NBER WORKING PAPER SERIES

UNIONS AND INEQUALITY OVER THE TWENTIETH CENTURY: NEW EVIDENCE FROM SURVEY DATA

Henry S. Farber
Daniel Herbst
Ilyana Kuziemko
Suresh Naidu

Working Paper 24587
http://www.nber.org/papers/w24587

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2018, Revised April 2021

We thank our research assistants Obaid Haque, Chitra Marti, Brendan Moore, Tamsin Kantor, AmyWickett, and Jon Zytnick and especially Fabiola Alba, Divyansh Devnani, Elisa Jacome, Elena Marchetti-Bowick, Amitis Oskoui, Paola Gabriela Villa Paro, Ahna Pearson, Shreya Tandon, and Maryam Rostoum. We have benefited from comments by seminar participants at Berkeley, Columbia, Georgetown, Harvard, INSEAD, SOLE, the NBER Development of the American Economy, Income Distribution and Macroeconomics, and Labor Studies meetings, McGill University, Princeton, Rutgers, Sciences Po, UMass Amherst, UC Davis, Universitat Pompeu Fabra, Stanford, and Vanderbilt. We are indebted to Devin Caughey and Eric Schickler for answering questions on the early Gallup data. We thank John Bakija, Gillian Brunet, Bill Collins, Angus Deaton, Arindrajit Dube, Barry Eidlin, Nicole Fortin, John Grigsby, Ethan Kaplan, Thomas Lemieux, Gregory Niemesh, John Schmitt, Stefanie Stantcheva, Bill Spriggs, and Gabriel Zucman for data and comments. All remaining errors are our own. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Unions and Inequality Over the Twentieth Century: New Evidence from Survey Data
Henry S. Farber, Daniel Herbst, Ilyana Kuziemko, and Suresh Naidu
NBER Working Paper No. 24587
May 2018, Revised April 2021
JEL No. J51,N32

ABSTRACT

U.S. income inequality has varied inversely with union density over the past hundred years. But moving beyond this aggregate relationship has proven difficult, in part because of limited microdata on union membership prior to 1973.We develop a new source of microdata on union membership dating back to 1936, survey data primarily from Gallup (N ≈ 980,000), to examine the long-run relationship between unions and inequality. We document dramatic changes in the demographics of union members: when density was at its mid-century peak, union households were much less educated and more non-white than other households, whereas pre-World-War-II and today they are more similar to non-union households on these dimensions. However, despite large changes in composition and density since 1936, the household union premium holds relatively steady between ten and twenty log points. We then use our data to examine the effect of unions on income inequality. Using distributional decompositions, time-series regressions, state-year regressions, as well as a new instrumental-variable strategy based on the 1935 legalization of unions and the World-War- II era War Labor Board, we find consistent evidence that unions reduce inequality, explaining a significant share of the dramatic fall in inequality between the mid-1930s and late 1940s.

Henry S. Farber
Industrial Relations Section
Simpson International Building
Princeton University
Princeton, NJ 08544-2098
and NBER
farber@princeton.edu

Daniel Herbst
Department of Economics
Eller College of Management
University of Arizona
Tucson, AZ 85721
dherbst@arizona.edu

Ilyana Kuziemko
Department of Economics
Princeton University
239 J.R. Rabinowitz Building
Princeton, NJ 08544
and NBER
kuziemko@princeton.edu

Suresh Naidu
Columbia University
420 West 118th Street
New York, NY 10027
and NBER
sn2430@columbia.edu

I. Introduction

Understanding the determinants of the U-shaped pattern of U.S. income inequality over the twentieth century has become a central goal among economists over the past few decades. Over the past one-hundred years, measures of inequality have moved inversely with union density (Figure I), and many scholars have posited a causal relationship between the two trends. But especially in the historical period, moving beyond this aggregate relationship toward more demanding tests of the causal effect of unions on inequality has proven difficult due to data limitations. While aggregate measures of union density date back to the early twentieth century, it is not until the Current Population Survey (CPS) introduces a question about union membership in 1973 that labor economists have had a consistent source of microdata that includes union status. Put differently, it is not until unions are in steady decline that they can be studied with representative U.S. microdata.

In this paper we bring a new source of household-level data to the study of unions and inequality. While the Census Bureau did not ask about union membership until the 1973 CPS, public opinion polls regularly asked about household union membership, together with extensive questions on demographics, socio-economic status and political views. We harmonize these surveys, primarily Gallup public opinion polls, going back to 1936. Our new dataset draws from over 500 surveys over the period from 1936-1986 and has over 980,000 observations, each providing union status at the household level. We combine these data with more familiar microdata sources (e.g., the CPS) to extend the analysis into the present day.

We use these new data to document a number of novel results consistent with a causal impact of unions on inequality. We begin by documenting the pattern of selection into unions from 1936 onward. We document a U-shape with respect to the education of union members. Before World War II and in recent decades, the education levels of non-union households and union households are similar. However, during peak-density years (1940s through 1960s), union households were substantially less educated than other households. During these peak-density years, union households were also more likely to be non-white than either before or after.

Second, we find that union households have 10-20% higher family income than non-union households, controlling for standard determinants of wages, and that these returns are higher for non-white and less-educated workers. Interestingly, the magnitude of the union premium and its patterns of heterogeneity by education and race remain relatively constant over our long sample period, despite the large swings in density and composition of union members that we document. Third, residual income inequality is lower for union households than non-union, consistent with Freeman (1980).

These first three results—that unions during their peak drew in disadvantaged groups such as the less-educated and non-white households; that over our full sample period they confer a large family-income premia, especially for disadvantaged groups, and their compression of residual income inequality—are consistent with unions' reducing inequality and that the high levels of union density at mid-century may help explain that era's low levels of inequality. Our remaining results focus directly on measures of inequality as the outcome of interest. First, following DiNardo, Fortin, and Lemieux (1996), we conduct a reweighting exercise, where we measure inequality of a counterfactual income distribution where all union households are paid their predicted *non-union* income, we find that the rise in unionization explains over one-fourth in the 1936-1968 decline in the Gini coefficient and, conversely, its decline explains over one-tenth of the rise in the Gini coefficient after 1968.

But these microeconomic estimates do not account for any effects of union density on the wages of non-union workers, and as such may underestimate the effect of unions on inequality. As an upper bound on the macroeconomic effect of unions on inequality, we follow and extend Katz and Murphy (1992) and Goldin and Katz (2008), regressing measures of inequality on skill-shares and union density over the 20th century. For a more conservative estimate, we take advantage of the fact that our microdata has state identifiers and regress state-year union density on inequality, controlling for state and year fixed effects. Both of these exercises yield robust negative correlations of union density with a variety of measures of income inequality.

Finally, we develop an instrumental-variables strategy that allows us to examine the effects of the sharp increase in union density in the 1930s through 1940s. We use the legalization of union organizing (via the 1935 Wagner Act and the 5-4 Supreme Court decision upholding its constitutionality in 1937) and the establishment of the National War Labor Board, which promoted unionization in establishments receiving defense contracts during World War II, as two large, negative shocks to the cost of union organizing. Both of these national policies have differential effects across states due to pre-existing factors such as industry mix. We show that both these policy shocks permanently increase state-level union density and reduce state-level measures of inequality, with only transitory effects on labor demand such as industry mix. Importantly, states that experienced these policy shocks do not exhibit increases in density or decreases in inequality outside of the treatment period. In particular, we show that other episodes of war-related defense production that did not explicitly promote union organization (e.g., mobilization during the Korean War) did not increase density nor reduce inequality. While the LATE we estimate with the

Wagner and World-War-II related shocks is specific to the mid-century institutional environment, it is consistent with unions playing a causal role in reducing inequality during this key period.

These results contribute to the long-running "market forces versus institutions" debate on the causes of inequality, particularly the determinants of the mid-century "Great Compression." Of course, most economists agree that market forces and institutions both play important roles in shaping the income and wage distributions, so the debate is more a question of emphasis. A key advantage of the "market forces" side of the debate is its grounding in a competitive model focusing on the supply and demand for skilled workers, which offers hypotheses on the joint movement of relative wages and relative quantities. Given the increase in relative college wages since the 1960s, authors in this tradition (with a long pedigree stretching back to Douglas (1930), Tinbergen (1970), and Freeman (1976)) have focused on changes in demand resulting from technology (Katz and Murphy, 1992; Autor, 2014; Card and Lemieux, 2001; Katz and Autor, 1999; Autor, Katz, and Kearney, 2008; Autor, Goldin, and Katz, 2020) interacting with the rate of schooling increases. Adaptations of the relative skill model to account for recent patterns in wage inequality include Beaudry, Green, and Sand (2016), Acemoglu and Autor (2011), Autor, Levy, and Murnane (2003), and Deming (2017).

On the institutions side, the literature includes Bound and Johnson (1992), DiNardo, Fortin, and Lemieux (1996) and Lee (1999), with recent literature incorporating firms as important determinants of inequality (Song et al., 2015; Autor et al., 2020; Card, Heining, and Kline, 2013). Authors in this tradition have highlighted the potential role for unions in reducing inequality (Card, 2001; DiNardo, Fortin, and Lemieux, 1996; Western and Rosenfeld, 2011). Two recent contributions are especially relevant to our study of unions and inequality at mid-century. Callaway and Collins (2018) uses detailed microdata from a survey of six cities in 1951 to estimate a union premium comparable in magnitude to what we find during the same period. Another recent paper, Collins and Niemesh (2019), emphasizes the role of unions in the Great Compression. They use the industry measures of union density constructed by Troy (1965) and form proxies of union density using 1940 IPUMS industry allocations within state economic areas. Both this paper and our analysis in Section 5 suggest that unions played a large role in reducing inequality at mid-century. We build on Collins and Niemesh (2019) by providing direct measures of household union membership at the annual level over this period.

The remainder of the paper is organized as follows. In Section II , we describe our data sources, in particular the Gallup data. This section also presents our new time-series on household union membership. Section III analyzes selection into unions, focusing on education and race. Section IV estimates household union income premiums over much of the twentieth century, and Section V presents our evidence on the effect of unions on the shape of the overall income distribution. Section VI offers concluding thoughts and directions for future work. All appendix material referred to in the text can be found in the online appendix.

II. Household union status, 1936 to present

In this section, we briefly describe how we combine Gallup and other historical microdata sources with more modern data to create a measure of *household union status* going back to the 1930s.

II.A. Gallup data

Since 1937, Gallup has often asked respondents whether anyone in the household is a member of a labor union. This question not only allows us to plot household union density over a nine-decade period, as we do in this section, but also to examine the types of households that had union members and whether union membership conferred a family-income premium, as we do in subsequent sections. Before beginning this analysis, we highlight a few key points about the Gallup and other historical data sources that we use. A far more complete treatment can be found in Appendix B.[1]

Before the 1950s when it adopts more modern sampling techniques to reach a more representative population, Gallup data suffers from several important sampling biases that tend to over-sample the better-off. First, George Gallup sought to sample *voters*, meaning under-sampling the South (which had low turnout even among whites) and in particular Southern blacks (who were almost completely disenfranchised). Further, the focus on voters resulted in over-sampling of the educated (due to their higher turnout). Second, survey-takers in these early years were given only vague instructions (e.g., “get a good spread" for age) and often found it more pleasant working in nicer areas, further oversampling the well-off. Even after 1950, these biases remain, though become smaller. We compare the (unweighted) Gallup data to decennial Census data in each decade in Appendix Tables B.1 and B.2.

As we are interested in the full U.S. population, we seek to correct these sampling biases to the extent possible. We weight the Gallup data to match Census $region$ ×

1. Much of the information summarized here and presented in more detail in Appendix B comes from Berinsky (2006).

race cells before 1942 and $region \times race \times education$ cells from 1942 (when Gallup adds its education question) onward. Moreover, in Appendix D, we show that all of our key results are robust to various weighting schemes, including not weighting at all.

As we can only compare Gallup to the Census every ten years, we also seek some annual measures to check Gallup's reliability at higher frequencies. In Appendix Figure A.1, we show that our Gallup unemployment measure matches in changes (and often in levels) that of the official Historical Statistics of the United States (HSUS) from the 1930s onward, picking up the high unemployment of the "Roosevelt Recession" period. As another test of whether Gallup can pick up high-frequency changes in population demographics, Appendix Figure A.2 shows the "missing men" during World War II deployment: the average age of men increases nearly three years, as millions of young men were sent overseas and no longer available for Gallup to interview.

Beyond sampling, Gallup's standard union membership survey question deserves mention, as it differs from that used in the most widely used modern economic survey data, the CPS. Gallup typically asks whether you *or your spouse* is a member of a union, so we cannot consistently extract individual-level union membership as one could in the CPS.[2] In Appendix D, we compare our key results whenever possible using individual instead of household union measures—while occasionally levels shift, the changes over time are remarkably similar.

II.B. Additional Data Sources

While we rely heavily on the Gallup data, we supplement Gallup with a number of additional survey data sources from the 1930s onward. Gallup does not ask family income for much of the 1950s, but the American National Election Survey (ANES) asks both family income and union household status throughout that period, so we augment our Gallup data with the ANES in much of our analysis.[3]

We have found one survey that includes a union question that pre-dates our Gallup data. This 1935-36 survey was conducted by the Bureau of Home Economics

2. In some but not all cases they will then ask who (the respondent or the spouse) but to be consistent across as many surveys as possible, we create a harmonized *household* union variable.

3. The ANES has a relatively small sample size in any given year so that our ability to use the ANES to provide detailed breakdowns of union status and income by geography or demographics is limited.

(BHS) and Bureau of Labor Statistics (BLS) to measure household demographics, income, and expenditures across a broad range of U.S. households, and we will henceforth refer to it as the 1936 Expenditure Survey. The survey asks about *union dues* as an expenditure category, which is how we measure household union membership. Rather than sampling randomly from the whole population, the agencies chose respondents from 257 cities, towns, and rural counties within six geographic regions. In most communities, the sample was limited to native-white families with both a husband and wife, though blacks were sampled the Southeast and blacks a single individuals in some major Northern cities.[4] To mitigate the effects of this selective sampling on our estimates, we employ the same cell-weighting strategy as we do in our Gallup sample.

We further supplement our sample with a 1946 survey performed by the U.S. Psychological Corporation that includes state identifiers, family income, union status and standard demographics.[5] In 1947 and 1950 we use data from National Opinion Research Corporation (NORC) as a check on our union density estimates from Gallup, but, as these data do not have state identifiers, we do not use them in our regression analysis. We also use the Panel Survey of Income Dynamics (PSID) for the late 1960s and early 1970s. From 1977 onward, we can use the CPS to examine household measures of union membership.[6]

Summary statistics for the CPS, ANES, and these additional data sources appear in Appendix Table B.3. In general, at least along the dimensions on which Gallup appears most suspect in its early years (share residing in the South, share white, education level), these data sources appear more representative. The table shows all data sources unweighted, though we will use ANES and CPS weights in years they are provided, to follow past literature. We weight the 1936 Expenditure survey and the 1946 U.S. Psychological Corporation survey in the same manner that we do

4. Black families were included in New York City, Columbus, OH, and the Southeast, and single individuals were included in Providence, RI, Columbus, OH, Portland, OR, and Chicago, IL. Note that Hausman (2016) uses these data in studying the effects of the 1936 Veteran's Bonus.

5. The Psychological Corporation survey was a public opinion survey conducted in April 1946, in 125 cities with 5,000 respondents (plus an additional rural sample). See Link (1946) for a description of the survey and cross-tabulations.

6. Beginning in 1977, the CPS includes both the union-membership question *and* individual state-of-residence identifiers. As most of our analysis conditions on state of residence, we generally do not use CPS data from 1973–1976, which has the union variable but only identifies twelve of the most populous states plus DC, and groups the rest into ten state groups.

Gallup.

II.C. *The union share of households over time*

Figure II plots our weighted Gallup-based measure of the union share of households, by year, alongside several other series (Appendix Figure D.1 shows that the weighted and unweighted Gallup measures are very similar). The Gallup series bounces around between eleven and fifteen percent from 1937 to 1940. Between 1941 and 1945, the years the U.S. is involved in World War II, the household union-membership rate in our Gallup data roughly doubles. The union share of households continues to grow at a slower pace in the years immediately after the war, before enjoying a second spurt to reach its peak in the early 1950s. After that point, the union share of households in the Gallup data slowly but steadily declines.

Also presented in Figure II are our supplemental survey-based series. Note that each of these series generally has fewer observations per year than Gallup. The ANES sits very close to Gallup, though as expected is noisier. The 1936 expenditure survey is very close to our earliest Gallup observation, in 1937. The U.S. Psychological Corporation appears substantially lower than our Gallup measures in 1946, whereas the two NORC surveys (from 1947 and 1950) are very close to the Gallup estimates for those years.

To avoid clutter and to focus on the earlier data, we end our series in the 1980s and do not plot our CPS series in this figure, instead plotting the official CPS/BLS individual worker series, divided by the number of households, in blue for comparison. Appendix Figure A.3 shows the Gallup and CPS household-level series from 1970 until today, allowing readers to more easily assess their degree of concordance during their period of overlap (1977-1986). Reassuringly, in the years when Gallup and the CPS overlap, they are quite close.[7] As we emphasized in Section II.A, our measure of union density is based on whether a *household* has a union member, as the Gallup data do not always allow us to examine respondent-level membership. Appendix Figure D.2 shows how our household notion of density compares to the more traditional *individual* measure of density within the ANES and CPS, where both measures can be computed. The household measure is always above the individual measure, as we would expect. But in both datasets, the household and individual measures track each other in changes quite closely.

7. Given the labor-intensity of reading in the Gallup data, we do not continue past 1986 and beyond this point rely on the CPS. We cut off at 1986 in order to have a ten-year period where Gallup and CPS overlap, which allows us to check consistency of Gallup over a substantial period of time.

II.D. *Comparison to historical aggregate series*

Finally, Figure II plots two widely-used historical *aggregate* data series, the BLS series (based on union self-reports of membership) and the Troy series (compiled by Leo Troy for the NBER and based on union's self-reported revenue data).[8] While the Gallup measures do not always agree with the BLS and Troy series in levels, they are, for the most part, highly consistent in changes. We describe these existing historical data sources in greater detail in Appendix E, summarizing key points below.

The density measures based on existing historical aggregate sources are everywhere above our microdata-based series until the 1950s, at which point they converge. As we document in Appendix E, labor historians believe the union self-reports of their own membership (which the BLS series uses) are significantly biased upwards. Especially from 1937-1955, when organized labor in the US was split into two warring factions—the American Federation of Labor and the Congress of Industrial Organization—the two federations over-stated their membership in attempts to gain advantages over the other. Membership inflation became such an issue that the federations themselves did not know their own membership. The CIO felt the need to commission a 1942 internal investigation into membership inflation, privately concluding that its official membership tally was inflated by a factor of two.

Leo Troy was aware of the membership inflation issue, and thus where possible bases estimates on dues revenue (from which he can back out membership using dues formulae). But as we discuss in Appendix E, revenue reports are missing for much of the early CIO, and the same incentives likely led unions to inflate dues revenue as well.

That respondents polled by Gallup did not share these incentives to overstate union membership is an advantage of our data. However, there is an important reason why Gallup and other opinion surveys may *understate* true union membership: individuals can be in unions without knowing it, especially during certain historical moments. As we discuss in greater detail in Section V.D, during World War II, the government gave unions the authority to default-enroll workers when they started a job at any firm receiving war-related defense contracts and to automatically deduct dues payments from their paychecks. Thus, some workers during this period of rapid growth in density may not have known they were members and thus

8. These series give aggregate union *counts* of membership, so we divide by estimates of total U.S. households (geometrically interpolated between Census years) to make the numbers as comparable as possible to Gallup. This transformation will obviously overstate the union share of households if many households had multiple union members.

answered Gallup survey enumerators honestly (though incorrectly) that they were not in a union. It is not surprising that the Gallup data most undershoots the Troy and BLS numbers during the war years. Similarly, moments of high unemployment complicated calculations of union density. Until Congress mandated annual reporting in 1959, unions had great discretion in how to count a union member who became unemployed, whereas an unemployed respondent in Gallup, no longer paying his union dues, might honestly consider himself no longer a member.[9] Indeed, Figure II shows that Gallup shows essentially no net growth between 1937-1940, which includes the period after the upholding of the NLRA, but also includes the Roosevelt Recession, whereas the BLS and Troy show robust growth.[10]

In summary, while the microdata-based versions of household union density we develop and the more widely used measures based on aggregate data differ slightly in levels (in a manner consistent with their non-trivial differences in methodology), they in almost all years firmly agree in changes. Like the Troy and BLS series, the Gallup data exhibit the same inverted U-shape over the twentieth century. Moreover, as we will show in Section V, the relationship between aggregate union density and inequality is very similar whether we use our new, microdata-based measures of household unionization rates or the traditional, aggregate measures.[11]

An important advantage of our series, however, is that it is based on microdata, which allow us to examine *who* joined unions and how this selection changed over time. It is to this task we now turn.

9. As noted, Gallup and ANES did *not* skip over the unemployed or those otherwise out of the labor force when fielding their union question, and many unemployed and retired respondents in these surveys nonetheless identify as union members.

10. Indeed, it is well documented that at least among the largest locals where data are available, *dues payments* plummeted for CIO unions during the 1938 recession, as millions of workers were laid off (Lichtenstein, 2003). We speculate that unions continued to report these laid-off workers as members.

11. Of course, it is possible that Gallup's non-representative sampling contributes to the gap between it and the BLS and Troy series. We suspect non-random sampling is not an important factor. First, the sampling biases with respect to calculating average density go in both directions (e.g., Gallup's oversampling the well-off creates negative bias but undersampling the union-hostile South creates positive bias). Second, as noted, the weighted and unweighted versions of the Gallup union density series are very similar (see Appendix Figure D.1).

III. Selection Into Unions

Labor economists have long debated the nature of selection into unions. We focus on selection into unions by education and then by race. Less-educated and non-white households have on average lower income than other households, and thus selection along these margins into unions reveals whether or not unions historically excluded or included the relative less advantaged. Besides being of independent interest, the nature of selection into unions is an indirect test about whether union density was causally related to the Great Compression: if union members were, say, more educated and whiter than non-union members in mid-century, it would be difficult to argue that the increased union density was exercising equalizing pressure.

While we focus on selection on observables, there is likely selection on unobservables that bias our results. These unobserved traits could include uncredentialled trade skills or raw ability. Lewis (1986) wrote "I have strong priors on the direction of the bias....the Micro, OLS, and CS wage gap estimates are biased upward—the omitted quality variables are positively correlated with union status." Abowd and Farber (1982) and Farber (1983) enriched the model of selection into unions to include selection by union employers from among the pool of workers who would like a union job. They argue that, because unions confer a larger wage advantage to the less skilled, the the marginal cost of skill to union employers is lower than for nonunion employers. The result is that most skilled will not want a union job, and employers will want to hire the most highly skilled from among those workers who do desire a union job. Thus, low observed skill workers will be positively selected into union jobs by employers based on their unobservables and high observed skill workers will be negatively selected into union jobs by workers based on their unobservables. This two-sided selection results in the union sector being composed of the center of the (observed plus unobserved to the econometrician) skill distribution for a particular job. Card (1996) presents evidence consistent with this two-sided view of selection, and argues that the resulting biases cancel each other out resulting in a relatively unbiased cross-sectional union premium.

III.A. Selection into unions by education

We begin our analysis of who joined unions by estimating the following equation, separately by survey-source d (e.g., Gallup, ANES, CPS) and year y:

$$Union_{hst} = \beta_{dy} Educ_h^R + \gamma_1 Female_h^R + f(age_h^R) + \mu_s + \nu_t + e_{hst}. \tag{1}$$

In this equation, subscripts h, s, and t denote household, state and survey-date, respectively (our Gallup data provides many surveys per year, so survey date t will map to some unique y and survey-date fixed effects subsume year fixed effects). The superscript R serves to remind readers that in many cases, a variable refers specifically to the *respondent* (not necessarily the household head). $Union_{hst}$ is an indicator for whether anyone in the household is a union member (and is the underlying household-level variable we use to construct the aggregate time-series in the previous section). $Educ_h^R$ is the respondent's education in years.[12] $Female_h^R$ is a female dummy, $f(age_h^R)$ is a function of age of the respondent (age and its square when respondent's age is recorded in years, fixed effects for each category when it is recorded in categories), and μ_s and ν_t are vectors of state and survey-date fixed effects, respectively.

The vector of estimated β_{dy} values tells us, for a given year y and using data from a given survey source d, how own years of schooling predicts whether you live in a union household, conditional on basic demographics and state of residence.[13] Note here that we are not yet controlling for race.

Figure III shows these results across our key datasets. A clear U-shape emerges, with the year-specific point-estimates remarkably consistent across all data sources.[14] In the earliest years (1936 through approximately 1943) the coefficients suggest that an additional year of education reduces the likelihood of living in a union household by only two to three percentage points. At the trough of the U (around 1960), we estimate that an additional year of education reduces the likelihood of living in a union household by roughly *five* percentage points. Since the 1960s, the negative marginal effect of education on the probability of living in a union household declines steadily: it reaches zero around 2000 and is now positive and in some years statistically significant, though small.

The differential increase in education among union households in recent decades may reflect, in part, the substantial growth of relatively highly-educated public sec-

12. Where a specific survey does not collect information directly on years of schooling but reports specific ranges or credentials, we use simple rules to convert these measures to years of schooling. The note to Figure III describes how we impute years of schooling in these cases.

13. For the ANES, given the small sample sizes, we constrain the coefficients on education (β_{dy}) to be equal across six-year bins in order to reduce sampling error. For the Gallup and other surveys, we estimate the coefficients on education (β_{dy}) by estimating separate regressions for each $survey\,source \times year$ combination.

14. This pattern holds when other education measures are used instead of years of schooling. Appendix Figures A.4, A.5, and A.6 show similar patterns when, respectively, a high-school dummy, college dummy and log years schooling serve as the education measure.

tor labor unions since the 1960s. Indeed, as we show in Appendix Figure A.7, before President Kennedy's 1962 executive order giving federal employees the right to organize, the share of union members in the public sector was nearly negligible, hovering around five percent, while today one in every two union members works in the public sector.[15] While we do not know sector for the Gallup, Psych. Corp., and 1936 expenditure surveys, we can compare our baseline selection patterns from the ANES and CPS to those when we drop any household with a public sector worker. As Appendix Figure A.8 shows, while the *levels* of the selection effect change slightly for this sample, the *increase* in the education of union households from 1970 onward is unchanged. While we do not have data from before 1950, any effect of public-sector unions is likely to be tiny, as both the public sector workforce was smaller and public-sector unions were essentially nonexistent.

Another possible explanation for the relative up-skilling of union households is the steep decline since the 1960s in the share of union members in manufacturing employment—also depicted in Appendix Figure A.7. The manufacturing share of union members is the rough inverse of the public-sector share, falling from nearly fifty percent in the 1950s to less than ten percent today. Appendix Figure A.8 also shows the education selection patterns after dropping households with *either* a public-sector or a manufacturing worker. A large majority of the up-skilling effect remains.[16] We return to this pattern in the conclusion when we discuss questions for future work.

As noted in Section II, we use a household and not an individual concept of union membership. In the discussion above, we have implicitly assumed that the selection patterns over time reflect less-educated workers joining unions in the middle decades of the 1900s, but in principle they could instead reflect changes in marriage patterns whereby union members, for whatever reason, became more likely to marry less-educated spouses during this period.

We address this concern in two ways. First, we reproduce the selection-by-education analysis (Figure III) after excluding observations where the respondent is female. In this sample we do *not* rely on the education of the spouse as a proxy for the education of the likely union member. Appendix Figure D.4 shows that selection into unions by years of schooling for the male-only sample yields the same U-shape as we saw with the full sample. Second, in the CPS era, we can directly compare results using

15. Over the period from 1973-2016, tabulation of CPS data indicates that 5.3 percent of college graduates employed in the private sector were members of labor unions. In contrast, fully 39.7 percent of college graduates employed in the public sector are union members.

16. These results use our standard weights as described in Section II and B, but Appendix Table D.1 shows robustness to other weighting schemes, including not weighting.

the household- and individual-based union membership concept. While we can only examine more recent years with our CPS data, both the individual and household selection series (plotted in Appendix Figure D.3) show the same marked increase in terms of selection by years of schooling from the 1970s until today.

All of this evidence suggests that union members were substantially less educated than non-members until quite recently and especially so in the 1950s and 1960s. While "skill" is multi-dimensional and has unobserved components, so long as unobserved dimensions of skill correlate with education, then the historical data from mid-century challenges Lewis' conjecture that "omitted quality variables are positively correlated with union status."

III.B. Selection into unions by race

We next examine selection by race, which is important for at least two reasons. First, given that school *quality* is an often unobserved dimension of skill (Card and Krueger, 1992) and blacks have always attended lower-quality schools than whites, race may serve as another proxy for skill and thus further inform the selection evidence in the previous subsection. Second, selection of union members by race over time is an important (and unresolved) historical question. Historians disagree on the degree to which unions discriminated against black workers over the twentieth century (Ashenfelter, 1972, Northrup, 1971; Foner, 1976; King Jr, 1986; Katznelson, 2013).

We analyze selection by race in the same manner as selection by years of schooling and simply replace $Educ_h^R$ with $White_h^R$ in equation (1).[17] The estimated coefficients on $White$ across time and data sources are presented in Figure IV. Again, a U-shape emerges, though it is noisier than that in the selection-by-education analysis. In the beginning of our sample period, whites are (conditional on our covariates) more likely to be in union households than non-whites. This advantage diminishes during the war years and continues to grow more negative until about the 1960s. While noisy, at this point, whites are about ten percentage points *less* likely to be in a union household than are other respondents. Since then, whites gain on non-white households and the differential attenuates toward zero as we reach the modern day.

While not quite as consistent as for education, selection by race again agrees for the most part across data sources. There is some disagreement between Gallup and CPS, whereby Gallup shows minimal selection with respect to race by the early

17. Results are essentially exactly the inverse when instead of $White$ we use a black dummy. We use $White$ instead because sometimes Gallup uses "negro" and sometimes "non-white" and thus $White$ would appear, in principle, a more stable marker.

1980s, whereas CPS shows that whites are still somewhat less likely to live in union households. However, by the end of the sample period, there is no remaining selection by race in the CPS either. As we noted in the previous Section, Gallup's sampling of the South changes over time, so in Appendix Figure A.9 we replicate the analysis dropping all observations from the South, finding very similar results.

We believe it is an important contribution to show that, at least with respect to *membership*, blacks were not underrepresented in unions throughout most of the twentieth century after conditioning on state of residence. But this result must be viewed in context. First, controlling for state in Figure IV means we partial out the massive under-representation of unions in the South, where blacks disproportionately lived at mid-century. There are many reasons why the Jim-Crow-era South was difficult to organize (e.g., less industrial employment), but the extreme hostility of white elites to unionization of black workers was certainly one of them (Friedman, 2000).

Second, outside of the South, part of the over-representation of blacks in unions is merely a byproduct of unions organizing lower-skilled areas of the economy, which were disproportionately non-white. Appendix Figure A.10 shows that controlling for years of schooling reduces the negative effect of the *White* coefficient in most years, though the basic U-shape remains.[18]

Third, membership rates alone do not fully capture non-white workers' experience in unions. While the mid-century leaders of the industrial unions of the CIO committed themselves publicly to policies of racial equality (Schickler, 2016), leadership roles remained overwhelmingly white, and U.S. labor history is littered with ugly examples of the white rank-and-file walking off the job in reaction to integration. By the early 1960s, over 100 locals of AFL-CIO unions (mostly in the South) remained explicitly segregated (Minchin, 2017). The 1964 Civil Rights Act led to large unions, even ones with Black leaders such as the UAW, being sued for discriminatory practices under Title VII. The AFL-CIO did not have a black officer until 2007.

Nonetheless, at mid-century, unions were organizing groups that were disproportionately non-white. Moreover, during most of the twentieth century the *non*-unionized sector practiced *de facto* or *de jure* racial discrimination, a topic we explore in the next section when we examine the union premium and in particular the premium by race.

18. For completeness, we also show (in Appendix Figure A.11) that the pattern of selection by education we see in Figure III barely changes if we simultaneously control for race.

IV. THE UNION FAMILY INCOME PREMIUM OVER THE TWENTIETH CENTURY

Estimating the union premium—the wage differential between union and otherwise-similar non-unions workers—is at the core of the modern empirical neoclassical approach towards measuring the effect of labor unions, pioneered by Lewis (1963). The early analysis by Lewis generally focused on industry-level differences, as consistent sources of microdata were not yet available. Freeman and Medoff (1984) were among the first to use CPS microdata to estimate determinants of union membership and the union premium with individual-level data. They find a union premium of roughly sixteen percent, averaging across studies in the 1970s. In general, a ten to twenty log-point union premium—controlling for Mincer-type covariates and estimated on cross-sectional wage data such as the CPS—has been found consistently in the literature. As noted in the introduction and in the Lewis (1986) review of the literature, there is almost no microdata-based estimates of the union premium prior to the 1968 PSID.[19]

A key challenge in this literature is separating any causal effect of union membership on wages from non-random selection into unions. On the one hand, if higher union wages create excess demand for union jobs, then union-sector employers have their pick of queueing workers and unobserved skill could be higher in the union sector, overstating the union premium. On the other hand, a higher union wage premium for less-skilled workers and union protections against firing might differentially attract workers with unobservably less skill and motivation. Naturally, researchers have turned to panel-data estimation to address this selection bias, though Freeman (1984) and Lewis (1986) warn about attenuation bias due to misreported union status, which fixed-effects regressions exacerbate. Card (1996) uses CPS ORG data to examine workers as they switch between the union and non-union sectors (using the 1977 CPS linkage to employer data to correct for measurement error), showing that the union premium remains significant even after accounting for negative selection at the top and positive selection at the bottom.[20]

19. While cross-sectional estimates of the union premium go back at least to the 1960s (see Johnson (1975) for a summary of research from that period), many are based on ecological regressions (e.g. Rosen (1970)) between union density and average wages at the industry or occupation (often not labor market) level. These macro estimates are summarized and critiqued in Lewis (1983). The one pre-PSID exception to our knowledge is Stafford (1968) who estimates a union premium of 16% in the 1966 Survey of Consumer Finance.

20. Lemieux (1998) performs a similar exercise using Canadian data, with the added advantage that he can focus on *involuntary* switchers. He finds estimates that are in fact quite close to OLS estimates of the union premium. Other scholars (e.g., Raphael, 2000 and Kulka-

IV.A. Baseline results

To construct a union premium series back to 1936, we use all the datasets employed in the selection analysis so long as they contain family income, which excludes most of Gallup data from the 1940s and 1950s. We also drop surveys with severe income top-coding (which we defined as more than 30 percent of observations in the top category), which results in losing some Gallup data from the 1970s.

Across all these surveys, we estimate the following regression equation separately by data source d and year y:

$$\begin{aligned} ln(y_{hst}) = &\beta_{dy} Union_h + \gamma_1 Female_h^R + \gamma_2 Race_h^R + f(age_h^R) + \\ &g(Employed_h) + \lambda_h^{eduR} + \nu_t + \mu_s + e_{hst}. \end{aligned} \tag{2}$$

While we are estimating a *household income* function, we do our best to mimic classic Mincerian controls. In the above equation, y_{hst} is household income of household h from survey date t in state s; $Union_h$ is an indicator for whether anyone in the household is a union member; $Female_h^R$ and $Race_h^R$ are, respectively, indicators for gender and fixed effects for racial categories *of the respondent*; $f(age_h^R)$ is a function of age of the respondent (age and its square when respondent's age is recorded in years, fixed effects for each category when it is recorded in categories); $g(Employed_h)$ is a flexible function controlling for the number of workers in the household; λ_h^{educR} is a vector of fixed effects for the educational attainment of the respondent; and μ_s and ν_t are vectors of state and survey-date fixed effects, respectively. Note that for the 1946 U.S. Psychological Corporation and for the Gallup surveys from 1961 onward, we cannot control for the number of workers per household, but we show later that this bias should be small.

As with our selection results in the previous section, Figure V shows our union premium results separately by survey source and year. While not a perfectly flat

rni and Hirsch, 2019) have used the Displaced Workers Survey (which records many *involuntary* separations thus lessening concerns about endogenous switching and is known to have limited mis-measurement of union status) to estimate worker-level panel regressions, again finding premiums close to cross-sectional OLS estimates (about 15 percent). Jakubson (1991) estimates longitudinal union premia in the PSID, getting estimates of around 5-8%, but does not account for measurement error. De Chaisemartin and d'Haultfoeuille (2020) show that once heterogeneous treatment effects are allowed for, it is difficult to find evidence of a fixed-effects union premium in the NLSY and show significant pre-trends in earnings.

line, the premium holds relatively stable. Of the more than sixty point estimates we report, only a handful are greater than 0.20 or less than 0.10. Not a single estimate has a confidence interval intersecting zero. Given the standard errors around each estimate, the family union premium does not appear to follow any discernible pattern over time.[21]

While the majority of our estimates are from cross-sectional data, there is a unique three-wave panel survey of the ANES (1956, 1958 and 1960) that allows us to estimate household union premium controlling for respondent fixed effects. The union premium estimated in this specification is almost identical to the cross-sectional estimate from the ANES in the same period, and statistically significant at the five-percent level despite a small sample. We provide more details and specifications in Appendix Table A.2. To our knowledge, this analysis yields the earliest panel-based estimate of the union premium, at least from U.S. data.a

Card (2001), using CPS data, noted as a puzzle that the union wage premium was surprisingly stable between 1973 and 1993, even as private-sector union density declined by half. Our results, if anything, deepen this puzzle, as we show that the premium remains somewhere between ten and twenty log points over a *nine-decade* period that saw density (as well as the degree of negative selection by skill) both increase and then decrease.[22] We have no clear resolution of this puzzle and indeed find it hard to write down a model of collective bargaining outcomes with standard union and firm objective functions that yields a steady premium in the face of increasing then declining density. One simple explanation is that the union premium is bounded below by some minimum, say five percent, below which workers will not pay dues and attend meetings. It may also be bounded above by some amount of product market (or other input market) competition on the firm side.[23] We flag this question and the testing of this hypothesis as a potentially fruitful area for future research.

21. In Appendix Table A.3 we check for heterogeneity by macroeconomic conditions, as in Blanchflower and Bryson (2004), but find little.

22. While the unions literature is mostly empirical, the few theory papers on unions that do exist do not help rationalize the surprising pattern of declining density alongside steady premiums. Existing models in which SBTC determines union density rates predict that the premium should dwindle as density declines. This result is also hard to rationalize with models that assume a union objective function that is a positive function of both union wages and membership, such as Dinlersoz and Greenwood (2016).

23. Rios-Avila and Hirsch (2014) offer this explanation for the steady nature of the union premium, between ten and twenty points, across time and countries.

IV.B. Robustness and Related Results

As a family union premium is a departure from the more familiar individual earnings premium estimated in past papers, Appendix Table A.1 shows the coefficients on the Mincer equation covariates in equation (2), so readers can compare it to standard earnings equations. In all cases, the coefficients on the covariates have the same signs and similar magnitudes as we typically see from an individual earnings regression.

As another check on whether the household nature of our inquiry creates biases, in Appendix Figure D.5 we use the CPS to compare our premium results with (a) the traditional worker-level earnings premium, where individual earnings are regressed on individual union membership and (b) a worker-level family income premium, where family income is regressed on individual union membership. Our premium results—family income regressed on household union membership—generally fall between these two other estimates. In almost all years, they agree in changes.

In Appendix Figure A.12, we show results after controlling for occupation of the household head. As noted, occupation categories vary considerably across survey sources so our attempts to harmonize will be imperfect, which is why we relegate this figure to the Appendix. The appendix figure reports coefficients that are somewhat larger than in the main Figure V, consistent with unions differentially drawing from households where the head has a lower-paid occupation.

As noted earlier, we cannot control for the employment status of household members in the Gallup and the Psychological Corporation data. Appendix Figure A.13 shows that any bias is likely very small: in the ANES, *not* controlling for employment status increases the estimated union premium only slightly, relative to the baseline results where these controls are included.[24]

The family income premium may not fully capture changes in the household's economic well-being. Union families may benefit from other forms of compensation such as health benefits or vacation, as has been documented in the CPS-era (see Freeman, 1981 and Buchmueller, DiNardo, and Valletta, 2004 among others). Unfortunately, Gallup and our other sources do not consistently ask about benefits. One exception is from a 1949 Gallup survey that asked about paid vacation. As we show

24. Union households are more likely to have at least one person employed (likely the union member himself), which explains why controlling for household employment has a (slight) negative effect on the estimated union household premia. However, living with a union member is a *negative* predictor of own employment (results available upon request), which likely accounts for the fact that controlling for total number of workers in the household has only a small effect on the estimated premium.

in Appendix Table A.4, Gallup respondents in union households are over twenty percentage points (about forty percent) more likely to report receiving paid vacation as a benefit.

On the other hand, the union premium may also reflect compensating differentials for workplace dis-amenities, which would suggest that our estimated premia are *overstating* the differential well-being of union households. Some evidence against this claim comes from another Gallup survey in 1939 that asks respondents how easily they could find a job "as good" as their current one. As we show in Appendix Table A.5, union households are significantly more likely to say it would be hard for them to find a job just as good. Similar to the union premium, this tendency is similar to that in the modern day (the same table shows these results using the 1977-2018 GSS). To the extent respondents considered non-wage job characteristics (safety, working conditions, benefits, etc.) this result is an additional piece of evidence that union members, even in the early days of the labor movement, felt their jobs were better—in a broad sense—than non-union members.

Our estimates of a sizable union premium contrast with recent papers using regression discontinuities in close NLRB representation elections to estimate the causal effect of unionization on firm-level outcomes (DiNardo and Lee, 2004; Lee and Mas, 2012; Frandsen, 2020). These papers have found little evidence of positive union wage premia, although some have found effects on non-wage benefits such as pensions (Knepper, 2020). What explains the discrepancy? A possibility is that the LATE identified by the RD papers is not informative about the average treatment effect of unions. Importantly, most existing union workplaces were organized earlier and most elections are not very close. It is reasonable that a clear (sizeable) union victory in an election reflects workers' expectations of substantial advantage while a very close election reflects workers' expectations of more limited advantage. As such, the LATE identified by the RD papers is likely not informative (and likely understates) the average advantage of unionization. We do not mean to imply that we have identified the true average causal effect of unions on wages, but neither is it the case that the small effects found in the close-election RD analyses are appropriate when applied broadly.

IV.C. Heterogeneous Union Household Income Effects

We have so far assumed that unions confer the same family income premium regardless of the characteristics of the respondent. We now explore heterogeneity by years of schooling and race.

We begin by augmenting our family income equation (2) by adding an interaction term between years of schooling and the household union dummy. Figure VI presents the coefficient on this interaction term, as usual, separately by survey-source and year. The results are consistent throughout the period and show that less-educated households enjoyed a larger union family income premium. Over the nine decades of our sample period, this differential effect appears relatively stable. For each additional year of education, the household union premium declines by roughly four log points.

The analogous results from adding $White_h^R \times Union_h$ to equation (2) instead of $Yearsofeduc_h^R \times Union_h$ are shown in Figure VII. The interactions are not statistically significant in the earliest surveys (the 1936 BLS Expenditure Survey and the 1942 Gallup Survey), though their signs suggest that white workers enjoyed *larger* premiums. However, in the 1946 Psychological Corporation survey and in succeeding Gallup, ANES and CPS surveys, there is consistent evidence of a larger union family income premium for nonwhites over the next five decades. This racial differential in the union effect on household income has declined somewhat since the 1990s and in the most recent CPS data it cannot be distinguished from zero.

We saw in our selection analysis that some of the disproportionate membership of non-white households was merely driven by disproportionate membership of the less-educated, so we check whether the differential premium to non-whites is similarly explained. In Appendix Figure A.14 we reproduce the analysis in Figure VII but include $Yearsofeduc_h^R \times Union_h$ in all regressions.[25] The results barely change, suggesting that even for households with the same level of education, black households enjoyed higher union premiums. Of course, the union premium equation is only identified by comparing family income for unionized versus non-unionized households, so this result does not mean that non-white union workers were paid more than white union workers, just that the white pay advantage was significantly smaller in the union sector. Returning to our discussion at the end of Section III, this result suggests that despite the many ways that the U.S. labor movement discriminated against non-whites, such discrimination appeared worse in the non-organized sector.

Our conclusion from the heterogeneity analysis is that, at least for most of our sample period, disadvantaged households (i.e., those with respondents who are non-white or less educated) are those most benefited (in terms of family income) by having a household member in a union. Ignoring this differential effect would tend to

25. For completeness, we also reproduce the heterogeneity by years of schooling analysis in Figure VI after adding $White_h \times Union_h$ interaction. The results barely change (see Appendix Figure A.15).

underestimate the effect of unions on inequality, especially from 1940–1990, when the differential premium for black households appears largest. We return to this point in Section V.D.

IV.D. Effects on Residual Income Dispersion

An influential view of unions is that they lower the return paid not only to observed skill, as we document above, but also to unobserved skill. Supporting this view is the fact that, at least in the CPS era, the union wage distribution is compressed even after conditioning on observable measures of human capital (e.g., Freeman and Medoff, 1984 and Card, 2001).

We implement an analogous analysis at the household level to determine if unions performed a similar function in earlier decades. Separately for union and non-union households, we regress log family income on all the covariates (except union) in equation (2). As before, we perform this analysis separately by survey-source and year. We then calculate residuals for each sector and compute the ratio of variances between the union and non-union residuals (which has an F-distribution with degrees of freedom given by the two sample sizes, allowing us to construct confidence intervals). If unions compress the distribution of unobserved skill, then this ratio should be less than one.

Figure VIII shows, over our sample period, the ratio of variance of residual log family income between the union and non-union sector, together with 95% confidence intervals. The ratio is uniformly below one, and often below 0.5, with confidence intervals that always exclude equality of the variances. Like the union premium estimates, there does not seem to be a strong pattern over time in the union-nonunion difference in residual income inequality. Instead, it appears that the CPS-era pattern of unions compressing residual inequality holds in a very similar manner throughout the post-1936 period.[26]

V. THE EFFECT OF UNIONS ON INEQUALITY

Empirically, we have so far documented that, in their effect on household income, unions have exhibited remarkable stability over the past eighty years. During our long sample period, the union premium has remained between ten and twenty log

26. For example, Card (2001) estimates a union-non-union variance ratio of around 0.61 in 1973 using individual male earnings, very similar to what we find in the 1970s for household income.

points, with the less-educated receiving an especially large premium. Moreover, the negative effect of unions on residual income variance is large and also relatively stable over time. By contrast, selection into unions varies considerably. From the 1940s to 1960s, when unions were at their peak and inequality at its nadir, *disadvantaged* households were much more likely to be union members than either before or since. These results support, at least indirectly, the hypothesis that unions compress the income distribution.

In this section, we explore in a more direct manner the relationship between unions and income inequality, joining an extensive empirical literature examining how unions shape the income distribution. It is helpful to separate this literature into two conceptual categories. First, assume that unions affect the wages of *only* their members and that estimates of the union premium can recover this causal effect, putting aside selection and spillover issues discussed earlier. Then, simple variance decompositions can estimate the counterfactual no-union income distribution and thus the effect of unions on inequality. For example, so long as unions draw from the bottom part of the counterfactual non-union wage distribution, then their conferring a union premium to this otherwise low-earning group reduces inequality. Moreover, residual wage inequality also appears to be lower among union workers, suggesting that unions reduce inequality with respect to unobservable traits as well (Card, 2001). DiNardo, Fortin, and Lemieux (1996) and Firpo, Fortin, and Lemieux (2009) take this approach and find that unions substantially reduce wage inequality, especially for men.

A second category of papers argues that unions affect *non-union* workers as well (so-called "unions spillover effects"). Unions can raise non-union wages via union "threat" effects (Farber, 2005; Taschereau-Dumouchel, 2020;Fortin, Lemieux, and Lloyd, 2018) or by the setting of wage standards throughout an industry (Western and Rosenfeld, 2011). Conversely, unions can lower non-union wages by creating surplus labor supply for uncovered firms (Lewis, 1963). Unions might also affect the compensation of *management* (Pischke, DiNardo, and Hallock, 2000; Frydman and Saks, 2010) and the returns to capital (Abowd, 1989; Lee and Mas, 2012; Dinardo and Hallock, 2002), thus reducing inequality by lowering compensation in the right tail of the income distribution. Finally, as an organized lobby for redistributive taxes and regulation, unions might affect the income distribution via political-economy mechanisms (Leighley and Nagler, 2007; Acemoglu and Robinson, 2013).

In this section, we add several new results to this literature. First, and most directly related to the results in the previous two sections, we conduct distributional decompositions following DiNardo, Fortin, and Lemieux (1996), where we show how measures of inequality change with the level and composition of union membership. This exercise jointly accounts for where union households are in the income distribution as well as the effect of union membership on a household's position in the income distribution. The identifying assumptions are as follows: first, that, conditional on our controls, union membership is not otherwise correlated with determinants of income and, second, that union membership affects *only* the income of union households (i.e., no "spillovers" to other workers or households). We show robustness to weakened versions of these assumptions, in particular showing evidence of spillovers using extensions to the reweighting methodology proposed by Fortin, Lemieux, and Lloyd (2018).

Second, we turn to more aggregate analysis. We follow some of the canonical work on the effect of skills shares on the college premium, adding union density to these standard, aggregate, time-series estimations. Note here that aggregate analysis does not rule out spillovers, but instead rests on the (strong) identifying assumption that conditional on our time-series controls, union density is exogenous. Next, we use the state identifiers in the Gallup data to conduct a parallel analysis at the state-year level. Finally, we leverage the historical cross-state variation in union density generated by the Wagner Act and World War II to obtain instrumental variables estimates of the effect of union density on inequality.

V.A. Distributional Decompositions

In this section we present the historical impact of unions on inequality using distributional decompositions, following DiNardo, Fortin, and Lemieux, 1996 (henceforth DFL). First, we compare observed inequality in each year to what inequality would look like without any union members. The difference provides a measure of unions' impact on inequality within a given year. Second, we use differences in this measure across key years in our data to identify the total contribution of unions to changes in inequality over time. In other words, we estimate how much of the fall and rise in inequality can be explained by unions.

Both of these exercises require estimating a counterfactual income distribution that would have existed had selection into unions been different than what was observed. Assuming union membership is conditionally independent of household income, we can simulate this counterfactual using reweighting procedures. Specifically, we will construct "deunionized" counterfactuals in each year by reweighting the non-union population so that their distribution of observables matches that of the general population.[27]

27. While the DFL methodology is by now standard, we provide a more complete review of

In our first exercise, we consider the income distribution under the counterfactual that nobody joins a union and compare it to the unweighted income distribution in each year. The top panel of Figure IX plots differences in Gini coefficients for true and reweighted populations over time, $\text{Gini}\left(F_{Y_t}\right) - \text{Gini}(F_{Y_t^{C_0}})$. Unsurprisingly, this within-year impact of unions tracks both the pattern of union density and negative selection into unions documented earlier. During the period of peak union density, unions reduced the Gini coefficient by 0.025 relative to the non-unionized counterfactual. More surprising is that even though union members are positively selected on education today, unions still exert a small equalizing force, suggesting that the within-union compression effect still dominates the union-non-union difference.

The bottom panel of Figure IX shows differences in log income percentiles between true and deunionized counterfactual distributions for the three years where we have continuous income data (1936 consumption survey, PSID, and CPS). In 1936 and 2014, the differences in these distributions are small, but in 1968 there is a large compressing effect of unions. We show the densities themselves in Appendix Figure F.1. In addition to true and deunionized density plots, the bottom panel of Figure IX shows dashed lines corresponding to a deunionized counterfactual that also accounts for potential spillover effects of unions. We construct these spillover-adjusted distributions following Fortin, Lemieux, and Lloyd (2018), who augment the standard DFL reweighting procedure to allow for labor-market-level union density effects on the household income distribution. This procedure consists of predicting wage distributions (flexibly using an ordered logit) for non-union workers as a function of labor-market level union density, and then imposing the counterfactual zero union density to obtain a non-union income distribution purged of union spillover effects.[28]

The time series and percentile plots tell a similar story: unions had a small impact on overall income inequality during the pre-war and modern eras, when density was low, but significantly compressed income inequality during the period in-between, when density was high. How much of the absolute change in inequality can we attribute to this differential impact from unions? To answer this question,

DFL reweighting methods in Appendix F.

28. Specifically, spillover-adjustment weights are constructed to remove the predicted impact of state-year-industry (in CPS) or state-year (in 1968 PSID) union density throughout the income distribution. Predictions are formed from an ordered probit of non-union household income against state-year-industry (in CPS) or state-year (in 1968 PSID) union densities. These labor market densities are only directly available in the CPS and PSID, and hence dashed lines are omitted for 1936, although we present results with predicted state-year shares (along with additional details) in Appendix F.

we decompose the absolute change in inequality into its "total union effect," the difference between observed changes in inequality and the change in inequality that would have occurred in the absence of unions. For the time period t_B to t, this total union effect is computed as the difference in within-year union effects,

$$\Delta^U = \left[\text{Gini}(F_{Y_t}) - \text{Gini}(F_{Y_{t_B}})\right] - \left[\text{Gini}(F_{Y_t^{C_0}}) - \text{Gini}(F_{Y_{t_B}^{C_0}})\right] \tag{3}$$

$$= \left[\text{Gini}(F_{Y_t}) - \text{Gini}(F_{Y_t^{C_0}})\right] - \left[\text{Gini}(F_{Y_{t_B}}) - \text{Gini}(F_{Y_{t_B}^{C_0}})\right]. \tag{4}$$

Table I reports the total union effect over different periods. The contribution of unions to the change in household inequality between 1936 and 1968 is considerable, with unions explaining 23% of the change in the Gini, 46% of the change in the 90/10, 18% of the change in the 90/50, and 80% of the change in the 10/50 (note that these are ratios of household income, not individual earnings). The contribution of unions to the change in household inequality since 1968 is smaller but not insignificant, with unions explaining about 10% of the increase in the gini, and between 12-18 percent of the change in the percentile ratios. With respect to skill premia, unions explain roughly 17% of the fall in the college premium between 1936 and 1968, but around 80% of the increase between 1968 and 2014.

In the left columns of Table I, we further decompose the total union effect into the portion attributable to changes in union membership (a "unionization effect") and the portion attributable to changes in union wages (a "union wage effect"). Note, however, that estimating these subcomponents requires predicting union membership in one year using estimates of union selection from another, which comes with considerable caveats in our mixed-dataset setting.[29]

In sum, the pure "micro" effect of the union-density *growth* on household inequality from 1936 to 1968 is considerable, even without accounting for spillovers, and typically larger than the effect of union-density *decline* on the recent rise in inequality. Further, even during periods of positively selected union members and low density, such as 1936 and today, unions are still an equalizing force, although nowhere as quantitatively important as during the period of peak union density, where union density was high and union members considerably less educated than non-union members.

29. Details on our detailed decomposition into unionization and union wage effects is provided in Appendix F.

V.B. Time-series Regressions

While the distributional decompositions capture the effect of union density on household income inequality, they require a strong assumption that there are no spillovers, threat effects, or political economy mechanisms that alter wages for non-union workers. The plausibility of these more macro mechanisms warrants an aggregate analysis, complementing the individual household regressions estimated above. Further, our household survey data is binned and misses inequality across individuals, as well as inequality at the bottom and the top of the distribution, which can be addressed with more standard inequality measures constructed from other sources.

Our aggregate analysis of the effect of unions on inequality is motivated by the literature on the college wage premium. Following Katz and Murphy (1992) as well as Goldin and Margo (1992) and using a mix of data from the Decennial Census, the CPS and a 1915 survey from Iowa, Goldin and Katz (2008) show that the evolution of the college premium between 1915 and 2005 is well-explained by the relative supply of college workers, controlling for flexible functions of time. Autor, Katz, and Kearney, 2008 confirm this analysis using data from the CPS in the 1963-2005 period and adding more covariates.

The analysis in this section (and the next) attempt to "horse race" institutional and market forces in ecological regressions over time (and across states), following a literature that has attempted to disentangle these two forces across countries (Blau and Kahn, 1996; Jaumotte and Osorio Buitron, 2020), albeit with limited identifying variation.

We begin by simply adding union density to the specifications estimated in these papers:

$$\log\left(\frac{wage_t^{Col}}{wage_t^{HS}}\right) = \beta UnionDensity_t + \gamma log\left(\frac{N_t^{Col}}{N_t^{HS}}\right) + f(t) + \lambda X_t + \epsilon_t. \tag{5}$$

The dependent variable is the log college wage premium, which we specify as a function of the supply of skilled workers, $log(N_t^{Col}/N_t^{HS})$, a polynomial in time, $f(t)$, other time-series controls X_t, which we vary to probe robustness, and, importantly, $UnionDensity_t$.[30]

30. As we do not have a strong view regarding whether, at the aggregate level, our Gallup-based estimate of early union density is better than the traditional BLS estimate, we take a simple average of the two, dividing the BLS estimate of union membership from Freeman et al. (1998) by total population for closer comparability.

We choose time-series controls X_t both to follow past literature as well as to capture the most obvious confounds in estimating the effect of unions on inequality. Specifically, following Autor, Katz, and Kearney (2008) we include the real value of the federal minimum wage and the civilian unemployment rate and following Piketty, Saez, and Stantcheva (2014) we include the top marginal tax rate in the federal individual income tax schedule. As unions historically push for full employment, higher minimum wages and higher top tax rates, these might be "bad controls" and their inclusion would understate the full effect of union density on inequality. We adjust for heteroskedasticity and AR(1) serial correlation in the error ϵ_t using Newey-West standard errors.[31]

The first two columns of Table II show the results from this exercise. Col. (1) does not include additional controls X_t, whereas col. (2) does. The coefficient on union density is negative and highly significant (and very similar to each other in magnitude), and we discuss specific magnitudes below.

We also find a significant and negative coefficient on skill shares and in fact (despite somewhat different sample periods) recover a coefficient very close to those in Goldin and Katz (2008), Autor, Katz, and Kearney (2008), and Autor, Goldin, and Katz (2020). Interestingly, as we show in Appendix Table A.6, union density and the skill-shares measure negatively co-vary at both the annual and state-year level (though this negative covariance is small and insignificant once we condition on our usual regression controls). Thus, controlling for skill shares tends to increase the significance of union density, and vice versa. This point is important because going forward we will sometimes use noisy measures of skill share (e.g., interpolations between Census years), but as skill shares and density both tend to decrease inequality and negatively covary, noisy measurement of this control variable should generally yield conservative coefficient estimates on density.

While the canonical analysis in Goldin and Katz (2008) and related work focuses on the college premium, we extend our analysis in Table II by using the same specifications as in cols. (1) and (2) but using other measures of inequality as outcomes. Cols. (3) - (4) of Table II are identical to Cols. (1) - (2) except that the 90/10 log wage ratio for men (also taken from the IPUMS Census and CPS) is used as the outcome

31. These regressions can be seen as following Katz and Autor (1999), who decompose group-level wages into their "latent competitive wage" (i.e., relative skill shares and technological trends, augmented with measures of institutions, such as union density). However, we do not model group level density as having group-level effects, as in Card and Lemieux (2001), who put relative union shares (college union density divided by HS grad union density) as a regressor in the relative wage equation; rather, we consider overall density as affecting the relative wage.

variable. The results are quite similar, with union density again having a negative and significant association with inequality that is robust to adding our vector of controls. Cols. (5) - (8) examine the 90/50 and the 10/50 ratios, showing that the effect we find on the 90/10 comes from the *bottom* half of the distribution, as the coefficients on density, while negative, are insignificant for the 90/50.

The rest of Table II examines annual data.[32] These additional years not only give us more observations, but also allow us to use inter-Census variation (e.g., during World War II). Cols. (9) and (10) use the Gini coefficient constructed by Kopczuk, Saez, and Song (2010) from Social Security data. The next two columns use the top-ten-percent income share from Piketty, Saez, and Zucman (2018).[33] The final two columns use the labor share of national income from Piketty, Saez, and Zucman (2018). For all three of these outcomes, the union density coefficient suggests a significant decrease in inequality (a negative coefficient for the Gini and top-ten share, and a positive one for labor share), robust to controls.

Appendix Tables A.7 and A.8 shows a series of robustness tests for each of the outcomes in Table II. We show results are robust to using the Gallup series alone or the BLS series alone to calculate $UnionDensity_t$ (instead of averaging the two together) and to substituting either a quartic or a quadratic for the cubic time polynomial. They also report more of the coefficients, which we suppress in the main tables in the interest of space.

Our estimate magnitudes are generally sensible yet economically significant. Table II implies that a ten percentage point increase in union density results in a 12-15 percent fall in the college premium, 2-1.7 percent falls in 90-10 wage ratios for men, small and insignificant effects on 90-50 male wage ratio, and 1.5 to 1.8 percent increase in the 50-10 wage ratio. We further find that the same size increase leads to a 0.016 to 0.014 decrease in the Gini, roughly 3% of the mean, and 2.3 to 3.5 percentage points in the top ten share and 4.5-4.8 percentage points in the labor share. These are large effects, and we view them as an upper bound on the true effects of

32. As noted earlier, a small complication in using these annual outcomes is that our pre-CPS estimates of the skill shares $log(N_t^{Col}/N_t^{HS})$ in equation (5) come from the Census and thus in principle are only available every ten years. To circumvent this issue, we include two separate education controls: (i) skill shares as measured (annually) in our Gallup data and an annual measure of skill shares equal to that from the CPS when it is available; and (ii) interpolating between Census years in the earlier period. In this sense, we treat education as a nuisance variable and simply try to control flexibly for it, allowing us to continue to estimate the conditional effect of union density.

33. Results are qualitatively similar, with smaller coefficients, if we instead use the top 10 share from Piketty and Saez (2003).

unions on inequality, and inclusive of a variety of economic and non-economic mechanisms by which unions could reduce inequality (e.g., both direct effects on wage and income distributions, but also indirect effects via politics, norms, and policies).

The magnitudes implied by the time-series analysis are clearly larger than those implied by the micro-effect of unions on union members, even including the spillover effects. There are clear limitations to the time-series analysis—perhaps most obviously, concerns about endogeneity of union density and suspect inference due to small samples. Moreover, unlike the analysis of skill shares in Goldin and Katz (2008) and similar papers, the inclusion of union density is not theoretically motivated.

To examine the role of spillovers more rigorously, we draw on the counterfactual distributions we estimated in the previous section. In Appendix F we use the difference between the actual Gini (constructed here from our survey data, not the SSA data) and the DFL counterfactual Gini coefficient from Section V.A as an outcome in the time-series regression, again controlling for skill shares and time polynomials. The coefficient on union density in this regression isolates the effect of union density on inequality that is solely due to the effect of unions on the incomes of union households. This could be called the pure "micro" effect of unions. The effect here is roughly between -0.04 and -.06, so that a 10 percentage point increase in union density reduces the Gini via the micro effect by roughly 0.005 points. But the effect of union density on the overall Gini itself is -0.3, where a 10 percentage point increase in density reduces the Gini by 0.03. This table suggests much of the effect of unions on inequality would be through the effects on non-union workers, but there are good reasons to think our selection equation is mis-specified (no controls for industry or occupation, for example, which Appendix Figure A.12 suggests increases the union premium) and use of binned income data implies we are underestimating the micro-effect of unions on inequality.

In the next section, we take an intermediate position on the scope of spillovers and the endogeneity of union density by estimating similar aggregate regressions at the state-year level, which allows a much richer set of controls, including state and year fixed effects.

V.C. State-Year Panel Regressions

While the time-series analysis generates summary accounts of the aggregate association of unions on the U.S. economy, a major limitation are the many unobserved factors (e.g., technology, macroeconomic policy, trade, outsourcing, industry structure) potentially correlated with both inequality and union density and not necessarily absorbed by our controls. In this section we replicate the analysis at the state-year level, controlling for state and year fixed effects, which can absorb a considerable amount of unobserved heterogeneity.

The Gallup data always contain state identifiers, so we can construct continuous state-year measures of union density throughout the pre-CPS period, something that was not possible with previous data.[34] Although we do not attempt to isolate exogenous variation in union density in this section, we can determine whether the inverse inequality-density relationship in the aggregate time series also holds at the state-year level, conditional on year and state fixed effects.[35] Importantly, as all states have access to the same national technology, the vector of year fixed effects in this design controls for simple variants of SBTC that affect all states the same way.

We combine our Gallup state-year measures with household state-year measures calculated from the CPS. We take a weighted average of Gallup-generated state-year union densities and CPS-generated state-year union densities, with weights proportional to the number of observations in each sample (so the CPS gets a much larger weight). This procedure results in a panel of annual state-year union density measures going back to 1937. Note that such a high-frequency panel was impossible to construct before the Gallup data, as in most years the BLS and Troy series did not break down their aggregate counts geographically, and when they did, it was generally only for a few years (Troy) or by coarse regions (BLS).

To examine the effect of unions on inequality, we closely follow equation (5) and estimate specifications of the form:

$$y_{st} = \beta UnionDensity_{st} + \gamma log\left(\frac{N_{st}^{Col}}{N_{st}^{HS}}\right) + \lambda X_{st} + \mu_{t,r(s)} + \delta_s + \epsilon_{st} \tag{6}$$

where y_{st} is a measure of inequality, for example the college-HS wage gap or the percent of total income accruing to the top ten percent, in state s and year t. A contribution of our paper that we use in this analysis (as well as in the next subsection) is the construction of historical state-year measures of the labor share of net income, following Piketty, Saez, and Zucman (2018). We present details and validation in Appendix H.

34. Troy (1965) presents state breakdowns for 1939 and 1956, and Hirsch, Macpherson, and Vroman (2001) use BLS reports to construct state-year measures of density from 1964 onwards.

35. Similar regressions are estimated at the cross-country level by Jaumotte and Osorio Buitron (2020), though their sample period of 1980-2010 is far shorter than ours.

As before, we control for skill-shares $\log(N^{Col}_{st}/N^{HS}_{st})$ in all specifications.[36] We include state fixed effects (δ_s) and a vector of year fixed effects that allow each year to have a different effect for the South ($\mu_{t,r(s)}$). Note that we include South-by-year fixed effects because, as discussed earlier, Gallup's sampling of the South improves over time and we want to flexibly control for this evolution. We cluster the standard errors at the state level.

As before, we will show results with and without X_{st}, a vector of additional state-year controls. We try our best to capture the same covariates as in equation (5), though in some cases controls that are available at the annual level in the historical period are not available at the state-year level. To control for economic expansions and contractions, we include in X_{st} state-year log income per capita and state-year measures of the share of households subject to the federal income tax. We include these measures as proxies for relative local economic prosperity, as annual state-level unemployment rates are not consistently available until the 1963 CPS. We include top marginal income tax rates by state, and to more fully capture the political-economy climate, we also control for a Democratic governor indicator variable as well as a state-year level "policy liberalism" index developed by Caughey and Warshaw (2016).[37] Manufacturing moving from the unionized Northeast and Midwest to the South and West is often cited as a reason for the decline in density, so we include in X_{st} the one-digit industry employment shares at the state-year level.

Because our Gallup sample size will become small for less populous states, our coefficients may be attenuated due to finite-sample bias in our state-year level union density measures. To address this concern, we use a "split-sample" IV strategy.[38] For each state-year, we split the Gallup observations into two random samples s_0 and s_1, and use the union density calculated from s_1 to instrument the union density calculated from s_0. This procedures yields the following first-stage equation:

36. The top-ten-percent and labor shares of income are available at the *annual* level, so just as in the time-series regressions we include the interpolated IPUMS-CPS education measure (at the state-year level) as well as the Gallup measure of education for these outcomes (at the state-year level).

37. We are indebted to Jon Bakija, Stefanie Stantcheva and John Grigsby for facilitating our access to the state-level income tax data.

38. See Angrist and Krueger (1995) for an early description of the methodology. Inoue and Solon (2010) and Aydemir and Borjas (2011) provide further exposition and applications, respectively.

$$(7)\quad UnionDensity^0_{st} = \eta UnionDensity^1_{st} - \iota \log\left(\frac{N^{Col}_{st}}{N^{HS}_{st}}\right) + \lambda^f X_{st} + \mu_{t\times South} + \delta_s + \nu_{st}.$$

The second-stage equation in the split-sample IV is merely equation (6) with $UnionDensity_{st}$ replaced by $\widehat{UnionDensity}^0_{st}$, the prediction generated from the first stage. Since $UnionDensity^1$ and $UnionDensity^0$ are calculated from a *random* split of the data, the sampling errors in the two measures will be orthogonal. Omitted-variable bias aside, if the only issue is measurement error, the IV estimator β^{IV} will yield a consistent, unattenuated estimate of β. We repeat this procedure 200 times and report bootstrapped estimates and standard errors, clustered by state.

Table III shows results from the specification in equation (6) across the state-year analogues of the inequality outcomes used in Table II. As in the previous subsection, the odd-numbered columns do not include the additional controls X_{st}, while the even-numbered columns do.

Cols. (1) and (2) show results when the college premium is the outcome variable. The coefficient on state-year union density is negative and significant, and the magnitude barely changes whether or not additional controls are included. Indeed, across the male percentile ratios and the Gini coefficient (cols. 3 to 10), the coefficient on state-year density is consistently signed, significant and quite robust to adding additional controls.

We now turn to regressions where state-year measures of top-ten and labor share of income are the outcomes. The first two columns for the top-ten share (cols. 11 and 12) and labor share (cols. 14 and 15) are analogous to all of the earlier outcomes and show a significant, robust negative (positive) coefficient when top-ten (labor) share is the outcome (though the point-estimate for the labor-share regressions is somewhat more sensitive to controls than our other outcomes). Unlike the earlier outcomes, which rely on Census income data and thus cannot extend earlier than 1940, these outcomes allow us to go back further in time, which we do in the third column for each outcome (cols. 13 and 16). Not only can we extend back to 1937 using Gallup density data, but we can also use the 1929 Handbook of American Trade Unions to develop a measure of state-level union density for 1929.[39] While

39. This measure is based on the distribution of union locals across states in 1929. Cohen, Malloy, and Nguyen (2016) construct a similar measure and validate it for a number of states. We provide more details on its construction in Appendix C. The next time the Handbook is available is 1937. We already have our Gallup data from that year, so the Handbook only provides one additional year of data (i.e., 1929).

we require microdata for much of the previous analysis in the paper, in this section, we need only a state-level measure, so can include this 1929 measure. Adding 1929 is especially useful because it pre-dates the New Deal and Great Depression, two events potentially linked to both inequality and union density. Cols. (13) and (16) replicate, respectively, cols. (11) and (14) and if anything adding this additional year slightly increases the magnitudes on the density variable.[40]

In Appendix Tables A.9 and A.10 we show a variety of specifications that add intermediate sets of controls between the odd and even columns reported in Table III. Further, we deal with possible unobserved but smooth state-specific changes in technology or other unobservables that may be confounding the estimated relationship by including state-specific linear and quadratic trends. These tables also contain a set of estimates (column 1) that do not use the split-sample IV for state-year union density. These estimates verify the presence of attenuation bias, with the split-sample IV coefficients roughly fifty percent larger than the OLS coefficients.

A natural concern is that unions' compression of state-level income distributions comes at the cost of slowed economic growth (e.g., lowered net business entry or capital flight). In fact, union density shows consistently positive, but sometimes insignificant, effects on log state income per capita, and we can rule out even small negative effects of unions on state-level economic activity (see Appendix Table A.11).

While the magnitudes across the three methodologies vary, they are not implausibly far apart. We can examine the share of the "Great Compression", the fall in inequality between 1936 and 1968 explained by the 12% increase in union density between those two years. Symmetrically, we can ask how much of the increase in inequality between 1968 and 2014 is explained by the 12 percentage point fall in union density. Focusing on the Gini coefficient, Table I shows that pure "micro" changes in unionization (without any spillovers) account for 24% of the fall in the Gini between 1936 and 1968, and further can account for 10% of the increase between 1968 and 2014. The time-series results imply much larger effects, with union density accounting for 35% of the mid-century fall in the Gini, and 21% of the recent increase, while the state-year results are smaller, implying that unions account for 14-17% of the mid-century fall in inequality and between 12-15% of the recent increase. The symmetry of the fall and rise of inequality explained by the rise and fall of union density is further suggestive of a true causal effect, rather than a purely spurious correlation.

40. While not all of our controls go back to 1929, we construct skill shares in 1929 by projecting backwards educational attainment for older ages in the 1940 Census using the reported state of residence in 1935. See Appendix C for more information and validation.

V.D. *Isolating exogenous policy variation*

While quite robust, our state panel analysis so far makes no attempt to isolate plausibly exogenous variation in union density. It is not hard to conceive of plausible bias stories. On the one hand, state union density may grow because of favorable local economic or political factors that themselves reduce inequality, a bias would overstate the role of unions in reducing inequality. On the other hand, reverse causality could mask any negative effect of unions on inequality if unions tends to organize in reaction to high or growing levels of inequality.

In this final exercise, we attempt to isolate exogenous components of the variation in state-level union density, focusing on a period highlighted by Goldin and Katz (2008). They note that in the years around World War II, particularly in the 1940s, the decline in inequality "went far beyond what can be accounted for by market forces alone," and they suggest that unions played a role. As Figure I shows, almost all the rise in U.S. density takes place during two short windows of time: immediately upon the legalization of labor organizing itself (via the 1935 Wagner Act and the 5-4 Supreme Court decision upholding it in 1937) and during the massive increase in demand for U.S. industrial production during World War II, when the federal government enforces pro-union policies at firms receiving defense contracts. We construct two measures that capture the incidence of these two policy shocks across states. First, we define our *Wagner shock*$_s$ as the number of new members added via NLRB elections and large recognition strikes between 1935 and 1938 in state s. This measure isolates the increase in union density driven by worker take-up of the new federal procedures created by the Wagner act, rather than changes due to, say, local variation in the 1938 recession, selective exits of union versus non-union firms, union-friendly state governments, or unionization occurring outside the NLRA process.[41] Second, we define our *War-spending shock*$_s$ as the value of defense production contracts from 1940-1945 received by state s. Both terms are defined per capita and then standardized (mean subtracted out and then divided by standard deviation).[42]

These two events provide hope for identification because they both have the fol-

41. Note that the NLRA exempted sectors such as government, railroads, and airlines which also experienced a modest increase in union density (Troy, 1965), so this instrument is not *mechanically* correlated with all increased unionization during this period.

42. Gillezeau (2017) looks at state-year persistence in union density over time, also using Gallup data to measure union density in 1939 and 1945 along with data from Troy, and uses state-level war contracts as a cross-sectional instrument. He does not look at inequality nor does he consider a panel specification as we do.

lowing three characteristics: (1) the source of the shock was a *national* policy and thus was not driven by local economic or political factors; (2) despite being driven by the federal government, these two shocks had differential effects across states, providing geographic variation; (3) these differential effects across states do not appear to stem from endogenous variation, as outside of the period of these two national policy shocks, more intensely treated states do not trend differently with respect to union density or inequality measures. Put differently, while we do not claim that these shocks hit a *random* set of states, the pre-existing differences across states do not correlate with differential changes in density or inequality outside of the treatment period. For example, in Appendix G we show that states with larger IV values had greater strike activity since at least 1914, suggesting they indeed may have had greater *latent* demand for unions long before the Wagner Act, and we use pre-1929 strikes interacted with post-Wagner Act as an alternative instrument in the Appendix. However, we show that these strikes were generally *unsuccessful*, and only during about a ten-year window beginning in 1935 (when the federal government briefly takes a pro-union stance) does this latent demand for unions translate into actual growth of union density. We show many more results and robustness checks as well as provide additional historical context in Appendix G.[43]

While, in Appendix G, we provide extensive evidence consistent with this policy-driven variation being exogenous, we acknowledge it is difficult to conclusively rule out alternative stories given the sweeping nature of the New Deal and World War II. Similarly, the uniqueness of the period suggests extreme caution in extrapolating these results to other periods in history. For these reasons, we view these results as complementary to the results shown above, and not definitive on their own.

We begin by displaying the underlying state-level variation in simple scatter plots. We plot the 1938-1929 changes in union density and our outcome variables separately on the Wagner shock and 1947-1938 changes on the War-spending shock. Using nine-year intervals may seem odd, but it is done intentionally. It allows us to avoid the worst years of the Great Depression and our period of missing data for state-year density (1930-1936), as well as avoid any year with war-related wage controls (1942-1946), as the Depression and the wage controls likely have their own effects on inequality beyond changes in union density. Beyond the union-friendly policies that we use as identification, defense production may have also increased demand for low-skilled workers, which might itself temporarily lower inequality and

43. In a previous working-paper version of this paper, we also experimented with so-called "Right-to-Work" laws as an alternative instrumental variable, but found no sufficiently robust effect of Right-to-Work on union density.

is another reason to avoid the war years. In the IV analysis it is especially important to include 1929, as it gives us a pre-Wagner Act datapoint, so the intervals 1929-1938 for the Wagner shock and 1938-1947 for the war shock present the natural starting points to our analysis.

The first-stage relationships in sub-figure (a) of Figure X show that both IVs have a significant and positive relationship with changes in state-level union density, with or without 1930 population weights. The remaining subfigures show the reduced-form relationships between the outcome variables and each IV. Again, we see that the expected relationship holds for both outcome variables and both IVs (though the relationship between the Wagner shock and top-ten share is noisier than the other three).

In Table IV we show the results from 2SLS estimations, separately for each IV. We add region fixed effects, the change in estimated skill shares, and the change in manufacturing employment share as controls, but otherwise these regressions are estimated using the same variation depicted in the raw scatter plots. Cols. (1) and (2) suggest a negative effect of an increase in union density (as instrumented by the Wagner shock and War shock, respectively) on a state's change in top-ten share, with the latter effect quite a bit larger. With only 47 observations, our first-stage F-statistics are naturally small (marginally above and below the rule-of-thumb cut-off value of ten for the first and second shocks, respectively). We therefore report weak-instrument robust Anderson-Rubin confidence intervals at the bottom of the tables. Cols (1) and (2) show that with weak-instrument robust confidence intervals we are unable to reject a 0 effect of union density with the Wagner Act instrument, but while the war spending instrument confidence intervals are unbounded below, they do exclude 0 and are consistent with negative effects of union density on top income shares. We find similar results (cols. 3 and 4) when state labor share is the outcome.

In the final columns, we pool the two shocks and also add placebo periods (other nine-year intervals that fall after the two treatment periods, i.e., 1947-1956, 1956-1965, etc.). We thus estimate a first-stage equation that uses $\textit{Wagner shock}_s \times \mathbb{I}_t^{t=1938}$ and $\textit{War-spending shock}_s \times \mathbb{I}_t^{t=1947}$ as instruments, and then controls for the main effects of $\textit{War-spending shock}_s$ and $\textit{Wagner shock}_s$ in the second stage. This estimation serves two purposes. First, pooling the shocks and adding control periods gives us more precision, as reflected in the higher F-statistics and the bounded weak-instrument confidence intervals (based on conditional-likelihood ratios, instead of Anderson-Rubin, to adjust for multiple instruments) that exclude zero. Second, finding effects of our IV variables *outside* of the treatment period would cast doubt on our identifying assumptions. Indeed, the main effects of the $\textit{War-spending shock}_s$

and *Wagner shock*$_s$ are small and insignificant in the final two columns of the table and the F statistic on the excluded instruments is now larger. These estimations suggest that a ten percentage point increase in union density reduces the state top-ten share by 6.2 percentage points; that same increase in density would increase the labor share by 3.3 percentage points. As we are identified via two state-level shocks, and for both Michigan is the most intensely treated state, in cols. (6) and (8) we show robustness to dropping Michigan. The first-stage relationship is in fact stronger; the coefficients of interest in the second-stage become somewhat smaller in magnitude but remain highly significant.

We show myriad other robustness tests in the Appendix, which we summarize briefly here. We pay special attention to changed in industrial mix as a potential confounder, with tests that include manufacturing employment share and other related variables on both the right- and left-hand side of regressions. We treat state-level policy and political changes (e.g., minimum wage, state income tax rates, and Democratic governorships) similarly. We use the microdata to show our first-stage is not driven by ecological bias.

In the Appendix, we analyze the Korean War (1950-1953) as an important placebo event. Though a smaller engagement than World War II, the conflict involved over five million U.S. service personnel, a major industrial mobilization effort, and domestic wage and price controls to address inflation concerns. Moreover, the same states tended to enjoy defense contracts as in World War II (the correlation in defense dollars per capita is above 0.8). Importantly, however, the federal government did *not* attach pro-union conditions to firms receiving defense contracts during Korea.[44] In the Appendix, we show the analogue of Figure X for the Korean War, finding no correlation between Korean-War defense spending and changes in state union density or inequality measures.

One might naturally worry, especially for the war-spending shock, that certain aspects of war production were sticky and would have facilitated a more egalitarian wage structure even absent the rise in density. However, we show in the Appendix that the there is no lasting effect on manufacturing share of employment in more heavily treated states, so at least industry-mix stickiness appears minimal. It also seems an unlikely moment for wage structures themselves to be sticky, given the historical level of labor-market churn immediately after V-J day as well as elevated inflation—which should erode any nominal wage stickiness—over the next two years.[45] Moreover, while it is often speculated that egalitarian social norms de-

44. See Stieber (1980) on the reduced status of labor during the Korean War relative to World War II. In 1951, the CIO walked out of the Wage Stabilization Board in protest.

45. With the end of defense production, non-farm payroll contracted by two million (or 4.9

veloped during the war and endured for a period thereafter, in Appendix G we use Gallup data to show that by 1945, survey respondents said that labor had gained more than its fair share during the war years and that in fact businessmen deserved more credit for their sacrifices, hardly a moment of pro-worker sentiment.

How could unions reduce inequality so drastically in this period? First, during our treatment period, unions organized the "superstar" firms (Autor et al., 2020) of their day (e.g., General Motors, Ford, U.S. Steel, and AT&T). Appendix Figure G.5 shows the number of the four largest companies with major union contracts, both by employment and market capitalization. The increase in union coverage among the largest firms over the treatment period is far more dramatic than the overall rise in union density (as displayed in Figure II). The resulting decrease in inequality (as measured by top ten share) could well be disproportionate: for example, large firms may exercise standard-setting influence in their sectors or have, by dint of their scale, low non-supervisory labor share and high payments to shareholders and CEOs (consistent with Frydman and Molloy, 2012, who argue unionization was the primary restraint on CEO pay in this period). This explanation is also consistent with the smaller effects when Michigan is dropped, as the large auto companies based in that state were the largest employers in the country and became unionized in our treatment period.

Moreover, while we show in the section that the policy shocks have large effects on state-level density, in Appendix Figure G.6 we show that they have disproporionately large effects on non-white union membership. Thus, the LATE that our policy variables estimate come from organizing the largest employers and at the same time some of the least advantaged workers. While the absence of matched firm-worker data from this period makes it difficult to distinguish precise mechanisms, we find these results intriguing and worthy of future work.

VI. Conclusion

We leverage historical polling data, allowing us to provide a systematic, representative study of unions' effects on the income distribution over a much longer period than existing work. A combination of low-skill composition, compression, and

percent) in the single month of September 1945, a record that would stand in absolute and percentage terms until the Covid-19 layoffs in April 2020. See on contraction of the labor force in 1945. At the same time, American military personnel shrunk by more than 10 million between 1945 and 1947, drastically expanding the civilian labor supply. See Acemoglu, Autor, and Lyle (2004) on military demobilization.

a large union income premium made mid-century unions a powerful force for equalizing the income distribution. We show that unions were a major force in the Great Compression, above and beyond what can be accounted for by the direct effect of unions on union members. We leverage cross-state instruments from the two policy shocks that explain almost all the increase in 20th century union density, and find that they have large effects on inequality as measured by the labor share or the top income share, further providing evidence that unions affect moments of the income distribution beyond what can explained by their effects on union members alone.

The famous U-shape in U.S. economic inequality over the twentieth century has been the object of a large and distinguished literature adjudicating the roles of supply-and-demand of skilled labor versus changes in labor-market institutions such as unions. Our results push the body of evidence towards the conclusion that institutions can have substantial and lasting effects on the income distribution, while also confirming a significant role for relative skill supplies. We believe the large and immediate effects of the Wagner Act and War Labor Board we find are hard to attribute to more secular and slower-moving changes like skill shares, but an important question would be how the subsequent rise in education triggered by the GI Bill helped sustain these low levels of inequality.

Looking forward, recent events suggest a spurt of grass-roots organizing activity, from the Covid-related mass walkouts at Amazon distribution centers and wildcat strikes at Tyson and other meat-processing plants to the wave of teachers strikes in 2018 and 2019. The configuration of crisis and mobilization targeting the country's largest firms recalls the 1930s, though our results suggest that without legal or other institutional changes at the federal level, translating this activity into growth in union density or coverage will be difficult.

We welcome future work that develops theoretical models explaining the joint evolution of union density, skill composition, premia, and overall inequality that we have documented. More work on the effect of unions, perhaps in light of the recent literature documenting pervasive labor market power (Manning, 2020), would inform whether unions could be an important part of a feasible policy package to lower inequality.

References

Abowd, John M (1989). "The effect of wage bargains on the stock market value of the firm". *The American Economic Review*, pp. 774–800.

Abowd, John M. and Henry S. Farber (1982). "Job Queues and the Union Status of Workers". *Industrial and Labor Relations Review* 35.3, pp. 354–367.

Acemoglu, Daron and David Autor (2011). “Skills, tasks and technologies: Implications for employment and earnings”. *Handbook of labor economics*. Vol. 4. Elsevier, pp. 1043–1171.

Acemoglu, Daron, David H Autor, and David Lyle (2004). “Women, war, and wages: The effect of female labor supply on the wage structure at midcentury”. *Journal of political Economy* 112.3, pp. 497–551.

Acemoglu, Daron and James A Robinson (2013). “Economics versus politics: Pitfalls of policy advice”. *The Journal of Economic Perspectives* 27.2, pp. 173–192.

Angrist, Joshua D and Alan B Krueger (1995). “Split-sample instrumental variables estimates of the return to schooling”. *Journal of Business & Economic Statistics* 13.2, pp. 225–235.

Ashenfelter, Orley (1972). “Racial discrimination and trade unionism”. *Journal of political economy* 80.3, Part 1, pp. 435–464.

Autor, David, Claudia Goldin, and Lawrence Katz (2020). “Extending the Race between Education and Technology”. *AEA Papers and Proceedings*. Vol. 110, pp. 347–51.

Autor, David et al. (2020). “The fall of the labor share and the rise of superstar firms”. *The Quarterly Journal of Economics* 135.2, pp. 645–709.

Autor, David H (2014). “Skills, education, and the rise of earnings inequality among the “other 99 percent"”. *Science* 344.6186, pp. 843–851.

Autor, David H, Lawrence F Katz, and Melissa S Kearney (2008). “Trends in US wage inequality: Revising the revisionists”. *The Review of economics and statistics* 90.2, pp. 300–323.

Autor, David H, Frank Levy, and Richard J Murnane (2003). “The skill content of recent technological change: An empirical exploration”. *The Quarterly journal of economics* 118.4, pp. 1279–1333.

Aydemir, Abdurrahman and George J Borjas (2011). “Attenuation bias in measuring the wage impact of immigration”. *Journal of Labor Economics* 29.1, pp. 69–112.

Beaudry, Paul, David A Green, and Benjamin M Sand (2016). “The great reversal in the demand for skill and cognitive tasks”. *Journal of Labor Economics* 34.S1, S199–S247.

Berinsky, Adam J (2006). “American public opinion in the 1930s and 1940s: The analysis of quota-controlled sample survey data”. *International Journal of Public Opinion Quarterly* 70.4, pp. 499–529.

Blanchflower, David G and Alex Bryson (2004). *The Union Wage Premium in the US and the UK*. 612. Centre for Economic Performance, London School of Economics and Political Science.

Blau, Francine D and Lawrence M Kahn (1996). “International differences in male wage inequality: institutions versus market forces”. *Journal of Political Economy* 104.4, pp. 791–837.

Bound, John and George Johnson (1992). “Changes in the Structure of Wages in the 1980’s: An Evaluation of Alternative Explanations”. *American Economic Review* 82.3, pp. 371–92.

Buchmueller, Thomas C, John DiNardo, and Robert Valletta (2004). “A submerging labor market institution? Unions and the nonwage aspects of work”. *Emerging labor market institutions for the twenty-first century*. University of Chicago Press, pp. 231–264.

Callaway, Brantly and William J Collins (2018). “Unions, workers, and wages at the peak of the American labor movement”. *Explorations in Economic History* 68, pp. 95–118.

Card, David (1996). “The Effect of Unions on the Structure of Wages: A Longitudinal Analysis”. *Econometrica* 64.4, pp. 957–979.

— (2001). “The effect of unions on wage inequality in the US labor market”. *ILR Review* 54.2, pp. 296–315.

Card, David, Jörg Heining, and Patrick Kline (2013). “Workplace heterogeneity and the rise of West German wage inequality”. *The Quarterly journal of economics* 128.3, pp. 967–1015.

Card, David and Alan B Krueger (1992). “Does school quality matter? Returns to education and the characteristics of public schools in the United States”. *Journal of political Economy* 100.1, pp. 1–40.

Card, David and Thomas Lemieux (2001). “Can falling supply explain the rising return to college for younger men? A cohort-based analysis”. *The Quarterly Journal of Economics* 116.2, pp. 705–746.

Caughey, Devin and Christopher Warshaw (2016). “The dynamics of state policy liberalism, 1936–2014”. *American Journal of Political Science* 60.4, pp. 899–913.

Cohen, Lauren, Christopher J Malloy, and Quoc Nguyen (2016). “The impact of forced migration on modern cities: Evidence from 1930s crop failures”. *Available at SSRN 2767564*.

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”. *The Economic History Review* 72.2, pp. 691–715.

De Chaisemartin, Clément and Xavier d’Haultfoeuille (2020). “Two-way fixed effects estimators with heterogeneous treatment effects”. *American Economic Review* 110.9, pp. 2964–96.

Deming, David J (2017). “The growing importance of social skills in the labor market”. *The Quarterly Journal of Economics* 132.4, pp. 1593–1640.
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, pp. 1001–1044.
Dinardo, John and Kevin F Hallock (2002). “When unions “mattered”: the impact of strikes on financial markets, 1925–1937”. *ILR Review* 55.2, pp. 219–233.
DiNardo, John and David S Lee (2004). “Economic impacts of new unionization on private sector employers: 1984–2001”. *The Quarterly Journal of Economics* 119.4, pp. 1383–1441.
Dinlersoz, Emin and Jeremy Greenwood (2016). “The rise and fall of unions in the United States”. *Journal of Monetary Economics* 83, pp. 129–146.
Douglas, Paul H (1930). *Real wages in the United States, 1890-1926*. Houghton Mifflin Company, New York.
Farber, Henry (1983). “The Determination of the Union Status of Workers”. *Econometrica* 51.
Farber, Henrys (2005). “Nonunion wage rates and the threat of unionization”. *ILR Review* 58.3, pp. 335–352.
Firpo, Sergio, Nicole M Fortin, and Thomas Lemieux (2009). “Unconditional quantile regressions”. *Econometrica* 77.3, pp. 953–973.
Foner, Philip Sheldon (1976). *Organized labor and the black worker: 1619-1973*. Vol. 475. International publishers.
Fortin, Nicole, Thomas Lemieux, and Neil Lloyd (2018). *Labor Market Institutions and the Distribution of Wages: The Role of Spillover Effects*. Tech. rep. Working Paper.
Frandsen, Brigham (2020). “The surprising impacts of unionization on establishments: Accounting for selection in close union representation elections”. *Journal of Labor Economics*.
Freeman, Richard (1976). “The overeducated american”.
— (1984). “Longitudinal Analyses of the Effects of Trade Unions”. *Journal of Labor Economics* 2.1, pp. 1–26.
Freeman, Richard and James Medoff (1984). *What do unions do*.
Freeman, Richard B (1980). “Unionism and the Dispersion of Wages”. *ILR Review* 34.1, pp. 3–23.
Freeman, Richard B et al. (1998). “Spurts in Union Growth: Defining Moments and Social Processes”. *NBER Chapters*, pp. 265–296.
Freeman, Richard B. (1981). “The Effect of Unionism on Fringe Benefits”. *Industrial and Labor Relations Review* 34.4, pp. 489–509.

Friedman, Gerald (2000). "The Political Economy of Early Southern Unionism: Race, Politics, and Labor in the South, 1880-1953". *Journal of Economic History*, pp. 384–413.

Frydman, Carola and Raven Molloy (2012). "Pay cuts for the boss: Executive compensation in the 1940s". *The Journal of Economic History* 72.1, pp. 225–251.

Frydman, Carola and Raven E Saks (2010). "Executive compensation: A new view from a long-term perspective, 1936–2005". *Review of Financial Studies*, p. 120.

Gillezeau, Rob (2017). *Labour Law Enforcement during World War II and the Growth of the US Trade Union Movement*. Tech. rep. Working Paper.

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

Goldin, Claudia Dale and Lawrence F Katz (2008). *The race between education and technology*. Harvard University Press.

Hausman, Joshua K. (Apr. 2016). "Fiscal Policy and Economic Recovery: The Case of the 1936 Veterans' Bonus". *American Economic Review* 106.4, pp. 1100–1143.

Hirsch, Barry T, David A Macpherson, and Wayne G Vroman (2001). "Estimates of union density by state". *Monthly Labor Review* 124.7, pp. 51–55.

Inoue, Atsushi and Gary Solon (2010). "Two-sample instrumental variables estimators". *The Review of Economics and Statistics* 92.3, pp. 557–561.

Jakubson, George (1991). "Estimation and testing of the union wage effect using panel data". *The Review of Economic Studies* 58.5, pp. 971–991.

Jaumotte, Florence and Carolina Osorio Buitron (2020). "Inequality: traditional drivers and the role of union power". *Oxford Economic Papers* 72.1, pp. 25–58.

Johnson, George E (1975). "Economic analysis of trade unionism". *The American Economic Review*, pp. 23–28.

Katz, Lawrence F and DH Autor (1999). "Changes in the Wage Structure and Earnings Inequality". 1987. Ed. by Orley Ashenfelter and D. Card, pp. 35–78.

Katz, Lawrence F and Kevin M Murphy (1992). "Changes in relative wages, 1963–1987: supply and demand factors". *The Quarterly Journal of Economics* 107.1, pp. 35–78.

Katznelson, Ira (2013). *Fear Itself: The New Deal and the origins of our time*. WW Norton & Company.

King Jr, Martin Luther (1986). "If the Negro wins, Labor wins". *A Testament of Hope: The Essential Writings of Martin Luther King, Jr*, pp. 201–207.

Knepper, Matthew (2020). "From the Fringe to the Fore: Labor Unions and Employee Compensation". *Review of Economics and Statistics* 102.1, pp. 98–112.

Kopczuk, Wojciech, Emmanuel Saez, and Jae Song (2010). "Earnings inequality and mobility in the United States: evidence from social security data since 1937". *The Quarterly Journal of Economics* 125.1, pp. 91–128.

Kulkarni, Abhir and Barry T Hirsch (2019). "Revisiting Union Wage and Job Loss Effects Using the Displaced Worker Surveys". *ILR Review*, p. 0019793920912728.

Lee, David S (1999). "Wage inequality in the United States during the 1980s: Rising dispersion or falling minimum wage?" *The Quarterly Journal of Economics* 114.3, pp. 977–1023.

Lee, David S and Alexandre Mas (2012). "Long-Run Impacts of Unions on Firms: New Evidence from Financial Markets, 1961–1999". *The Quarterly Journal of Economics* 127.1, pp. 333–378.

Leighley, Jan E and Jonathan Nagler (2007). "Unions, voter turnout, and class bias in the US electorate, 1964–2004". *Journal of Politics* 69.2, pp. 430–441.

Lemieux, Thomas (1998). "Estimating the Effects of Unions on Wage Inequality in a Panel Data Model with Comparative Advantage and Nonrandom Selection". *Journal of Labor Economics* 16.2, pp. 261–291.

Lewis, H Gregg (1963). *Unionism and Relative Wages in the United States*. University of Chicago Press.

— (1983). "Union relative wage effects: A survey of macro estimates". *Journal of Labor Economics* 1.1, pp. 1–27.

— (1986). "Union relative wage effects". *Handbook of labor economics* 2, pp. 1139–1181.

Lichtenstein, Nelson (2003). *Labor's War at Home: The CIO in World War II*. Temple University Press.

Link, Henry C (1946). "The Psychological Corporation's index of public opinion." *Journal of Applied Psychology* 30.4, p. 297.

Manning, Alan (2020). "Monopsony in labor markets: a review". *ILR Review*, p. 0019793920922499.

Minchin, Timothy J (2017). *Labor under fire: A history of the AFL-CIO since 1979*. UNC Press Books.

Northrup, Herbert Roof (1971). *Organized labor and the Negro*. Periodicals Service Co.

Piketty, Thomas and Emmanuel Saez (2003). "Income inequality in the United States, 1913–1998". *The Quarterly journal of economics* 118.1, pp. 1–41.

Piketty, Thomas, Emmanuel Saez, and Stefanie Stantcheva (2014). "Optimal taxation of top labor incomes: A tale of three elasticities". *American economic journal: economic policy* 6.1, pp. 230–271.

Piketty, Thomas, Emmanuel Saez, and Gabriel Zucman (2018). "Distributional national accounts: methods and estimates for the United States". *The Quarterly Journal of Economics* 133.2, pp. 553–609.

Pischke, Jörn-Steffen, John E DiNardo, and Kevin F Hallock (2000). "Unions and the labor market for managers".

Raphael, Steven (2000). "Estimating the Union Earnings Effect Using a Sample of Displaced Workers". *Industrial and Labor Relations Review* 53.3, pp. 503–521.

Rios-Avila, Fernando and Barry T. Hirsch (2014). "Unions, Wage Gaps, and Wage Dispersion: New Evidence from the Americas". *Industrial Relations: A Journal of Economy and Society* 53.1, pp. 1–27. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/irel.12044.

Rosen, Sherwin (1970). "Unionism and the occupational wage structure in the United States". *International Economic Review* 11.2, pp. 269–286.

Schickler, Eric (2016). *Racial realignment: The transformation of American liberalism, 1932–1965*. Princeton University Press.

Song, Jae et al. (2015). *Firming up inequality*. Tech. rep. National Bureau of Economic Research.

Stafford, Frank P (1968). "Concentration and labor earnings: comment". *The American Economic Review* 58.1, pp. 174–181.

Stieber, Jack (1980). "Labor's Walkout from the Korean War Wage Stabilization Board". *Labor History* 21.2, pp. 239–260.

Taschereau-Dumouchel, Mathieu (2020). "The union threat".

Tinbergen, Jan (1970). "A Positive and a Normative Theory of Income Distribution". *Review of Income and wealth* 16.3, pp. 221–265.

Troy, Leo (1965). *Trade Union Membership, 1897–1962*. NBER.

Western, Bruce and Jake Rosenfeld (2011). "Unions, norms, and the rise in US wage inequality". *American Sociological Review* 76.4, pp. 513–537.

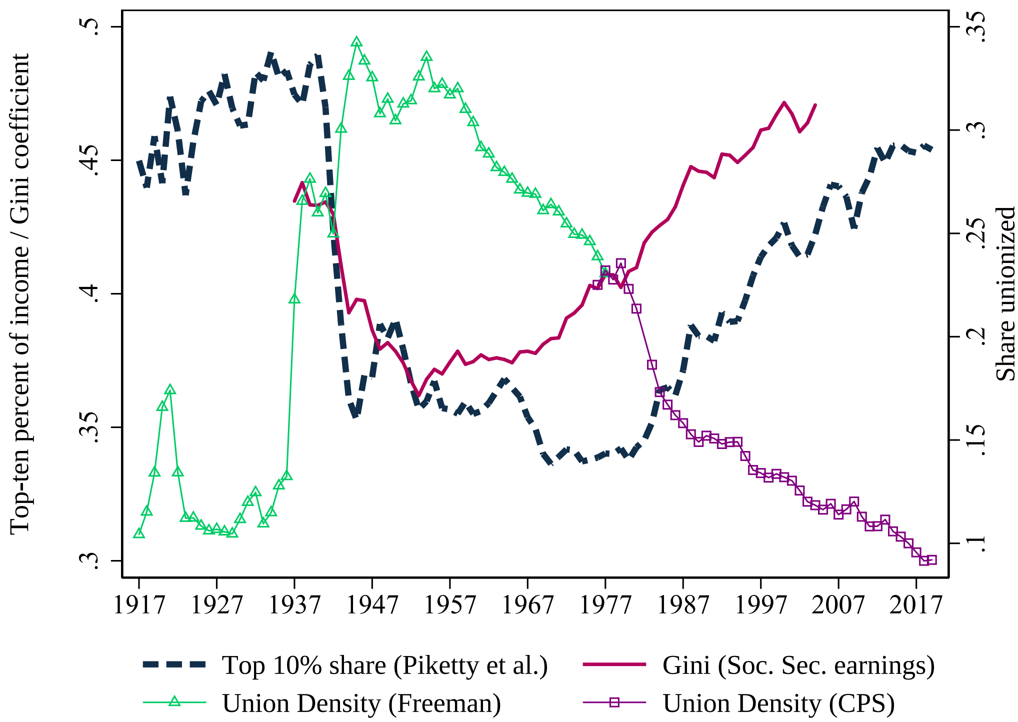

FIGURE I: UNION DENSITY AND INEQUALITY MEASURES, 1917-2019

Data sources: Top share individual income inequality from Piketty, Saez, and Zucman (2018). Union density is number of unionized workers as a share of non-agricultural workforce from Historical Statistics of the United States, together with individual union density as a share of employed civilian workers ages 16 to 65 from the Current Population Survey. We discuss these data sources in detail in Section II.B and Appendix E.

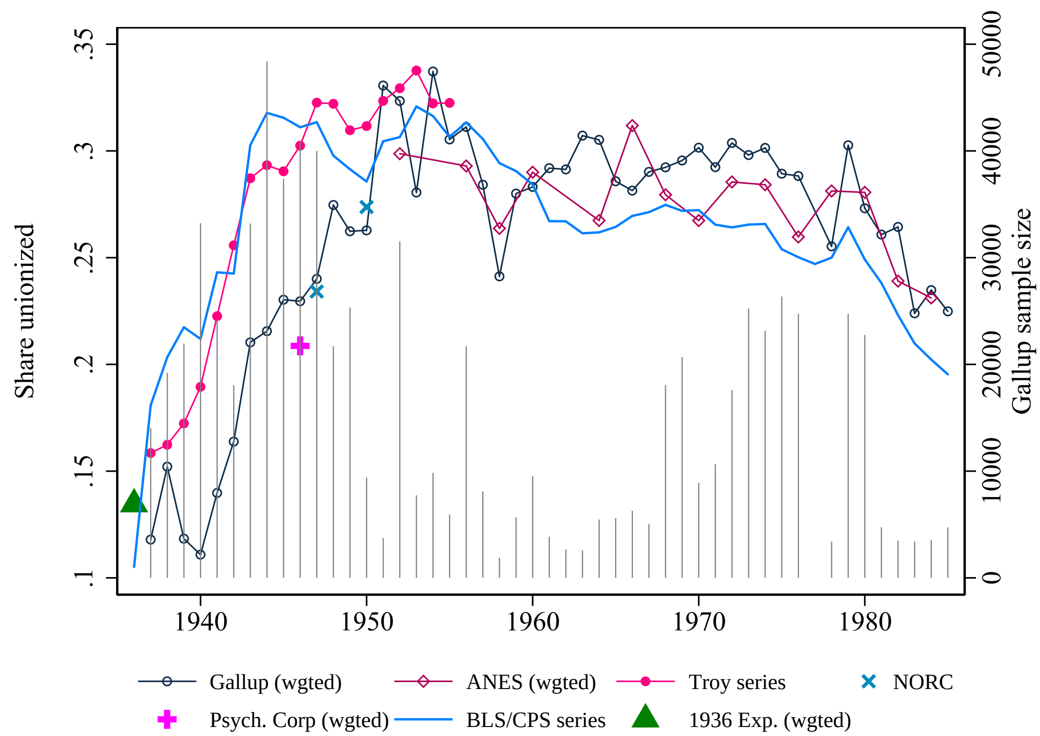

FIGURE II: THE SHARE OF HOUSEHOLDS WITH A UNION MEMBER, COMPARING OUR SURVEY-BASED MEASURES TO EXISTING TIME SERIES, 1936-1985

Data sources: Microdata sources used in this graph are from Gallup data, 1937–1986; CPS, 1977–2014; BLS Expenditure Survey, 1936; ANES, 1952–1996, U.S. Psych. Corporation, 1946. The historical data sources are the Leo Troy series (Troy, 1965) and the Bureau of Labor Statistics series (Freeman et al., 1998). See Sections II.B and Appendices B and E.

Notes: For our microdata sources, we include individuals age 18-65 whenever possible (for the Psych Corp and BLS Expenditure surveys, the sample is ages 21-65). The vertical spikes indicate the number of Gallup observations per year *that include the union variable* (plotted on the right-hand-side axis). The existing time series (the BLS and Troy measures) are *counts* of union members, so we divide them by Census estimates of the number of households (geometrically interpolated between Census years) to make them as comparable as possible to our household membership series. The Gallup, ANES, 1936 Expenditure and Psychological Corporation are all weighted, either with survey-provided weights or to match Census demographics as described in Section II.B and Appendix B.

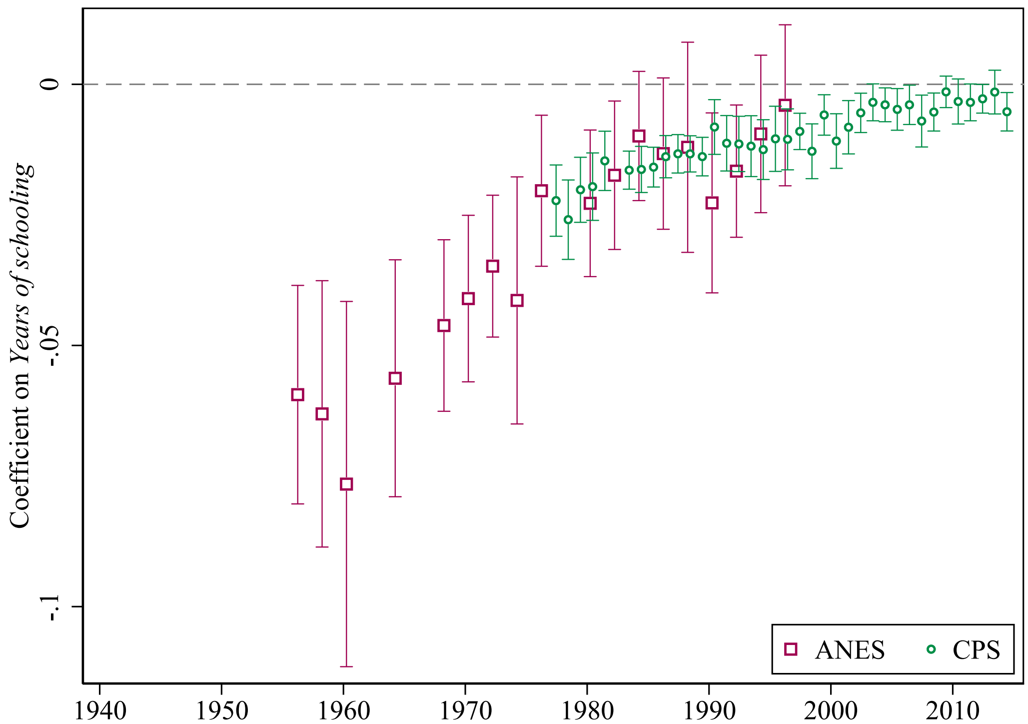

FIGURE III: HOW DOES YEARS OF SCHOOLING PREDICT UNION HOUSEHOLD STATUS?

Data sources: Gallup data, 1942–1986; CPS, 1977–2014; BLS Expenditure Survey, 1936; ANES, 1952–1996, U.S. Psych. Corporation, 1946; Panel Study of Income Dynamics, 1968, 1976. See Section II.B for a description of each data source.

Notes: We regress household union status on $Years of education$, state s and survey-date t fixed effects, age and its square, and gender (all demographics refer to the survey *respondent*); for the CPS we also control for number of *employed* household members (because in the CPS the union question is only asked of those who are employed). We estimate this equation separately by survey source and by year. We harmonize years of schooling in the following manner: ten years for "less than high school;" twelve years for "high school;" fourteen years for "some college" or "vocational training;" sixteen years for "college" or "more than college." The figure plots the coefficient on *Years of schooling*. For the ANES, because the samples are smaller, we group surveys into six-year bins. The plotted confidence intervals are based on standard errors clustered by state. Note that Gallup does not consistently ask respondent education until 1942, which is why the Gallup analysis here begins later than in some other analyses.

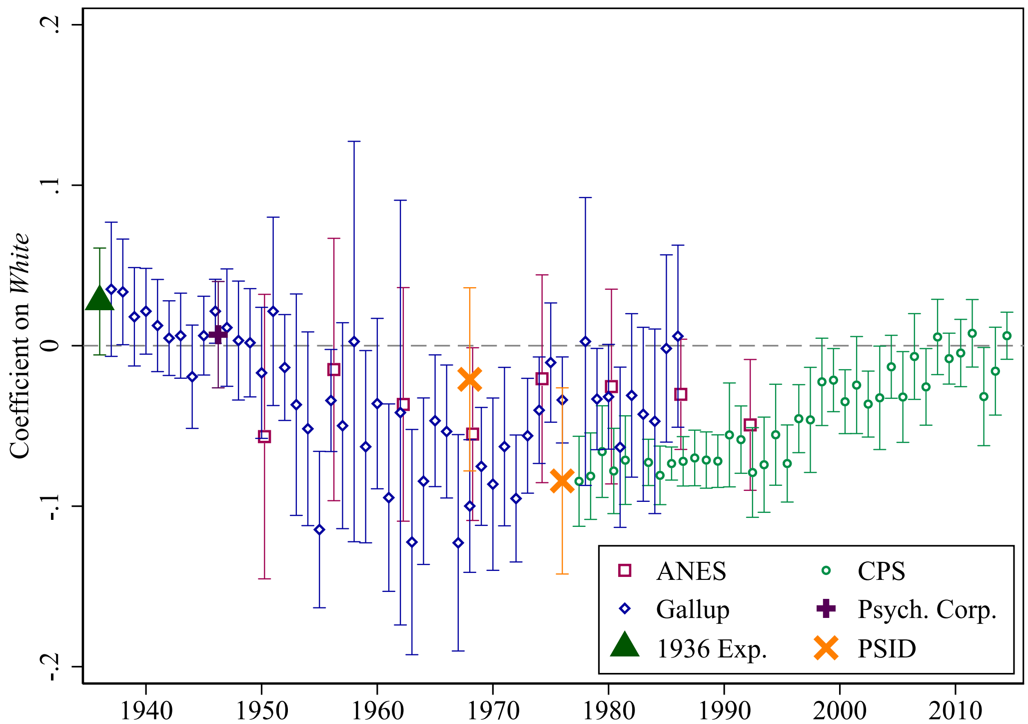

FIGURE IV: HOW DOES RACE PREDICT UNION HOUSEHOLD STATUS?

Data sources: Gallup data, 1937–1986; CPS, 1978–2016; BLS Expenditure Survey, 1936; ANES, 1952–1996, U.S. Psych. Corporation, 1946; Panel Study of Income Dynamics, 1968, 1976. See Section II.B for a description of each data source.

Notes: For each data source, we estimate (separately by year if a data source has multiply years), household union status on a *White* dummy variable, state s and survey-date t fixed effects, age and its square, and gender (all demographics refer to the survey *respondent*); for the CPS we also control for number of *employed* household members (because in the CPS the union question is only asked of those who are employed). We plot in this graph the coefficients on *White* from each of these estimations. For the ANES, because the samples are smaller, we group surveys into six-year bins. Confidence intervals are based on standard errors clustered by state.

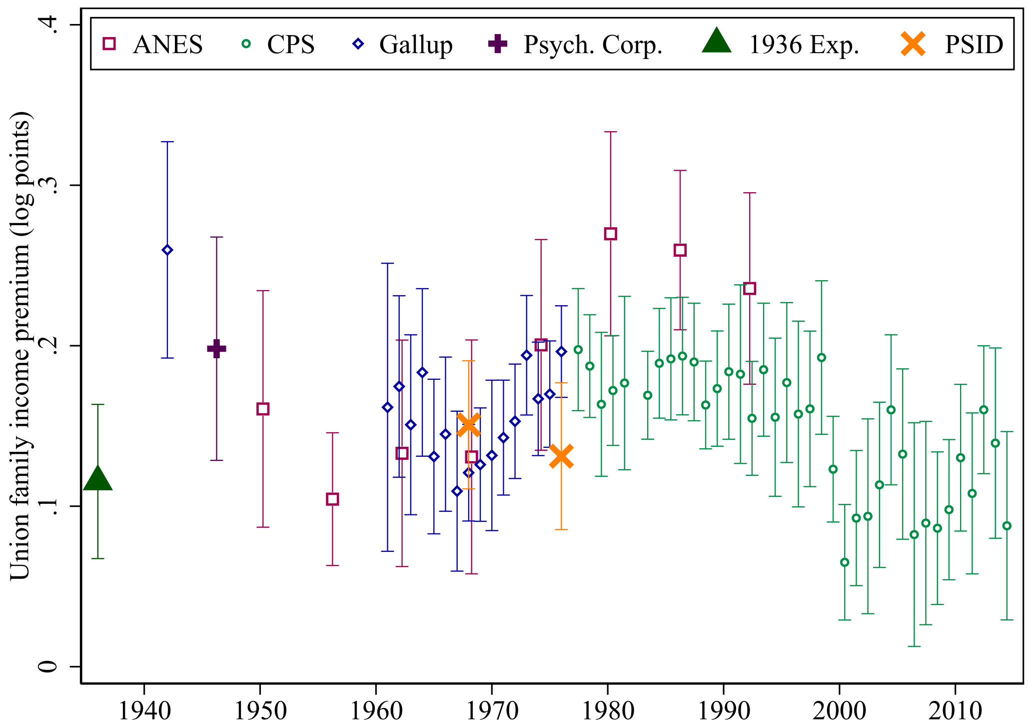

FIGURE V: ESTIMATES OF THE UNION FAMILY INCOME PREMIUM

Data sources: Gallup data, 1942, 1961–1976; CPS, 1977–2014; BLS Expenditure Survey, 1936; ANES, 1952–1996, U.S. Psych. Corporation, 1946; Panel Study of Income Dynamics, 1968, 1976. See Section II.B for a description of each data source. See Appendix C for details on CPS family-income variable construction.

Notes: Each plotted point comes from estimating equation (2), which regresses log family income on *household union status*, with controls for years of schooling (harmonized into four categories corresponding to 10, 12, 14 and 16 years), age, gender, race, state and survey-date fixed effects. Whenever possible we also include controls for employment status of households members. Occupation controls are not included as they are not consistent across data sources or within data sources across time. We estimate a separate regression for each survey source and year. For the ANES, because the samples are smaller, we group surveys into six-year bins. The plotted confidence intervals are based on standard errors clustered by state.

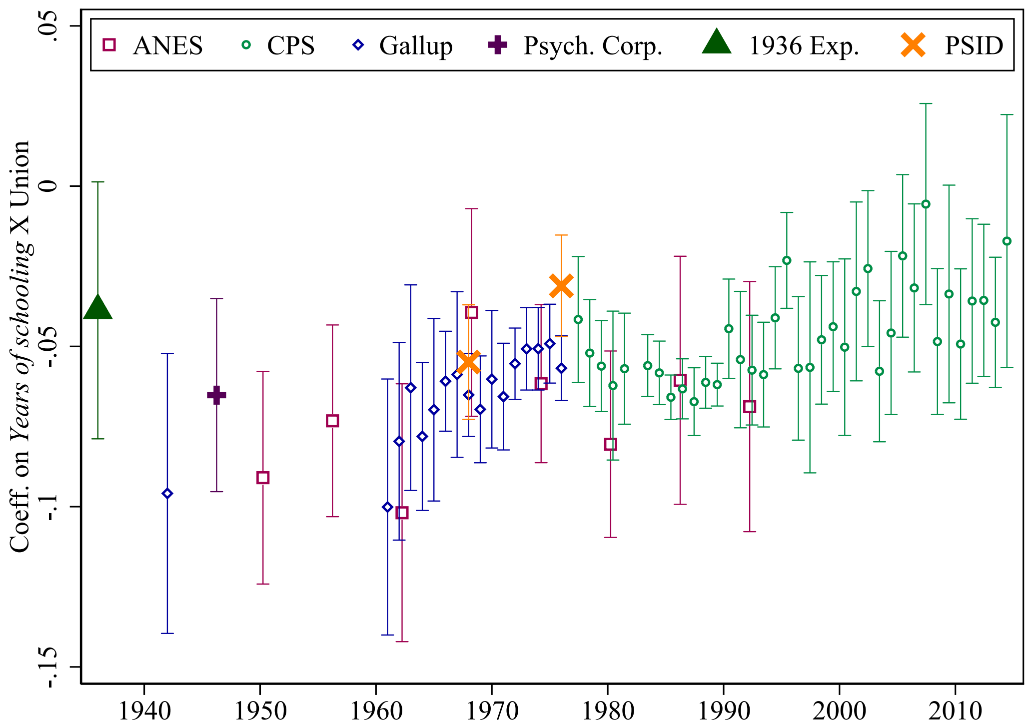

FIGURE VI: DIFFERENTIAL FAMILY UNION PREMIUM BY RESPONDENT'S YEARS OF SCHOOLING

Data sources: Gallup data, 1942, 1961–1976; CPS, 1977–2014; BLS Expenditure Survey, 1936; ANES, 1952–1996, U.S. Psych. Corporation, 1946; Panel Study of Income Dynamics, 1968, 1976. See Section II.B for a description of each data source. See Appendix C for details on CPS family income variable construction.

Notes: Each plotted point comes from estimating an equation regressing log family income on household union status, its interaction with respondents' years of schooling, and all other controls in the union-premium equation (2). We estimate this equation separately by survey source and by year. The *Years of schooling* variable is harmonized across surveys into four categories (10, 12, 14 and 16 years). The figure plots the coefficient on the interaction *Years of schooling* × *Union*. For the ANES, because the samples are smaller, we group surveys into six-year bins. The plotted confidence intervals are based on standard errors clustered by state.

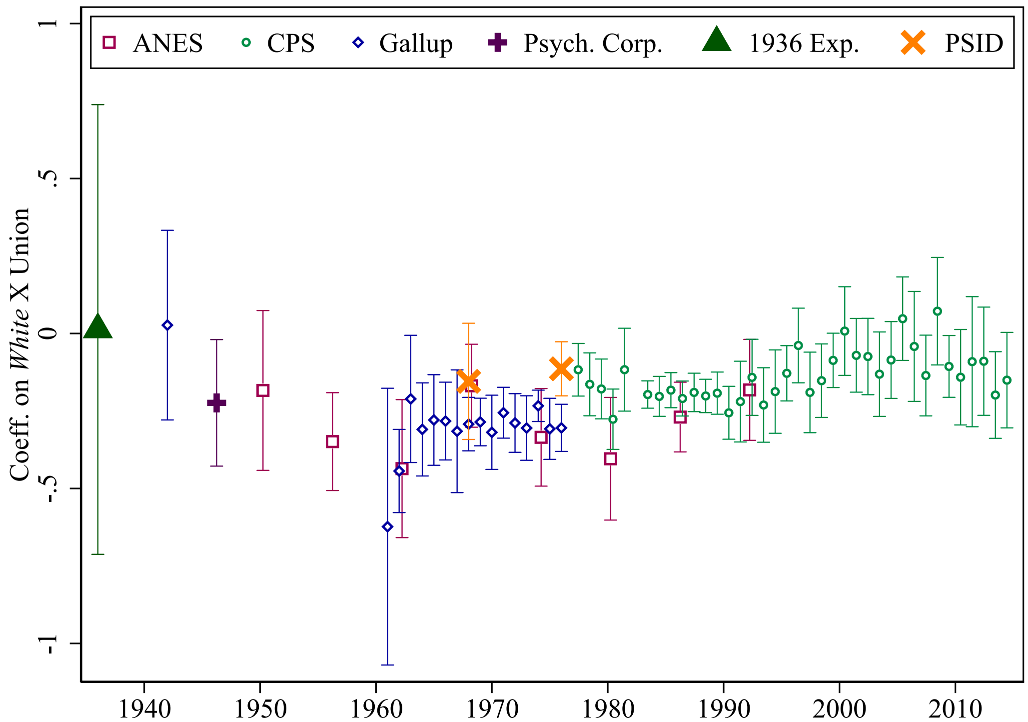

FIGURE VII: DIFFERENTIAL FAMILY UNION PREMIUM FOR WHITES RELATIVE TO MINORITIES

Data sources: Gallup data, 1942, 1961–1976; CPS, 1977–2014; BLS Expenditure Survey, 1936; ANES, 1952–1996, U.S. Psych. Corporation, 1946; Panel Study of Income Dynamics, 1968, 1976. See Section II.B for a description of each data source. See Appendix C for details on CPS family income variable construction.

Notes: Each plotted point comes from estimating an equation regressing log family income on household union status, its interaction with a White dummy variable, and all other controls in the union-premium equation (2). We estimate this equation separately by survey source and by year. The figure plots the coefficient on the interaction $White \times Union$. For the ANES, because the samples are smaller, we group surveys into six-year bins. The plotted confidence intervals are based on standard errors clustered by state.

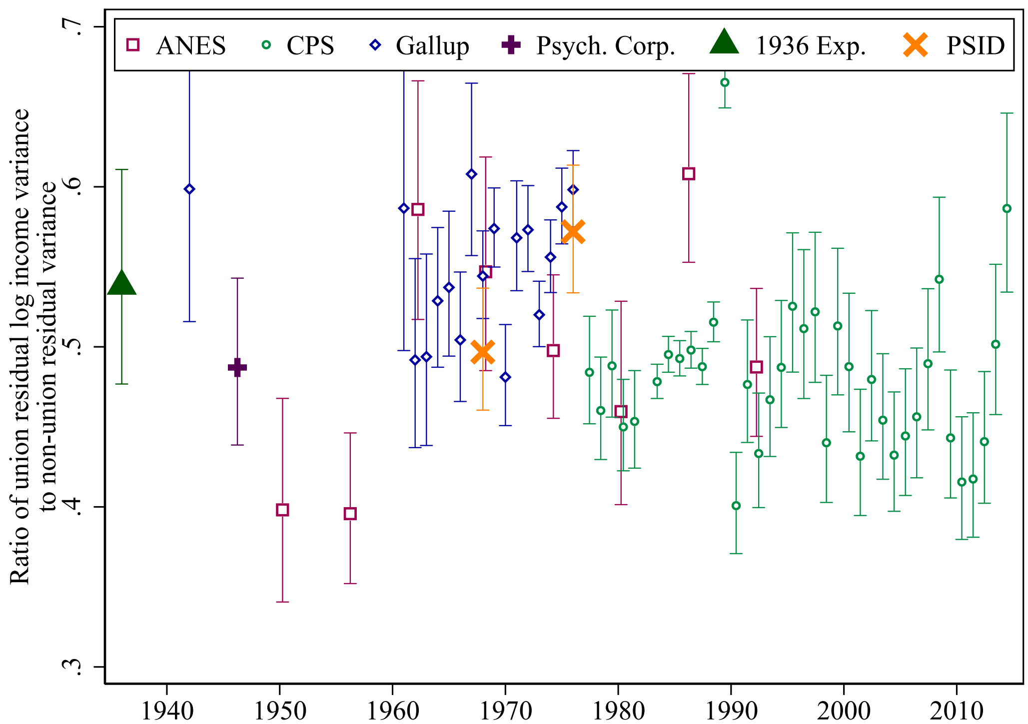

FIGURE VIII: RATIO OF RESIDUAL VARIANCE BETWEEN UNION AND NON-UNION SECTORS

Data sources: Gallup data, 1942, 1961–1976; CPS, 1977–2014; BLS Expenditure Survey, 1936; ANES, 1952–1996, U.S. Psych. Corporation, 1946; Panel Study of Income Dynamics, 1968, 1976. See Section II.B for a description of each data source. See Appendix C for details on CPS family income variable construction.

Notes: Each plotted point is the ratio of variance of residuals from regressing log family income on the controls in equation (2) separately for union and non-union households. As usual, we perform this analysis separately by survey source and year. See Section IV.D for more detail. The figure plots the ratio of the variance of residuals in the union sector to that of the non-union sector (so ratios less than one suggest that residual variance in the union sector is more compressed than in the non-union). The plotted confidence intervals are based on inverting the F-statistic testing the null that the ratio is equal to 1. For the ANES, because the samples are smaller, we group surveys into six-year bins.

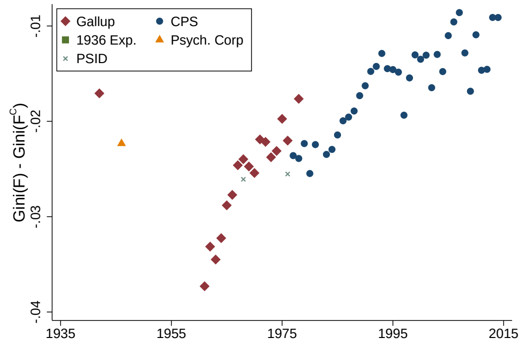

(A) Yearly union impact (assuming no spillovers to non-union households)

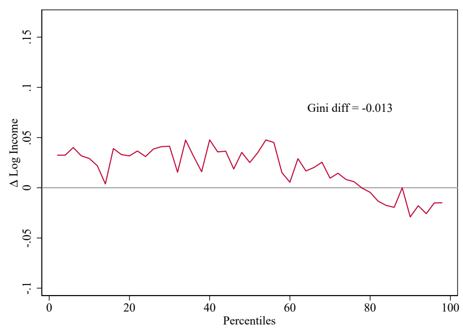

(B) 1936 (Exp. survey)

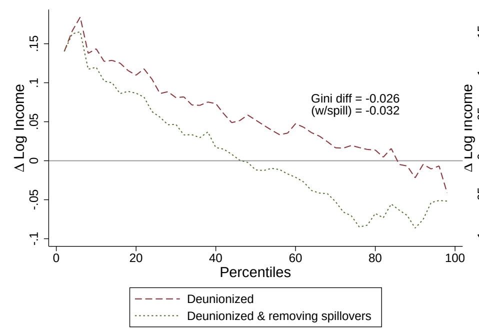

(C) 1968 (PSID)

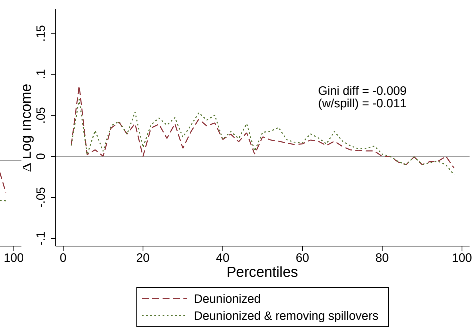

(D) 2014 (CPS)

Figure IX: Actual vs. "no-unions" counterfactual income distributions

Data sources: Gallup data, 1942, 1961–1976; CPS, 1978–2016; BLS Expenditure Survey, 1936; ANES, 1952–1996, U.S. Psych. Corporation, 1946. See Section II.B for a description of each data source. See Appendix C for details on CPS family income variable construction. See section V.A and appendix F for DFL reweighting factor construction.

Notes: This figure compares the observed population (F_Y) and the counterfactual population without unions (F_{Y_n}) in selected years. The counterfactual population's income distribution is calculated by upweighting the non-union observations by the inverse of the predicted probability of being union, estimated using a logistic regression of union household on race, age, age-squared, education dummies, and state indicators. Panel (a) plots yearly differences in true and counterfactual Gini coefficients. Panels (b) through (c) plot differences in true and counterfactual log-family-income percentiles for 1936, 1968, and 2014, respectively. Income is denominated in 2014 dollars using CPI.

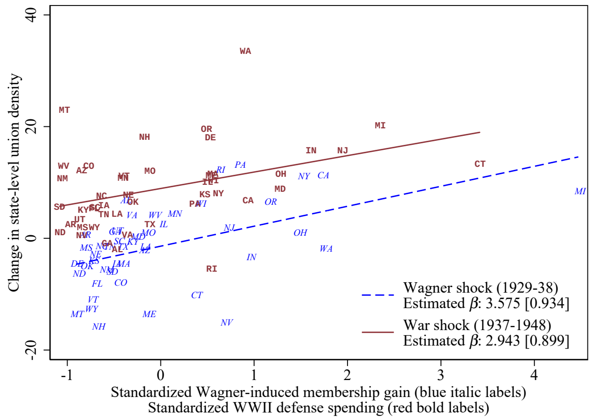

(A) Union Density (first stage)

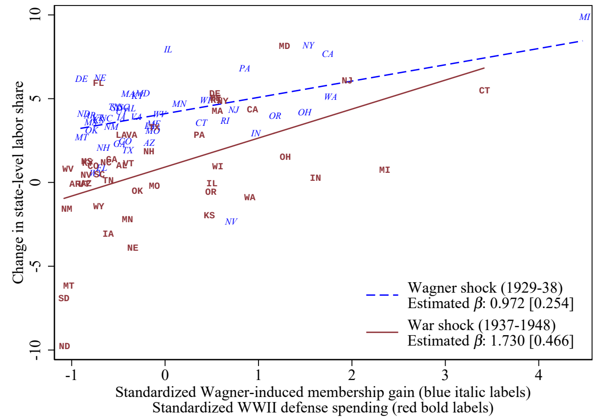

(B) Labor Share (reduced form)

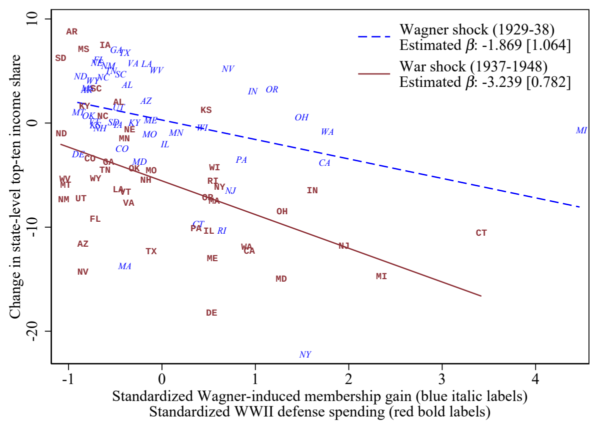

(C) Top Ten (reduced form)

Figure X: Union density, inequality measures regressed on Wagner-Act and WW-II-spending policy-shock variables

Data sources: The outcome variable for panel (a) comes from Gallup data for 1947 and 1938 and from the 1929 Handbook of American Trade Unions (see Appendix C and Cohen, Malloy, and Nguyen (2016) for details and validation of the 1929 measure). The outcome variable for panel (b) comes from our estimate of historical state-year labor shares, detailed in Appendix H. The outcome variable for panel (c) are top-one-percent shares of state income, taken from Frank, 2015.

Notes: Each subfigure shows two scatter plots: the outcome variable against the Wagner shock (states labeled in blue, italic Times font); and the outcome variable against the war shock (red, bold, Courier font). In all cases, the outcome variables are in nine-year changes (the effect of the Wagner shock is estimated from 1929-1938 and the War-spending shock from 1938-1947) and plotted for all 47 states in our data. Both shocks are standardized and plotted on the same x-axis. Except for standardizing the x-axis variables, we plot *the raw data* (not residualized). We display the β and robust standard errors from the (bivariate) OLS regressions of the outcome variable against each shock.

	Time Period (1)	Total Change in Statistic (2)	Change Attributable to: Δ Union Wages (3)	Δ Unionization (4)	Total Union Effect (5)
Panel A: Gini Coefficient	1936 to 1968	-0.0526	0.00169	-0.0149	-0.0132
			(-3.223)	(28.37)	(25.14)
	1968 to 2014	0.144	0.0111	0.00587	0.0170
			(7.702)	(4.075)	(11.78)
Panel B: Log 90/10	1936 to 1968	-0.188	-0.00911	-0.0980	-0.107
			(4.846)	(52.17)	(57.01)
	1968 to 2014	0.817	0.106	0.0494	0.155
			(12.94)	(6.041)	(18.98)
Panel C: Log 90/50	1936 to 1968	-0.102	0.0129	-0.0328	-0.0198
			(-12.63)	(31.99)	(19.36)
	1968 to 2014	0.360	0.0120	0.0281	0.0401
			(3.327)	(7.818)	(11.14)
Panel D: Log 10/50	1936 to 1968	0.0855	0.0220	0.0653	0.0873
			(25.78)	(76.33)	(102.1)
	1968 to 2014	-0.458	-0.0938	-0.0213	-0.115
			(20.50)	(4.644)	(25.14)
Panel E: Log College Premium	1936 to 1968	-0.231	-0.00415	-0.0417	-0.0459
			(1.797)	(18.06)	(19.86)
	1968 to 2014	0.0688	0.0386	0.0219	0.0605
			(56.10)	(31.86)	(87.96)

TABLE I: DECOMPOSITION OF CHANGE IN GINI

Data sources: Data for years 1936, 1968, and 2014 are taken from the 1936 Expenditure Survey, PSID, and CPS, respectively. Gini coefficient, log income ratios, and college premium are calculated using household-level income in the labeled, with weights applied according to DiNardo, Fortin, and Lemieux, 1996. See section V.A and appendix F for reweighting factor construction.

Notes: This table reports the union-related components of DFL decompositions of changes in Gini coefficient, log 90/10, log 90/50, and log 10/50 income ratios, and log college premium over time. Each panel represents a different inequality measure and each row represents a separate decomposition. Column 1 specifies the beginning and end years of the decomposition. Column 2 reports the total change in computed inequality measure, and columns 3-5 report components of that change from a DFL decomposition. Column 3 reports the change in inequality measure attributable to changes in union versus non-union incomes. Column 4 reports the change in inequality attributable to changes in the conditional unionization rate. Column 5 reports the total effect of both union wage changes and unionization (Column 3 + Column 4). Numbers in parentheses report components as a percentage of total change in the inequality measure. Each component is calculated using true and counterfactual inequality measures, where counterfactuals are constructed by reweighting households according to their relative predicted probabilities of union membership in beginning and end years. Predicted union membership is estimated using logistic regressions of household union status against education, race, a quadratic in respondent age, and state fixed effects. See section V.A and appendix F for reweighting details and formal definitions of components.

	Dependent variable:													
	Coll. premium		Log 90/10		Log 90/50		Log 10/50		Gini coeff.		Top 10 share		Labor share	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)
Union density	-1.090**	-1.115	-2.189***	-1.936***	-0.450	-0.489	1.739***	1.447**	-0.168***	-0.160***	-69.16***	-61.97***	43.12***	39.43***
	[0.477]	[0.693]	[0.415]	[0.688]	[0.332]	[0.366]	[0.420]	[0.629]	[0.0386]	[0.0422]	[18.10]	[18.08]	[10.71]	[13.21]
Skill share	-0.586***	-0.572***	-0.158	0.179	-0.329***	-0.232***	-0.172	-0.411***						
	[0.0996]	[0.125]	[0.0986]	[0.119]	[0.0882]	[0.0669]	[0.125]	[0.121]						
Mean, dept. var	0.476	0.476	1.423	1.423	0.662	0.662	-0.762	-0.762	0.410	0.410	38.304	38.304	73.144	73.144
Annual edu. controls?	No	No	No	No	No	No	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Addit. controls?	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes
Cubic polynomial?	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Min. Year	1940	1940	1940	1940	1940	1940	1940	1940	1940	1940	1940	1940	1940	1940
Max. Year	2014	2014	2014	2014	2014	2014	2014	2014	2004	2004	2014	2014	2014	2014
Observations	54	54	54	54	54	54	54	54	65	65	75	75	75	75

TABLE II: AGGREGATE INEQUALITY AS A FUNCTION OF UNION DENSITY

Data sources: For cols. (1) - (8), outcome variables generated from Census IPUMS and CPS; for cols (9) and (10) from Kopczuk, Saez, and Song (2010); for cols (11)-(14) from Piketty, Saez, and Zucman (2018). The union density explanatory variable is the simple average between the Gallup- and BLS-based density measures (see Section V.B for detail).

Notes: Note that union density is out of one (not 100) to conserve table space by avoiding coefficients with zeros after the decimal. All regressions include controls for the log share of college versus high-school educated workers, calculated in the early years from Census IPUMS and for later years from the CPS. The first eight columns use outcome variables calculated from the source (so are only available in Census years until the CPS), but the last eight columns use *annual* measures as outcomes , calculated from administrative data. For these measures, we have to control *annually* for skill shares. We include two annual controls: annual skills shares as measured in Gallup and annual skills shares as measured in the Census IPUMS and the CPS (interpolated between Census years in the pre-CPS years). As these two measures are correlated, we do not report their coefficients because they are hard to interpret (and are not the variables of interest). For each outcome variable, the first specification has *parsimonious* controls (only a time cubic and the skill shares controls) and the second has *additional* controls (federal minimum wage, the national unemployment rate, and the top marginal tax rate in the federal income tax schedule). Appendix Tables A.7 and A.8 provides additional specifications using the college premium, the log percentile ratios, the Gini coefficient, the top-ten share and the labor share as outcomes. Note that the log 90/10, 90/50, and 10/50 ratios are for men only, but all other inequality measures pool both men and women. Standard errors are robust to heteroskedasticity and AR(1) serial correlation. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$

	Dependent Variable:															
	Coll. prem.		log 90/10		log 90/50		log 10/50		Gini coeff.		Top 10			Labor share		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
Household union share	-0.187	-0.214*	-0.345**	-0.307**	-0.140	-0.122	0.205*	0.184*	-0.055**	-0.054**	-4.192**	-3.479**	-4.704**	5.567***	3.972**	5.861***
	[0.136]	[0.128]	[0.168]	[0.149]	[0.088]	[0.086]	[0.113]	[0.102]	[0.027]	[0.022]	[1.917]	[1.693]	[1.990]	[1.870]	[1.789]	[1.884]
Mean	0.462	0.462	1.408	1.408	0.666	0.666	-0.742	-0.742	0.394	0.394	37.123	37.123	37.151	74.559	74.559	74.532
Controls?	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	No	Yes	No
Min Year	1940	1940	1940	1940	1940	1940	1940	1940	1940	1940	1937	1937	1929	1937	1937	1929
Max. Year	2014	2014	2014	2014	2014	2014	2014	2014	2014	2014	2014	2014	2014	2014	2014	2014
Observations	1,960	1,960	1,960	1,960	1,960	1,960	1,960	1,960	1,960	1,960	3,537	3,537	3,584	3,537	3,537	3,584

TABLE III: STATE-YEAR INEQUALITY AS A FUNCTION OF UNION DENSITY

Data sources: For cols. (1) through (10), dependent variables created using Census and CPS data. Note that the Gini coefficient used in Table II is not available at the state level, so in cols. (9) and (10) we calculate a state-level annual Gini from the Census and CPS. For cols. (11) through (13) outcome variables are taken from Frank, 2015; for cols (14) through (16) we construct a state-level labor share of net income (see Appendix H for details and validation). The key explanatory variable comes from state-year average household union share generated from Gallup in the earlier years and the CPS in later years. Cols. (13) and (16) add a 1929 measure of state-year density based on data from the Handbook of American Trade Unions (see Appendix C and Cohen, Malloy, and Nguyen (2016) for details and validation) and a 1929 measure of skill shares based on the 1940 Census with age and migration adjustment (see Appendix C for details and validation).

Notes: Note that union density is out of one (not 100) to conserve table space by avoiding coefficients with zeros after the decimal. All estimates are from split-sample-IV regressions (see Section V.C for estimating equations), repeated 200 times (bootstrapped estimates and standard errors, clustered by state, reported). All regressions include state and year fixed effects; $South \times Year$ fixed effects; and state-year education controls (both from Gallup and CPS at the annual level, and interpolated from the IPUMS Census at the decade level). "Controls" include state-year share of employment in all one-digit industry categories, state-year log income, state-year share of households filing taxes, state-year minimum wage, state top marginal income tax rate, a "policy liberalism" index (from Caughey and Warshaw, 2016), a dummy for Democratic governor, and state-year top marginal tax rates. Sample size is larger for the top 10 and labor share outcomes because they are available at the annual level and go back further in time; for the other outcomes, until the CPS in the 1970s, we only have data from the decadal Census beginning in 1940. $^{*}p < 0.1,^{**}p < 0.05,^{***}p < 0.01$

	Top 10		Labor share		Top 10		Labor share	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Change in density	-0.289**	-1.154**	0.221**	0.563**	-0.621**	-0.555**	0.325**	0.230**
	[0.122]	[0.397]	[0.0774]	[0.244]	[0.118]	[0.136]	[0.0613]	[0.0505]
Change in skill share	28.83	-10.06	-9.580	1.670	-5.516**	-5.174**	1.450	1.143
	[26.18]	[6.214]	[11.81]	[3.167]	[1.844]	[1.804]	[1.057]	[0.878]
Change in manuf. share	11.71	19.60	19.44	-6.201	0.891	-2.062	6.405	9.923
	[31.31]	[65.65]	[15.43]	[37.14]	[13.09]	[12.12]	[7.331]	[6.220]
Wagner shock					0.345	0.395	0.0872	0.176
					[0.234]	[0.312]	[0.0960]	[0.143]
War shock					-0.307	-0.346	0.00390	0.0185
					[0.349]	[0.360]	[0.136]	[0.136]
Dept. var. mean	0.292	-5.554	4.107	0.920	0.643	0.643	0.0320	0.0206
First-stage F-stat	12.68	8.237	12.68	8.237	16.22	24.30	16.22	24.30
Top CI	-.593175	.	-.064511	.244108	-.943238	-.936806	.177024	.072219
Bottom CI	.005594	-.604125	.346269	.	-.402154	-.366746	.385776	.27621
Interval	1929-38	1938-47	1929-38	1938-47	All	All	All	All
Ex. Mich	No	No	No	No	No	Yes	No	Yes
Observations	47	47	47	47	409	400	409	400

TABLE IV: IV ESTIMATION OF CHANGES IN STATE INEQUALITY ON CHANGES IN STATE DENSITY

Data sources: Data on state-year density Gallup data from 1938-1977, from Gallup and CPS from 1978 onward. State-year density data from 1929 is from the 1929 Handbook of American Trade Unions (see Appendix C and Cohen, Malloy, and Nguyen (2016) for details and validation of the 1929 measure). The top-one-percent shares of state-year income are taken from Frank (2015). The labor-share measures come from our estimate of historical state-year labor shares, detailed in Appendix H.

Notes: Cols. (1) and (2) display IV regression results when the nine-year change in top-ten share is the outcome variable. Col. (1) models the change between 1929 and 1938, using the *Wagner shock* as the excluded instrument; col. (2) models the change between 1939 and 1947, using the *War-spending shock* as the excluded instrument. Cols. (3) and (4) are analogous to cols. (1) and (2) except that the change in the labor share is the outcome variable. The remaining columns include *placebo* intervals (1947-1956, 1956-1965, etc...). Col. (5) models nine-year changes in the top-ten share, with $\textit{Wagner shock}_s \times \mathbb{I}_t^{t=1938}$ and $\textit{War-spending shock}_s \times \mathbb{I}_t^{t=1947}$ are the excluded instruments, and the main effects of *Wagner shock* and *War-spending shock* as controls. Col. (6) replicates col. (5) after dropping Michigan (which has the largest value for both policy shock variables). Col. (7) and (8) are analogous to cols (5) and (6) but with nine-year changes in labor share as the outcome. Cols. (1) through (4) include Census-region fixed effects, and cols. (5) through (8) region fixed effects interacted with year. The "Top CI" and "Bottom CI" reported in the table footer in each column refer to confidence intervals robust to weak instruments. They are based on Anderson-Rubin tests (cols. 1-4) or conditional-likelihood ratio tests (cols. 5-8). A missing value indicates negative or positive infinity.