-experience cells, state fixed effects, and the 1979 coverage rate and its square. Standard errors are clustered by state. Since the regressions include state-by-year and industry-by-year fixed effects and control for the 1979 coverage rate of the market, the effect of density is identified from differences in the evolution of union density across industries within the same state and year.

To summarize the estimates, we follow Fortin, Lemieux and Lloyd (2021) and convert the effect of a 1 percentage point increase in union density on the distribution into its effect on each quantile. Figure E.3 shows that higher union density raises wages at the bottom of the distribution and lowers them at the top, for both uncovered and covered workers. For uncovered workers, a 1 percentage point increase in union density raises wages by 0.4 percent at the

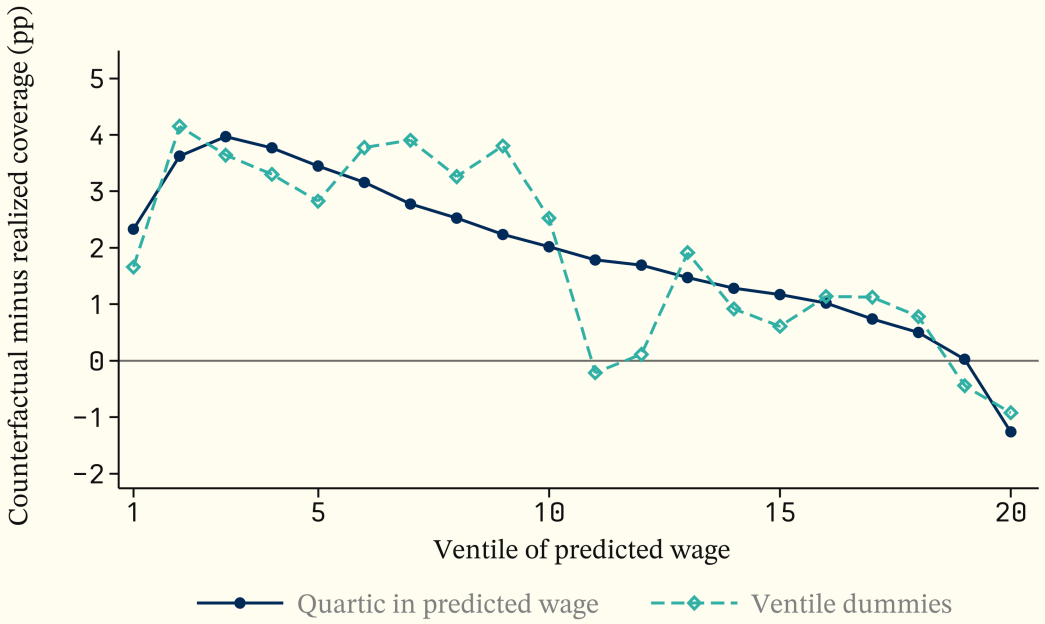


FIGURE E.2
Effect of Severe Exposure on Union Coverage by Predicted Wage

Notes: This figure displays the average difference between the counterfactual probability of coverage $\pi^c(x)$ and the observed probability of coverage $p_1(x)$ for severely exposed workers in 1988, in percentage points, by ventile of a predicted wage index. The predicted wage index is the fitted value from a regression of 1979 log wages on schooling, a quartic in experience, their interaction, education-by-experience cells, marital status, race, metropolitan area, and region. The solid line allows the effect of severe exposure on the log odds of coverage to vary with a quartic in the predicted wage index, and the dashed line allows it to vary freely across ventiles. See Appendix E.3 for details.

10th percentile and 0.25 percent at the median, and lowers them by 0.2 percent at the 90th percentile. This means that union density compresses the distribution of wages, as in Fortin, Lemieux and Lloyd (2021), reducing the 90–50 gap by 0.004 log points and the 90–10 gap by 0.006 log points for each percentage point of density. The decline in union density in more exposed markets therefore widened the wage distribution.

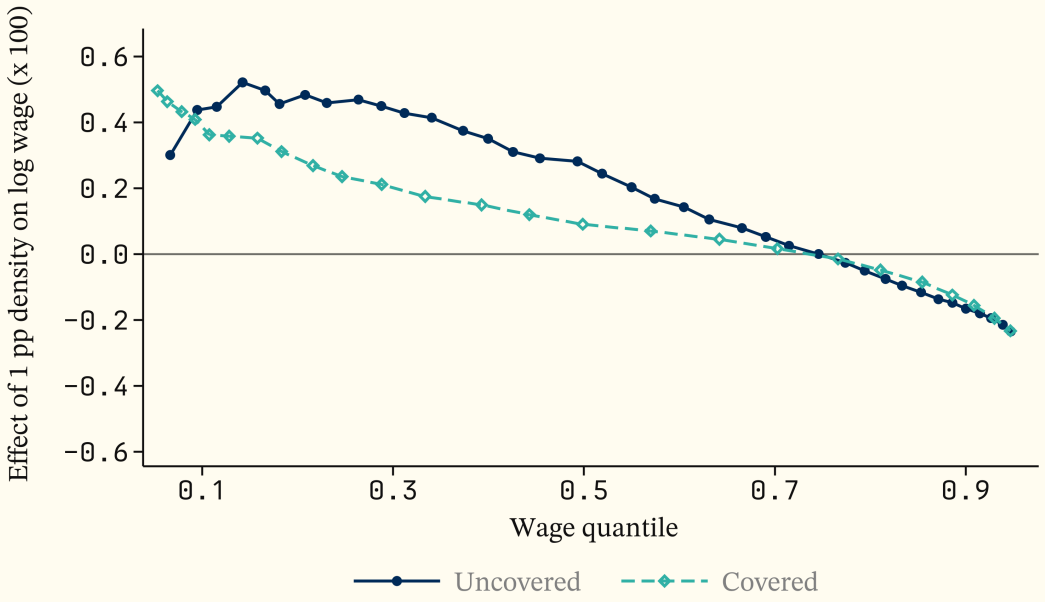


FIGURE E.3
Effect of Union Density on the Distribution of Wages

Notes: This figure displays the effect of a 1 percentage point increase in union density on each quantile of the wage distribution, multiplied by 100, for uncovered and covered workers in severely exposed states between 1984 and 1988. The effects are computed from the stacked distribution regressions described in Appendix E.3 by increasing the union density of every market by 1 percentage point and converting the change in the distribution into its effect on each quantile following Fortin, Lemieux and Lloyd (2021). Each point corresponds to one wage bin.

Figure E.4 shows that this pattern is robust to the specification of the density model for uncovered workers. The alternative specifications replace the state-by-year fixed effects with state-specific linear trends as in Fortin, Lemieux and Lloyd (2021), either with their controls or with all of our controls, allow the 1979 coverage rate to enter linearly in the bin position or not at all, and add interactions of the bin position with industry. In every other specification, higher density raises wages at the bottom of the distribution relative to the top, and the effect on the 90–50 gap ranges from −0.004 to −0.006 log points per percentage point of density. The exception is the specification that adds industry-by-position interactions together with state-by-year fixed effects, which reduces the effect on the 90–50 gap to −0.002. This specification

removes the variation in density across industries within each state and year, which is the variation that identifies the effect of density once we condition on state-by-year fixed effects. The effect on the 50–10 gap is small and changes sign across specifications, so the spillover mainly affects the top half of the distribution.

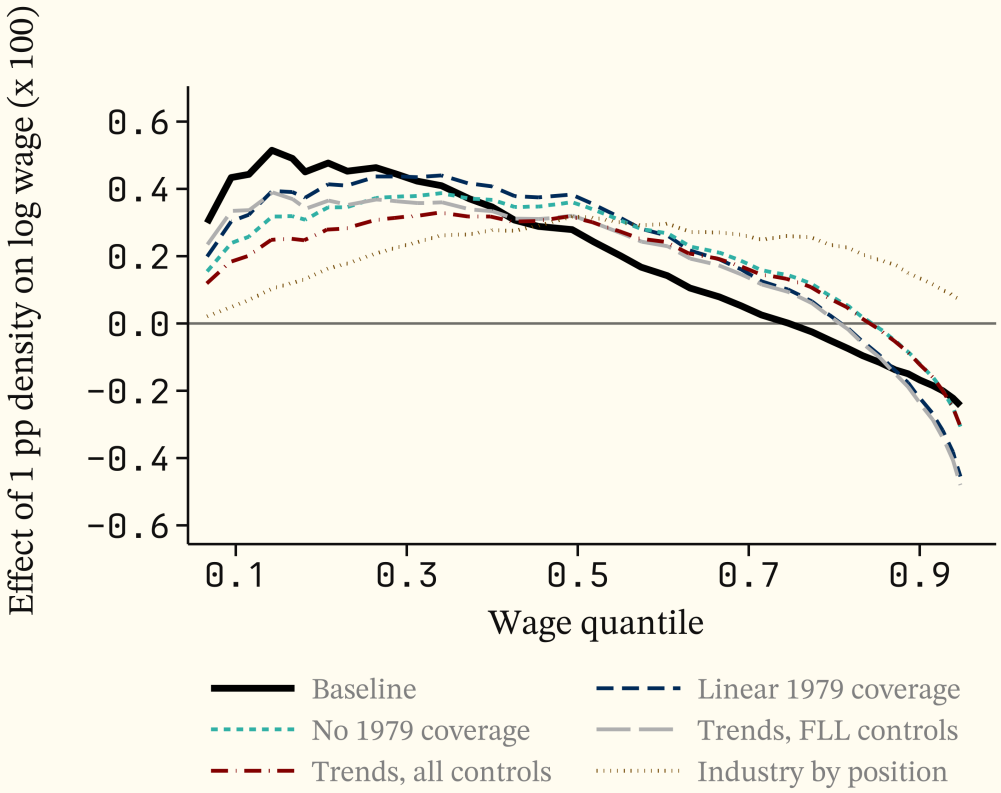


FIGURE E.4
Effect of Union Density on the Distribution of Wages — Alternative Specifications

Notes: This figure displays the effect of a 1 percentage point increase in union density on each quantile of the wage distribution, multiplied by 100, for uncovered workers in severely exposed states between 1984 and 1988, under alternative specifications of the stacked distribution regression. Baseline is the specification described in Appendix E.3. Linear 1979 coverage interacts the bin position with the 1979 coverage rate linearly rather than as a quadratic. No 1979 coverage omits the interactions of the bin position with the 1979 coverage rate. Trends, FLL controls replaces the state-by-year fixed effects with state fixed effects and state-specific linear trends and uses the controls in Fortin, Lemieux and Lloyd (2021), and Trends, all controls adds our additional controls and the interactions of the bin position with education and experience, the 1979 coverage rate, and industry. Industry by position adds interactions of the bin position with industry to the specification with state-by-year fixed effects and a linear 1979 coverage rate. Standard errors are not shown.

The counterfactual distributions by coverage. We estimate G_0 using the distribution regressions in Appendix E.1 with coverage added to the model. At each threshold, we add coverage and its interactions with year, S, and $S{\times}T$, as well as interactions of coverage with schooling, experience, and education-by-experience cells, so that every term of the difference-in-differences

differs by coverage. We estimate Ψ_2 by setting both $S \times T$ terms to zero and averaging the fitted values over severely exposed workers in 1988 at their observed coverage, which weights $G_0(y \mid 1, x)$ and $G_0(y \mid 0, x)$ by $p_1(x)$.

Components of the decomposition. We rearrange and invert each of Ψ_0–Ψ_3 and smooth the resulting quantile functions as described in Appendix E.1. Since we smooth each quantile function before taking differences, the three components add up exactly to the total effect on each quantile and on each measure of inequality. The share of the total effect attributed to each component in Table 8 divides the effect of that component by the total effect.

E.4 CAUSAL INTERPRETATION OF THE DECOMPOSITION

The intermediate distributions in Appendix E.3 are identified under the assumptions in that section, and the decomposition of the total effect into their differences holds as an accounting identity. This section describes the causal effect that each component aims to measure and the additional assumptions required for that interpretation.

Causal parameters

The decomposition asks how much of the recession's effect on the wage distribution operated through its effects on unions. To make this precise, we need to describe what wages would have been if the recession had affected only some aspects of unionization. We therefore extend our potential outcomes. Let $Y(s, d, u)$ denote the wage a worker would earn under exposure s, if the union density of their market were d and their own coverage were u. Let $U(s)$ and $D(s)$ denote the worker's coverage and the density of their market under exposure s, so that $D(1) = d_1(M)$ and $D(0) = d^c(M)$ for severely exposed workers in 1988. The wages we observe for severely exposed workers in 1988 are $Y(1, D(1), U(1))$, and the wages they would have earned under the less severe counterfactual are $Y(0, D(0), U(0))$. The total effect compares the distributions of these two wages, which we denote T_0 and T_3.

The recession lowered both the density of unions in each market and the coverage of individual workers, and each channel corresponds to undoing one of these effects at a time. We start from the observed distribution T_0. First, consider what wages would have been if the density of each worker's market had followed its counterfactual path, but every worker kept their observed coverage and everything else about the recession were unchanged. This is the distribution of $Y(1, D(0), U(1))$, which we denote T_1. The difference $T_0 - T_1$ is the union spillover, which captures how the decline in union density changed the wages of workers, regardless of whether they themselves were covered.

Second, consider what wages would have been if, in addition, the wage structure were at its counterfactual value, but coverage remained at its observed post-recession rate. To define this, we need the wages that coverage would produce for a typical worker rather than for the particular workers with that coverage status. For coverage status u and characteristics x, let

$$P_u(y, x) = \int \Pr\big(Y(0, d, u) \le y \mid S = 1, T = 1, X = x\big)\, dF_{D(0)|S,T,X}(d \mid 1, 1, x)$$

denote the probability that a severely exposed worker with characteristics x would earn less than y under the less severe counterfactual if their coverage were set to u, averaging over the counterfactual density of their markets. We then assign coverage within each group of workers

with the same characteristics at its observed rate $p_1(x)$, which gives

$$T_2(y) = \int \Big[p_1(x)\, P_1(y,x) + \big(1 - p_1(x)\big)\, P_0(y,x) \Big]\, dF_{X|S,T}(x \mid 1,1).$$

Note that T_2 specifies how many workers are covered, but not which ones. The difference $T_1 - T_2$ is the non-union wage structure component, which captures all of the remaining changes in how wages are set that do not operate through unions. Finally, the difference $T_2 - T_3$ is the union composition component, which captures the effect of the decline in the rate of coverage, holding wages at their counterfactual values. These three components add up to the total effect $T_0 - T_3$.

Assumptions

The observed distribution Ψ_0 equals T_0, and by Lemma E1 the counterfactual distribution Ψ_3 equals T_3. The total effect therefore only requires the assumptions in Appendix E.1. To interpret the intermediate distributions Ψ_1 and Ψ_2 as T_1 and T_2, we need three additional assumptions, each among severely exposed workers in 1988.

Assumption A2 (Union density unconfoundedness). *Among workers with the same coverage and characteristics, the realized and counterfactual density of a worker's market are independent of their potential wages:* $\{Y(1,d,u)\}_{d,u} \perp \big(D(1), D(0)\big) \mid S = 1, T = 1, U(1), X$.

Assumption A2 rules out that unions survived in markets whose workers would have earned more regardless. Since we control for each market's 1979 coverage rate, it only restricts how density changed, not its level, and it places no restrictions on which workers are covered. Under Assumption A2, the wage distribution of workers with the same coverage and characteristics in markets with density d describes what these workers would earn at density d. This means we can evaluate it at the counterfactual density of each worker's market, which is identified under the assumptions in Appendix E.3, so Ψ_1 equals T_1 and the union spillover $\Psi_0 - \Psi_1$ measures the effect of the decline in union density.

Assumption A3 (Counterfactual coverage unconfoundedness). *Among workers with the same characteristics, counterfactual coverage and market density are independent of potential wages:* $\{Y(0,d,u)\}_{d,u} \perp \big(U(0), D(0)\big) \mid S = 1, T = 1, X$.

Assumption A4 (Coverage and markets). *Among workers with the same characteristics, coverage is independent of the worker's market:* $\big(U(0), U(1)\big) \perp M \mid S = 1, T = 1, X$.

Assumption A3 rules out that the workers who would be covered absent the recession differ in their wage prospects from other workers with the same characteristics. Assumption A4

ensures that covered and uncovered workers with the same characteristics are spread across markets in the same way, so they face the same distribution of union densities. Our counterfactual distributions $G_0(y \mid u, x)$ describe the wages of the particular workers who would have coverage u absent the recession. Under Assumptions A3 and A4, these are also the wages that coverage u would produce for any worker with characteristics x, $G_0(y \mid u, x) = P_u(y, x)$. This means that Ψ_2 equals T_2, so the non-union wage structure and union composition components measure the effects described above.

Each conclusion therefore rests on a different set of assumptions. The total effect $\Psi_0 - \Psi_3$ relies only on the assumptions in Appendix E.1, and the union spillover $\Psi_0 - \Psi_1$ additionally relies on Assumption A2. Assumptions A3 and A4 are only needed to split the remainder $\Psi_1 - \Psi_3$ into its two components. If they fail, Ψ_2 differs from T_2, and the error $\Psi_2 - T_2$ is added to the union composition component and subtracted from the non-union wage structure component, leaving the total effect and the union spillover unaffected.

Discussion of Assumptions

We believe that Assumption A3 is the strongest of these assumptions. It requires that, among workers with the same characteristics, the workers who would be covered absent the recession have the same wage prospects as other workers, both when covered and when not covered. Under this assumption, the counterfactual wage distributions of covered and uncovered workers, $G_0(y \mid 1, x)$ and $G_0(y \mid 0, x)$, describe the wages that coverage and non-coverage would produce for any worker with characteristics x. The union composition component multiplies the coverage lost in the recession by the difference between these two distributions. If covered and uncovered workers differ in ways that our characteristics do not capture, part of this difference reflects who is covered rather than the effect of coverage, and this part is attributed to the composition component.

Note that Assumption A3 concerns selection into coverage absent the recession. It compares the workers who would be union members under the less severe counterfactual to other workers with the same characteristics in that same scenario. It places no restrictions on which workers lost coverage in the recession, since T_2 assigns coverage at the observed post-recession rate without specifying which workers keep it. The union composition component therefore measures the effect of the lower rate of coverage, valued at what coverage does for a typical worker, rather than the wage losses of the particular workers who lost coverage. These could differ if, for instance, the workers who lost coverage were employed at less productive union employers.

Selection into coverage could take several forms, which have different implications for the

sign of the bias. We focus on the case where workers who would be covered absent the recession have better wage prospects than other workers with the same characteristics and coverage does not lower wages. In other words, union members are positively selected relative to similar non-members. This is consistent with union employers selecting from queues of applicants for union jobs, or with unions organizing establishments with more productive workers.

To see how selection affects our estimates, consider workers with characteristics x and let $C_u(y)$ and $N_u(y)$ denote the probability of earning less than y under the less severe counterfactual with coverage set to u, for the workers who would and would not be covered absent the recession. We observe $G_0(y \mid 1, x) = C_1(y)$ and $G_0(y \mid 0, x) = N_0(y)$, but not $C_0(y)$ or $N_1(y)$. Under Assumption A4, the distribution that coverage would produce for all workers is $P_u(y, x) = \pi^c(x)C_u(y) + \big(1 - \pi^c(x)\big)N_u(y)$, and substituting into T_2 shows that the error in the union composition component is

$$\Psi_2(y) - T_2(y) = \int \Big[\underbrace{p_1(x)\big(1 - \pi^c(x)\big)\big(C_1(y) - N_1(y)\big)}_{\text{covered wages}} + \underbrace{\big(1 - p_1(x)\big)\pi^c(x)\big(N_0(y) - C_0(y)\big)}_{\text{uncovered wages}} \Big] dF_{X|S,T}(x \mid 1, 1) \tag{E.6}$$

Under Assumption A3, $C_u(y) = N_u(y)$ for each u and both terms are zero. Positive selection means that the workers who would be covered have higher wages, so they are less likely to earn less than any y: $C_1(y) \leq N_1(y)$ and $C_0(y) \leq N_0(y)$. The first term is then negative and the second positive. Since the coverage lost in the recession lowers wages at the bottom of the distribution, raising the distribution function there, a positive error means we overstate the composition component and a negative error means we understate it. The sign of the bias therefore depends on which of the two terms dominates. We consider two forms of selection.

Level selection. Workers who would be covered earn more than other workers with the same characteristics whether or not they are covered, for instance because of higher unobserved productivity, so that $N_1(y) - C_1(y) = N_0(y) - C_0(y) = s\,\Delta_G(y, x)$ for some $s \in [0, 1]$, where $\Delta_G(y, x) = G_0(y \mid 0, x) - G_0(y \mid 1, x)$ is the observed difference between the two distributions. Substituting into equation (E.6),

$$\Psi_2(y) - T_2(y) = s \int \big[\pi^c(x) - p_1(x)\big]\,\Delta_G(y, x)\, dF_{X|S,T}(x \mid 1, 1) \;\geq\; 0 \tag{E.7}$$

since severe exposure lowered coverage, $\pi^c(x) > p_1(x)$. Intuitively, the difference between the wage distributions of covered and uncovered workers overstates the effect of coverage on wages, since part of it reflects the higher wages these workers would earn anyway. Our estimates would therefore overstate the effect of the coverage lost in the recession, so the union composition component is overstated and the non-union wage structure component is under-

stated by the same amount.

Gains selection. Workers who would be covered are those whose wages coverage raises the most, while they would earn the same as other workers without coverage, so that $C_0(y) = N_0(y)$. This is consistent with workers sorting into union jobs when the union premium for them is largest. Only the first term in equation (E.6) remains, and the error is negative. The wage distribution of covered workers then overstates the effect of coverage for a typical worker. Since we use it to describe the wages of the smaller share of workers who remain covered after the recession, our estimate of the post-recession distribution Ψ_2 overstates the wages these workers would receive from coverage, which makes it look more like the counterfactual distribution with higher coverage. The union composition component is therefore understated and the non-union wage structure component overstated.

We can bound the size of these biases using the estimated distributions, since under positive selection and a non-negative effect of coverage, the bias is a known function of the difference between $G_0(y \mid 1, x)$ and $G_0(y \mid 0, x)$ and of the observed and counterfactual coverage rates. Under level selection, the bias is largest when all of the difference between covered and uncovered workers reflects selection. In this extreme case, the bias would account for 55 percent of our estimate of the composition component for the 90–50 gap and around 40 percent for the 90–10 and 80–20 gaps. Under gains selection, the true composition component for the 90–50 gap could be up to around four times our estimate. The bias in the direction that would overstate the composition component is limited because, as equation (E.7) shows, under level selection it is proportional to the 1.9 percentage points of coverage lost in the recession, $\pi^c(x) - p_1(x)$. Under gains selection, it is instead proportional to $p_1(x)(1 - \pi^c(x))$, which is around 0.18 on average. The data do not tell us which form of selection holds. However, even under level selection, which works against our results, the composition component remains positive and is at least around half of our estimate.

Selection could also be negative. Card (1996) finds that union workers are positively selected among workers with lower observed skills but negatively selected among workers with higher observed skills. Negative selection is outside the range our bounds consider, and it reverses their implications since negative level selection would lead us to understate the composition component. However, the coverage lost in the recession was concentrated among workers with lower predicted wages (Figure E.2), where Card (1996) finds positive selection, so we view the bounds above as the relevant case. Finally, our characteristics X include detailed measures of education, experience, industry, and occupation, as well as the 1979 coverage rate of the worker's market, which limits the scope for selection on the characteristics that typically predict union coverage.

F. Model Appendix

This appendix discusses the model from Section III formally along with proofs for each result. We build on Manning (2003) and Card et al. (2018) for wage setting, Autor, Dube and McGrew (2023) for the link between competition and the shape of the wage distribution, Dixit (1989) for exit, and Rosen (1969) and Taschereau-Dumouchel (2020) for the threat of unionization, and we model the union sector as a wage floor based on Freeman (1980).

F.1 ENVIRONMENT

Throughout, $c \equiv \varepsilon/(1+\varepsilon)$ and $\Delta \equiv \psi T\kappa/\varepsilon$, and we denote post-recession objects with primes.

Firms

There are M active firms, where M is large and finite. Firm j has productivity p_j drawn independently from a continuous distribution Γ with support $[\underline{p}, \overline{p}]$. Each firm produces $y_j = p_j \ln l_j$, so $MRPL_j = p_j/l_j$, and pays a fixed cost $F > 0$ per period. A fraction ψ of firms are unionized, which is also local union density. We choose worker units so that $l_j > 1$ at every active firm, which makes $\ln l_j > 0$.

Workers

There are N potential workers and $L \leq N$ participate. Conditional on participating, worker i derives utility $\varepsilon \ln w_j + e_{ij}$ from firm j, where e_{ij} is i.i.d. Type I extreme value. The probability that worker i chooses firm j is the logit share $s_j = w_j^\varepsilon / \sum_k w_k^\varepsilon$, so $l_j = Ls_j$ and the elasticity of firm j's labor supply with respect to its own wage is $\varepsilon(1-s_j)$.

(F1) Monopsonistic Competition. Each firm's share s_j is small and firms take $A \equiv L/\sum_k w_k^\varepsilon$ as given. Firm j faces labor supply $l_j = Aw_j^\varepsilon$ with elasticity ε.

Since raising all wages together leaves every s_j unchanged, ε is a firm-level elasticity and participation determines total labor supply. Participation depends on the average wage $w^a \equiv B/L$, where $B \equiv \sum_k w_k l_k$ is the total wage bill

$$L = N\left(\frac{w^a}{b}\right)^\theta \tag{F.1}$$

with $\theta > 0$ and b the value of non-employment.[47]

[47] This is the participation curve implied by reservation wages distributed on $[0, b]$ with cdf $(r/b)^\theta$, where a worker participates if the average wage exceeds her reservation wage. We use the average wage rather than the logit inclusive value $(\sum_k w_k^\varepsilon)^{1/\varepsilon}$ because the latter rises mechanically when ε falls.

Unions

(F2) Union Wage Floor. Union firms must pay at least a standard rate $\bar{w}$. This rate is set outside the local market and does not respond to local conditions.

Workers at a non-union firm paying w organize if their gain from doing so, $\ln(\bar{w}/w)$, exceeds the cost of organizing. The cost is random from the firm's perspective and is lower where unions already have a presence.

(F3) Organizing Probability. The cost of organizing is uniformly distributed on $[0, 1/(\psi\kappa)]$, so a non-union firm paying w is organized with probability

$$\Pi(w, \psi) = \begin{cases} \min\left\{1, \psi\kappa \ln\left(\frac{\bar{w}}{w}\right)\right\} & \text{if } w < \bar{w} \\ 0 & \text{if } w \geq \bar{w} \end{cases} \tag{F.2}$$

and $\psi\kappa \ln(\bar{w}/w) < 1$ at the lowest paying active non-union firm. If organized, the firm pays $\bar{w}$ and loses profit $T > 0$, which is constant across firms.

The recession

Productivity falls from p_j to λp_j with $\lambda < 1$ and then returns to p_j.

(F4) Irreversible Exit. Exit is irreversible. Re-entry requires a sunk cost that exceeds the present value of post-recovery profits for exiting firms.

(F5) Labor Supply Elasticity Falls. The labor supply elasticity depends on the number of active firms, $\varepsilon = \varepsilon(M)$, with $\varepsilon'(M) > 0$.

F.2 WAGE SETTING

Lemma F1 (Benchmark wage). *A non-union firm facing no threat of organization sets $w_n^0 = c \cdot MRPL$, with*

$$w_n^0(p) = \left(\frac{cp}{A}\right)^{\frac{1}{1+\varepsilon}}, \qquad w_n^0 l_n^0 = cp$$

For any two such firms with $p_L < p_H$, $w_n^0(p_L)/w_n^0(p_H) = (p_L/p_H)^{1/(1+\varepsilon)}$. The wage is below the firm's competitive wage $w_c(p) \equiv (p/A)^{1/(1+\varepsilon)}$, at which $MRPL = w$.

Proof. The firm maximizes $f_0(w) = p\ln(Aw^\varepsilon) - Aw^{1+\varepsilon} - F$. The first order condition $\varepsilon p/w = (1+\varepsilon)Aw^\varepsilon$ gives $wl = cp$, or $w = c \cdot p/l$. Since $f_0'' < 0$, this is the unique optimum. Substituting $l = Aw^\varepsilon$ gives the closed form and the wage ratio. Solving $p/(Aw^\varepsilon) = w$ gives $w_c(p)$, and $w_n^0 < w_c$ since $c < 1$. Since $MRPL > w_n^0$, the firm hires every willing worker. □

Lemma F2 (Threatened wage). *A non-union firm facing Assumption F3 and paying $w < \bar{w}$ sets*

$$w^\tau(p) = \left(\frac{c(p+\Delta)}{A}\right)^{\frac{1}{1+\varepsilon}}, \qquad w^\tau l^\tau = c(p+\Delta)$$

and pays a premium over the benchmark wage of

$$\ln w^\tau(p) - \ln w_n^0(p) = \frac{1}{1+\varepsilon} \ln\left(1 + \frac{\Delta}{p}\right) \tag{F.3}$$

which is positive and decreasing in p. If $\psi T\kappa < p$, then $w^\tau(p) < w_c(p)$ and the firm hires every willing worker.

Proof. The firm maximizes $f_\tau(w) = f_0(w) - \psi\kappa T \ln(\bar{w}/w)$. The first order condition is $(\varepsilon p + \psi T\kappa)/w = (1+\varepsilon)Aw^\varepsilon$, or $(1+\varepsilon)Aw^{1+\varepsilon} = \varepsilon(p+\Delta)$, and $f_\tau'' < 0$. Dividing by the benchmark wage gives equation (F.3). Finally, $w^\tau < w_c$ if and only if $c(p+\Delta) < p$, or $\psi T\kappa < p$. If the firm instead paid above w_c, it would ration workers and its objective would be decreasing in w when $\psi T\kappa < p$. □

The threat acts like a common increase Δ in productivity. Let $p_2 \equiv A\bar{w}^{1+\varepsilon}/c$ denote the productivity at which $w_n^0 = \bar{w}$.

Lemma F3 (Non-union firms). *A non-union firm (a) with $p < p_2 - \Delta$ pays $w^\tau(p) < \bar{w}$ and faces a positive probability of organization, (b) with $p_2 - \Delta \le p < p_2$ pays exactly $\bar{w}$ and faces no risk, and (c) with $p \ge p_2$ pays $w_n^0(p) \ge \bar{w}$ and faces no risk. The wage is continuous and increasing in p.*

Proof. The firm's objective is f_τ for $w < \bar{w}$ and f_0 for $w \ge \bar{w}$. It is continuous at $\bar{w}$ since $\Pi(\bar{w}, \psi) = 0$, and both pieces are strictly concave. From Lemma F2, $w^\tau = \bar{w}$ at $p = p_2 - \Delta$. In case (a), $w^\tau < \bar{w}$ is feasible and f_0 is decreasing above $\bar{w}$ since $w_n^0 < w^\tau < \bar{w}$. In case (b), f_τ is increasing below $\bar{w}$ and f_0 is decreasing above $\bar{w}$, so the optimum is $\bar{w}$. In case (c), $w_n^0 \ge \bar{w}$ is feasible and $f_\tau < f_0$ for all $w < \bar{w}$. □

Note. The uniform cost distribution is what makes the premium in equation (F.3) decreasing in p. With an exponential cost distribution with mean $1/(\psi\kappa)$, the probability of organization is $1 - (w/\bar{w})^{\psi\kappa}$. The premium is then small at the very bottom, since a firm far below $\bar{w}$ is likely to be organized regardless, largest in the lower middle of the distribution, and zero at $\bar{w}$. This is closer to the shape Fortin, Lemieux and Lloyd (2021) estimate. The results below only require that the premium is positive where the floor would bind, increasing in ψ holding ε fixed, and concentrated in the bottom of the non-union distribution.

F.3 UNION FIRMS AND EXIT

Let $p_1 \equiv A\bar{w}^{1+\varepsilon} = cp_2$ denote the productivity at which $w_c = \bar{w}$. Since $c < 1$, $p_1 < p_2$.

Lemma F4 (Union firms). *A union firm (i) with $p \geq p_2$ pays $w_n^0(p)$ and hires $l_n^0(p)$, (ii) with $p_1 \leq p < p_2$ pays $\bar{w}$ and hires every willing worker, $l_u = A\bar{w}^{\varepsilon} > l_n^0(p)$, and (iii) with $p < p_1$ pays $\bar{w}$ and hires along its demand curve, $l_u = p/\bar{w} < A\bar{w}^{\varepsilon}$. The union premium $\ln \bar{w} - \ln w_n^0(p)$ is positive and decreasing in p for $p < p_2$ and zero otherwise.*

Proof. If $w_n^0(p) \geq \bar{w}$, or $p \geq p_2$, the floor does not bind. Otherwise, the firm pays $\bar{w}$, since any higher wage raises costs and the firm can hire the same workers at $\bar{w}$. The firm then chooses $l \leq A\bar{w}^{\varepsilon}$ to maximize $p \ln l - \bar{w}l$, whose unconstrained optimum is $l = p/\bar{w}$, so it hires $\min\{A\bar{w}^{\varepsilon}, p/\bar{w}\}$. The two are equal at $p = p_1$. The union premium is positive and decreasing since $w_n^0(p)$ is increasing and equals $\bar{w}$ at p_2. □

Lemma F5 (Profit and exit). *Union profit $\pi_u(p)$, threatened non-union profit $\pi_n(p)$, and benchmark non-union profit $\pi_n^0(p)$ are continuous and strictly increasing in p. If each is negative at $\underline{p}$ and positive at $\overline{p}$, there are unique viability thresholds p_u^* and p_n^*, and firms with productivity below the threshold exit.*

Proof. By the envelope theorem, the derivative of each with respect to p is $\ln l > 0$ at the optimum. □

In case (iii), union profit is $p \ln(p/\bar{w}) - p - F$, so the union viability threshold satisfies

$$p_u^* \ln \frac{p_u^*}{\bar{w}} = p_u^* + F, \qquad \frac{\partial p_u^*}{\partial \bar{w}} = \frac{p_u^*}{\bar{w} \ln(p_u^*/\bar{w})} > 0 \tag{F.4}$$

where the derivative follows from total differentiation. This applies when the marginal union firm is in case (iii), or $\pi_u(p_1) > 0$, which holds when $F < A\bar{w}^{1+\varepsilon}[\ln(A\bar{w}^{\varepsilon}) - 1]$.

Proposition F1 (Differential profitability). *For all p, (a) $\pi_u(p) \leq \pi_n^0(p)$, strictly if $p < p_2$, (b) $\pi_n^0(p) - T \leq \pi_n(p) \leq \pi_n^0(p)$, and (c) if $T \leq \pi_n^0(p) - \pi_u(p)$, then $\pi_u(p) \leq \pi_n(p)$.*

Proof. (a) The union firm chooses (w, l) from $\{w \geq \bar{w}, l \leq Aw^{\varepsilon}\}$, which is a subset of the benchmark firm's choice set $\{l \leq Aw^{\varepsilon}\}$, with the same objective. The inequality is strict when the benchmark optimum is not feasible, or $w_n^0 < \bar{w}$. (b) The threat is a cost, which gives the upper bound. The firm can always pay w_n^0 and earn $\pi_n^0 - \Pi(w_n^0, \psi)T \geq \pi_n^0 - T$, which gives the lower bound. (c) follows from (b). □

The condition in (c) means that the loss from being organized is smaller than the cost of the union constraint.

Note. A natural alternative to a constant T is the firm's own loss in profit from being organized, $\pi(w)-\pi_u(p)$, in which case the non-union firm maximizes $(1-\Pi)\pi(w)+\Pi\pi_u(p)$. The premium is still positive where the floor would bind and decreasing in p, but it is zero at exactly p_2, so non-union firms no longer match the union scale. The condition in Proposition F1(c) also holds automatically, since the firm can always pay $\bar{w}$ and earn $\pi_u(p)$. However, we lose the closed form in Lemma F2.

(F6) Floor Binds at the Exit Margin. The marginal non-union firm is threatened, $p_n^* < p_2 - \Delta$, and $T \leq \pi_n^0(p_n^*) - \pi_u(p_n^*)$.

Proposition F2 (Differential exit). *Under Assumption F6, $p_u^* > p_n^*$.*

Proof. By Proposition F1(b), $\pi_n(p_n^*) \geq \pi_n^0(p_n^*) - \Pi(w_n^0, \psi)T > \pi_n^0(p_n^*) - T$, since $\Pi < 1$ under Assumption F3. By Assumption F6, $\pi_n^0(p_n^*) - T \geq \pi_u(p_n^*)$, so $\pi_u(p_n^*) < \pi_n(p_n^*) = 0$. Since π_u is increasing, $p_u^* > p_n^*$. □

F.4 THE RECESSION

Proposition F3 (The downturn). *During the downturn, holding $\bar{w}$, Δ, and A fixed, (a) threatened non-union log wages change by $\frac{1}{1+\varepsilon}\ln\frac{\lambda p+\Delta}{p+\Delta} < 0$, (b) a non-union firm exits if $p < p_n^*/\lambda$ and a union firm exits if $p < p_u^*/\lambda$, and (c) a case (iii) union firm cuts employment to $\lambda p/\bar{w}$ while its wage is unchanged.*

Proof. Apply Lemmas F2, F4 and F5 with productivity λp. □

Proposition F4 (Deunionization). *The fraction of union firms that exit is $\Gamma(p_u^*/\lambda)$, which is larger than the fraction of non-union firms that exit, $\Gamma(p_n^*/\lambda)$. It is increasing in $\bar{w}$ and decreasing in λ. The post-recession union share is*

$$\psi' = \frac{\psi[1-\Gamma(p_u^*/\lambda)]}{\psi[1-\Gamma(p_u^*/\lambda)] + (1-\psi)[1-\Gamma(p_n^*/\lambda)]} < \psi$$

and does not recover.

Proof. The fractions follow from independent productivity draws with M large, and the ranking from Proposition F2. The comparative statics follow from equation (F.4). The union share follows from Bayes' rule, and it does not recover by Assumption F4. □

With a common Γ, union firms are more likely to exit only because $p_u^* > p_n^*$. If union firms also have lower productivity, both lead to more union exit. We can distinguish the two by whether union firms are more likely to exit than non-union firms with the same productivity.

After productivity recovers, $M' < M$, so $\varepsilon' = \varepsilon(M') < \varepsilon$ by Assumption F5 and $c' < c$. Union density falls to $\psi' < \psi$ by Proposition F4, and $\Delta' = \psi' T\kappa/\varepsilon'$.

F.5 HETEROGENEITY AND INEQUALITY

From Lemma F2, threatened non-union log wages are

$$\ln w^{\tau}(p) = \kappa_0 + \frac{1}{1+\varepsilon}\ln(p+\Delta), \qquad \kappa_0 = \frac{1}{1+\varepsilon}\ln\frac{c}{A} \tag{F.5}$$

where κ_0 is common to all firms.

Proposition F5 (Fan-out). *For threatened firms with $p_L < p_H$*

$$\frac{\partial[\ln w_L - \ln w_H]}{\partial\varepsilon} = \frac{\ln(p_H+\Delta) - \ln(p_L+\Delta)}{(1+\varepsilon)^2} > 0, \qquad \frac{\partial[\ln w_L - \ln w_H]}{\partial\Delta} = \frac{1}{1+\varepsilon}\left[\frac{1}{p_L+\Delta} - \frac{1}{p_H+\Delta}\right] > 0$$

with $\partial\Delta/\partial\psi = T\kappa/\varepsilon > 0$. A fall in ε and a fall in ψ holding ε fixed both widen the wage gap between low- and high-productivity firms.

Proof. Differentiate equation (F.5). Since $\ln w_L - \ln w_H < 0$, a positive derivative means that a fall in the parameter widens the gap. □

Note. Δ falls with ψ but rises when ε falls. In the first order condition in Lemma F2, the incentive to raise wages to hire more workers, $\varepsilon p/w$, scales with ε while the threat incentive, $\psi T\kappa/w$, does not. The net change in Δ is only negative if $\psi'/\psi < \varepsilon'/\varepsilon$, so the model only signs the effect of ψ on Δ holding ε fixed.

Lemma F6 (The threat premium rises when competition weakens). *Holding ψ fixed, the threat premium $P(\varepsilon) = \frac{1}{1+\varepsilon}\ln(1+\Delta/p)$ satisfies*

$$\frac{dP}{d\varepsilon} = -\frac{\ln(1+\Delta/p)}{(1+\varepsilon)^2} - \frac{1}{1+\varepsilon}\cdot\frac{1}{p+\Delta}\cdot\frac{\psi T\kappa}{\varepsilon^2} < 0$$

Proof. Differentiate using $d\Delta/d\varepsilon = -\psi T\kappa/\varepsilon^2$. □

Proposition F6 (Composition). *Let s denote the employment share of workers earning exactly $\bar{w}$ (union cases (ii) and (iii) and non-union firms matching the union scale) and G_n the distribution of log wages among all other workers. Then $G = s\cdot\mathbb{1}[w \geq \ln\bar{w}] + (1-s)G_n$ and, with $\bar{w}$ fixed,*

$$G' - G = (1-s)[G_n' - G_n] + (s-s')[G_n' - \mathbb{1}[w \geq \ln\bar{w}]] \tag{F.6}$$

Proof. Add and subtract $(1-s)G_n'$ in $G' - G$. □

The first term in equation (F.6) is the change in the non-union wage distribution (Proposition F5) holding the union share fixed. The second term moves a share $s - s'$ of workers

from the union wage $\bar{w}$ into G'_n. Displaced union workers come from firms with $p < p^*_u/\lambda$. Those re-employed at a surviving non-union firm earn equation (F.5), which is increasing in the firm's productivity, so those absorbed by low-productivity firms enter the bottom of G'_n. Those not re-employed leave the labor force and are not in G'. The model does not determine which surviving firms hire displaced workers.

F.6 DISCUSSION OF ASSUMPTIONS

The labor supply elasticity and the number of firms

Assumption F5 follows from two standard models. First, in Burdett and Mortensen (1998), a firm paying w loses workers to better outside offers, and the elasticity of its labor supply is increasing in the arrival rate of offers relative to the job destruction rate (Manning, 2003). If offers arrive in proportion to the number of active employers, exit lowers the arrival rate and the elasticity. Second, in Berger, Herkenhoff and Mongey (2022), workers choose among firms within a local market with elasticity η and across markets with elasticity $\theta_m < \eta$, and firms take their market share s_j into account. The firm's elasticity is $\varepsilon_j = (1 - s_j)\eta + s_j\theta_m$, which in a symmetric market is $(1-1/M)\eta+\theta_m/M$ and is increasing in M. However, with heterogeneous firms, the elasticity of large firms falls the most when M falls. Therefore, this delivers the fall in the level of ε in Assumption F5 but not the common ε we use for the fan-out.

The threat of unionization

Union density enters Assumption F3 through the scale of the cost distribution, which corresponds to the union density term in Freeman and Medoff (1981). Any Π that is increasing in ψ, decreasing in w, and zero at $\bar{w}$ gives a premium that is positive and decreasing in p, and the log form makes it an additive shift in productivity. The threat of unionization is consistent with evidence that non-union wages fall when right-to-work laws pass, which lowers the credibility of organizing without changing union density or union wages (Farber, 2005; Fortin, Lemieux and Lloyd, 2022). In contrast, the labor supply spillover from union to non-union employers runs in the opposite direction (Neumark and Wachter, 1995). Taschereau-Dumouchel (2020) develops a general equilibrium model of the threat where workers organize through a majority vote.