WORKING PAPER #620
PRINCETON UNIVERSITY
INDUSTRIAL RELATIONS SECTION
May 2018
Version: May 2, 2018
<http://arks.princeton.edu/ark:/88435/dsp01gx41mm54w>

Unions and Inequality Over the Twentieth Century: New Evidence from Survey Data*

Henry S. Farber, Daniel Herbst, Ilyana Kuziemko, Suresh Naidu

Abstract

It is well-documented that, since at least the early twentieth century, U.S. income inequality has varied inversely with union density. But moving beyond this aggregate relationship has proven difficult, in part because of the absence of micro-level data on union membership prior to 1973. We develop a new source of micro-data on union membership, opinion polls primarily from Gallup ($N \approx 980,000$), to look at the effects of unions on inequality from 1936 to the present. First, we present a new time series of household union membership from this period. Second, we use these data to show that, throughout this period, union density is inversely correlated with the relative skill of union members. When density was at its peak in the 1950s and 1960s, union members were relatively less-skilled, whereas today and in the pre-World War II period, union members are equally skilled as non-members. Third, we estimate union household income premiums over this same period, finding that despite large changes in union density and selection, the premium holds steady, at roughly 15–20 log points, over the past eighty years. Finally, we present a number of direct results that, across a variety of identifying assumptions, suggest unions have had a significant, equalizing effect on the income distribution over our long sample period.

*We thank our research assistants Obaid Haque, Elena Marchetti-Bowick, Chitra Marti, Brendan Moore, Tamsin Kantor, Amy Wickett, and Jon Zytnick and especially Fabiola Alba, Divyansh Devnani, Elisa Jacome, Amitis Oskoui, and Maryam Rostoum. We have benefited from comments by seminar participants at Columbia, Georgetown, Harvard, INSEAD, the NBER Development of the American Economy meetings, McGill University, Princeton, Rutgers, Sciences Po, Universitat Pompeu Fabra, Stanford, and Vanderbilt. We are indebted to Devin Caughey and Eric Schickler for answering questions on the early Gallup data. All remaining errors are our own. Financial support was provided by the Industrial Relations Section at Princeton University. Farber: Princeton University and NBER, farber@princeton.edu. Herbst: Princeton University, dherbst@princeton.edu. Kuziemko: Princeton University and NBER, kuziemko@princeton.edu. Naidu: Columbia University and NBER, sn2430@columbia.edu.

1 Introduction

Understanding the determinants of the U -shaped pattern of U.S. income inequality over the 20th century has become a central goal among economists over the past few decades. While there is a substantial literature in labor economics and sociology that argues for a causal relationship from redistributive institutions such as labor unions to lowered inequality (Card, 2001; DiNardo *et al.*, 1996; Western and Rosenfeld, 2011), the consensus among economists is that more fundamental drivers, namely technological developments that increase the demand for skilled labor, better explain the time-series variation in inequality (Acemoglu *et al.*, 2001).¹

In the aggregate, there is a well-documented inverse relationship between income inequality and union membership in the US (see Figure 1). But moving beyond this aggregate relationship has proven difficult. 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 U.S. micro data. It is perhaps not surprising that, instead, economists focus on models of inequality that emphasize the supply and demand for skill, as the U.S. Census has tracked education and wages consistently since 1940 (see, e.g., the seminal work by Goldin and Katz, 2009). Indeed, to the extent recent papers have tried to explicitly model the relationship between unions and inequality, they have argued that declining union density is largely an artifact of rising skill-biased technological change (SBTC), and thus unions themselves have little if any direct effect on inequality.²

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 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

¹As Acemoglu *et al.* (2001) write of the post-1970 period: “Most economists, however, discount the role of unions in the increase in inequality.” Indeed, literature searches and polls of actual economists support the idea that economists favor alternative explanations of the variation in inequality over time. For example, a search for articles from 1990 onward on “inequality and technology” in Google Scholar yields 1.75 million hits, compared to 183,000 for “inequality and unions.” In 1997, the Federal Reserve conducted a poll of economists, asking them their view on the key driver on increased inequality. Forty-four percent said “technological change,” fifteen percent said “other,” twelve percent said “international trade,” nine percent said “decline in real minimum wage,” nine percent said “decline in unionization,” and the rest said “rising immigration.” See 1997 Economic Report of the President, Figure 5-3.

²We will discuss these papers in greater detail in the next section.

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 micro-data sources (e.g., the CPS) to extend the analysis into the present day.

These new data sources allow us to revisit the role of unions in shaping the income distribution and contribute to the long-running “institutions versus market forces” debate on the causes of inequality. Bringing new micro data to the study of unions allows us to perform a number of novel tests on the plausibility of unions playing a significant role in reducing income inequality. For example, suppose we find that during the “Great Compression” period (roughly 1940–1970, when unions were at their peak and inequality at its lowest), union members were relatively educated or disproportionately white.³ Then this pattern of selection, with generally low skill and low wage workers being underrepresented in the union sector during the twentieth century’s most egalitarian years, would argue against unions playing an important role in reducing inequality. By contrast, we find the opposite pattern of selection, and the results we do find push against the idea that unions are simply an artifact of larger market forces and instead favor the idea that they may have their own causal effect in reducing inequality. We support this hypothesis with a number of new empirical results, which fall into two broad sets.

Our first set of results makes heavy use of the fact that our data sources are all *microdata* at the household or individual level, allowing us to examine over an eighty-year period how household union status covaries with characteristics like education, race, and household income. We consider many of the stylized facts about unions established with CPS data, and extend them to earlier decades. We begin by showing that patterns of selection into unions has varied substantially over time: the education of union members relative to non-union members has followed a marked *U*-shape pattern, mirroring the pattern of inequality itself and the sharp inverse of union density. That is, during the Great Compression period, when density was the highest, unions were drawing in the least skilled workers. Today, as in the 1930s, unions are smaller and their members at least (if not slightly more) educated than the rest of the workforce. A similar pattern emerges for minorities: unions were relatively less white during the Great Compression period than either before or after.

A key stylized fact about CPS-era unions is that members enjoy a wage premium, but did this advantage exist as union density was growing in the 1930s and 1940s or during unions’ Great Compression heyday? We show that the income advantage accruing to union

³While a widely used term, there does not appear to be a consensus on the exact years that the “Great Compression” begins and ends. The term appears to originate with Goldin and Margo (1992) and in their paper refers more specifically to the 1940s. The Wikipedia entry, however, suggests the “Great Compression” ends around 1970. In this paper, we will use it to refer to, roughly, 1940–1970, but nowhere does this exact definition matter for any result.

households relative to non-union households with the same demographics and skill proxies is roughly constant (between fifteen and twenty log points) over our eighty-year period, despite the huge swings in union density and composition. While extrapolating selection patterns from the CPS era back to the more distant past would have been problematic, remarkably, the union premium from CPS data is a decent estimate for the entire post-1935 period. The union premium is larger for the less-educated (for every additional year of education, the union premium falls by roughly four log points), an effect that is also remarkably stable over our entire sample period. Finally, we observe a similarly constant pattern for the ratio of *residual* income variance in the union sector to that in the non-union sector (roughly 0.60 over our entire sample period), suggesting that not only have unions consistently reduced the differentials paid to observed characteristics like education, but unions have also had a similar (and stable) effect on the return to unobservable traits.

Together, the *U*-shape in selection by skill and the constant patterns in union premia suggest that during the Great Compression period, unions were conferring a substantial advantage to what would otherwise have been low-income households, thus compressing the income distribution. Importantly, these stylized facts also relate to the potentially endogenous relationship between unions and SBTC posited in recent papers. These papers model union density as an explicit function of SBTC and predict certain co-movements in density, the union premium and selection into unions by skill. For example, Acemoglu *et al.* (2001) predict that skilled workers will differentially exit unions as SBTC increases, as their non-union market wages are bid up, whereas our data show the opposite pattern. Similarly, this model also implies that the union premium falls as SBTC rises, as the scope for rent-sharing between high- and low-skill workers narrows, a prediction that is also rejected by our data. While union membership undoubtedly responds to market forces, we argue that, in light of these new stylized facts, such forces alone cannot explain patterns in union membership over the 20th century.

In our second set of results, we move beyond documenting stylized facts consistent with the role of unions depressing inequality, and instead more explicitly model inequality as a function of union density. We begin by estimating distributional regressions following Firpo *et al.* (2009), where we show how measures of inequality vary as union density is perturbed. This exercise is most directly related to the stylized facts we document on selection and the union premium. Roughly speaking, for each year of our data, we estimate the effect of unions on the unconditional household income distribution, accounting for the changing position of union households in the counterfactual non-union income distribution as well as the union income premium. Across our eighty-year sample period, we find a consistent negative effect of reweighting the full income distribution toward the union income distribution on both the

Gini coefficient and the 90/10 ratio. As would be expected given the changes in selection documented earlier, the negative effect of unions on inequality is especially large in the Great Compression period, when unions were organizing the most negatively-selected workers.

Next, we turn to regression analysis where instead of micro data we employ annual aggregated data and include union density as an explanatory variable. We begin by simply adding union density to the canonical regressions estimated by Katz and Murphy (1992), Autor *et al.* (2008) and Goldin and Katz (2009), who use aggregate time-series regressions to show that the supply of skilled workers is a strong, negative predictor of the college-wage premium. We then refine our time series analysis by introducing variation at the state level. This parallel analysis, aggregated at the state-year level, is carried out controlling for state and year fixed effects. We are able to implement this analysis only because the Gallup data always include a state-of-residence variable and thus allow us to calculate union density at the state-year level (existing measures of union density before the CPS are only consistently available annually at the national level). In both the annual and the state-year regression analyses, union density has a negative effect on standard measures of inequality such as the skill premium, the 90/10 ratio, the Gini coefficient, and the top-ten-percent income share. While each of these exercises depends on a different set of (admittedly strong) identifying assumptions, they each yield a negative and significant effect of union density on measures of income inequality, in many cases comparable in magnitude to and independent of the effect of skill shares.

The rest of the paper is organized as follows. In the next section, we discuss the existing literature on inequality and unions and our contribution. Section 3 describes our data sources, in particular the Gallup data. Section 4 presents our new time-series on household union membership. Section 5 analyzes selection into unions, focusing on education and race. Section 6 estimates household union income premiums over much of the 20th century, and Section 7 presents our evidence on the effect of unions on the shape of the overall income distribution. Section 8 offers concluding thoughts and directions for future work.

2 The Economics of Unions and Inequality

Economists' interest in the effects of unions on workers goes back at least to Adam Smith, who noted asymmetries between combinations of employers versus workers, and John Stuart Mill, who claimed that in the presence of imperfections "trade unions...are the necessary instrumentality of [the] free market."⁴ The literature on unions is vast, and below we provide a selective review of those papers most directly related to the relationship between unions

⁴ See Mill (2008) [p. 319 (original publication year 1848)].

and inequality.

The modern empirical neoclassical approach towards measuring the effect of labor unions was pioneered by Lewis (1963). This literature has largely focused on estimating the union premium: the differential in wages between union and (otherwise similar) non-union workers. 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 0.10 to 0.20 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.⁵

The underlying market structure is important for interpreting the welfare impacts of a significant union premium. A classical view is that labor unions are monopoly sellers of labor to the firm, lowering employment in exchange for higher wages at firms that face otherwise competitive conditions (Rees, 1963). Unions thus create deadweight losses, and their decline should be associated with greater efficiency. Another view is that unions induce firms to share with workers the rents from product market power. Globalization, technological change, and deregulation all may have lowered rents available to be redistributed, but there could still remain a role for unions to change the split of the remaining rents between workers and firms (as in Abowd and Lemieux (1993), Rose (1987) and Abowd (1989)), potentially mitigating increases in inequality induced by market forces, and with ambiguous welfare implications. Consistent with this older rent-sharing literature, a recent literature has documented both pervasive firm- and establishment-specific wage premia (Song *et al.*, 2015; Card *et al.*, 2013; Barth *et al.*, 2016), with unionized firms having larger premia (Hirsch and Müller, 2018). Further, if, as a recent and rapidly growing literature has suggested (see Manning (2011) and Naidu *et al.* (2018) for recent surveys), labor markets are monopsonistic, then union bargaining could raise worker wages to their efficient level (as the quote above from Mill suggests).⁶ While we lack the data to distinguish these interpretations of the union premium, they should be kept in mind throughout our paper.

A key challenge in this literature is separating any causal effect of union membership

⁵As noted, there is very limited work examining the pre-CPS era. Callaway and Collins (2016) is a related recent paper that uses detailed microdata from a survey of six cities in 1951 to estimate a union premium comparable in magnitude to what we find in Section 6. Lewis (1986) is an earlier survey of the literature on union premia, with the earliest micro estimates from the late 1960s.

⁶A recent paper by Benmelech *et al.* (2018) shows that local employer concentration has increased considerably across labor markets over the last forty years. Moreover, this increased concentration is associated with significantly lower wages *except* in highly unionized labor markets.

on wages from non-random selection into unions based on counterfactual non-union wages. 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. Card (1996) uses panel data (allowing him to examine workers as they switch between the union and non-union sectors) to show that the union premium remains significant even after accounting for negative selection at the top and positive selection at the bottom.⁷ Recent identification of the effect of unions on firms and workers has used discontinuities in union-recognition election results in U.S. firms (DiNardo and Lee, 2004; Lee and Mas, 2012; Frandsen, 2013).⁸

A related empirical literature focuses more directly on the connection between unions and the shape of 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 (so, put aside the selection issues raised in the previous paragraph and potential spillovers to the wages of non-union workers). 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 no-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 *et al.* (1996) and Firpo *et al.* (2009) take this approach and find that unions substantially reduce wage inequality, especially for men.

A second category of papers argues that unions affect the wages of *non-union* workers as well. Unions can raise non-union wages via union “threat” effects (Farber, 2005; Taschereau-Dumouchel, 2015) or by the setting of wage fairness norms 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 *et al.*, 2000; Frydman and Saks, 2010) and the returns to capital (Abowd, 1989; Lee and Mas, 2012; Dinardo and Hallock, 2002), thus reducing inequality by

⁷Lemieux (1998) performs a similar exercise using Canadian data, with the added advantage that he can focus on *involuntary* switchers. He finds that estimates that account for non-random selection are in fact quite close to OLS estimates of the union premium.

⁸Although see Frandsen (2017), who cautions against validity of this IV by showing that unions disproportionately lose (win) “close elections” in years when Republicans (Democrats) have a majority on the National Labor Relations Board.

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). Given these diverse mechanisms, the effect of union membership on the income distribution might be larger or smaller than that implied by micro-data analysis of the union wage premium and selection into unions, a point we return to in Section 7.

A final group of papers of particular interest to us is a small macro-theoretical literature on the relationship between unionization, technology, and inequality. These papers argue that any empirical relationship between unions and inequality is spurious, driven by variation of an omitted variable—skill-biased-technological change—that simultaneously reduces union density and increases inequality. For example, in Acemoglu *et al.* (2001), skilled workers earn more in the non-union sector relative to the compressed union sector, and this wedge grows with SBTC. In Açıkgöz and Kaymak (2014) both this mechanism as well as selection of firms into the union sector are at play, and increases in SBTC contribute to both high-skill workers not wanting to be in unions as well as firms being increasingly unwilling to hire low-skill workers at the union wage, as in Farber (1983). These papers both predict that skilled workers should opt out in increasing numbers as the skill-premium increases, lowering union density. While the opting out of skilled workers is an equilibrium result in Acemoglu *et al.* and Açıkgöz and Kaymak, Dinlersoz and Greenwood (2012) rule out the unionization of skilled workers by assumption (they argue that skilled workers are too diverse for unions to organize). Thus, union density in their model declines as SBTC rises because unskilled unionized workers become relatively more expensive to employ and unions have, by construction, no option to organize the more skilled.

We see three key contributions in extending micro-data analyses of unions back to the 1930s. First, economists’ understanding of the basic economics of U.S. labor unions—the size or stability of the union premium, selection into unions by education or other proxies for non-union wages, differences in residual wage variance between the covered- and non-covered sectors—relies almost entirely on CPS data and is thus limited to 1973 and later. We use our new micro-data to examine these stylized facts going back to 1936 (Sections 5 and 6). Importantly, tracing out how selection and the union premium varies during the decline, at the nadir, and then during the rise of U.S. income inequality sheds light on whether unions are a plausible factor in explaining the time-series pattern of inequality. Second and related, the recent macro-theoretical models described above have distinct predictions for the co-movement of union density, skill composition, and premia that could only be tested in a limited way with CPS-era data. In fact, our data reject many of these predictions, suggesting that union density is not merely an artifact of skill-biased technological change. This finding

helps motivate our third contribution, in which we model inequality as an explicit function of union density (Section 7) in regression analyses.

3 Newly Available Data on Union Membership

As noted in the introduction, one of the main contributions of the paper is the introduction of newly available household-level data that include information on union membership. We draw much of these data from public opinion polls, which have recently been posted online by the Roper Center at Cornell.⁹

Polling has a long history in American life. The earliest systematic polls were conducted by magazines, in particular *Literary Digest*, which would include a returnable postcard with opinion questions to conduct “straw polls” on the issues of the day (Igo, 2007).¹⁰ Beginning in the late 1930s, George Gallup, Elmo Roper, and Archibald Crossley began importing techniques from market research into the domain of public opinion polling.

Gallup established the American Institute of Public Opinion (AIPO) and set out to conduct nationwide surveys of American opinions on a range of social and political issues.¹¹ Gallup was scrupulously non-partisan, never running polls on behalf of a particular party. AIPO also devoted considerable efforts to develop neutral, easy to understand question wording. By 1940, about eight million people had encountered Gallup’s tri-weekly polling report, *America Speaks!* Gallup and other pollsters made money by selling their results to businesses for consumer research and newspapers for public opinion.

3.1 Gallup Methodology Before and After 1950

Before 1950, Gallup used so-called “quota-based” sampling. Survey-takers had to fill quotas for each pre-determined strata thought to capture distinct political views. Enumerators were given both hard (e.g., gender, must have one-third female) and soft (e.g., age, “get a good spread”) quotas, but within each quota, interviewers had a lot of discretion. As Berinsky (2006a) notes, “interviewers preferred to work in safer areas and tended to question approachable respondents,” which likely led to Gallup over-sampling, within each quota strata,

⁹See <https://ropercenter.cornell.edu/ipoll-database/>.

¹⁰The Earliest *Literary Digest* poll we could find referenced was a poll to predict the outcome of the 1916 presidential election.

¹¹Similar organizations were formed at roughly the same time: Roper’s company was steadily employed by Fortune magazine starting in 1935, Henry Cantril started the Organization of Public Opinion Research (OPOR) in 1940, and the University of Chicago’s National Opinion Research Center (NORC) was founded in 1941.

more prosperous and well-off respondents.¹²

Gallup once noted that the “the voting public...is the universe of the opinion researcher,” suggesting his aim was to be representative of *voters*, which implies substantial underrepresentation of certain segments of the population. Presumably because the South had low turnout (given many of its elections during this time did not even manage a Republican challenger), it was under-sampled. Southern blacks were differentially underrepresented among Southerners, consistent with their near total disenfranchisement during this period. Gallup purposely over-sampled men because of a belief that women merely adopted their husbands’ opinions on Election Day.¹³

Consistent with discretion within the quota-based sampling leading to oversampling of the well-to-do, Gallup over-predicts the Republican vote share in 1940 and 1944, though in both cases he still correctly predicts Roosevelt victories. In 1948, this over-sampling of Republican voters leads him to incorrectly call the election.

From 1950 onward, Gallup uses modern-day probabilistic sampling procedures. Weights are often provided, but their documentation is not consistent. As a result, in our analyses of the Gallup data we use weights that we generate from the Census, as detailed in Appendix C. Before 1942, our weights adjust Gallup demographics at the *region* $\times$ *race* level. From 1942 onward, we adjust at the *region* $\times$ *race* $\times$ *education* level (because, before 1942, education variables are not consistently available in Gallup).

Appendix Table A.1 and A.2 compare *unweighted* Gallup data with their closest Census year. We provide greater detail for the 1940 comparison (when Gallup is thought to be at its worst in terms of representation of the full adult population). As expected, Gallup over-samples the educated and substantially undersamples women, though these biases are far smaller by 1950. As we only have two surveys in 1940 that include education, we also show detailed occupational breakdowns (doing our best to match Gallup and Census categories). Reassuringly, major categories like farming appear to be well represented in Gallup in 1940. Again as expected, the South is undersampled in 1940, though outside the South, the remaining regions are all represented roughly equal to their population and importantly the black share *outside* the South is very similar in Gallup and the Census.

In the coming sections, we focus on results using *weighted* Gallup data, but we show some

¹²Berinsky (2006a) provides great detail on Gallup’s quota-based sampling procedures, from which we draw much of the information in this subsection.

¹³It is worth noting that any oversampling of men is not a substantial problem for our purposes since we are interested in measures of union status and income at the *household* level. Since most men and women are combined in households, particularly in the earlier years, reports of “any union members in the household” and “household income” should not be affected by whether the surveyed individual in the household was male or female.

unweighted results for the sake of comparison (and all other unweighted results are available from the authors). In general, weighting makes little difference to our regression results.

3.2 Additional Checks on Pre-1950 Gallup Data

Given readers may have natural concerns about the earlier Gallup data, we present a few more plausibility checks that we can assess in *non*-Census years. 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). The concordance is reassuring, given that a concern voiced by pollsters at the time was that the surveys were prone to sample the relatively well-off (both due to enumerator discretion and quotas tilted towards voters). We in fact see large shares of Gallup respondents during the so-called “Roosevelt recession” of 1937-38 report that the household head is unemployed (note that levels are not directly comparable as Gallup and HSUS use different definitions of unemployment in the 1930s).¹⁴ Note that, as with most of our results, weighting the Gallup data has little effect.

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.

3.3 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 often we will we augment our Gallup data with the ANES.¹⁵ We also found a handful of additional survey datasets that ask union status as well as the other variables we need to estimate a family income premium (i.e., education, family income, state of residence and basic demographics). The first is an expenditure survey that the BLS conducted in 1936–1937, which we will henceforth refer

¹⁴Gallup unemployment is measured as the share of households (generally the household head but sometimes the respondent) who report being unemployed (including, sometimes, those who are in fact out of the labor force), while the HSUS reports unemployed workers as those who had no employment, were available for work, and had made “specific efforts” to find employment. The HSUS unemployment rate is the share unemployed out of the employed plus unemployed population.

¹⁵The 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.

to as the 1936 Expenditure Survey.¹⁶ The second is a 1946 survey performed by the U.S. Psychological Corporation. 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 micro-data regression analysis. From 1977 onward, we can use the CPS to examine household measures of union membership.¹⁷

Summary statistics for the CPS, ANES as well as for these additional data sources appear in Appendix Table A.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 do not weight the other additional surveys.

3.4 Existing Time Series of Aggregate Union Density

As noted in the introduction, aggregate counts of union members exist going back to the early 20th century. We describe the two main time series—the BLS and Troy series—in greater detail in Appendix D and give only a very brief summary here.

The BLS series is based on unions’ reports of their aggregate membership counts. A key issue, which we describe in detail in the Appendix, is that individual unions had an incentive to overstate membership. Especially from 1937–1955, when the main U.S. labor umbrella organization had split into two warring factions—the American Federation of Labor and the Congress of Industrial Organizations—the two federations over-stated their membership in attempt to gain advantages over the other.

The Troy series was compiled by Leo Troy under the auspices of the NBER. He attempts to correct some of these over-reporting issues in the BLS, but in many years he is forced to rely on the same membership statistics, while in other years he derives membership by dividing total revenue by average dues rates. However, as we discuss in Appendix D, revenue data were also believed to be vastly overstated.

In fact, unions had no obligation to report membership to the government until 1959, when every union involved in interstate commerce had to file reports. From this point onward, these aggregate counts are believed to be more accurate (Troy, 1965).

¹⁶Note that interviews conducted in 1937 refer to consumption and income in 1936. The survey asks about *union dues* as an expenditure category, and thus can be used to measure household union membership.

¹⁷We could in principle go back to 1973 using the CPS, but only beginning in 1977 does the CPS include both the union membership question and state (as opposed to merely state groups or regions) identifiers, and as we are generally controlling for state in all our results, we want to remain consistent with analysis from our other data sets.

4 A New Time-series Measure of Household Union Membership

The typical Gallup union question is “Are you (or is your husband) a member of a labor union?”, with the choices most often being: “neither,” “yes, I am,” “yes, he is,” “yes, both are.” In 1959, “husband” changes to “husband/wife.” In some years, however, the question does not ask which member or members of the household is or are in a union, so we cannot, for example, always measure individual union status. We harmonize these questions to form a measure of *household* union status, where we code a household as union if either household head or spouse is a union member. While technically the implied unit of observation is *couple* (or individual if the respondent is not part of a couple), we will generally refer to this measure as *household union status*. Importantly, Gallup asks this question of *all* respondents, not skipping those in, say, agricultural occupations or who are unemployed.

4.1 The Union Share of Households Over Time

Figure 2 plots our (weighted) Gallup-based measure of the union share of households, by year, alongside several other series.¹⁸ The Gallup series bounces around between eleven and thirteen percent from 1937 to 1940. Between 1941 and 1945, the years the U.S. is involved in World War II, union membership in our Gallup data roughly doubles. Union membership 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, union membership in the Gallup data slowly but steadily declines.

Also presented in Figure 2 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) sit somewhat higher than Gallup estimates for those years.

Finally, Figure 2 plots the widely-used historical data series described in Section 3.4, the BLS and the Troy series.¹⁹ To avoid clutter, we do not plot our CPS series in this figure. Appendix Figure A.4 shows the Gallup and CPS series, zoomed in on the period of their overlap (1977-1986), so that readers can more easily assess their degree of concordance.

¹⁸Appendix Figure A.3 shows this series alongside its unweighted analogue. In most years, the two series look very similar.

¹⁹Recall that 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.

Reassuringly, in the years when Gallup and the CPS overlap, they are quite close.²⁰

4.2 Differences Between Gallup and the Historical Aggregate Series

While the Gallup measures do not always agree with the BLS and Troy series in levels, they are, for the most part, consistent in changes.

The largest differences exist in the early years of the series, with much more agreement in the 1950s, when both union reporting and Gallup sampling improve. As noted in Appendix D, we have reason to believe that both the Troy and the BLS series over-stated membership. While unions had incentives to overstate membership, respondents themselves had no incentive to tell Gallup survey-takers they were union members when they were not, so this bias is unlikely to affect the Gallup numbers. However, there are reasons that Gallup and other opinion surveys may *under-state* union membership relative to these historical series based on union reports.

First, there is a legitimate possibility that individuals are union members without knowing it, especially during certain historical moments, meaning union reports would accurately count them as members but they would (truthfully, but potentially inaccurately) tell Gallup that they were not currently members. During World War II, some unions default-enrolled all new workers and automatically collected dues from workers' paychecks (workers would have to actively take steps to dis-enroll). Workers could thus be members of unions without knowing it. During the war, a period of rapid union growth, the National War Labor Board (NWLB) gave unions this privilege—default enrollment of new workers—in war-related plants, in exchange for a no-strike pledge (Lichtenstein, 2003). It is thus perhaps not surprising that the increase in membership during the war, while large in the Gallup data, is even larger in the union-membership-reports-based series.

Second, and related, 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 (who is no longer paying dues) might well consider himself no longer a member.²¹ Indeed, Figure 2 shows that the Gallup estimate diverges from the BLS and Troy estimates

²⁰Give 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, so we can check consistency of Gallup over a substantial period of time.

²¹As 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.

the most in 1937-38 recession period. Gallup shows essentially no growth during this period, whereas the BLS and Troy show robust growth, despite soaring unemployment.²²

Third, as noted, Gallup over-samples the well-to-do, especially before 1950, which likely biased union membership toward zero. But offsetting sampling biases—most notably the under-representation of the South, a region historically hostile to unions—likely overstate density estimates. While we can never fully discount the possibility that non-representative sampling is causing Gallup to understate density, that we only find a marginal changes in density after applying weights suggests it is not a major factor (see Appendix Figure A.3).

In summary, we are not surprised that the union density measures based on opinion surveys differ slightly in levels from the more widely used measures in the literature, given non-trivial differences in methodology. However, in almost all cases they firmly agree in changes, and, like the Troy and BLS series, Gallup exhibits the same inverted *U*-shape over the twentieth century. Moreover, as we will show in Section 7, the relationship between aggregate union density and inequality is very similar whether we use our new, micro-data-based measures of household unionization rates or the traditional, aggregate measures.

An important advantage of our series, however, is that it is based on micro data, which allow us far greater scope to sort among theories relating inequality to unionization. It is to this task we now turn.

5 Selection Into Unions

The analysis in this section exploits a key feature of our microdata sources: they allow us to examine the joint distribution of unions and other covariates over time. Of particular interest is the relative skill of union members relative to non-union workers. First, how selection into unions varies across time is informative about the plausibility of the claim that unions have played an important role in compressing the income distribution. For example, suppose we find that the expansion of unions from, say, 1940 to the 1960s was all among relatively *high*-skill workers. Then it would cast doubt on the claim that unions contributed to the mid-century Great Compression in wages, as they would be conferring an advantage to already highly-paid workers. Second, as discussed in Section 2, some existing models of unionization, technology and inequality predict that the relative skill of union members will decline as union density declines and rise as union density rises. Plotting selection by skill across time allows us to test these predictions.

²²Indeed, it is well documented that at least among the largest locals where data are available, *dues payments* plummeted for CIO unions during this period, as millions of workers were laid off (Lichtenstein, 2003).

We take two approaches to examining the relationship between selection into unions and union density. We begin by estimating a basic union selection model (regressing a union indicator on demographics and skill proxies) by year, and reporting the coefficient on skill proxies across time. This approach allows the data to “speak for themselves” in a transparent manner. If there is indeed a positive (negative) relationship between union density and the relative skill of union members, then these coefficient estimates across time should take an (inverted) *U*-shape, as union density across time is itself inverted-*U*-shaped. We then take a more parametric approach and include density as an explanatory variable in a union selection regression model.

5.1 Selection Into Union Across Time

We begin our education selection analysis by estimating the following equation by survey-source d and year y :

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

In this equation, subscripts h , s , and t denote household, state and survey-date, respectively. The superscript R serves to remind readers that in many cases, a variable refers specifically to the *respondent* (not necessarily the household head, though via assortative mating by race, age and education, the respondent, if she is not the head herself, will often have similar demographic characteristics). $Union_{hst}$ 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.²³ $Female_h^R$ 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.²⁴

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 education predicts whether you live in a union household, conditional on basic demographics and state of residence.²⁵ Note here that we are not yet

²³Where a specific survey does not collect information directly on years of education but reports specific ranges or credentials, we use simple rules to convert these measures to years of education. The note to Figure 3 describes how we impute years of education in these cases. We also carry out some analyses with alternative measures of education (e.g., a high school dummy) to probe robustness.

²⁴Year fixed effects are effectively included in the model in the form of the *survey-date* fixed effects (the ν_t).

²⁵For the ANES, given the small sample sizes, we constrain the coefficients on education (β_{dy})

controlling for race.

Figure 3 shows these results across our key datasets. A clear *U*-shape emerges, with the year-specific point-estimates remarkably consistent across all data sources.²⁶ In the earliest years (1936 through the mid-1940s, say) the coefficients suggest that an additional year of education reduces the likelihood of living in a union household by about 2 to 2.5 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 four 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 statistically significant, though small.

This change may reflect, in part, the substantial growth of public sector labor unions since the 1960s. Since at least the early 1970s, the public sector workforce has had higher education than the private sector workforce. And, importantly, unions have been much more successful organizing workers in the public sector than in the private sector.²⁷ 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 skill of union households from 1970 onward is unchanged.

Another possibility is that the steep decline in manufacturing employment since the 1960s drives the most recent half of the *U*-shape in our selection results. Appendix Figure A.8 also shows selection patterns after dropping households with *either* a public-sector or a manufacturing worker. A large majority of the up-skilling effect remains.

As a final robustness check on our education selection results, we experiment with different weighting schemes. As noted in Section 3, sampling before 1950 in the Gallup data may be suspect (and as we showed in Appendix Table A.2, especially suspect in terms of over-sampling the educated, at least when we compare them to Census estimates in 1940). However, Appendix Table A.4 shows that selection by education in Gallup before 1950 barely changes no matter how we weight (including not weighting at all).

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.

²⁶This pattern holds when other education measures are used instead of years of schooling. Appendix Figures A.5, A.6, and A.7 show similar patterns when log years schooling, a high-school dummy or college dummy serve as the education measure.

²⁷Over 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.

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. 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 during the pre-World War II and Great Compression period, with some arguing that unions exacerbated white-black wage gaps and others arguing that, at the very least, unions were less discriminatory than the non-union labor market more generally during this period (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).²⁸ We start by examining selection by race *unconditional* on education, because Gallup always asks about race but only beginning in 1942 does it consistently ask about education and we want to maximize the time period over which we can estimate selection by race. We return to selection by race conditional on education later in the section.

The estimated coefficients on *White* across time and data sources are presented in Figure 4. Again, a *U*-shape emerges. In the beginning of our sample periods, whites are (conditional on our covariates) more likely to be in union households than non-whites. This advantage disappears and in fact reverses during the war years and continues to grow more negative until about the 1960s. Since then, whites gain on non-white households until today there is essentially no differential.

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 the ANES and Gallup in the 1950s and 1960s, though the large standard errors around the ANES estimates certainly do not allow us to reject equality. There is some disagreement between Gallup and CPS, whereby Gallup shows minimal negative selection with respect to education by the early 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.

Appendix Figure A.9 replicates this analysis, but excludes the South, given that Gallup significantly under-sampled blacks in the region at least through 1950 and its correction of this bias over time might make comparisons of union selection by race across time difficult.

²⁸Results 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.

In fact, the results are very similar.

5.2 Modeling Selection as a Parametric Function of Density

So far, we have shown that, graphically, selection into unions by proxies for skill takes an U shape from 1936 to today, the inverse of the shape of the union density time-series. In this section, we pool our data and include union density itself as a regressor. Relative to the graphical evidence, this analysis allows an easier summary of the correlation between density and selection and also facilitates robustness checks. Robustness checks are especially important because other key trends over the twentieth century are U -shaped, so any relationship between density (inverted- U -shaped) and negative selection (U -shaped) could easily be spurious and driven by some third factor.²⁹

We begin by examining selection by education, modifying equation (1). Here, we pool all years of CPS and Gallup data (results are unchanged if we also add ANES and our one-off datasets) and estimate variants of the following equation:

$$\begin{aligned} Union_{hst} = & \beta_1 Educ_{hst}^R \times Density_t + \beta_2 Educ_h^R + \\ & \beta_3 Educ_h^R \times Year_t + \beta_4 Educ_h^R \times Year_t^2 + \\ & \gamma_1 Female_h^R + \gamma_2 White_h^R + f(age_h^R) + \nu_t + \mu_s + e_{hst}, \end{aligned} \quad (2)$$

where $Density_t$ is aggregate union density in year t as estimated in our Gallup/CPS data and all other notation follows naturally. Including the $Educ_h^R \times Year_t$ and $Educ_h^R \times Year_t^2$ terms allows selection by education to take an arbitrary U -shape of its own. Note that, given the evidence we just documented that selection into unions by race is also U -shaped, we include a control for race in all specifications, as well as show separate specifications with race interacted with year and year squared. As noted, a surprising number of trends appear to take a U or inverted- U shape over the period we study. This specification allows us to examine whether there actually is a direct relationship between union density and the skill and demographic composition of union members, or is composition following, perhaps coincidentally, its own inverted U pattern or being driven by some other U -shaped trend.

To provide a baseline, col. (1) of Table 1 estimates equation (2) with the restriction that $\beta_3 = \beta_4 = 0$ (so, not allowing selection by education to follow its own U -shape). The coefficient of interest on $Educ_{hst}^R \times Density_t$ is large and negative, as we would expect from

²⁹Two important examples of U -shaped trends over twentieth-century America are the share foreign born and political polarization in Congress.

the graphical evidence in the previous subsection. Col. (2) allows selection by education to follow its own U -shape, and indeed both of these terms are significant and of the expected sign. As we would expect given the collinearity, the coefficient on $Educ_{hst}^R \times Density_t$ falls (by just over half) but remains large, negative and highly significant. As a final check, in col. (3) we drop any household in the CPS that has a public-sector worker, and the coefficient of interest is largely unaffected.³⁰ Col. (4) shows that higher union density is associated not only with more unskilled workers joining unions, but also with more non-whites joining (even allowing selection by race to follow its own, arbitrary U -shape), though the final column of this table shows that this relationship is sensitive to dropping public-sector households. As a further robustness check, we confirm in Appendix Table A.5 that *state-year* union density also predicts greater unionization rates among the less-educated and minorities, controlling for both state and year fixed effects.

We conclude from this analysis that, as unions expand or contract, the marginal union member is relatively low-skilled. This is counter to the predictions of some SBTC-based models of deunionization driven by increased self-selection of high-skilled workers out of unions. Finally, we reiterate that much of the prominent U -shape in Figures 3 and 4 appears correlated with variation in union density (which itself has an inverted U -shape) and is not merely following some arbitrary U -shape of its own.

6 The Union Family Income Premium Over the Twentieth Century

The analysis in Section 5 shows that union households were much lower skilled during the period of peak union density than either before World War II or in recent periods. We now use the survey microdata to estimate union premia, both on average and by skill level over over nine-decade period. Given that almost all evidence on the size of the union premium is from the CPS era, estimating union premia from earlier periods is interesting in its own right.

Measuring the premium over longer periods of time is important for at least two other reasons. First, it helps us assess the plausibility of the argument that unions reduce inequality: for example, if the union premium was much smaller in the immediate post-war period, then it would be harder to argue that unions were conferring advantages to otherwise low-wage

³⁰We cannot drop public-sector workers in Gallup, as Gallup does not designate sector. However, given the rise of public sector unions is not until the 1970s and in these regressions we only use Gallup observations until 1976, the vast majority of unionized households in our Gallup data are working in the private sector.

workers during the Great Compression period and beyond, as we claim. Second, some models of SBTC-driven union density imply that the union premium should increase as density grows and decline as density falls, a prediction we can test with our data.

As in the previous two sections, we draw from a variety of data sources. From Gallup, we can estimate a premium in 1942 and then each year beginning in 1961 (though we drop any year in which the top-coded category contains more than thirty percent of observations). Gallup occasionally asked household income before 1961, but only in 1942 do we have a survey that also includes the union question and education. Our other major data source is the ANES. We can estimate a household union income premium in the ANES in 1952, and then every other year after 1956.³¹

As noted in Section 3, we supplement these surveys with any other survey we are able to find that asks union membership, household income and education.³² Recall that these supplemental surveys are the 1936 Expenditure Survey and the 1946 U.S. Psychological Corporation survey.³³

6.1 Estimating the Union Family Income Premium Over Time

Across all these surveys, we are able to estimate the following regression equation (as with selection, we estimate it separately by data source d and year y):

$$\begin{aligned} \ln(y_{hst}) = & \beta_{dy} Union_h + \gamma_1 Female_h^R + \gamma_2 Race_h^R + f(age_h^R) + \\ & g(Employed_h) + \lambda_h^{eduR} + \nu_t + \mu_s + e_{hst}. \end{aligned} \quad (3)$$

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

³¹As seen in Figure 2, we can estimate union density in the ANES in 1948, but that survey does not include household income.

³²In fact, the real constraint are surveys that ask both household income and union status. Almost all surveys we have found that ask for household income also ask respondents' education.

³³The 1936 Expenditure Survey is the only one in which income is not binned. To make it comparable to our other data sources, we winsorize at 95% and 5% percent and then create 15 equally sized income bins. We discuss the details of the binned income variables in Appendix B.

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 below we show that this bias should be small.

As with our selection results in the previous section, we show premium results separately by survey source and year in Figure 5. Remarkably, the household income premium has not appreciably changed from 1936 to today, holding stable at roughly fifteen to twenty log points. While not a perfectly flat line, given the standard errors around each estimate, the family union premium does not appear to follow any discernible pattern over time. Of the sixty-some point estimates, only a handful are greater than 0.20 or less than 0.10. For none of these estimates do the confidence interval intersect zero.

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 is constant over a *nine-decade* period that saw density (as well as the degree of negative selection) both increase and then decrease. This stability of the premium is not predicted by the existing models in which SBTC determines union density rates, as in those models 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 (2012). 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 constant premium in the face of declining density. We flag this question as a potentially fruitful area for future research.

6.2 Robustness and Related Results

While we do not focus on them, Appendix Table A.6 shows the coefficients on the Mincer equation covariates in equation (3). For ANES (since it covers a long time period), we split the sample in two, so that coefficients can be compared across time. In all cases, the coefficients on the covariates have the same signs as we typically see from an individual earnings regression.

In Appendix Figure A.10, we show results after controlling for occupation of the household head. As noted, occupation categories vary considerably across survey sources, which is why we relegate this figure to the Appendix. The appendix figure reports coefficients that are somewhat larger than in the main Figure 5, 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.11 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.³⁴

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 *et al.*, 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 in Appendix Table A.7, 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.8, union households are far more likely to say it would be hard for them to find a job just as good. 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 felt their jobs were better—in a broad sense—than non-union members.

Our estimates of a sizeable and relatively stable union premium contrast with recent papers using NLRB regression discontinuities in close representation elections to estimate the causal effect of unionization on firm-level outcomes (DiNardo and Lee, 2004; Frandsen, 2013). These papers have found little evidence of positive union premia. 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

³⁴Union 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.

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 identify are appropriate when applied broadly.

6.3 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 education and race (whites versus non-whites).

We begin by augmenting our family income equation (3) by adding an interaction term between years of schooling and the household union dummy. Figure 6 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 when instead of adding $Years\ of\ educ_h^R \times Union_h$ to equation (3) we add $White_h^R \times Union_h$ are shown in Figure 7. 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.

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 underestimate the effect of unions on inequality, especially during the Great Compression period and beyond (1940-1990 or so) when the differential premium for black households appears largest.

6.4 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 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 (3). 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 8 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 in fact holds in a very similar manner throughout the post-1936 period.³⁵

7 The Effect of Unions on Inequality

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 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 is not constant across time. In the Great Compression period, 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 pieces of evidence suggest, at least indirectly, that unions were a powerful force pushing to lower income inequality during the heyday of the labor movement. In this section, we explore in a more direct manner the relationship between unions and income inequality.

³⁵For 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.

First, and most directly related to the results in the previous two sections, we estimate distributional regressions following Firpo *et al.* (2009), where we show how measures of inequality change with small increases in union density. This regression 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, 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).

Second, we turn to more aggregate analysis. We follow some of the canonical work on the effect of skills shares on the college premium, and add 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. Finally, as the Gallup data always has state identifiers (thus allowing us to create, for the first time, state-year level density measures going back to the 1930s), we also conduct a parallel analysis at the state-year level.

The next three subsections describes the implementation of and results from each of these exercises. Note that each asks a slightly different question, so we would expect magnitudes to differ. We defer a more detailed comparison of magnitudes to Section 7.4.

7.1 Distributional Regressions

In this section we present the effects of union membership on inequality using recentered influence functions (RIF) as in Firpo *et al.* (2009). We provide a brief overview of the methodology below, with greater detail provided in Appendix E. In our case, the RIF is a transformation of the dependent variable that allows the coefficient on the household union dummy to recover the effect of union density on some aggregate statistic of the family income distribution. We present effects of union density on the 90-10 ratio of log wages as well as on the Gini coefficient.

The virtue of RIF regressions is that they provide a tractable method to estimate the marginal effect of a variable on a statistic of the whole distribution. Let the joint distribution of family income y and covariates X be $F(y, X)$, and let the probability of a household being a union household be represented by $p \in [0, 1]$, so that $\mathbb{E}(\text{union}) = p$, where *union* is our usual household union dummy. Then, under the assumption of no spillovers, we can decompose the family income distribution into the $\text{union} = 1$ and $\text{union} = 0$ distributions so that

$$F(y, X) = p \cdot F(y, X \mid \text{union} = 1) + (1 - p) \cdot F(y, X \mid \text{union} = 0). \quad (4)$$

Note that the assumption of no spillovers is incorporated into this expression by making the union and non-union joint distribution of income and characteristics independent of the share of the population that is unionized.

Consider some distributional statistic $\nu(F)$, such as the difference between the 90th and 10th percentile, $Q_{90} - Q_{10} = F^{-1}(\frac{90}{100}) - F^{-1}(\frac{10}{100})$ or the Gini coefficient $Gini(F)$. We can write the effect of a small change in union density on this statistic as $\frac{d\nu(F)}{dp}$. Firpo *et al.* (2009) show that the value of this derivative is obtained from a regression of the RIF of the statistic ν on the independent variable of interest. As with the selection and union premium results in the previous two sections, we run this estimation separately by survey source d and year y , using a specification parallel to equation (3) but with the RIF of a given inequality statistic ν instead of family income as the dependent variable:

$$RIF(y_{hst}, \nu(F)) = \beta Union_h + \gamma_1 Female_h^R + \gamma_2 Race_h^R + f(age_h^R) + g(Employed_h) + \lambda_h^{eduR} + \nu_t + \mu_s + e_{hst}. \quad (5)$$

In our setting, the RIF of a given statistic is, roughly speaking, the derivative of ν as the distribution in expression (4) is slightly perturbed toward the $union = 1$ distribution (i.e., as p increases slightly), an expression that must be derived for each statistic of interest.³⁶

Figure 9 shows the effect of moving the family income distribution toward that of union families on the difference between the 90th and 10th percentile of log family income. Interestingly, while noisy, these estimates also display a *U*-shape over the 20th century. Given the constant income premium and *U*-shaped selection documented in the previous two sections, this pattern is unsurprising; when unions drew from the lower skilled, a constant wage premium would compress the wage distribution more. The implied effect on the 90-10 ratio varies from roughly -0.5 in the 1960s to -0.2 in earlier years as well as the CPS period.

In Appendix Figure A.12 we show the effects separately by the tenth, fiftieth and ninetieth percentile. Effects are uniformly positive for the 10th percentile, and negative or close to zero for the 90th percentile. The effect on the median is always positive, but depending on the data source and time period is sometimes above and sometimes below the effect on the tenth percentile. As we discuss in Appendix E, some of these differences may arise because of variation across data sources and across time in how coarsely the income data are binned, and we show robustness to changing bin sizes.

We conduct a parallel analysis for the Gini coefficient, which summarizes changes in inequality coming from all parts of the distribution, though, relative to the 90/10 ratio, it

³⁶For the sake of illustration, we derive the RIF of the 90/10 ratio in Appendix E.

is especially sensitive to changes in the middle part of the distribution. Figure 10 shows the effect of union density on the Gini coefficient in each year. For the bulk of the sample period, the effect is quite stable, roughly -0.075, with magnitude falling as we approach the present, again, consistent with the up-skilling of union households since the 1960s.

7.2 Time-series Regressions

While the distributional regressions capture the effect of union density on 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.

We begin our aggregate analysis of the effect of unions on inequality by adapting regressions from the time-series literature on the college wage premium, such as the analysis in Goldin and Katz (2009), which spans the whole 20th century. Following Katz and Murphy (1992) and 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 (2009) 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 *et al.* (2008) confirm this analysis using data from the CPS in the 1963-2005 period and adding more covariates, and argue that the non-market factors stressed by Card and DiNardo (2002), Lee (1999), and Lemieux (2006) have limited explanatory power in explaining the rise of inequality, measured as the 90-50 or 50-10 ratios. However, these papers do not consider unions as a potential factor in their analysis.

Motivated by these earlier papers, we begin by estimating the following equation:

$$\log\left(\frac{wage_t^{Col}}{wage_t^{HS}}\right) = \beta UnionDensity_t + \gamma \log\left(\frac{N_t^{Col}}{N_t^{HS}}\right) + f(t) + \lambda X_t + \epsilon_t. \quad (6)$$

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, $f(t)$, (which is meant in the SBTC literature to pick up changes in relative demand for skill, although in our context it could also pick up unmeasured institutional factors driving the college premium), other time-series controls X_t , which we vary to probe robustness, and, importantly, $UnionDensity_t$.³⁷

We choose time-series controls X_t both to follow past literature as well as to capture

³⁷As 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.

the most obvious confounds in estimating the effect of unions on inequality. Specifically, following Autor *et al.* (2008) we include the real value of the federal minimum wage and the civilian unemployment rate and following Piketty *et al.* (2014) we include the top marginal tax rate in the federal individual income tax schedule. We adjust for heteroskedasticity and AR(1) serial correlation in the error ϵ_t using Newey-West standard errors.

The first two columns of Table 2 show the results from this exercise. Col. (1) does not include additional controls X_t , whereas col. (2) does. While we will defer a more detailed discussion of magnitudes until Section 7.4, we note here that in both specifications, the coefficient on union density is negative and highly significant (and very similar to each other in magnitude). Appendix Table A.9 shows a variety of alternative specifications (e.g., using the Gallup series alone or the BLS series alone to calculate $UnionDensity_t$, instead of averaging the two together, substituting either a quartic or a quadratic for the cubic time polynomial) and also reports more of the coefficients, which we suppress in the main tables in the interest of space.

While the canonical analysis in Goldin and Katz (2009) and related work focuses on the college premium, we extend our analysis in Table 2 by using the same specifications as in cols. (1) and (2) but using other measures of inequality as outcomes. Cols. (3) - (4) of Table 2 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 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. Appendix Table A.10 is the analogue to A.9 and shows the result is robust across alternative specifications.

Cols. (5) - (8) of Table 2 extend this analysis to inequality measures constructed from administrative data (rather than surveys). These have the advantage of being available annually, instead of just every ten years from the Census in the pre-CPS era.³⁸ 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. (5) and (6) use the Gini coefficient constructed by Kopczuk *et al.* (2010) from Social Security data. Without controls, the relationship is negative and highly significant, whereas adding controls decreases the magnitude by roughly one-fourth and significance is

³⁸A small complication in using these annual outcomes is that our estimates of the skill shares $\log(N_t^{Col}/N_t^{HS})$ in equation (6) come from survey data, and thus in principle are only available every ten years in the pre-CPS era. 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.

reduced to the ten-percent level. The final two columns uses the top-ten-percent income share from Piketty and Saez (2003). Without controls, the union coefficient is negative and highly significant. But this result is the most sensitive to controls, with the coefficient falling by half and losing its significant. The usual robustness checks for the Gini and the top-ten-percent-share outcomes are in Appendix Table A.11 and A.12, respectively.

While there are clear limitations to the time-series analysis (e.g. no exogenous variation in union density, and small samples make inference suspect), all specifications in this section control for a cubic in time (and we show results using a quadratic or a quartic in the Appendix Tables), ensuring that the robust negative relationships we find between unions and our measures of inequality are not simply mirroring the *U*-shape present in many time-series over the 20th century.

7.3 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 is that there are many unobserved factors (e.g. macroeconomic policy, trade, outsourcing, industry structure) that are likely correlated with both inequality and union density that are not 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.³⁹ One limitation of our survey-based data is that small states get small samples, resulting in noisy estimates for annual variation. We use both winsorized measures as well as a split-sample instrumental variables strategy to mitigate this problem. Although we do not have exogenous variation in union density, we can see if the inverse inequality-density relationship that holds in the aggregate time series hold at the state-year level, conditional on year and state fixed effects.⁴⁰ Importantly, as all states have access to the same national technology, the 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 densities, with weights proportional to the number

³⁹Troy (1965) presents state breakdowns for 1939 and 1956, and Hirsch *et al.* (2001) use BLS reports to construct state-year measures of density from 1964 onwards.

⁴⁰Similar regressions are estimated at the cross-country level by Jaumotte and Osorio (2015), though their sample period of 1980-2010 is far shorter than ours.

of observations in each sample (so the CPS gets a much larger weight). This procedure results in a panel of state-year union density measures going back to 1937. Note that such a 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.

To examine the effect of unions on inequality, we closely follow equation 6 and estimate specifications of the form:

$$y_{st} = \beta \text{UnionDensity}_{st} + \gamma \log\left(\frac{N_{st}^{Col}}{N_{st}^{HS}}\right) + \lambda X_{st} + \mu_{t \times \text{South}} + \delta_s + \epsilon_{st} \quad (7)$$

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 . As before, we control for skill-shares $\log(N_{st}^{Col}/N_{st}^{HS})$ in all specifications.⁴¹ We include South-by-year and state fixed effects in all regressions, $\mu_{t \times \text{South}}$ and δ_s , respectively. Note that we include South-by-year fixed effects because, as we showed earlier, Gallup’s sampling of the South improves over time. 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 (6), 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 GDP per capita and state-year measures of the share of households subject to the federal income tax, as annual state-level unemployment rates are not consistently available until the 1963 CPS. We do not have information on state tax schedules going back to the 1930s, so to capture the political-economy climate, we control for a state-year level “policy liberalism” index developed by Caughey and Warshaw (2016). 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. Finally, to deal with possible unobserved but smooth state-specific changes in technology or other unobservables that may be confounding the estimated relationship, X_{st} also includes state-specific linear and quadratic trends.

As mentioned above, 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.⁴² For

⁴¹The top-ten-percent share of income is available at the *annual* level, so just as in the time-series regressions we include both the interpolated IPUMS-CPS education measure (at the state-year level) as well as the Gallup measure of education for that outcome (at the state-year level).

⁴²See Angrist and Krueger (1995) for an early description of the methodology, and Inoue and Solon (2010) and Aydemir and Borjas (2011) for more recent exposition and applications, respec-

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 procedure yields the following first-stage equation:

$$UnionDensity_{st}^0 = \eta UnionDensity_{st}^1 - \iota \log \left(\frac{N_{st}^{Col}}{N_{st}^{HS}} \right) + \lambda^f X_{st} + \mu_{t \times South} + \delta_s + \nu_{st}. \quad (8)$$

The second-stage equation in the split-sample IV is merely equation (7) with $UnionDensity_{st}$ replaced by $\widehat{UnionDensity}_{st}^0$, the prediction generated from the first stage. Since $UnionDensity_{st}^1$ and $UnionDensity_{st}^0$ are calculated from a *random* split of the data, the sampling errors in the two measures will be orthogonal. Omitted variable issues aside, if the only issue is measurement error, the IV estimator β^{IV} will yield a consistent, unattenuated estimate of β .

Table 3 shows results from the specification in equation (7) across the state-year analogues of the four main inequality outcomes used in Table 2.⁴³ 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 our three other outcomes (the male 90/10 ratio, the Gini coefficient, and the top-ten-percent shares), the coefficient on state-year density is negative and quite robust to adding additional controls.

In Appendix Tables A.13 to A.17 we show a variety of specifications that add intermediate sets of controls between the odd and even columns reported in Table 3.⁴⁴ 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

tively.

⁴³For cols. (1) to (4), we simply calculate the skill premium and the 90/10 ratio at the state-year level instead of the year level from the same survey data sources. For cols. (5) and (6), as the Social Security data are not disaggregated to the state level, we instead just calculate the household-income Gini from our survey data sources. For cols. (7) and (8), we use recently constructed top income shares at the state-year level from Frank (2015), based on internal IRS data (downloaded from the World Wealth and Income Database Alvaredo *et al.*, 2016). These data are not adjusted for capital gains. Frank (2015) show that when the same methodology is applied to national US data the results are still quite close to Piketty and Saez (2003).

⁴⁴An additional robustness check we perform but do not report is including state-year variation in the passage of right-to-work laws. Including such controls has no effect on our results, and we are pursuing the independent effect of right-to-work laws on union composition in ongoing work. See Feigenbaum *et al.* (2018) for recent work on right-to-work laws, mainly using CPS-era data.

the cost of slowed economic growth (e.g., lowered net business entry). In fact, union density shows consistently positive, but sometimes insignificant, effects on state GDP per capita, and we can rule out even small negative effects of unions on state-level economic activity (see Appendix Table A.17).

7.4 Discussion of Estimate Magnitudes

While each of our three approaches involves strong identification assumptions, all regression results point to negative effects of unions on inequality. How large are these effects and why might they differ across approaches?

In Table 4 we show results from a simple exercise: taking our various coefficients, how much of the *decline* in inequality from 1940-1960 can be explained by the eleven-percentage-point rise in union density, and how much of the *increase* in inequality between 1970 and 2004 can be explained by the twelve-percentage-point decline in density?⁴⁵ We begin with coefficients from the RIF regressions. As we noted, the RIF regression coefficients are larger in magnitude during unions’ heyday relative to the more recent period, because unions were organizing less advantaged households during their Great Compression peak years. We thus consider a “high” and “low” RIF effect, the “high” using the simple average of all the RIF coefficients between 1950 and 1970 (i.e., those plotted in Figure 9 and 10 for the 90/10 and the Gini, respectively) and the “low” using the simple average of coefficients from 2005-2015. Multiplying the “high” coefficients by the rise (decline) in union density between 1940-1960 (1970-2004) can explain a sizeable share of the decline (rise) in the 90/10 ratio and Gini coefficients, roughly between ten and fifteen percent, comparable with the results in DiNardo *et al.* (1996) and Card (2001). By contrast, using coefficient estimates from the recent period would suggest that changes in union density explain less than ten (and in some cases less than five) percent of the changes. An important point that the RIF analysis highlights is that the question “how would inequality change today if union density were to somehow increase by x percent” is not, strictly speaking, well-defined: one would want to know from where in the counterfactual non-union income distribution this hypothetical increase in density would come.

Notwithstanding this point, the rest of the table will assume a constant effect of union density across time. For the aggregate time-series and state-year regressions, we always use point-estimates that include the additional covariates (i.e., the even-numbered columns of

⁴⁵We end the recent period in 2004 as it is the last year for which the Social Security earnings Gini coefficient measures is available. While somewhat arbitrary, we choose 1940 and 1970 as starting points as they are Census years and thus we have college premium and 90/10 outcomes for those years.

Tables 2 and 3). The implied explanatory power of union density is considerable using the time-series coefficients: we in fact substantially “over-explain” the decline in the skill premium from 1940-1960 (potentially consistent with the observation in Goldin and Katz, 2009 that skill shares *under*-explain the fall in the skill premium between 1940 and 1950). Roughly speaking, explanatory power when we use the state-year coefficients is about one-fifth to one-half as large as that from the aggregate time-series. Like the RIF results, results from the time-series and state-year regressions imply that changes in union density have similar power to explain *both* the decline and then the rise of our inequality measures: in general, the magnitudes of the share-explained estimates are similar in the top and bottom panels of the table.

Why does the implied explanatory power of union density vary across our three econometric approaches? As we have emphasized throughout, each of these approaches requires strong identifying assumptions, and to us the most obvious explanation for why they differ is that omitted-variable and other forms of endogeneity affect some of these estimates more than others (in a manner that is impossible to measure or even sign). But even if each of these exercises were somehow perfectly identified, we would *not* expect them to give the same coefficient estimates. For example, the smaller effects yielded by the RIF compared to the time-series could be due to the fact that only the latter picks up spillover effects such as fairness norms or political economy mechanisms. That the state-year estimates are smaller than the national estimates could suggest that these spillover effects are more powerful at the national level or that some countervailing spillover effects (e.g., creating excess labor supply in uncovered sectors) operate at the more local level. Similarly, note that inequality effects at lower levels of geography do not generally “aggregate up” to the national effect, so we should not expect the results from the aggregate time-series and the state-year panel regressions to be similar even absent spillovers. For example, unions’ reducing the top-ten share in the richest states will have a much larger effect on the *national* top-ten share than their reducing the top-ten share in a poor state like Mississippi (whose 90th percentile during much of our sample period is not much higher than the national median).

8 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. We show that a combination of low-skill composition, compression, and a large union income premium made mid-century unions a powerful force for equalizing the income distribution. As unions have receded, it is perhaps surprising—and counter to the predictions of some

SBTC models of endogenous unionization—that relatively *skilled* workers are the ones that remain. This pattern mimics the pre-World War II era, when unions were both small and their members relatively skilled. Our results show that over the last nine decades, when unions expand, whether at the national level or the state level, they tend to draw in unskilled workers and raise their relative wages, with significant impacts on inequality.

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

References

- ABOWD, J. A. and LEMIEUX, T. (1993). The effects of product market competition on collective bargaining agreements: The case of foreign competition in Canada. *The Quarterly Journal of Economics*, **108** (4), 983–1014.
- ABOWD, J. M. (1989). The effect of wage bargains on the stock market value of the firm. *The American Economic Review*, pp. 774–800.
- ACEMOGLU, D., AGHION, P. and VIOLANTE, G. L. (2001). Deunionization, technical change and inequality. In *Carnegie-Rochester conference series on public policy*, Elsevier, vol. 55, pp. 229–264.
- and ROBINSON, J. A. (2013). Economics versus politics: Pitfalls of policy advice. *The Journal of Economic Perspectives*, **27** (2), 173–192.
- AÇIKGÖZ, Ö. T. and KAYMAK, B. (2014). The rising skill premium and deunionization. *Journal of Monetary Economics*, **63**, 37–50.
- ALVAREDO, F., ATKINSON, A. B., PIKETTY, T., SAEZ, E. and ZUCMAN, G. (2016). The world wealth and income database. *Website: <http://www.wid.world>*.
- ANGRIST, J. D. and KRUEGER, A. B. (1995). Split-sample instrumental variables estimates of the return to schooling. *Journal of Business & Economic Statistics*, **13** (2), 225–235.
- ASHENFELTER, O. (1972). Racial discrimination and trade unionism. *Journal of political economy*, **80** (3, Part 1), 435–464.
- AUTOR, D. H., KATZ, L. F. and KEARNEY, M. S. (2008). Trends in US wage inequality: Revising the revisionists. *The Review of Economics and Statistics*, **90** (2), 300–323.
- AYDEMİR, A. and BORJAS, G. J. (2011). Attenuation bias in measuring the wage impact of immigration. *Journal of Labor Economics*, **29** (1), 69–112.
- BARTH, E., BRYSON, A., DAVIS, J. C. and FREEMAN, R. (2016). Its where you work: Increases in the dispersion of earnings across establishments and individuals in the United States. *Journal of Labor Economics*, **34** (S2), S67–S97.
- BENMELECH, E., BERGMAN, N. and KIM, H. (2018). *Strong Employers and Weak Em-*

- ployees: How Does Employer Concentration Affect Wages?* Tech. rep., National Bureau of Economic Research.
- BERINSKY, A. J. (2006a). American public opinion in the 1930s and 1940s: The analysis of quota-controlled sample survey data. *International Journal of Public Opinion Quarterly*, **70** (4), 499–529.
- (2006b). American public opinion in the 1930s and 1940s: The analysis of quota-controlled sample survey data. *International Journal of Public Opinion Quarterly*, **70** (4), 499–529.
- BUCHMUELLER, T. C., DiNARDO, J. and VALLETTA, R. (2004). A submerging labor market institution? unions and the nonwage aspects of work. In *Emerging labor market institutions for the twenty-first century*, University of Chicago Press, pp. 231–264.
- BUREAU OF LABOR STATISTICS (1945). *Extent of Collective Bargaining and Union Status, January 1945*. Tech. rep., bulletin # 829.
- CALLAWAY, B. and COLLINS, W. J. (2016). Unions, workers, and wages at the peak of the american labor movement.
- CARD, D. (1996). The effect of unions on the structure of wages: A longitudinal analysis. *Econometrica*, **64** (4), 957–979.
- (2001). The effect of unions on wage inequality in the us labor market. *ILR Review*, **54** (2), 296–315.
- and DiNARDO, J. E. (2002). Skill-biased technological change and rising wage inequality: some problems and puzzles. *Journal of Labor Economics*, **20** (4), 733–783.
- , HEINING, J. and KLINE, P. (2013). Workplace heterogeneity and the rise of west german wage inequality. *The Quarterly journal of economics*, **128** (3), 967–1015.
- and KRUEGER, A. B. (1992). Does school quality matter? returns to education and the characteristics of public schools in the united states. *Journal of political Economy*, **100** (1), 1–40.
- CAUGHEY, D. and WARSHAW, C. (2016). The dynamics of state policy liberalism, 1936–2014. *American Journal of Political Science*, **60** (4), 899–913.
- DEVILLE, J.-C., SÄRNDAL, C.-E. and SAUTORY, O. (1993). Generalized raking procedures in survey sampling. *Journal of the American Statistical Association*, **88** (423), 1013–1020.
- DiNARDO, J., FORTIN, N. M. and LEMIEUX, T. (1996). Labor market institutions and the distribution of wages, 1973-1992: A semiparametric approach. *Econometrica*, **64** (5), 1001–1044.
- DINARDO, J. and HALLOCK, K. F. (2002). When unions mattered: the impact of strikes on financial markets, 1925–1937. *ILR Review*, **55** (2), 219–233.
- DiNARDO, J. and LEE, D. S. (2004). Economic impacts of new unionization on private sector employers: 1984–2001. *The Quarterly Journal of Economics*, **119** (4), 1383–1441.
- DINLERSOZ, E. M. and GREENWOOD, J. (2012). *The Rise and Fall of Unions in the US*. Tech. rep., National Bureau of Economic Research.
- DUBE, A. (2017). Minimum wages and the distribution of family incomes.
- FARBER, H. (1983). The determination of the union status of workers. *Econometrica*, **51**.
- (2005). Nonunion wage rates and the threat of unionization. *ILR Review*, **58** (3), 335–352.

- FEIGENBAUM, J., HERTEL-FERNANDEZ, A. and WILLIAMSON, V. (2018). *From the bargaining table to the ballot box: political effects of right to work laws*. Tech. rep., National Bureau of Economic Research.
- FIRPO, S., FORTIN, N. M. and LEMIEUX, T. (2009). Unconditional quantile regressions. *Econometrica*, **77** (3), 953–973.
- FONER, P. S. (1976). *Organized labor and the black worker: 1619-1973*, vol. 475. International publishers.
- FRANDSEN, B. (2013). The surprising impacts of unionization on establishments: Accounting for selection in close union representation elections. *Manuscript. Department of Economics, Brigham Young University*.
- FRANDSEN, B. R. (2017). Party bias in union representation elections: Testing for manipulation in the regression discontinuity design when the running variable is discrete. In *Regression Discontinuity Designs: Theory and Applications*, Emerald Publishing Limited, pp. 281–315.
- FRANK, M. (2015). Frank-sommeiller-price series for top income shares by us states since 1917. *WTID Methodological Notes*.
- FREEMAN, R. and MEDOFF, J. (1984). Medoff. what do unions do.
- FREEMAN, R. B. (1981). The effect of unionism on fringe benefits. *Industrial and Labor Relations Review*, **34** (4), 489–509.
- *et al.* (1998). Spurts in union growth: Defining moments and social processes. *NBER Chapters*, pp. 265–296.
- FRYDMAN, C. and SAKS, R. E. (2010). Executive compensation: A new view from a long-term perspective, 1936–2005. *Review of Financial Studies*, p. 120.
- GALENSON, W. (1960). *The CIO challenge to the AFL: a history of the American labor movement, 1935-1941*. Harvard University Press.
- GOLDIN, C. and MARGO, R. (1992). The great compression: The wage structure in the united states at mid-century. *The Quarterly Journal of Economics*, **107**, 1–34.
- GOLDIN, C. D. and KATZ, L. F. (2009). *The race between education and technology*. Harvard University Press.
- HAVNES, T. and MOGSTAD, M. (2015). Is universal child care leveling the playing field? *Journal of Public Economics*, **127**, 100–114.
- HIRSCH, B. and MÜLLER, S. (2018). Firm wage premia, industrial relations, and rent sharing in germany.
- HIRSCH, B. T., MACPHERSON, D. A. and VROMAN, W. G. (2001). Estimates of union density by state. *Monthly Labor Review*, **124** (7), 51–55.
- IGO, S. E. (2007). *The averaged American: Surveys, citizens, and the making of a mass public*. Harvard University Press.
- INOUE, A. and SOLON, G. (2010). Two-sample instrumental variables estimators. *The Review of Economics and Statistics*, **92** (3), 557–561.
- JAUMOTTE, F. and OSORIO, C. (2015). *Inequality and labor market institutions*. International Monetary Fund.

- KATZ, L. F. and MURPHY, K. M. (1992). Changes in relative wages, 1963–1987: supply and demand factors. *The Quarterly Journal of Economics*, **107** (1), 35–78.
- KATZNELSON, I. (2013). *Fear itself: The new deal and the origins of our time*. WW Norton & Company.
- KING JR, M. L. (1986). If the negro wins, labor wins. *A Testament of Hope: The Essential Writings of Martin Luther King, Jr*, pp. 201–207.
- KOPCZUK, W., SAEZ, E. and SONG, J. (2010). Earnings inequality and mobility in the united states: evidence from social security data since 1937. *The Quarterly Journal of Economics*, **125** (1), 91–128.
- LEE, D. S. (1999). Wage inequality in the united states during the 1980s: Rising dispersion or falling minimum wage? *The Quarterly Journal of Economics*, **114** (3), 977–1023.
- and MAS, A. (2012). Long-run impacts of unions on firms: New evidence from financial markets, 1961–1999. *The Quarterly Journal of Economics*, **127** (1), 333–378.
- LEIGHLEY, J. E. and NAGLER, J. (2007). Unions, voter turnout, and class bias in the us electorate, 1964–2004. *Journal of Politics*, **69** (2), 430–441.
- LEMIEUX, T. (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), 261–291.
- (2006). Increasing residual wage inequality: Composition effects, noisy data, or rising demand for skill? *The American Economic Review*, **96** (3), 461–498.
- LEWIS, H. G. (1963). *Unionism and Relative Wages in the United States*. University of Chicago Press.
- (1986). Union relative wage effects. *Handbook of labor economics*, **2**, 1139–1181.
- LICHTENSTEIN, N. (2003). *Labor’s War at Home: The CIO in World War II*. Temple University Press.
- MANNING, A. (2011). Imperfect competition in the labor market. In *Handbook of labor economics*, vol. 4, pp. 973–1041.
- MILL, J. S. (2008). *Principles of Political Economy*. Oxford University Press.
- NAIDU, S., POSNER, E. A. and WEYL, E. G. (2018). Antitrust remedies for labor market power. *Harvard Law Review*.
- NORTHROP, H. R. (1971). *Organized labor and the Negro*. Periodicals Service Co.
- PIKETTY, T. and SAEZ, E. (2003). Income inequality in the united states, 1913–1998. *The Quarterly journal of economics*, **118** (1), 1–41.
- , — and STANTCHEVA, S. (2014). Optimal taxation of top labor incomes: A tale of three elasticities. *American economic journal: economic policy*, **6** (1), 230–271.
- PISCHKE, J.-S., DINARDO, J. E. and HALLOCK, K. F. (2000). Unions and the labor market for managers.
- REES, A. (1963). The effects of unions on resource allocation. *The Journal of Law and Economics*, **6**, 69–78.
- ROSE, N. L. (1987). Labor rent sharing and regulation: Evidence from the trucking industry.

- Journal of Political Economy*, **95** (6), 1146–1178.
- SCHICKLER, E. and CAUGHEY, D. (2011). Public opinion, organized labor, and the limits of new deal liberalism, 1936–1945. *Studies in American Political Development*, **25** (02), 162–189.
- SONG, J., PRICE, D. J., GUVENEN, F., BLOOM, N. and VON WACHTER, T. (2015). *Firming up inequality*. Tech. rep., National Bureau of Economic Research.
- TASCHEREAU-DUMOUCHEL, M. (2015). The union threat.
- TROY, L. (1965). *Trade Union Membership, 1897–1962*. NBER.
- and SHEFLIN, N. (1985). Union sourcebook: Membership, finances, structure, directory. *West Orange, NJ: Industrial Relations Data and Information Services*.
- WESTERN, B. and ROSENFELD, J. (2011). Unions, norms, and the rise in us wage inequality. *American Sociological Review*, **76** (4), 513–537.
- WOLMAN, L. (1924). *The Growth of American Trade Unions, 1880-1923*. NBER.

Figure 1: Union density and inequality measures, 1917-2011

Data sources: Top share inequality from Piketty and Saez (2003, updated 2016). Union density data from Historical Statistics of the United States and the Current Population Survey. We discuss these data sources in detail in Section 3 and Appendix D.

Figure 2: Union density over the twentieth century: Comparing our survey-based measures to existing time-series

Data sources: See Sections 3 and Appendix D.

Notes: No sample restrictions are imposed (so farmers and those over age 65 are included in this graph). The vertical spikes indicate the number of Gallup observations per year that include the union variable (plotted on the right-hand-side axis).

Figure 3: How does educational attainment 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. See Section 3 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. (The notes to Figure 6 describe how we impute years of education if the survey source only gives us categories of educational attainment.) We estimate this equation separately by survey source and by year. The figure plots the coefficient on *Year of education*. 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 4: 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. See Section 3 for a description of each data source.

Notes: For each data source, we estimate, separately by data source and year, household union status on a *White* dummy variable, state s and survey-date t fixed effects, age and its square, and gender. 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 5: Estimates of the union family income premium

Data source: Gallup data, 1942, 1961–1974; CPS, 1978–2016; BLS Expenditure Survey, 1936; ANES, 1952–1996, U.S. Psych. Corporation, 1946. See Section 3 for a description of each data source. See Appendix B for details on CPS family income variable construction.

Notes: Each plotted point comes from estimating equation (3), which regresses log family income on controls for age, gender, race, state and survey-date fixed effects. Occupation controls are not included. 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 6: Differential family union premium by respondent's years of schooling

Data source: Gallup data, 1942, 1961–1974; CPS, 1978–2016; BLS Expenditure Survey, 1936; ANES, 1952–1996, U.S. Psych. Corporation, 1946. See Section 3 for a description of each data source. See Appendix B for details on CPS family income variable construction.

Notes: Each plotted point comes from estimating equation regressing log family income on household union status, its interaction with respondents' log years of schooling, and all other controls in equation (3). We estimate this equation separately by survey source and by year. Some survey sources give actual years of schooling. For those that do not, we impute in the following manner: six years for “less than middle school;” eight years for “middle school;” ten years for “some high school;” twelve years for “high school;” fourteen years for “some college” or “vocational training;” sixteen years for “college,” eighteen years for “more than college.” The figure plots the coefficient on the interaction $Years\ of\ schooling \times Union$. 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 7: Differential family union premium for whites relative to minorities

Data source: Gallup data, 1942, 1961–1974; CPS, 1978–2016; BLS Expenditure Survey, 1936; ANES, 1952–1996, U.S. Psych. Corporation, 1946. See Section 3 for a description of each data source. See Appendix B for details on CPS family income variable construction.

Notes: Each plotted point comes from estimating equation regressing log family income on household union status, its interaction with a White dummy variable, and all other controls in equation (3). 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 8: Ratio of residual variance between union and non-union sectors

Data source: Gallup data, 1942, 1961–1974; CPS, 1978–2016; BLS Expenditure Survey, 1936; ANES, 1952–1996, U.S. Psych. Corporation, 1946. See Section 3 for a description of each data source. See Appendix B 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 (3) separately for union and non-union households, separately by survey source and by year. See Section 6.4 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.

Figure 9: Effects of union density on the 90/10 ratio of family income

Data source: Gallup data, 1942, 1961–1974; CPS, 1978–2016; BLS Expenditure Survey, 1936; ANES, 1952–1996, U.S. Psych. Corporation, 1946. See Section 3 for a description of each data source. See Appendix B for details on CPS family income variable construction.

Notes: Each plotted point comes from estimating equation (5), which regresses the recentered influence function (RIF) for the difference between the 90th and 10th quantile on controls for age, gender, race, educational attainment fixed effects, household employment status controls, state and survey-date fixed effects. Occupation controls are not included. For the ANES, because the samples are smaller, we group surveys into six-year bins. The plotted confidence intervals are robust to heteroskedasticity.

Figure 10: Effects of union density on the Gini coefficient of family income

Data source: See Section 3 for a description of each data source. See Appendix B for details on CPS family income variable construction.

Notes: Each plotted point comes from estimating equation (5), with regressed the recentered influence function (RIF) for the Gini coefficient on controls for age, gender, race, educational attainment fixed effects, household employment status controls, state and survey-date fixed effects. Occupation controls are not included. For the ANES, because the samples are smaller, we group surveys into six-year bins. The plotted confidence intervals are robust to heteroskedasticity.

Table 1: Is selection into unions a function of union density?

	Dep't var.: At least one household member is in a union				
	(1)	(2)	(3)	(4)	(5)
Yrs. of educ x Aggregate union density	-0.280*** [0.0259]	-0.124*** [0.0235]	-0.104*** [0.0231]	-0.122*** [0.0234]	-0.102*** [0.0227]
Yrs. of educ x Year		-0.449*** [0.0272]	-0.266*** [0.0281]	-0.451*** [0.0280]	-0.266*** [0.0292]
Yrs. of educ x Year sq.		0.230*** [0.0136]	0.135*** [0.0138]	0.231*** [0.0140]	0.135*** [0.0143]
White x Aggregate union density				-0.183* [0.0956]	-0.0923 [0.113]
White x Year				0.0842 [0.274]	-0.0303 [0.276]
White x Year sq.				-0.0381 [0.135]	0.0165 [0.136]
Dep't var. mean	0.203	0.203	0.178	0.203	0.178
Year FE?	Yes	Yes	Yes	Yes	Yes
Drop pub. sect. HH?	No	No	Yes	No	Yes
Observations	1143694	1143694	1005363	1143694	1005363

Sources: Gallup, 1937–1976, CPS, 1977-2015.

Notes: These regressions test whether selection into union households by education and by race is a function of union density. All regressions include state and year fixed effects, and controls for age (and its square), race, gender and education. Variables that are interacted with *Year* ($Year^2$) are divided by 1,000 (1,000,000) for ease of presentation. Note that *Year* refers to calendar year (e.g., 1942, 1943, etc). Thus, interacting a variable with *Year* and *Year sq.* allows selection into unions with respect to that variable to follow an arbitrary *U*-shape in calendar time. Dropping public sector households is only done in the 1977-2015 CPS sample. Standard errors are clustered by state. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 2: Aggregate inequality as a function of union density

	Dependent variable:							
	Coll. premium		90/10 ratio		Gini coeff.		Top 10 share	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Union Density	-1.338** [0.535]	-1.556** [0.720]	-1.407*** [0.421]	-1.456** [0.636]	-0.132** [0.0536]	-0.106** [0.0525]	-33.64** [14.41]	-15.57 [12.50]
Mean, dept. var	0.512	0.512	1.376	1.376	0.410	0.410	35.848	35.848
Gallup edu. control?	No	No	No	No	Yes	Yes	Yes	Yes
Addit. controls?	No	Yes	No	Yes	No	Yes	No	Yes
Cubic polynomial?	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	49	49	49	49	65	65	70	70

Sources: For cols. (1) - (4), outcome variables generated from Census IPUMS and CPS; for cols (5) and (6) from Kopczuk *et al.* (2010) and for cols (7) and (8) from Piketty and Saez (2003, updated 2016). The union density explanatory variable is the simple average between the Gallup and BLS-based density measures (see Section C for detail).

Notes: 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 four columns use outcome variables calculated from the source (so are only available in Census years until the CPS), but the last four columns use as outcomes *annual* measures, 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). 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.9, A.10, A.11, A.12 provide additional specifications using the skill premium, the log 90/10 ratio, the Gini coefficient and the top-ten share, respectively, as outcomes. Note that to the log 90/10 is for men only, but all other inequality measures pool both men and women. Standard errors are robust to heteroskedasticity and AR(1) serial correlation. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 3: State-year inequality as a function of union density

	Dependent variable:							
	Coll. premium		90/10 ratio		Gini coeff.		Top 10 share	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Household union share	-0.442*** [0.118]	-0.434*** [0.104]	-0.296** [0.118]	-0.244** [0.0953]	-0.0635*** [0.0235]	-0.0738*** [0.0235]	-5.342*** [2.065]	-3.147** [1.399]
Mean, dept. var.	0.490	0.497	1.386	1.398	0.376	0.378	36.61	36.96
Industry shares	No	Yes	No	Yes	No	Yes	No	Yes
State-spec. quad.	No	Yes	No	Yes	No	Yes	No	Yes
Income covars.	No	Yes	No	Yes	No	Yes	No	Yes
Policy covars.	No	Yes	No	Yes	No	Yes	No	Yes
Observations	1640	1505	1640	1505	1640	1505	3107	2723

Sources: For cols. (1) through (6), dependent variables created using Census and CPS data; for (7) and (8) they are taken from Frank (2015). See Appendix for variable construction. 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.

Notes: All estimates are from split-sample-IV regressions (see Section 7.3 for estimating equations). 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). “Industry shares” controls for state-year share of employment in all one-digit industry categories. “State-spec. quad.” indicates that state-specific quadratic time trends are included. “Income covars.” indicate that state-year GDP and state-year share of households filing taxes are included. “Policy covars.” indicate that state-year minimum wage and a “policy liberalism” index (from Caughey and Warshaw, 2016) are included. Sample size is larger for the top 10 outcome because it is available at the annual level in all years; for the other outcomes, until the CPS in the 1970s, we only have data from the decadal Census.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 4: How much of the variation in inequality is explained by changes in union density?

		Share explained by Δ union density, using β values from...			
	Total Δ	RIF (high)	RIF (low)	Time-series	State panel
—1940-1960					
Union density	0.113				
Skill premium	-0.0926			1.894	0.449
Male 90/10 ratio	-0.326	0.135	0.0379	0.504	0.200
Gini coefficient	-0.0585	0.167	0.0890	0.204	0.116
Top ten percent	-12.77			0.137	0.0188
—1970-2004					
Union density	-0.123				
Skill premium	0.212			0.904	0.214
Male 90/10 ratio	0.465	0.103	0.0290	0.385	0.153
Gini coefficient	0.0875	0.122	0.0649	0.149	0.0843
Top ten percent	12.13			0.158	0.0216

Sources: For cols. (1) through (6), dependent variables created using Census and CPS data; for (5) and (6) they are taken from Frank (2015). See Appendix for variable construction. 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.

Notes: Col. (1) is the change in each outcome variable j (skill premium, male 90/10 ratio, gini coefficient, and top 10 percent) as well as union density between 1940 and 1960 (1970 and 2004 in lower panel). Values in cols. (2)-(5) are calculated, for column i , row j , as $(\Delta_{unionDensity} * \beta_{ij}) / \Delta_j$ where β_{2j} is the average of all RIF coefficients between 1950 and 1970 for outcome j (Figures 9 and 10), β_{3j} is the average of all RIF coefficients between 2005 and 2015 (Figures 9 and 10), β_{4j} is the point-estimate in the even-numbered column of Table 2 corresponding to outcome j , and β_{5j} is the point-estimate in the even-numbered column of Table 3 corresponding to outcome j . Note that 1950-1970 is roughly the period when the RIF coefficients are largest in magnitude (which we term “RIF (high)”), while in the more recent years they are smaller in magnitude (“RIF (low)”).

Appendix A. Supplementary Figures and Tables Noted in the Text

Appendix Figure A.1: Comparing unemployment rates in Gallup and the HSUS

Data sources: Gallup and Historical Statistics of the United States (HSUS)

Notes: Sample in Gallup includes farmers

Appendix Figure A.2: Age distribution in Gallup, by gender, 1937-1952

Data sources: Gallup microdata.

Appendix Figure A.3: Union share of households in the Gallup data (weighted vs. un-weighted)

Data sources: Gallup. See Section 3 and Appendix C for more detail on data and weight construction.

Appendix Figure A.4: Comparing union density in Gallup and CPS, 1970–present

Data sources: Gallup and Current Population Survey

Appendix Figure A.5: Selection of union households by log years schooling

Sources: Gallup, 1937–1986. CPS, 1978–2016; BLS Expenditure Survey, 1936; ANES, 1952–1996, U.S. Psych. Corporation, 1946.

Notes: For each data source, we estimate, separately by data source and year, household union status on *Log Years Education*, state s and survey-date t fixed effects, age and its square, and gender. We plot in this graph the coefficients on *Log Years Education* from each of these estimations. We control for survey-date and state fixed effects, age and its square, and gender. For the ANES, because the samples are smaller, we group surveys into six-year bins. Standard errors are clustered by state.

Appendix Figure A.6: Selection of union households by high-school graduation

Sources: Gallup, 1937–1986. CPS, 1978–2016; BLS Expenditure Survey, 1936; ANES, 1952–1996, U.S. Psych. Corporation, 1946.

Notes: For each data source, we estimate, separately by data source and year, household union status on a *High School Grad* dummy variable, state s and survey-date t fixed effects, age and its square, and gender. We plot in this graph the coefficients on *High School Grad* from each of these estimations. We control for survey-date and state fixed effects, age and its square, and gender. For the ANES, because the samples are smaller, we group surveys into six-year bins. Standard errors are clustered by state.

Appendix Figure A.7: Selection of union households by college graduation

Sources: Gallup, 1937–1986. CPS, 1978–2016; BLS Expenditure Survey, 1936; ANES, 1952–1996, U.S. Psych. Corporation, 1946.

Notes: For each data source, we estimate, separately by data source and year, household union status on a *College Grad* dummy variable, state s and survey-date t fixed effects, age and its square, and gender. We plot in this graph the coefficients on *College Grad* from each of these estimations. We control for survey-date and state fixed effects, age and its square, and gender. For the ANES, because the samples are smaller, we group surveys into six-year bins. Standard errors are clustered by state.

Appendix Figure A.8: Selection of union households by education in the ANES and CPS (dropping households with a public- or manufacturing-sector worker)

Sources: CPS, 1978–2016; ANES, 1952–1996.

Notes: For each data source, we estimate, separately by data source and year, household union status on a *Years of education* variable, state s and survey-date t fixed effects, age and its square, and gender. We plot in this graph the coefficients on *Years of education* from each of these estimations. We control for survey-date and state fixed effects, age and its square, and gender. For the ANES, because the samples are smaller, we group surveys into six-year bins. Note that we only include ANES and CPS in this graph, because other data sources do not allow us to identify industrial sectors of workers in the household.

Appendix Figure A.9: Selection of union households by race (dropping Southern states)

Data sources: Gallup data, 1937–1986; CPS, 1978–2016; BLS Expenditure Survey, 1936; ANES, 1952–1996. See Section C for a description of each data source.

Notes: For each data source, we estimate, separately by data source and year, household union status on a *White* dummy variable, state s and survey-date t fixed effects, age and its square, and gender. 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. Note that we cannot use the U.S. Psychological Corporation survey in this figure because, while it has state identifiers (thus we can thus control for state fixed effects) the codebook does *not* provide the state names that correspond to the codes (so we cannot drop the South). Confidence intervals are based on standard errors clustered by state.

Appendix Figure A.10: Estimates of the union family income premium (including occupation controls when available)

Data source: Gallup data, 1942, 1961–1974; CPS, 1978–2016; BLS Expenditure Survey, 1936; ANES, 1952–1996, U.S. Psych. Corporation, 1946. See Section C for a description of each data source. See Appendix B for details on family income variable construction.

Notes: Each plotted point comes from estimating equation (3), which regressed log family income on controls for age, gender, race, state and survey-date fixed effects and (in most cases) fixed effects for the occupation of the head. 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.

Appendix Figure A.11: Estimates of the union family income premium from ANES (with and without employment status controls)

Data source: See Section 3 for a description of ANES data.

Notes: Each plotted point comes from estimating equation (3), which regresses log family income on controls for age, gender, race, state and survey-date fixed effects. In addition, the first series also includes an indicator for the household head being employed and a separate indicator for the respondent being employed. See Section 6.1 for more detail.

Appendix Figure A.12: Effects of union density on family income quantiles

Data source: Gallup data, 1942, 1961–1974; CPS, 1978–2016; BLS Expenditure Survey, 1936; ANES, 1952–1996, U.S. Psych. Corporation, 1946. See Section C for a description of each data source. See Appendix B for details on family income variable construction.

Notes: Each plotted point comes from estimating equation (5), which regresses the recentered influence function (RIF) for the specified quantile on controls for age, gender, race, educational attainment fixed effects, household employment status controls, state and survey-date fixed effects. Occupation controls are not included. The plotted confidence intervals are robust to heteroskedasticity.

Appendix Figure A.13: Effects of union density on family income quantiles (using *unbinned* CPS family income)

Data source: Gallup data, 1942, 1961–1974; CPS, 1978–2016; BLS Expenditure Survey, 1936; ANES, 1952–1996, U.S. Psych. Corporation, 1946. See Section C for a description of each data source. See Appendix B for details on family income variable construction.

Notes: Each plotted point comes from estimating equation (5), which regresses the recentered influence function (RIF) for the specified quantile on controls for age, gender, race, educational attainment fixed effects, household employment status controls, state and survey-date fixed effects. Occupation controls are not included. The plotted confidence intervals are robust to heteroskedasticity.

Appendix Table A.1: Comparing Gallup and IPUMS, 1950–1980

	1950		1960		1970		1980	
	Census	Gallup	Census	Gallup	Census	Gallup	Census	Gallup
South Share	0.242	0.117	0.259	0.138	0.271	0.247	0.296	0.256
— <i>South</i>								
Female	0.516	0.505	0.521	0.518	0.529	0.507	0.529	0.503
Age	44.61	44.31	45.07	47.64	45.94	46.35	45.20	46.13
Black	0.200	0.0849	0.182	0.147	0.160	0.129	0.159	0.160
HS grad.	0.294	0.373	0.366	0.372	0.473	0.529	0.619	0.635
— <i>Non-South</i>								
Female	0.515	0.504	0.517	0.512	0.528	0.506	0.528	0.503
Age	46.67	43.75	45.96	45.87	46.27	45.38	45.28	44.10
Black	0.0530	0.0454	0.0611	0.0586	0.0709	0.0614	0.0782	0.0874
HS grad.	0.385	0.473	0.450	0.531	0.579	0.659	0.710	0.755
Observ.	296223	182171	5388972	95064	2444218	138098	7475162	128507

Sources: Gallup surveys and 1950–1980 IPUMS.

Notes: We use the Gallup definition of the “South”: all eleven states of the former Confederacy plus Oklahoma. All Census results use IPUMS person weights.

Appendix Table A.2: Comparing Gallup and IPUMS in 1940

	Gallup	Census	Census	Gallup	Census
<i>–Demographics</i>					
Black	0.0290	0.0895	0.0906	0.0325	0.0357
Female	0.338	0.505	0.344	0.341	0.343
Age	40.45	39.61	40.06	40.40	40.55
HS Graduate	0.493	0.278	0.266	0.494	0.290
College Graduate	0.0720	0.0472	0.0499	0.0709	0.0543
<i>–Geography</i>					
Northeast	0.0835	0.0660	0.0629	0.0946	0.0854
Mid Atlantic	0.262	0.253	0.241	0.297	0.327
East Central	0.207	0.187	0.186	0.235	0.252
West Central	0.176	0.127	0.129	0.200	0.175
South	0.118	0.258	0.263	- -	- -
Rocky Mountain	0.0751	0.0284	0.0308	0.0851	0.0418
Pacific Coast	0.0784	0.0754	0.0818	0.0888	0.111
<i>–Occupation</i>					
Professional	0.0780	0.113	0.122	0.0793	0.129
Farmer	0.209	0.156	0.159	0.185	0.109
Proprietors, managers, officials	0.0104	0.0928	0.0875	0.0106	0.0933
Clerks (white collar)	0.294	0.0535	0.0539	0.301	0.0609
Skilled workmen and foremen	0.0906	- -	- -	0.0953	- -
Sales workers	- -	0.0462	0.0457	- -	0.0499
Craftsmen	- -	0.142	0.139	- -	0.153
Operatives	- -	0.146	0.147	- -	0.159
Unskilled or semi-skilled labor	0.190	- -	- -	0.200	- -
Laborers	- -	0.0932	0.0973	- -	0.0944
Service workers (priv. HH)	- -	0.0103	0.0105	- -	0.00626
Other service workers	- -	0.0477	0.0468	- -	0.0508
No answer, N/A, etc.	0.0826	0.0999	0.0920	0.0836	0.0949
HH/gender adjustment	N/A	N/A	Yes	N/A	Yes
Ex. S/SW?	No	No	No	Yes	Yes
Observations	See notes	736832	736832	See notes	544375

Sources: Gallup surveys and 1940 IPUMS.

Notes: The Gallup sample size varies substantially by variable during this period. For the col. (1) sample, all demographics except for education and all geographic variables have a sample size around 159,000 (with small variations due to missing observations). The occupation codes have a sample size of roughly 21,000. The high school completion indicator has a sample size of 5,700. In col. (4) each sample size is roughly twelve percent smaller. “HH / gender adjustment” underweights women and people in large households in the IPUMS, to better match Gallup sampling (which only sampled one person per household and had a target female share of one-third). “Ex S/SW” excludes Southern and Southwestern states (all eleven states of the former Confederacy plus Oklahoma). Note that occupation categories are coarser in Gallup than in the Census (but unfortunately, Gallup categories do not nest Census categories). We do our best to match occupation across these different categorizations. All Census results use IPUMS person weights.

Appendix Table A.3: Summary statistics from supplementary data sets

	(1) ANES	(2) BLS exp. dataset	(3) U.S. Psych. Corp.	(4) NORC
Union household	0.254	0.141	0.184	0.274
Female	0.548	0.0346	0.507	0.514
White	0.858	0.920	0.879	0.903
Age	39.67	40.85	39.11	39.84
HS graduate	0.360	0.405	0.470	0.403
South	0.277	0.232		
Log fam. inc.	9.380	7.121	3.720	7.913
Sample period	1952-1988	1936	1946	1950
Observations	30757	4058	4956	1106

Notes: See Section 3 for details on the data sources.

Appendix Table A.4: Gallup selection results through 1950, robustness to weights

	Dependent variable: Union household				
	(1)	(2)	(3)	(4)	(5)
Years of education	-0.0250*** [0.00229]	-0.0293*** [0.00191]	-0.0265*** [0.00233]	-0.0267*** [0.00240]	-0.0237*** [0.00221]
Dept. var. mean	0.233	0.215	0.200	0.203	0.193
Weighting scheme	Baseline	None	White x Sth	Raking	Schickler
Observations	198257	208986	198257	195565	59275

Sources: See Section 3 and Appendix C for details.

Notes: All regressions include state and survey-date fixed effects and include ages 21–64. Baseline weights are those we use throughout the paper (weights to make Gallup match interpolated Census cells for *White* $\times$ *South* $\times$ *Education categories* (16 cells)). *White* $\times$ *Sth* are analogous, but match only on those four cells. Raking weights are constructed by matching yearly marginal mean population shares by *Black*, *Female*, and *Region* to interpolated census shares. See Deville *et al.* (1993) for more details. “Schickler weights” are taken from Schickler and Caughey (2011), match on *Black* and whether a residence has a phone, and are only available through 1945.

Standard errors in brackets, clustered by state. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Appendix Table A.5: Is selection into unions a function of union density (using state-year variation in union density, instead of annual variation as in Table 1)?

	Dep't var.: At least one household member is in a union			
	(1)	(2)	(3)	(4)
Yrs. educ x State-year union density	-0.0803*** [0.0147]	-0.0839*** [0.0123]	-0.0772*** [0.0145]	-0.0812*** [0.0123]
White x State-year union density			-0.293*** [0.0586]	-0.279*** [0.0664]
Dep't var. mean	0.203	0.178	0.203	0.178
Year FE?	Yes	Yes	Yes	Yes
Drop pub. sect. HH?	No	Yes	No	Yes
Observations	1134257	998146	1134257	998146

Sources: Gallup, 1937–1976, CPS, 1977-2015.

Notes: These regressions test whether selection into union households by education and by race is a function of union density. All regressions include state and year fixed effects, and controls for age (and its square), race, gender and education. Dropping public sector households is only done in the 1977-2015 CPS sample. Standard errors are clustered by state. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Appendix Table A.6: Estimating family union income premium and reporting coefficients on additional covariates, by data source and time period

	Dep't var: Logged family income					
	(1)	(2)	(3)	(4)	(5)	(6)
Union household	0.212*** [0.0372]	0.255*** [0.0348]	0.0836*** [0.0229]	0.163*** [0.00289]	0.150*** [0.0360]	0.234*** [0.0189]
Years of educ., respondent		0.132*** [0.00549]	0.108*** [0.00587]	0.0974*** [0.000612]	0.107*** [0.00700]	0.111*** [0.00465]
Years of educ., household head	0.0915*** [0.00766]					
White	0.632*** [0.0720]	0.406*** [0.0618]	0.255*** [0.0469]	0.385*** [0.00501]	0.394*** [0.0626]	0.327*** [0.0381]
Respondent is female		-0.112*** [0.0300]	-0.155*** [0.0254]	-0.113*** [0.00271]	-0.163*** [0.0329]	-0.122*** [0.0149]
Household head is female	-0.892*** [0.168]					
Age	0.0834*** [0.00820]	0.0589*** [0.00968]		0.0619*** [0.000796]	0.0597*** [0.00765]	0.0579*** [0.00396]
Age squared, divided by 1,000	-0.899*** [0.0917]	-0.671*** [0.120]		-0.741*** [0.00972]	-0.679*** [0.0927]	-0.587*** [0.0463]
Age 30-39			0.202*** [0.0412]			
Age 40-49			0.193*** [0.0295]			
Age 50-59			0.199*** [0.0426]			
Data source	Exp. survey	Gallup	U.S. Psych.	Gallup	ANES	ANES
Year(s) in sample	1937	1942	1946	1961-1975	1952-1970	1972-1990
Observations	4157	2524	2373	177099	2628	11777

Sources: See Section C for details.

Notes: All regressions include state fixed effects. For Gallup, survey date fixed effects are included and for ANES, year fixed effects. We control for number of employed individuals in the household, except in the Gallup and U.S. Psych. data, where this control is not available. For the U.S. Psych. survey, age is given in categories, not in years, and the omitted age category in the regression is “under 30” (and we drop any observation above age 60). Otherwise, all other samples include ages 21–64. Standard errors in brackets, clustered by state. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Appendix Table A.7: Paid vacation as a function of union status (Gallup, 1949)

	Dep't var: Do you (or husband) get paid vacation?				
	(1)	(2)	(3)	(4)	(5)
Union household	0.220*** [0.0332]	0.183*** [0.0308]	0.288** [0.126]	0.280* [0.143]	0.121*** [0.0312]
White x Union household			-0.111 [0.127]		
Years educ. x Union household				-0.00964 [0.0130]	
Low-skill labor x Union					0.149*** [0.0493]
Dept. var. mean	0.523	0.526	0.526	0.526	0.526
State FE?	Yes	Yes	Yes	Yes	Yes
Demographic controls?	Yes	Yes	Yes	Yes	Yes
Occupation FE?	No	Yes	Yes	Yes	Yes
Observations	1895	1864	1864	1864	1864

Data sources: Data from a Gallup survey in May 1949. For details on Gallup data, see Section 3 and Appendix C.

Notes: The dependent variable is a dummy. Demographic controls include respondent's age and its square, education (four fixed effects), gender and race. When occupation controls are added, they refer to the head of the household. Low-skill occupation dummy in the final column refer to the Gallup categories of "unskilled and semi-skilled labor." Standard errors are in brackets and clustered by state. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Appendix Table A.8: Ease of finding a job “just as good as” the one you have, as function of union status (Gallup, 1939)

	Would be easy to find another job just as good			
	(1)	(2)	(3)	(4)
Union household	-0.122*** [0.0278]	-0.123*** [0.0256]	-0.0951*** [0.0288]	-0.1000*** [0.0298]
Mean, dept. var.	0.499	0.499	0.499	0.495
State FE	Yes	Yes	Yes	Yes
Demogr. controls	No	Yes	Yes	Yes
Educ. controls	No	No	Yes	Yes
Occup. controls	No	No	Yes	Yes
Ex. South	No	No	No	Yes
Observations	1952	1952	1952	1686

Notes: Data from a Gallup survey in March 1939. The exact wording of the question from which we create the outcome variable is: “If you lost your present job (business, farm), how hard do you think it would be for you to get another job (business, farm) just as good?” We code the responses “impossible” and “quite hard” as zero and “fairly hard” and “easy” as one. Demographic controls include respondent’s age and its square, education (four fixed effects), gender and race. When occupation controls are added, they refer to the head of the household. Low-skill occupation dummy in the final column refer to the Gallup categories of “unskilled and semi-skilled labor.” Standard errors are in brackets and clustered by state. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Appendix Table A.9: Aggregate college-high-school premium as a function of union density (additional specifications)

	Dep't var: College High School Premium						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Educ. Share Ratio	-0.382** (0.178)	-0.287** (0.140)	-0.334** (0.163)	-0.307** (0.152)	-0.239 (0.195)	-0.218 (0.167)	-0.391* (0.227)
Union Density (Gallup)		-1.507*** (0.481)					
Union Density (BLS)			-0.973* (0.570)				
Union Density (Average)				-1.338** (0.535)	-1.556** (0.720)	-1.712*** (0.581)	-1.370** (0.575)
Mean, dept. var	0.512	0.512	0.512	0.512	0.512	0.512	0.512
Controls?	No	No	No	No	Yes	Yes	Yes
Time Polynomial?	Cubic	Cubic	Cubic	Cubic	Cubic	Quadratic	Quartic
Observations	49	49	49	49	49	49	49

Sources: The college premium was created using Census and CPS data. See Appendix B for details on variable construction.

Notes: This table shows companion specifications to the regressions reported in cols. (1) and (2) in the main Table 2. All regressions include “time polynomial” controls, either up to quadratic, cubic or quartic level. “Controls” include the federal minimum wage, unemployment rate for civilian men, and top marginal tax rates. “Union Density (Average)” uses the mean of BLS and Gallup Union Density Series. Standard errors are robust to heteroskedasticity and AR(1) serial correlation. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Appendix Table A.10: Aggregate male 90/10 ratio as a function of union density (additional specifications)

	Dep't var: Log Percentile 90-10 Men						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Educ. Share Ratio	-0.175** (0.076)	-0.100 (0.095)	-0.103 (0.087)	-0.096 (0.093)	-0.084 (0.127)	0.156 (0.211)	-0.189 (0.158)
Union Density (Gallup)		-1.185*** (0.416)					
Union Density (BLS)			-1.465*** (0.395)				
Union Density (Average)				-1.407*** (0.421)	-1.456** (0.636)	-3.242*** (0.582)	-1.328** (0.566)
Mean, dept. var	1.376	1.376	1.376	1.376	1.376	1.376	1.376
Controls?	No	No	No	No	Yes	Yes	Yes
Time Polynomial?	Cubic	Cubic	Cubic	Cubic	Cubic	Quadratic	Quartic
Observations	49	49	49	49	49	49	49

Sources: The Log 90-10 Percentile for Men was created using Census and CPS data. See Appendix B details on variable construction.

Notes: This table shows companion specifications to the regressions reported in cols. (3) and (4) in the main Table 2. All regressions include “Time Polynomial” controls, either up to quadratic, cubic or quartic level. “Controls” include the federal minimum wage, unemployment rate for civilian men, and top marginal tax rates. “Union Density (Average)” uses the mean of BLS and Gallup Union Density Series. Standard errors are robust to heteroskedasticity and AR(1) serial correlation. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Appendix Table A.11: Aggregate Gini coefficient as a function of union density (additional specifications)

	Dep't var: Gini Coefficient					
	(1)	(2)	(3)	(4)	(5)	(6)
Union Density (Gallup)	-0.083 (0.053)					
Union Density (BLS)		-0.120*** (0.040)				
Union Density (Average)			-0.132** (0.054)	-0.106** (0.053)	-0.216*** (0.046)	-0.127** (0.050)
Mean, dept. var	0.410	0.410	0.410	0.410	0.410	0.410
Educ. Control	Yes	Yes	Yes	Yes	Yes	Yes
Addit. controls?	No	No	No	Yes	Yes	Yes
Time Polynomial?	Cubic	Cubic	Cubic	Cubic	Quadratic	Quartic
Observations	65	65	65	65	65	65

Sources: The Gini Coefficient is calculated from social security data by Kopczuk *et al.* (2010).

Notes: This table shows companion specifications to the regressions reported in cols. (5) and (6) in the main Table 2. All regressions include “Time Polynomial” controls, either up to quadratic, cubic or quartic level. “Controls” include the federal minimum wage, unemployment rate for civilian men, and top marginal tax rates. “Union Density (Average)” uses the mean of BLS and Gallup Union Density Series. “Educ. Controls” include College-High School Share Ratio using CPS, Gallup Data and Census (interpolated) data. Standard errors are robust to heteroskedasticity and AR(1) serial correlation. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Appendix Table A.12: Aggregate top-ten income share as a function of union density (additional specifications)

	Dep't var: Top 10 Income Share					
	(1)	(2)	(3)	(4)	(5)	(6)
Union Density (Gallup)	-16.285 (10.664)					
Union Density (BLS)		-37.049** (14.165)				
Union Density (Average)			-33.638** (14.413)	-15.570 (12.497)	-35.340*** (12.292)	-13.626 (14.119)
Mean, dept. var	35.848	35.848	35.848	35.848	35.848	35.848
Educ. Control	Yes	Yes	Yes	Yes	Yes	Yes
Addit. controls?	No	No	No	Yes	Yes	Yes
Time Polynomial?	Cubic	Cubic	Cubic	Cubic	Quadratic	Quartic
Observations	70	70	70	70	70	70

Sources: The Top 10 percent income share is calculated from IRS data and updated by Piketty and Saez (2003).

Notes: This table shows companion specifications to the regressions reported in cols. (7) and (8) in the main Table 2. All regressions include “Time Polynomial” controls, either up to quadratic, cubic or quartic level. “Controls” include the federal minimum wage, unemployment rate for civilian men, and top marginal tax rates. “Union Density (Average)” uses the mean of BLS and Gallup Union Density Series. “Educ. Controls” include College-High School Share Ratio using CPS, Gallup Data and Census (interpolated) data. See Appendix for variable construction. Standard errors are robust to heteroskedasticity and AR(1) serial correlation.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Appendix Table A.13: State-Year College High School Premium as a function of union density (additional specification)

	Dep't var: College High School Premium					
	(1)	(2)	(3)	(4)	(5)	(6)
Household union share	-0.166*** 0.046	-0.442*** 0.118	-0.435*** 0.118	-0.510*** 0.122	-0.438*** 0.102	-0.434*** 0.104
Mean, dept. var.	0.490	0.490	0.490	0.497	0.497	0.497
R-squared	0.707	0.696	0.701	0.678	0.763	0.764
Education Control	Yes	Yes	Yes	Yes	Yes	Yes
Industry Shares	No	No	No	Yes	Yes	Yes
Income covars.	No	No	Yes	Yes	Yes	Yes
Policy covars.	No	No	No	No	No	Yes
Split-Sample IV	No	Yes	Yes	Yes	Yes	Yes
State-spec. quad.	No	No	No	No	Yes	Yes
Min. Year	1940	1940	1940	1940	1940	1940
Max. Year	2009	2009	2009	2009	2009	2009
Observations	1640	1640	1640	1505	1505	1505

Sources: College High School Premium was created using Census and CPS data. See Appendix B details on variable construction.

Notes: This table shows companion specifications to the regressions reported in cols. (1) and (2) in the main Table 3. IV estimates are from split-sample-IV regressions (see Section 7.3 for estimating equations). 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). “Industry shares” controls for state-year share of employment in all one-digit industry categories. “State-spec. quad.” indicates that state-specific quadratic time trends are included. “Income covars.” indicate that state-year GDP and state-year share of households filing taxes are included. “Policy covars.” indicate that state-year minimum wage and a “policy liberalism” index (from Caughey and Warshaw, 2016) are included. Standard errors are clustered at the state level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Appendix Table A.14: State-Year male 90/10 ratio as a function of union density (additional specification)

	Dep't var: Log Percentile 90-10 Men					
	(1)	(2)	(3)	(4)	(5)	(6)
Household union share	-0.133** 0.062	-0.296** 0.118	-0.272** 0.119	-0.361*** 0.116	-0.247*** 0.094	-0.244** 0.095
Mean, dept. var.	1.386	1.386	1.386	1.398	1.398	1.398
R-squared	0.671	0.669	0.685	0.681	0.790	0.790
Education Control	Yes	Yes	Yes	Yes	Yes	Yes
Industry Shares	No	No	No	Yes	Yes	Yes
Income covars.	No	No	Yes	Yes	Yes	Yes
Policy covars.	No	No	No	No	No	Yes
Split-Sample IV	No	Yes	Yes	Yes	Yes	Yes
State-spec. quad.	No	No	No	No	Yes	Yes
Min. Year	1940	1940	1940	1940	1940	1940
Max. Year	2009	2009	2009	2009	2009	2009
Observations	1640	1640	1640	1505	1505	1505

Sources: Log 90-10 Percentile for Men was created using Census and CPS data. See Appendix B details on variable construction.

Notes: This table shows companion specifications to the regressions reported in cols. (3) and (4) in the main Table 3. IV estimates are from split-sample-IV regressions (see Section 7.3 for estimating equations). 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). “Industry shares” controls for state-year share of employment in all one-digit industry categories. “State-spec. quad.” indicates that state-specific quadratic time trends are included. “Income covars.” indicate that state-year GDP and state-year share of households filing taxes are included. “Policy covars.” indicate that state-year minimum wage and a “policy liberalism” index (from Caughey and Warshaw, 2016) are included. Standard errors are clustered at the state level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Appendix Table A.15: State-Year Gini coefficient as a function of union density (additional specification)

	Dep't var: Gini Coefficient					
	(1)	(2)	(3)	(4)	(5)	(6)
Household union share	-0.027** 0.011	-0.063*** 0.024	-0.058** 0.023	-0.076*** 0.025	-0.074*** 0.023	-0.074*** 0.024
Mean, dept. var.	0.376	0.376	0.376	0.378	0.378	0.378
R-squared	0.712	0.710	0.720	0.718	0.783	0.783
Education Control	Yes	Yes	Yes	Yes	Yes	Yes
Industry Shares	No	No	No	Yes	Yes	Yes
Income covars.	No	No	Yes	Yes	Yes	Yes
Policy covars.	No	No	No	No	No	Yes
Split-Sample IV	No	Yes	Yes	Yes	Yes	Yes
State-spec. quad.	No	No	No	No	Yes	Yes
Min. Year	1940	1940	1940	1940	1940	1940
Max. Year	2009	2009	2009	2009	2009	2009
Observations	1640	1640	1640	1505	1505	1505

Sources: The Gini Coefficient was created using Census and CPS data. See Appendix B details on variable construction.

Notes: This table shows companion specifications to the regressions reported in cols. (5) and (6) in the main Table 3. IV estimates are from split-sample-IV regressions (see Section 7.3 for estimating equations). 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). “Industry shares” controls for state-year share of employment in all one-digit industry categories. “State-spec. quad.” indicates that state-specific quadratic time trends are included. “Income covars.” indicate that state-year GDP and state-year share of households filing taxes are included. “Policy covars.” indicate that state-year minimum wage and a “policy liberalism” index (from Caughey and Warshaw, 2016) are included. Standard errors are clustered at the state level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Appendix Table A.16: State-year top-ten income share as a function of union density (additional specification)

	Dep't var: Top 10p. Income					
	(1)	(2)	(3)	(4)	(5)	(6)
Household union share	-2.124** 0.960	-5.342*** 2.065	-6.158*** 2.103	-3.994** 1.727	-3.087** 1.410	-3.147** 1.399
Mean, dept. var.	36.606	36.608	36.608	36.964	36.964	36.964
R-squared	0.782	0.779	0.805	0.832	0.915	0.915
Education Control	Yes	Yes	Yes	Yes	Yes	Yes
Industry Shares	No	No	No	Yes	Yes	Yes
Income covars.	No	No	Yes	Yes	Yes	Yes
Policy covars.	No	No	No	No	No	Yes
Split-Sample IV	No	Yes	Yes	Yes	Yes	Yes
State-spec. quad.	No	No	No	No	Yes	Yes
Min. Year	1940	1940	1940	1940	1940	1940
Max. Year	2009	2009	2009	2009	2009	2009
Observations	3108	3107	3107	2723	2723	2723

Sources: The Top 10 percent share of income data comes from Frank (2015) and is discussed in the text.

Notes: This table shows companion specifications to the regressions reported in cols. (7) and (8) in the main Table 3. IV estimates are from split-sample-IV regressions (see Section 7.3 for estimating equations). 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). “Industry shares” controls for state-year share of employment in all one-digit industry categories. “State-spec. quad.” indicates that state-specific quadratic time trends are included. “Income covars.” indicate that state-year GDP and state-year share of households filing taxes are included. “Policy covars.” indicate that state-year minimum wage and a “policy liberalism” index (from Caughey and Warshaw, 2016) are included. Standard errors are clustered at the state level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Appendix Table A.17: Log state-year GDP per capita as a function of union density

	Dep't var: Log GDP/Cap					
	(1)	(2)	(3)	(4)	(5)	(6)
Household union share	0.071*** 0.023	0.130*** 0.050	0.142*** 0.050	0.140*** 0.050	0.018 0.039	0.020 0.040
Mean, dept. var.	-5.379	-5.378	-5.378	-5.322	-5.322	-5.322
R-squared	0.998	0.998	0.998	0.998	0.999	0.999
Education Control	Yes	Yes	Yes	Yes	Yes	Yes
Industry Shares	No	No	No	Yes	Yes	Yes
Share returns.	No	No	Yes	Yes	Yes	Yes
Other covars.	No	No	No	No	No	Yes
Split-Sample IV	No	Yes	Yes	Yes	Yes	Yes
State-spec. quad.	No	No	Yes	No	Yes	Yes
Min. Year	1940	1940	1940	1940	1940	1940
Max. Year	2009	2009	2009	2009	2009	2009
Observations	3108	3107	3107	2723	2723	2723

Sources: Log State GDP/Cap data comes from Frank (2015).

Notes: IV estimates are from split-sample-IV regressions (see Section 7.3 for estimating equations). 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). “Industry shares” controls for state-year share of employment in all one-digit industry categories. “State-spec. quad.” indicates that state-specific quadratic time trends are included. “Income covars.” indicate that state-year GDP and state-year share of households filing taxes are included. “Policy covars.” indicate that state-year minimum wage and a “policy liberalism” index (from Caughey and Warshaw, 2016) are included. Standard errors are clustered at the state level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Appendix B. Sample Selection and Construction of Key Variables

B.1. Sample Selection

To construct our main Gallup sample, we apply the following selection criteria to the population of recorded Gallup survey respondents from years 1937 through 1987. First, we eliminate respondents to surveys in which the union membership question was not asked. Second, we remove any respondents younger than 21 or older than 64. Third, we remove respondents who identify themselves as farmers. Fourth, we remove respondents who live in Alaska, Hawaii, or Washington D.C.

Our CPS sample is taken from the May supplements in years 1976 to 1981 and both May and March supplements in years 1983 to 2015. We excluded armed forces from our sample. For state-year measures, the state identifiers are only available starting in 1979, and so our CPS-based series begins there. We excluded Alaska, DC and Hawaii from time series, additionally we excluded Idaho from the state-year analysis to make it comparable with the Gallup sample.

B.2. Variable Construction

Union Density In both CPS and Gallup, union density is calculated as the number of households with at least one reported union member between 21 and 64 years old divided by the total number of households. We use the CPS to construct both household union density at the aggregate time-series level as well as the state-year level.

College-High School Share Ratio To construct college-high school share ratio in both CPS and Gallup samples, we first calculate the share of respondents aged 21 to 64 years in each of three categories: high school or less, some college, and college or more. The college-high school share ratio is calculated as the ratio of college share plus half of some college share relative to high school share plus half of some college share in each year:

$$CollegeHSShareRatio_t = \frac{ShareCollege_t + 0.5 * ShareSomeCollege_t}{ShareHighSchool_t + 0.5 * ShareSomeCollege_t} \quad (9)$$

Family Income Our Gallup measure of family income covers years 1942 and 1961 through 1974. Gallup family income is derived from the responses to survey questions of the following form: “Which best represents the total annual income, before taxes, of all the members of your immediate family living in your household?” Responses are coded into income bins which vary across surveys. We construct a harmonized income measure by calculating the midpoint of each interior binned response and multiplying bounds on top- and bottom-coded responses by factors of 1.25 and 0.75, respectively. Our CPS measure of family income is taken from the May and March supplements in years 1978 through 2015. This measures combines all reported income from household members 15 years and older. To construct this variable in early CPS years (May and March before 1990), we use the family income variable, which is binned into 12 categories. For the following years (CPS March only) we use the family income variable non-binned, which reports the total income for the respondent’s family.

College premium, 90-10 wage ratio and Gini coefficient While our RIF analysis constructs measures of aggregate income inequality using family income variable defined above, our time series analysis calculates the college premium, 90-10 wage ratio and Gini coefficient following Goldin and Katz (2009). Each measure uses the log weekly earnings of 18 to 64 year-old full-time and full-year wage and salary workers setting unemployed and NILF to zero. Estimation controls include a full-time dummy, a female dummy, a non-white dummy, a quartic in experience and the interaction of female with non-white and the quartic in experience.

Appendix C. Background on Gallup Data

Gallup Sampling

Documentation for Gallup surveys prior to 1950 describe the sampling procedure as follows:

Prior to 1950, the samples for all Gallup surveys, excluding special surveys, were a combination of what is known as a purposive design for the selection of cities, towns, and rural areas, and the quota method for the selection of individuals within such selected areas. The first step in obtaining the sample was to draw a national sample of places (cities, towns, and rural areas). These were distributed by six regions and five or six city size, urban rural groups or strata in proportion to the distribution of the population of voting age by these regional-city size strata. The distribution of cases between the non-south and south, however, was on the basis of the vote in presidential elections. Within each region the sample of such places was drawn separately for each of the larger states and for groups of smaller states. The places were selected to provide broad geographic distribution within states and at the same time in combination to be politically representative of the state or group of states in terms of three previous elections. Specifically they were selected so that in combination they matched the state vote for three previous elections within small tolerances. Great emphasis was placed on election data as a control in the era from 1935 to 1950. Within the civil divisions in the sample, respondents were selected on the basis of age, sex and socioeconomic quotas. Otherwise, interviewers were given considerable latitude within the sample areas, being permitted to draw their cases from households and from persons on the street anywhere in the community.

Beginning in 1950, Gallup's sampling procedure changed. The following excerpt is taken from post-1950 Gallup survey documentation on sampling:

All Gallup polls since 1950, excluding certain special surveys, have been based on a national probability sample of interviewing areas. Refinements in the sample design have been introduced at various points in time since then. However, over this period the design in its essentials has conformed to the current procedure, as follows:

1. The United States is divided into seven size-of-community strata: cities of population 1,000,000 and over; 250,000 to 999,999; and 50,000 to 249,999; with the urbanized areas of all these cities forming a single stratum; cities of 2,500 to 49,999; rural villages; and farm or open country rural areas.
2. Within each of these strata, the population is further divided into seven regions: New England, Middle Atlantic, East Central, West Central, South, Mountain, and Pacific Coast.
3. Within each size-of-community and regional stratum the population is arrayed in geographic order and zoned into equal-sized groups of sampling units.

4. In each zone, pairs of localities are selected with probability of selection proportional to the size of each locality's population—producing two replicated samples of localities.
5. Within selected cities for which population data are reported by census tracts or enumeration districts, these sample subdivisions are drawn with probability of selection proportional to the size of the population.
6. For other cities, minor civil divisions, and rural areas in the sample for which population data are not reported by census tracts or enumeration districts, small, definable geographic areas are drawn, with the probability of selection proportional to size where available data permit; otherwise with equal probability.
7. Within each subdivision selected for which block statistics are available, a block or block cluster is drawn with probability of selection proportional to the number of dwelling units.
8. In cities and towns for which block statistics are not available, blocks are drawn at random, that is, with equal probability.
9. In subdivisions that are rural or open country in character, segments approximately equal in size of population are delineated and drawn with equal probability.
10. In each cluster of blocks and each segment so selected, a randomly selected starting point is designated on the interviewer's map of the area. Starting at this point, interviewers are required to follow a given direction in the selection of households, taking households in sequence, until their assigned number of interviews has been completed.
11. Within each occupied dwelling unit or household reached, the interviewer asks to speak to the youngest man 18 or older at home, or if no man is at home, the oldest woman 18 or older. This method of selection within the household has been developed empirically to produce an age distribution by men and women separately which compares closely with the age distribution of the population. It increases the probability of selecting younger men, who are at home relatively infrequently, and the probability of reaching older women in the household who tend to be under-represented unless given a disproportionate chance of being drawn from among those at home. The method of selection among those at home within the household is not strictly random, but it is systematic and objective and eliminates interviewer judgement in the selection process.
12. Interviewing is conducted at times when adults are most likely to be at home, which means on weekends or if on weekdays, after 4:00 p.m. for women and after 6:00 p.m. for men.
13. Allowance for persons not at home is made by a "times-at-home" weighting procedure rather than by "call-backs." this procedure is a standard method for reducing the sample bias that would otherwise result from underrepresentation of persons who are difficult to find at home.

14. The pre-stratification by regions is routinely supplemented by fitting each obtained sample to the latest available census bureau estimates of the regional distribution of the population. Also, minor adjustments of the sample are made by educational attainment (by men and women separately), based on the annual estimates of the census bureau derived from their current population survey. The sampling procedure described is designed to produce an approximation of the adult civilian population living in the United States, except for those persons in institutions such as hospitals.

Gallup Weighting Procedure

To construct weights, we use post-stratification methods (i.e., cell-weighting). Specifically, we weight observations in the Gallup data so that the annual proportions of education-race-region cells in Gallup match the corresponding proportions in U.S. Census data. The process involves several steps: First, we construct comparable measures of education (less than high school, high school graduate, some college, college graduate), race (white, non-white), and region (South, non-South) in both Gallup and Census data. Second, we construct annual proportions of each education-race-region cell for each dataset. In the Census data, we apply representative household weights and linearly interpolate values for intercensal years to best approximate the “true” annual proportions of each cell. Third, we generate cell-specific weights w_{ct} by applying the following formula:

$$w_{ct} = \frac{\pi_{ct}^C}{\pi_{ct}^G} \quad (10)$$

where c denotes a particular education-race-region cell (e.g., white Southerners with a college degree), and π_{ct}^C and π_{ct}^G denote annual cell proportions for Census and Gallup, respectively. Finally, we let $w_{it} = w_{ct}$ for each respondent i in year t corresponding to cell c in the Gallup data and re-normalize so $\sum_i^{N_t} \frac{w_{it}}{N_t} = 1$ for each year t .

We repeat the procedure above for several alternative cell definitions (e.g., education-race-age-state, age-gender-region). Our preferred weights use education-race-region cells because we find this definition makes our sample as representative as possible without compromising comparability across surveys or creating excessively small or “empty” cells.⁴⁶ For surveys without education data, we use race-region weights.

Comparing Gallup to Census Microdata

We begin with Gallup data from 1950 onward, returning shortly to earlier data. Table A.1 compares Gallup data to 1950–1980 Census data. To summarize how the *actual* (unweighted) Gallup observations compare to the full U.S. adult population, we compare unweighted Gallup data to Census IPUMS tabulations. Given Gallup’s well-documented under-sampling of the South, we show results separately for Southern and non-Southern states.

⁴⁶For a more thorough discussion of post-stratification weighting, including optimal cell “finess,” see Berinsky (2006b)

In 1950 and 1960, Gallup exhibits some under-sampling of the South, but, by 1970, this bias had disappeared. From 1950 to 1960, Gallup under-sampled blacks in both the South and the Non-South. This bias continued in the South through 1970, to a smaller degree. These biases reflect the substantial disenfranchisement of blacks, particularly in the South during this period. Age and gender appear representative in Gallup in both regions in each decade.

Gallup respondents outside the South are more educated than their Census counterparts, with the largest gap being a high school completion difference of ten percentage points in 1950. In the South, except for 1950, Gallup and IPUMS show similar levels of education. Gallup Southern respondents have higher high school completion rates than those in the Census in 1950, not surprising as it was still under-sampling Southern blacks in that year. Later in the paper we will show results with the Gallup data both unweighted and weighted to match Census characteristics, but Table A.1 gives some sense of how much “work” the weights must do.

Table A.2 looks separately at 1940, given that Gallup’s sampling procedures were quite different during its earlier years. In fact, in 1940, very few Gallup surveys ask about education (the summary statistics we present for that variable are based on only 5,767 observations), so in this table we include occupation categories as supplemental proxies for socio-economic status. The first column shows, again, unweighted Gallup data. Col. (2) presents summary statistics for all adults in the 1940 IPUMS. Perhaps the most striking discrepancy is gender: consistent with their stated methodology at the time, Gallup over-samples men. Col. (3) adjusts the Census sampling so that men are sampled at the Gallup frequencies and also down-weights large households (since Gallup only interviews one person per household). Comparing col. (1) versus (3) shows, as expected, that Gallup significantly under-samples the South.

Consistent with concerns about Gallup over-sampling the affluent, Gallup respondents in 1940 are substantially more educated than their Census counterparts. Unfortunately, given that only in 1942 does Gallup begin to regularly include an education question, the Gallup sample for which we have an education measure in 1940 is quite small (about 5,700 individuals, relative to over 150,000 for the other Gallup variables in 1940). Given the small education sample in 1940, we use occupational categories to further explore socio-economic status in Gallup versus the 1940 Census. Gallup and IPUMS use different occupation categories—Gallup’s are much coarser and unfortunately IPUMS categories do not completely nest Gallup categories—so comparisons are not straightforward. Consistent with the concerns cited earlier that Gallup over-sampled the well-to-do, Gallup respondents appear to have slightly higher-status occupations relative to their Census counterparts. For example, “professionals” and “proprietors, managers, officials” appear more numerous in Gallup (these categories are especially useful because IPUMS categories fully nest these Gallup occupations). Reassuringly, farmers and farm laborers are similarly represented in both samples (these two Gallup categories are also fully nested in IPUMS categories, again easing comparisons across data sources).

For the most part, these patterns hold when we drop Southern states from both samples (the final two columns of Table A.2). Importantly, outside of the South, Gallup appears to sample blacks in proportion to their population, even in the very early years of its existence. Also, outside the South, Gallup appears to accurately sample the remaining six regions of

the US.⁴⁷

In general, we show results with Gallup data using weights to match (interpolated) Census IPUMS summary statistics, even though the need for weights is not obvious after 1950 or 1960. From 1937 until 1941, we weight so that Gallup matched the IPUMS in terms of *White* $\times$ *South* cells, given that the summary statistics show that Gallup sampling along these dimensions appears suspect in the early years. Beginning in 1942 (the first year in which Gallup surveys ask the union and education questions in the same survey) we weight by *White* $\times$ *Education* $\times$ *South*, where *Education* $\in$ {No high school degree, HS degree, Some college, College graduate}, thus giving us $2 \times 4 \times 2 = 16$ cells on which to match. In practice, however, our results are very similar with and without weights.

⁴⁷We use Gallup-defined geographic regions in this table.

Appendix D. Existing Measures of Union Density Pre-Dating the Current Population Survey

The CPS first asks respondents their union status in 1973, and then only in selected months until 1983 from which time information on union status was collected each month in the CPS as part of the outgoing rotation group supplement. Before this survey, the primary sources for union density are the BLS and Troy/NBER historical time series mentioned in the introduction. The data underlying these calculations are union reports of membership and dues revenue when available, and a variety of other sources when not available. Neither of these data sources ever used representative samples of individual workers to calculate union density.

In general, the data derived from union reports likely become more accurate by the 1960s. Post-1959 the BLS collected mandatory financial reports from unions as a condition of the Labor-Management Reporting and Disclosure (Landrum-Griffin) Act, and Troy and Sheflin (1985) incorporate these data into their estimates of union density. Beginning in 1964, the BLS disaggregates union membership counts by state, and Hirsch *et al.* (2001) splice these reports together with the CPS to form state-year union density panel beginning in 1964 and continuing through today.⁴⁸

Before the 1960s, however, union data were far less standardized. In the remainder of this section, we detail the methodology of the two most widely used data sources on aggregate union density: the BLS and Troy series.

D.1. The BLS Estimate of Early Union Density

The BLS series is based on union-reported membership figures starting in the late 1940s. Prior to 1948, the methodology for calculating union membership does not appear standardized. For example, the 1945 Monthly Labor Report notes as its sources: “This study is based on an analysis of approximately 15,000 employer-union agreements as well as employment, union membership, and *other data available to the Bureau of Labor Statistics* [emphasis ours]” (Bureau of Labor Statistics, 1945)⁴⁹

It is obviously hard to verify information from unspecified “sources available to the BLS” but even in instances where the BLS can rely on union membership reports, concerns arise. A key issue is that unions had important incentives to over-state their membership and until the late 1950s faced no penalty for doing so. In the early and mid-1930s, the main umbrella organization for local unions was the American Federation of Labor (AFL). They were often charged with over-stating their membership, presumably to inflate their political influence.

⁴⁸Freeman *et al.* (1998) constructs a time-series of union density from 1880 to 1995, splicing together the official series from the BLS with series constructed from the CPS. Freeman reports alternative series constructed by other scholars (Troy (1965), Troy and Sheflin (1985), Wolman (1924), and Galenson (1960)) in the Appendix to his paper.

⁴⁹For example, one alternative source the BLS used was convention representation formulas. “Convention formulas” specified the number of seats, as a function of membership, each union would have at the umbrella organization convention. Inverting this formula and using the convention records, rough estimates of union membership could be formed.

For example, a 1934 *New York Times* story casts doubt on the AFL’s claim to represent over six million workers, noting that “complete and authoritative data are lacking” and that the figures provided by the AFL “are not regarded as accurate.”⁵⁰ Individual unions also had an incentive to inflate the numbers they reported to the AFL. For example, the number of seats each union would receive at the annual convention was based on a formula to which membership was the main input.

If anything, these incentives to over-report likely grew after 1937, when the Committee on Industrial Organization broke away from the AFL to form a rival umbrella organization, the Congress of Industrial Organizations (CIO). Both federations of labor, the AFL and CIO, now competed for local unions to join their umbrella organizations, as well as for sympathies of government officials, tasks that were aided by a public perception that the federation was large and growing. Based on our read of *New York Times* articles on unions in the late 1930s and early 1940s, one of the most common if not the most common topic is the conflict between the two federations.⁵¹ Individual unions still had incentives to compete for influence within their given federation, and thus inflate membership.

Membership inflation became such an issue that the federations themselves may not have known how many actual members they had. In fact, the CIO commissioned an *internal* investigation into membership inflation, conducted by then-United Steelworkers of America president Philip Murray. Murray’s 1942 report concluded that actual CIO membership was less than fifty percent of the official number the federation was reporting. (Galenson, 1960)

D.2. The Troy Estimates of Early Union Density

In his NBER volumes estimating union density, Troy is well aware of the problems documented above with the BLS estimates. For this reason, he defines membership as “dues-paying members” and proceeds to estimate union membership using unions’ financial reports where available, presumably under the assumption that financial reports were less biased than membership reports. For each union, he divides aggregate union dues revenue by average full-time member dues to recover an estimate of union membership. While Troy is cognizant of the limitations of his data and methodology, he believes the biases are largely *understating* union membership (e.g. some groups, such as veterans, pay lower than average or no dues).

But union financial reports, like membership reports, are also not verified until the late 1950s. Nor is it obvious that union revenue data are not similarly inflated (in fact, the AFL accused the CIO of lying about their income data, as we mention in footnote 51). Moreover, revenue data are largely incomplete for the 1930s and 1940s. For example, in his

⁵⁰See, “*Organized Labor is Put at 6,700,000*”, *New York Times*, May 1935. reporting that “For one thing, complete and authoritative data are lacking, and this is especially true during times of depression, when some unions drop unemployed workers from the rolls and exempt them from paying dues. . . . The [AFL] reported an average membership of 2,609,011 for the year ended Aug. 31, 1934. These official figures, which are not regarded as an accurate measure of the movement, are far below the peak figure of 4,078,740 for 1920.”

⁵¹As just one example, a 1938 *NYT* headline and subtitles read: “Green Says Lewis Falsified Report; A.F.L. Head Alleges Statement on C.I.O. membership is an ‘Amazing Inflation; Questions Income Data,” referring to AFL head William Green and CIO head John Lewis, respectively.

1940 estimates, Troy (1965) notes that the sources for 54.4% of his total is *not* in fact from financial reports, but instead an “Other” category, which includes personal correspondence with unions, asking their membership.⁵² As such, for these early years, the Troy data in fact appears to face the same issue with membership-inflation as does the BLS data.⁵³

In addition, Troy imputes the membership of many CIO unions in the late 1930s and 1940s by assigning them the membership of their AFL counterpart in the same sector.⁵⁴ This procedure likely over-states CIO membership, given that the AFL was believed to be twice as large as the CIO during this period (we also find this 2:1 ratio in our Gallup data), though obviously that average ratio may vary by sector.

In summary, while a likely improvement over the BLS series, it is difficult to believe that Troy’s estimates (or Troy and Sheflin (1985)) are without extensive mismeasurement. Given the limitations of the existing pre-CPS data on union density, in the next section we introduce a new source: Gallup and other opinion surveys.

D.3. Other Pre-CPS State-Year Measures of Union Density

The only sources of state-year data on union density prior to the CPS we are aware of are measures created by Hirsch *et al.* (2001) from BLS reports (which begin disaggregating union membership regionally, often by state, in 1964) from 1964-1977, and measures created by Troy and Sheflin (1985) for the years 1939 and 1956. Our Gallup measure is quite highly correlated (correlation = .724) with the existing Hirsch-Macpherson measures (individual union density as a fraction of non-farm employment) for the 1964-1986 years, which are where there is overlap. This correlation increases to .75 when we restriction attention to the CPS years with state identifiers (1978-1986).

The historical Troy measures for 1939 and 1956 are constructed from even more fragmentary records than the annual series we discuss above (as many union reports did not disaggregate either revenue or membership by state), we are also correlated with these data in both cross-sections and changes (1939 correlation = 0.78, 1953 correlation = 0.75, correlation in changes = 0.5).

⁵²“Other” is down to 10% by 1960 (Troy (1965)).

⁵³Troy (1965) also only presents validation exercises for his post-1950 data, comparing reported measurement with that inferred from dues receipts for the Chemical and Rubber Workers in 1953, leaving it open whether the BLS or Troy (or neither) is correct for the pre-1950 series.

⁵⁴From Troy (1965) [pp. A53]: “The average membership per local industrial union is arbitrarily estimated to be 300, and this figure is multiplied each year by the number of such unions reported by the CIO. The estimate of an average membership of 300 is deemed a fair one since the average membership of the local trade and federal labor unions of the AFL, a class of unions similar to the local industrial unions of the CIO, varies from a low of 82 in 1937 to a high of 193 in 1948.”

Appendix E. Details on RIF Regressions

In this appendix section we outline the recentered influence function approach of Firpo *et al.* (2009), adapted for our outcome measures. The relevant property of a recentered influence function is that its expectation equals the distributional statistic of interest. For quantile τ denoted Q_τ , the quantile RIF is given by $RIF(y, Q_\tau) = Q_\tau + \frac{\tau - \mathbf{1}(y_h < Q_\tau)}{f_y(Q_\tau)}$, and taking expectations verifies that $E[RIF(y, Q_\tau)] = Q_\tau$.

Firpo *et al.* (2009) show that a regression of the RIF on covariates yields the approximate effect of the covariates on the distributional statistic of interest (applied to the unconditional distribution). The intuition is that the marginal effect of an increase in union density to a given statistic is given by the average effect of each union individual on that statistic (each observation’s “influence”). The estimate for β in equation (5) is therefore the effect of a change in union density on the probability that a household’s income is less than the value of the quantile τ , i.e., $\frac{dF(Q_\tau)}{dU_{union}}$, divided by the density of household income at Q_τ ($\frac{dF(Q_\tau)}{dQ_\tau}$). The resulting coefficient thus measures $\frac{dQ_\tau}{dU_{union}}$ the marginal change in the value of the quantile at τ in response to a small change in $Union$. The RIF of the Gini is not particularly illustrative and we omit it here. For further examples and exposition of distributional regressions see Havnes and Mogstad (2015) and Dube (2017).

Extending this to the 90-10 ratio immediately suggests the following RIF, and inspection confirms that $E[RIF(y_h, Q_{.9} - Q_{.1})] = Q_{.9} - Q_{.1}$.

$$RIF(y_h, Q_{.9} - Q_{.1}) = Q_{.9} - Q_{.1} + \frac{.9 - \mathbf{1}(y_h < Q_{.9})}{f_y(Q_{.9})} - \frac{.1 - \mathbf{1}(y_h < Q_{.1})}{f_y(Q_{.1})} \quad (11)$$

One limitation of our historical data is that the income distribution is binned, sometimes very coarsely, so the kernel density used to estimate f will be biased. We can assess this using the CPS where we have both the unbinned family income (which is still topcoded) as well as the binned version. While the exact magnitude and precision of the RIF estimates are sensitive to the binning procedure, the qualitative pattern in the 90-10 ratio is not. Appendix Figure A.13 shows the percentiles with unbinned data; while the effect on the median is unchanged, the effect on the 90th and the 10th percentiles are both more positive than the binned analogue in Appendix Figure A.12, which roughly cancels out in the corresponding 90-10 ratio.