Why Unions Still Matter: The Effects of Unionization on the Distribution of Employee Earnings

Brigham R. Frandsen*

Massachusetts Institute of Technology

January 30, 2012

Abstract

This paper estimates the causal effect of unionization on the distribution of employee earnings using a regression discontinuity design that links administrative records on individual earnings to union certification election results. The results suggest unions raise the lower end of the distribution by around 30 log points, with a much smaller effect on the upper tail, and a modest effect on average earnings. Estimates of average effects by baseline earnings quantile suggest the distributional effects correspond to individual-level earnings effects that vary by skill. Unionization also appears to reduce employment of the lowest skilled workers. These results are consistent with a model of union wage setting in which unions set wages so as to maximize the probability of certification, subject to a minimum profit constraint for the employer. The optimal union wage schedule pays low-skilled union members above marginal product but reduces the return to skill. The estimates suggest that about one quarter of the increase in the variance of log earnings from 1979 to 2009 can be accounted for by falling U.S. private sector unionization rates, a larger fraction than earlier studies have found.

*This research uses data from the Census Bureau's Longitudinal Employer Household Dynamics Program, which was partially supported by the following National Science Foundation Grants: SES-9978093, SES-0339191 and ITR-0427889; National Institute on Aging Grant AG018854; and grants from the Alfred P. Sloan Foundation. Any opinions and conclusions expressed herein are those of the author and do not necessarily represent the views of the U.S. Census Bureau. All results have been reviewed to ensure that no confidential information is disclosed. The author is indebted to JP Ferguson, Jim Davis, Javi Miranda, and Kristin McCue for help with the data. The author also gratefully acknowledges support from NSF Grant SES-0922355.

1 Introduction

How do unions affect workers' pay and employment? This question is at the heart of the debate over the economic consequences of unionization. The effects of unionization are important both for labor relations policy changes, such as card check rules that make it easier for workers to organize, and to explanations of increasing U.S. inequality. Unions affect pay both by changing average wages—a between groups effect—and by changing the distribution of wages within sectors. If, as many labor economists believe, the net effect of unionization is wage compression, then the persistent decline in private sector unionization may indeed account for a substantial portion of increasing inequality.[1]

Much of the long history of research on unions and worker pay has attempted to quantify the distributional impact of unionization. Most studies have focused on comparisons between unionized and nonunionized workers, controlling for observable characteristics, and have found a 10 to 20 percent average wage gap.[2] Comparisons across the distribution have found larger union wage gaps at the lower end of the distribution, consistent with compression in the overall earnings distribution (Chamberlain, 1994; Card, 1996; DiNardo, Fortin, and Lemieux, 1996; Card, Lemieux, and Riddell, 2004).

A question raised by these studies is whether the union-nonunion differences reflect causal effects or unobserved factors. Recent quasi-experimental studies focusing on *employer* outcomes have found small effects for unions on the margin of certifying, casting doubt on the interpretation of large average union-nonunion differences as causal (DiNardo and Lee, 2004; Lee and Mas, forthcoming).

Do these quasi-experimental findings mean unions have minimal economic effects? Perhaps not. Small employer-level average effects can mask substantial effects on individual workers through offsetting distributional effects or effects on worker composition. The economic consequences of unionization, especially the contribution of unionization to changes in inequality, depend crucially on individual-level effects, but these are missed by approaches looking only at employer-level outcomes, a gap this paper attempts to fill.

The purpose of this paper is to identify the causal effect of unionization on the distribution of employee earnings. The target of estimation can be understood in terms of a hypothetical

[1] See Card, Lemieux, and Riddell (2004) for a survey of the long literature on unions and inequality.

[2] Early studies along these lines were survey by Lewis (1986), and revisited using more recent data and methodologies by Blanchflower and Bryson (2003).

experiment where workers are randomly assigned by establishment to be unionized or not. The causal effect of interest is the difference between the subsequent distribution of earnings among employees at the unionized and non-unionized establishments. To approximate this hypothetical experiment, this paper uses a regression discontinuity (RD) design based on union certification elections, an approach introduced by DiNardo and Lee (2004). Here I go beyond DiNardo and Lee by looking at individual earnings instead of plant-level averages, by assessing employment effects, and by looking at wage distributions as well as means. The individual-level analysis is based on administrative records on individual earnings matched to establishment-level election results. A key feature of the research design is that *all* employees at an establishment at the time of a union election are followed, whether or not they are union members, and whether they stayed employed at the same establishment, changed to a new job, or became unemployed. Following individuals allows estimation of the union effect on employment status and job switching, something not possible using traditional cross-section approaches. Including all employees at the plant captures potentially important effects on workers excluded from the bargaining unit that would be missed by focusing only on union membership status.

Consistent with DiNardo and Lee (2004), the RD estimates reported here show little effect of union certification on average wages. At the same time, the results provide clear evidence of a distributional effect. Unionization raises the lower tail of the earnings distribution by around 30 log points, with much smaller, possibly negative, effects at the upper end. Further empirical results show unionization decreases retention among the lowest-skilled employees, but may increase retention among mid- to low-skilled workers, with little effect on workers in the upper half of the skill distribution. These results are consistent with movement along a downward sloping demand schedule for less-skilled workers, though they also suggest cross-skill substitution possibilities are limited.

The causal effects identified in the empirical work are specific to unions on the margin of winning or losing a certification election. Vote tally data show that most elections are decided by a relatively narrow margin, so in this sense the estimates reflect the typical case. Effects local to the threshold of union victory are also of policy interest since labor policy changes currently under debate would be most likely to affect marginal elections.[3]

[3]Two recently proposed policies in the private-sector are the Employee Free Choice Act (EFCA, also known as

On the methodological side, I use Frandsen, Frölich, and Melly's (2010) recently developed procedure for estimating quantile treatment effects in the regression discontinuity design. The quantile-based analysis is juxtaposed with results from a simpler strategy estimating average treatment effects conditional on baseline earnings rank, an approach that takes advantage of the panel structure of the earnings data. This juxtaposition provides a test of the rank invariance assumption needed to interpret the quantile treatment effects as informative about individual-level effects, a conclusion that could not be drawn looking at quantile treatment effects alone. The approach of conditioning on baseline earnings quantile also allows estimation of effects on other outcomes across the earnings distribution, something not possible with quantile regression methods.

The next section describes the U.S. private sector unionization process. Section 3 describes the data used in the empirical work. Section 4 lays out the research design and the econometric framework for identifying and estimating the effect of unionization on the distribution of earnings, and section 5 presents the estimation results. Section 6 interprets the results in terms of a model of union electoral competition and section 7 summarizes the findings and concludes.

2 Background

Since 1935, most U.S. private sector unionization has been governed by the National Labor Relations Act (NLRA), which specifies the rights of unionizing workers. While an employer may voluntarily bargain with the workers' chosen representative, or in some cases may be required to do so even without an election, the traditional process by which workers unionize is through a National Labor Relations Board (NLRB) secret ballot election.[4] Although in practice an organizing drive is often fraught with disputes and delays, the following steps describe the typical path a group of workers follows to form a union:[5]

1. Petition drive: Union organizers lobby workers, collect signatures expressing a desire to hold

the Card Check Rule), and a National Labor Relations Board representation case procedure change, both of which would likely increase the odds of union representation at establishments near the margin (Greenhouse, 2009; National Labor Relations Board, 2011).

[4]Secret ballot election has historically been the dominant form of new unionization, although in recent years voluntary recognition through neutrality agreements and card checks have become more common. (Brudney, 2005)

[5]The simple process laid out here follows the procedures described in NLRB (2010). See Ferguson (2008) and DiNardo and Lee (2004) for a more complete description of the possible complications and objections that can be raised at each step.

an election, and submit a petition to the NLRB to hold an election. If the petition is accepted, the NLRB ascertains the scope of the bargaining unit and sets the election time and place, usually the workplace.

2. Election: Eligible workers vote for or against the union, and the union wins if it receives a simple majority (50 percent + 1) of the votes cast.

3. Certification: If the union wins, the NLRB certifies it as the sole authorized representative of the workers in the bargaining unit, and requires the employer to bargain "in good faith" with the union.

4. Bargaining: The employer negotiates with union representatives over a collective bargaining agreement. If an agreement is reached—which does not always happen—the contract becomes binding for all employees in the unit.

NLRB certification elections may include two or more competing unions on the ballot. In the case of multiple competing unions, a simple majority is still required for certification. Elections may also be held to remove union representation altogether (decertification) or to replace one union with another. These cases, however, occur relatively infrequently, and the analysis focuses on certification elections.

3 Data

3.1 Union Elections and Contracts

The analysis uses a dataset on the universe of NLRB union representation election results from 1990 to 2007, which was compiled from data obtained from Hank Farber, J.P. Ferguson, and Thomas Holmes.[6]. Each record in this dataset represents a union certification election held at an establishment, and includes the number of votes cast for and against union representation, the date of the election, and the employer's name and address. Out of the universe of NLRB elections, those elections held in the states and years where individuals earnings data are available were retained and matched to Census earnings data (see below).

[6]The union data from Thomas Holmes is available on his website, http://www.econ.umn.edu/~holmes/data/geo_spill/

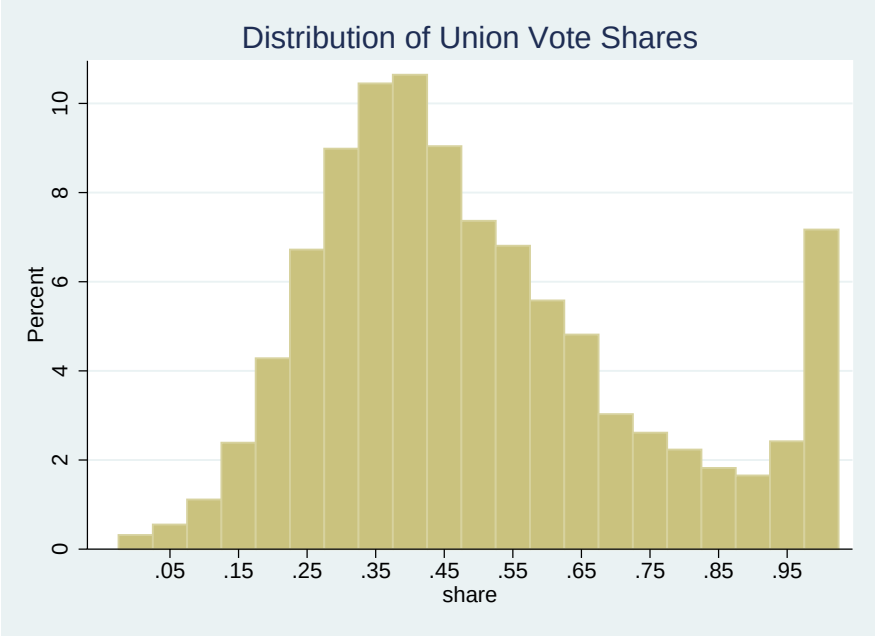

Figure 1: Votes-weighted histogram of the union vote share in representation elections from 1992 to 2001. Data are from NLRB election records, restricted to elections where 10 or more votes were cast.

The matched union elections dataset contains data on 12,894 representation elections, involving around 785,000 votes cast. The average union share was 56.5 percent and unions won 54.2 percent of the elections (52.9 percent employee-weighted). Figure 1 shows the distribution of the union vote share in the sample. The mode is around 40 percent, with a significant number of elections in which the union received all votes.[7] Figure 2 shows the distribution of the union margin of victory in terms of number of votes. This figure shows that close elections represent the typical case, a fact that is important for the interpretation of the estimation results.

Data on an important outcome of union representation—the presence of a collective bargaining agreement—was inferred from Federal Mediation and Conciliation Service (FMCS) records on contract expiry and renewal, obtained from the FMCS and Thomas Holmes.[8] In principle, employers and unions are required to provide notice to the FMCS within 60 days of a contract expiration or modification, although in practice there is significant undercount (DiNardo and Lee, 2004). Still, these data provide a lower bound on an important effect of union representation on collective bargaining. Each FMCS record contains the expiration date of the contract being reported, the

[7]The figure corresponds to Figure II in DiNardo and Lee (2004), who use similar data

[8]Recent FMCS records on contract expiry are available from the FMCS website: http://www.fmcs.gov/internet/downloadsList.asp?categoryID=275.

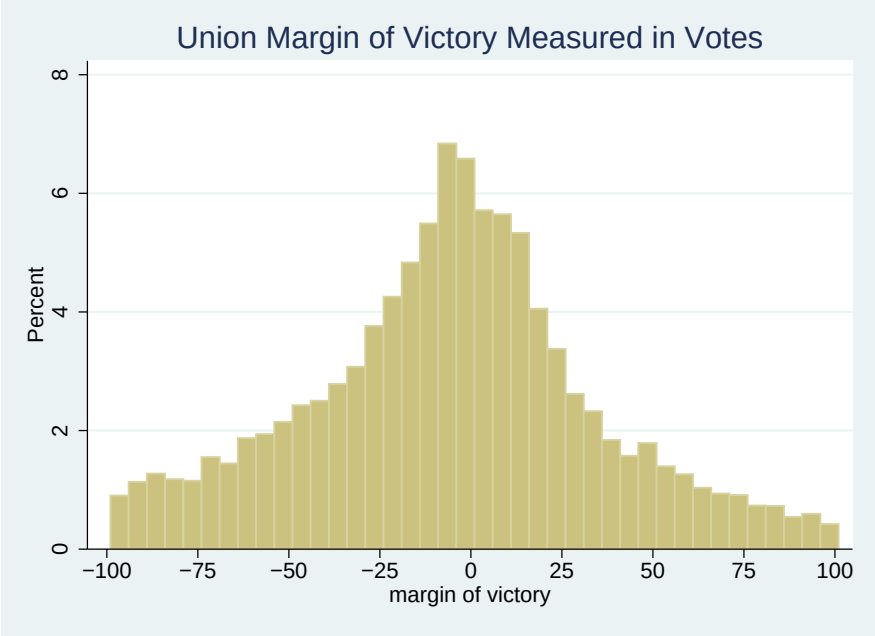

Figure 2: Votes-weighted histogram of the union margin of victory (in terms of number of votes) in representation elections from 1992 to 2001. Only elections decided by 100 or fewer votes are included in the histogram. Data are from NLRB election records, restricted to elections where 10 or more votes were cast.

employer's name and address, and the affected unit's location. These contracts data were matched via name and address to elections records to form an indicator for each election of whether a collective bargaining agreement was reported.

3.2 LEHD

Individual-level earnings were obtained from the U.S. Census Bureau's Longitudinal Employer-Household Dynamics (LEHD) database. The LEHD integrates the universe of unemployment insurance-covered (UI) earnings records held by participating state agencies into a cohesive data structure using person and employer identifiers, allowing linkages to other sources of data.[9]

The Employment History Files (EHF) within the LEHD contain quarterly records of individuals' UI-covered earnings. The EHF for each of the 30 covered states contains a record for each employee-employer combination—a job—that produced at least one dollar of wages in that state in each year. The data cover a period as wide as 1985 to 2008, although for most states the data only go back

[9]For more details on the construction and uses of the LEHD database, see McKinney and Vilhuber (2008), Lane (2008), Abowd, Haltiwanger, and Lane (2004), and Abowd, Stephens, Vilhuber, Andersson, McKinney, Roemer, and Woodcock (2009).

to the early 1990s. The EHF contains more than 2.8 billion records, although I focus on workers employed at election establishments just before the time of the election.

Workers employed at the time of the election in each establishment in the NLRB election dataset can be identified by matching employer name and address information from the election record with employer information in the LEHD. The data appendix describes in detail the procedure used for matching the two datasets. The matching procedure identified over 1.3 million individuals who were employed at establishments just before the time a union election was held from 1990-2007. The subsequent earnings and employment histories of these individuals constitute the main outcomes of interest. Individuals who leave the election establishment are followed and their UI-covered earnings elsewhere are included as outcomes. For supplementary analysis on the effect of unionization on earnings actually paid at election establishments the outcomes are the earnings of all workers employed at election establishments at a given time after the election.

Table 1 reports summary statistics on pre-election earnings, post-election earnings, indicators for positive earnings (a measure of employment status) and retention by union status for workers in the sample. Average (post-election) annual earnings—defined as the sum of the four quarterly earnings beginning six months after the election—in this sample is around $31,000 (deflated to 2000 dollars). Unionized workers earn on average slightly more than non-unionized workers, consistent with the large literature on union wage gaps, although the difference here is smaller than the union premium typically estimated. This wage gap, however, reflects both the causal effect of unionization, as well as a selection effect. The $1,100 pre-election earnings gap suggests the selection effect may be substantial. The table also reports statistics for the sample restricted to close elections (+/- 10 percent). Most of the union relative earnings effect disappears in the restricted sample, consistent with bias in the full-sample comparisons by union status, and with DiNardo and Lee's (2004) findings of a small average effect of unionization. However, this small average effect could be masking significant, but offsetting effects elsewhere in the distribution. The table also shows a statistically insignificant (p-value=.28) but noticeable difference in pre-election earnings in the restricted sample. While this difference likely occurred by chance (the restricted sample represents 1,330 elections), it suggests the need to control for pre-election earnings in the analysis. Post election employment is about 95 percent for all groups, although only around 60 percent of workers remained employed at the election establishment by six quarters after the election.

Table 1: Pre-election Earnings and Outcomes by Unionization Status

	Full Sample			Discontinuity Sample		
	All	Non-unionized	Unionized	All	Non-unionized	Unionized
Annual earnings	31,013	30,782	31,225	31,700	31,702	31,697
Pre-election earnings	31,238	30,692	31,739	32,175	32,596	31,564
Post-election employment	.950	.952	.947	.951	.953	.948
6-quarter retention	.604	.607	.600	.611	.595	.635
N	1,354,036	648,903	705,133	427,567	253,074	174,493

Notes: the table reports means of the variables at left by union status for workers employed at plants in LEHD-covered states where a union election was held between 1990 and 2007. Earnings variables are measured in 2000 dollars. Annual earnings are defined at the sum of the four quarterly earnings starting two quarters after the union election closed. Pre-election earnings are the sum of the four quarterly earnings prior to the union election closing date. Post-election employment status is defined as an indicator for positive annual earnings. Stayed employed at plant is an indicator for remaining employed at the plant 6 quarters after the election closed. The discontinuity sample restricts to elections that were decided by 10 percent or less.

4 Research Design and Econometric Framework

A fundamental obstacle to measuring the effect of unionization on earnings is selection bias: earnings within unionized plants may differ for reasons other than union representation. This study seeks to overcome selection issues by using the regression discontinuity (RD) research design, originally developed by Thistlethwaite and Campbell (1960). The motivation for the design, first used in this context by DiNardo and Lee (2004), is that if plants and workers where the union barely won and barely lost are comparable, then close union elections approximate a randomized experiment, and the resulting difference in the distribution of earnings provides a reliable estimate of the causal effect of unionization.[10] To formalize this idea, let $D = 1(R > 0)$ be an indicator for a union victory, where R is the union margin of victory (negative for losses). Unionization here is a plant-level treatment, which potentially affects all workers at the establishment, in or out of the bargaining unit. Let Y_1 be an individual's earnings if the union were to win, and let Y_0 be the earnings otherwise, so that observed earnings is $Y = Y_0 + (Y_1 - Y_0) D$.

The distributional effects of unionization are captured by the quantile treatment effect, or the

[10]See Lee (2008) for further discussion on the conditions under which a close election provides as-good-as-randomized variation.

difference between the quantiles of potential earnings at the margin of union victory:

$$QTE(\tau) \equiv Q_{Y_1|R=0}(\tau) - Q_{Y_0|R=0}(\tau). \tag{1}$$

While the quantile treatment effects correspond only to the effect of marginally certifying unions, Figure 2 implies that this reflects the typical case. This is also the relevant effect for analyzing the impact of recently proposed policy changes which would primarily impact relatively close certification elections (Greenhouse, 2009; National Labor Relations Board, 2011). Quantile treatment effects do not necessarily correspond to the effect on any particular individual, but they do capture the distributional effects for impacted workers, which is what traditional welfare analysis depends on (Atkinson, 1970).

A closer link to effects on particular individuals is captured by the average treatment effect conditional on pre-election earnings rank, given by

$$\beta_Y(\tau) \equiv E\left[Y_1 - Y_0 | F_{Y_{-1}|R=0}(Y_{-1}) = \tau, R = 0\right], \tag{2}$$

where $F_{Y_{-1}|R=0}$ is the cdf of pre-election earnings at the threshold. This gives the expected effect of unionization for a worker at a given rank in the earnings distribution prior to unionization. If unionization does not systematically affect workers' relative ranks in the earnings distribution and ranks are stable over time among a group of workers, then (2) will coincide with the quantile treatment effect, (1), a point made formally in the appendix. Comparing estimates of these two parameters is a test of whether this rank similarity holds, and therefore whether the quantile treatment effect can be interpreted as an individual-level effect.[11] The effect defined by (2) can be also be generalized to describe the effects on other outcomes, W, conditional on pre-election earnings rank: $\beta_W(\tau) = E\left[W_1 - W_0 | F_{Y_{-1}|R=0}(Y_{-1}) = \tau, R = 0\right]$, where W might be measures of employment status or retention at an establishment. Comparing these effects for different outcomes at the same earnings rank will help interpret the results in terms of economic responses of employers and employees to unionization.

[11] The importance of rank similarity restrictions is discussed in Heckman, Smith, and Clements (1997) and Chernozhukov and Hansen (2005). However, even if the quantile treatment effect cannot be interpreted as corresponding to individual-level effects, it is still of policy interest as summarizing the effect of unionization on the marginal distribution of earnings.

Treatment effects (1) and (2) provide comparisons of potential outcome distributions, but do not directly describe how the marginal distribution of earnings among affected workers would change if the fraction of unionized workers, $p \equiv \Pr(D = 1|R = 0)$ were to incrementally change. A parameter that captures this notion is the unconditional quantile partial effect (UQPE), described by Firpo, Fortin, and Lemieux (2009), defined as the derivative of the τ-quantile of realized earnings with respect to the probability of unionization, p:

$$UQPE(\tau) \equiv \frac{\partial Q_{Y|R=0}(\tau)}{\partial p}.$$

Noting that the marginal τ-quantile is defined by $F_{Y|R=0}\left(Q_{Y|R=0}(\tau)\right) = \tau$, and the marginal cdf of earnings at the threshold is $F_{Y|R=0} = pF_{Y_1|R=0} + (1-p)F_{Y_0|R=0}$, implicit differentiation immediately gives an expression for the UQPE:

$$UQPE(\tau) = \frac{F_{Y_0|R=0}(q_\tau) - F_{Y_1|R=0}(q_\tau)}{f_{Y|R=0}(q_\tau)}, \tag{3}$$

where $q_\tau = Q_{Y|R=0}(\tau)$ and $f_{Y|R=0}$ is the density of realized earnings at the threshold. The UQPE is therefore simply a rescaling of the cdf treatment effect. Besides its intuitive interpretation, an advantage of the UQPE is that since cdfs are expectations, UQPE can be easily estimated by regression, and incorporating covariates for efficiency is straightforward.

The key assumption for identifying the treatment effects above is that the conditional distribution of potential earnings as a function of the union vote share is smooth near the threshold of union victory, and thus any jumps in the observed distribution of earnings at the threshold are due to the treatment. Formally:

Assumption 1: Local Smoothness $F_{Y_d|R}(y|r)$ is continuous in r over an ε-neighborhood of zero, and is strictly increasing in y over the same neighborhood, for $d \in \{0, 1\}$.

This assumption is satisfied if, for example, unions, workers, and firms are *a priori* uncertain about the outcome of the election when it is close (see Lee, 2008 for a formal proof). The condition that the distribution be increasing in y ensures that quantiles are uniquely defined at the threshold. Given the sharp RD design setup and this local smoothness assumption, the identifying conditions

given in Frandsen, Frölich, and Melly (2010) are satisfied.[12]

I focus here on the effect of union representation at an establishment, which leads to the sharp RD design (Campbell, 1969). Union representation leads to a number of downstream channels affecting wages, including collective bargaining agreements. Considering a collective bargaining agreement as the treatment variable would lead to a "fuzzy" RD design since an agreement does not always follow from a union victory (Ferguson, 2008). While the econometric framework described in the text applies equally well to the sharp and fuzzy designs, I focus on representation because contracts are not always documented and because the resulting instrumental variables setup is not likely to satisfy the required exclusion restriction. FMCS records severely undercount union agreements (DiNardo and Lee, 2004), and there is typically a time lag of several months between a certification election and a collective bargaining agreement being reached. Behavioral responses to an election outcome that occur before an agreement is reached or without regard to an agreement (as observed) lead to violations of the exclusion restriction since a union victory may affect outcomes event absent a contract. However, I begin with estimates of the effect of union representation on collective bargaining as an outcome of interest.

Frandsen, Frölich, and Melly (2010) outlines a tractable procedure for estimating quantile treatment effects in this framework. In the sharp RD design considered here, the estimator becomes particularly simple: it is the difference between kernel-weighted local linear estimates of the quantiles of earnings approaching the threshold of union victory from the right and from the left. Formally, the estimator can be written:

$$\widehat{QTE}\left(\tau\right) = \hat{Q}_{Y_1|R=0}\left(\tau\right) - \hat{Q}_{Y_0|R=0}\left(\tau\right),$$

where $\hat{Q}_{Y_1|R=0}\left(\tau\right)$ and $\hat{Q}_{Y_0|R=0}\left(\tau\right)$ are locally linear, consistent estimates of the τ-quantile of earnings at the threshold from the right and the left, respectively. I use a uniform kernel with several different bandwidths to explore sensitivity.[13] Bandwidths between .025 and .1 yield similar results, with the larger bandwidths leading to more precise estimates. The estimator for the local quantile

[12]The other identification conditions in Frandsen, Frölich, and Melly (2010)—that the probability of treatment jumps discretely at the threshold, and that crossing the threshold has a montonic effect on treatment status—are automatically satisfied in the sharp RD setup considered here.

[13]Imbens and Kalyanaraman's (2009) optimal RD bandwidth algorithm yields extremely wide bandwidths of around .4. For the sake of conservativeness and to reduce bias I use narrower bandwidths in the range of .05 to .1.

treatment effect is consistent and asymptotically normally distributed. Inference is clustered at the establishment level.

Nonparametric estimates of effects conditional on pre-election earnings, (2), are obtained by grouping observations close to the threshold by percentile of pre-election earnings and estimating the union-nonunion difference in means for each percentile. Estimates of the unconditional quantile partial effect, (3), are regression coefficients on a dummy for union victory where the dependent variable is $-1\left(Y_i \leq q_\tau\right)/\hat{f}_{Y|R=0}\left(q_\tau\right)$ where q_τ is the observed τ-quantile of earnings near the threshold and $\hat{f}_{Y|R=0}$ is the kernel-smoothed density of earnings near the threshold. The UQPE regressions also control for covariates such as pre-election earnings, year, state, and industry dummies.

5 Results

5.1 Effects on Collective Bargaining

I look first at whether a union contract was reached as a check on the research design and the matching procedure, as well as an interesting outcome in its own right. Regression discontinuity estimates of the effect of union representation on collective bargaining show that a union victory significantly increases the likelihood of a collective bargaining agreement being reported following an election. Figure 3 plots raw fractions and fitted probabilities by union vote share that a collective bargaining agreement at an establishment was reported to the FMCS following a representation election. Among establishments where the union lost the election (50 percent or lower vote share), around 10 percent had a contract reported, mostly from other bargaining units at the establishment. At the threshold of union victory, the probability jumps to between 40 and 50 percent. The estimated effect is .271 (s.e.=.0157) for any contract reported to the FMCS. If contracts are under-reported to the FMCS then this effect is a lower bound for the true impact on the likelihood of a union contract. An undercount by a factor of 2—as suggested by DiNardo and Lee (2004)—would imply that close to 90 percent of union victories result in a union contract.

The figure also plots probabilities that a contract was reported prior to the election. Little difference can be seen on either side of the threshold, and the estimated effect is an insignificant -.013 (s.e.=.0108). This constitutes a test of the validity of the design: while prior union activity appears to be positively correlated with union success in the election (the probability of a prior

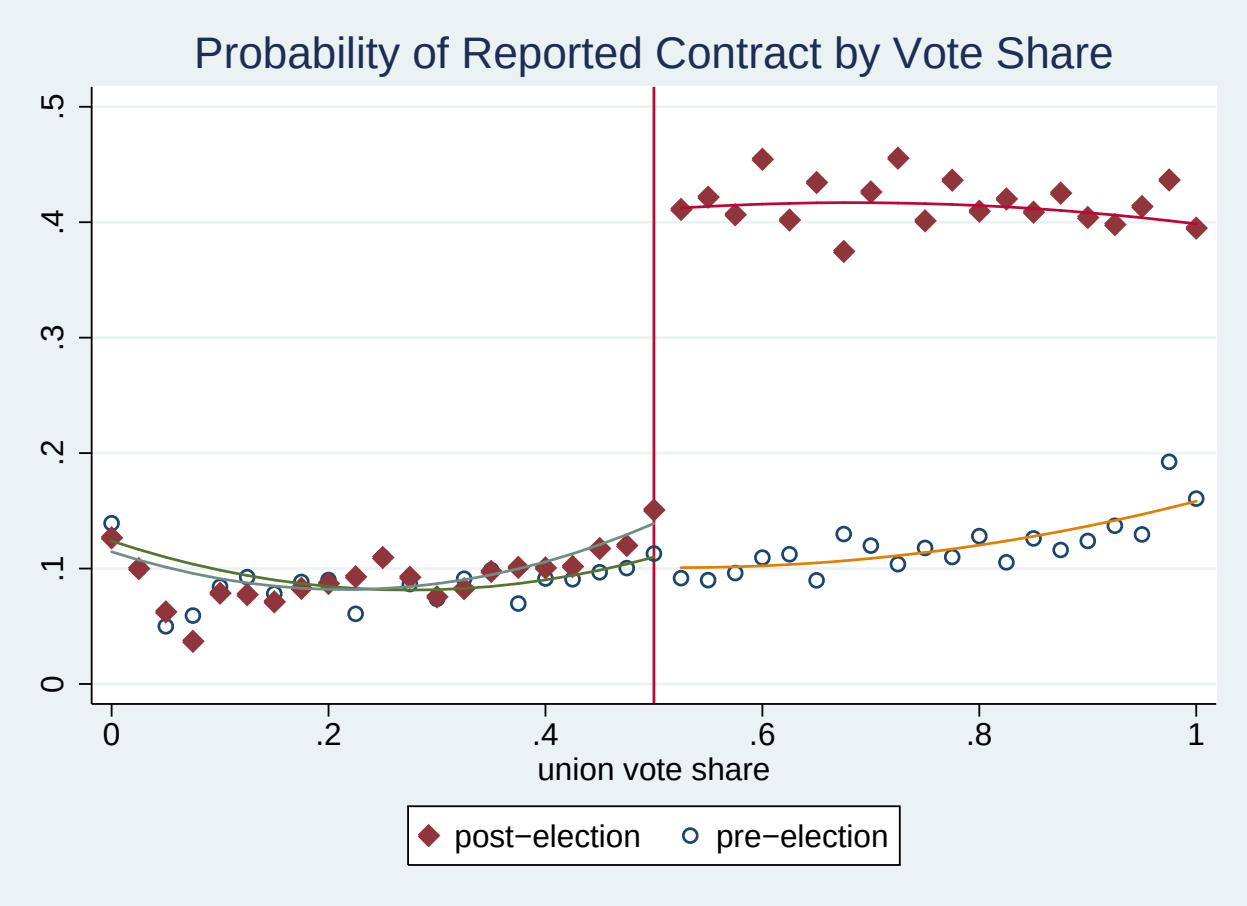

Figure 3: Probability of a contract expiration reported to the FMCS by union vote share. The points are raw fractions by non-overlapping vote share bins of width .025. The lines are predicted values from a piecewise-quadratic fit. Data are from FMCS Form F-7 filings and NLRB election records.

union contract is slightly upward sloping in vote share), there is no jump at the threshold and establishments on either side of the threshold appear comparable in terms of prior union presence. The figure also clarifies the economic content the "treatment" of union representation: a union victory in a representation election clearly leads to an increase in collective bargaining agreements, but it is far from the sole determining factor; a victorious union may fail to reach a contract with the employer, and even establishments where the union lost may at some point engage in collective bargaining.

5.2 Individual-level Earnings and Employment Effects

The next set of analyses estimates the effect of union representation on post-election earnings and employment outcomes of workers employed at establishments just before a union election. Regression discontinuity estimates support the notion that unionization has little effect on the average wages, consistent with DiNardo and Lee, but significantly compresses the distribution of employee earnings. These findings can be seen in Table 2, which reports RD estimates of overall average and distributional effects on log annual earnings during the year starting two quarters after the election, an indicator for positive annual earnings, and an indicator for remaining employed at the original plant six quarters after the election. Estimates in the first column are restricted to full-time, full-year (FTFY) workers (based on pre-election earnings history) in an effort to isolate a wage effect from an hours effect. The first column shows the effect on FTFY workers' log earnings is small: -.002 with a standard error of .011. The effect on employment status (defined as positive annual earnings) is likewise negligible for these workers, with a point estimate of -.001 (s.e.=.005). This result is consistent with DiNardo and Lee's (2004) finding of little union effect on average wages and employment. The effect on FTFY workers remaining employed at the same establishment six quarters after the election is a modestly positive .06, but imprecisely estimated.

The small average effect on log earnings and employment for FTFY workers may miss significant effects for less-attached workers or on other features of the distribution. Column (2) expands the sample to all workers, including less-attached workers with lower pre-election earnings. The effect on log annual earnings for all workers is .065 (s.e.=.031), which suggests significant effects for less-attached workers with low pre-election earnings. The average effect on employment for all workers in column (2) is still small, estimated at -.008 (s.e.=.008), and the effect on six-quarter retention is .042, but also imprecisely estimated. The difference in the earnings effect from column (1) to column (2) suggests that effects may be heterogeneous across the distribution. The remaining columns report estimates of quantile treatment effects, (1), and effects conditional on pre-election quantile, (2), for selected quantiles. Columns (3) and (4) show that both estimates suggest a large effect on log annual earnings at the 10th percentile, with significant point estimates of over .3 for the average effect at the .1-quantile of pre-election earnings and over .7 for the .1-quantile treatment effect, although the confidence intervals are wide. This effect quickly drops to near zero

for the median and 90th percentile, and the point estimates are negative, although not significantly different from zero. The effect on employment stays near zero for the reported quantiles. The effect on retention is larger for the 10th percentile, with a point estimate of .110, but this is significant only at the 10 percent level.

Table 2: Effects of Union Representation on Earnings and Employment

	Average		Quantiles					
Dependent variable	FTFY	All	0.1		0.5		0.9	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Log(annual earnings)	-0.002	0.065	0.320	0.732	-0.004	-0.021	-0.021	-0.018
	(0.011)	(0.031)	(0.130)	(0.315)	(0.018)	(0.049)	(0.016)	(0.072)
Employed	-0.001	-0.008	-0.007	-	0.000	-	-0.001	-
	(0.005)	(0.008)	(0.019)		(0.005)		(0.006)	
6-quarter retention	0.060	0.042	0.110	-	0.028	-	0.026	-
	(0.035)	(0.030)	(0.067)		(0.035)		(0.051)	
Method	LS	LS	LS	QR	LS	QR	LS	QR
N	228,796	427,567	427,567		427,567		427,567	

Notes: The table reports regression discontinuity estimates of the unionization effect on post-election earnings and employment outcomes listed in the left-hand column for workers employed at establishments at the time of a union election from 1990-2007. Earnings (in 2000 dollars) are defined as the sum of the four quarterly earnings starting two quarters after the union election closed. 6-quarter retention is an indicator for a worker still being employed at the election establishment 6 quarters after the election. The full-time, full-year (FTFY) sample in column (1) is defined as workers employed at the establishment for the four consecutive quarters before the election with earnings during that period of at least $20,000 (in 2000 dollars). All other columns use all workers. The quantile estimates in columns (3), (5), and (7) are least squares estimates conditional on pre-election earnings quantile. The estimates in columns (4), (6), and (8) are local linear quantile regression estimates. The bandwidth (in terms of union vote share) is .1. Standard errors are clustered at the establishment level.

As a whole, the results in Table 2 suggest that while average effects are small, there may be significant variation across the distribution. A more detailed analysis of distributional effects confirms that log earnings effects are large at the lower end of the distribution and small elsewhere. The union effects on the distribution of earnings are summarized graphically in Figure 4, which plots estimates and pointwise confidence intervals for quantile treatment effects (Panel A.) and unconditional quantile partial effects (Panel B.) for the .1 through the .99 quantile. Both sets of estimates show unionization had a large and significant effect at the lower end of the distribution, with a much smaller and possibly negative effect at the upper end. The quantile treatment effect on the .1-quantile is about .7, meaning the .1-quantile of unionized earnings at the threshold is about 70 log points higher than the .1-quantile of non-unionized earnings, although the confidence interval is very wide. The unconditional quantile partial effect at the .1-quantile is about .2, meaning a 10 percent increase in the unionization rate would increase the 10th percentile of observed earnings among this group of workers by about 2 percent. The unconditional quantile partial effects are

much more precisely estimated because they control for pre-election earnings, state, time, and industry dummies.

The estimates shown in Figure 4 correspond to effects on the distribution, not necessarily to effects on particular individuals. Nevertheless, estimated average effects for individuals at different percentiles of the pre-election earnings distribution show a very similar pattern of large, significantly positive effects at the lower end of the distribution, with much smaller effects elsewhere. Figure 5 plots point estimates and confidence intervals for the effect of unionization on average log annual earnings by percentile of the pre-election earnings distribution. The effect for individuals at the 5th percentile of pre-election earnings is around 62 log points and is significantly greater than zero at the 5 percent level. However, at about the 20th percentile and higher, the estimated effects are small and not significantly different from zero. That the pattern and magnitudes of the effects in Figure 5 rank appear very similar to the quantile treatment effects in Figure 4 is consistent with unionization not systematically affecting relative ranks and ranks remaining stable over time within a group of workers. If this holds, then the quantile treatment effects can be interpreted not only as effects on the marginal distribution of earnings, but also as an individual-level effect for a worker of a given rank in the earnings distribution.

The results thus far imply that unionization compresses and shifts the earnings distribution slightly to the right. This effect can be seen directly in Figure 6, which plots estimates of the counterfactual earnings densities by union status, conditional on a close election. The solid curve represents the density of potential earnings without unionization, and the dashed curve represents the density of potential earnings under unionization. The unionized density is lower in the tails, reflecting compression, and shifted slightly to the right, reflecting the modest positive effect throughout much of the distribution seen in Figure 4.

While earnings effects appear to vary considerably over the distribution, do employment or retention effects also vary? Economic intuition suggests two countervailing forces. If unionization leads to higher wages at the bottom of the distribution, low-skilled workers will be more likely to want to stay, while employers would like to respond by substituting away from lower-skilled labor, as in Card's (1996) two-sided selection model. Estimates of the effect of unionization on employment status suggest that for the lowest-skilled workers, the negative employer-driven effect dominates. Figure 7 plots point estimates and confidence intervals for the effect of unionization

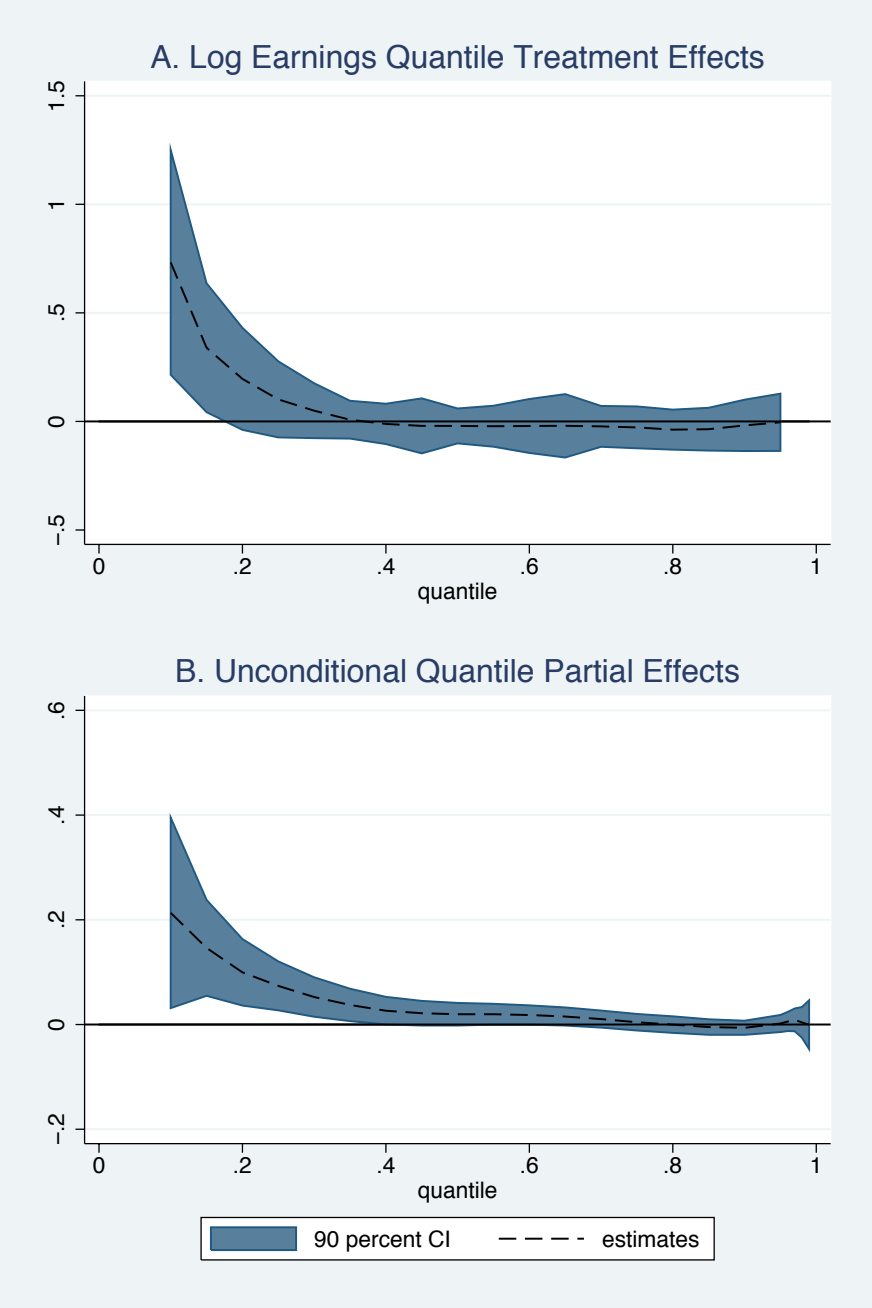

Figure 4: Estimates and 90-percent confidence intervals for the effect of unionization on the employees' log annual earnings by quantile. Panel A. reports quantile treatment effects. Panel B. reports unconditional quantile partial effects.

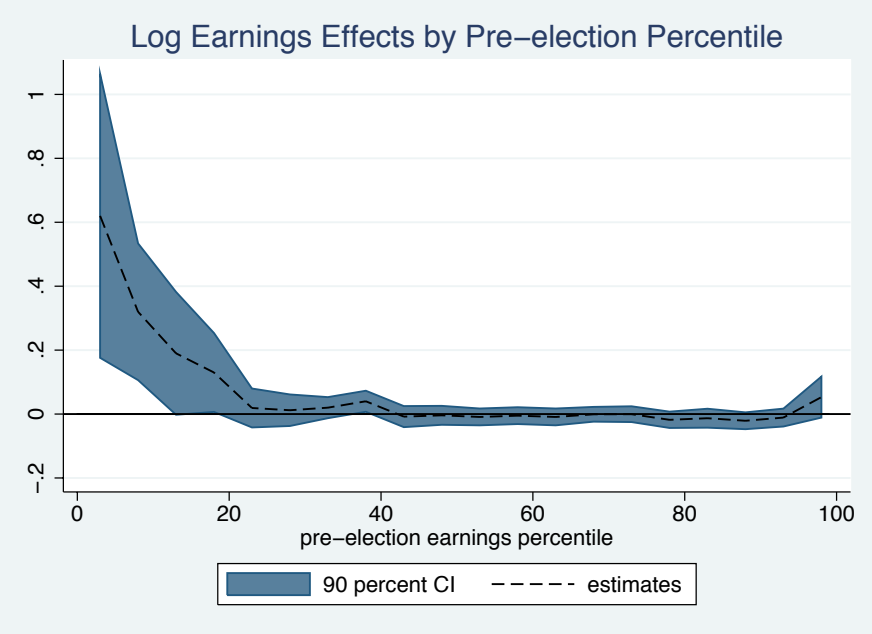

Figure 5: Estimates and 90-percent confidence intervals for the effect of unionization on the employees' log annual earnings by pre-election earnings percentile.

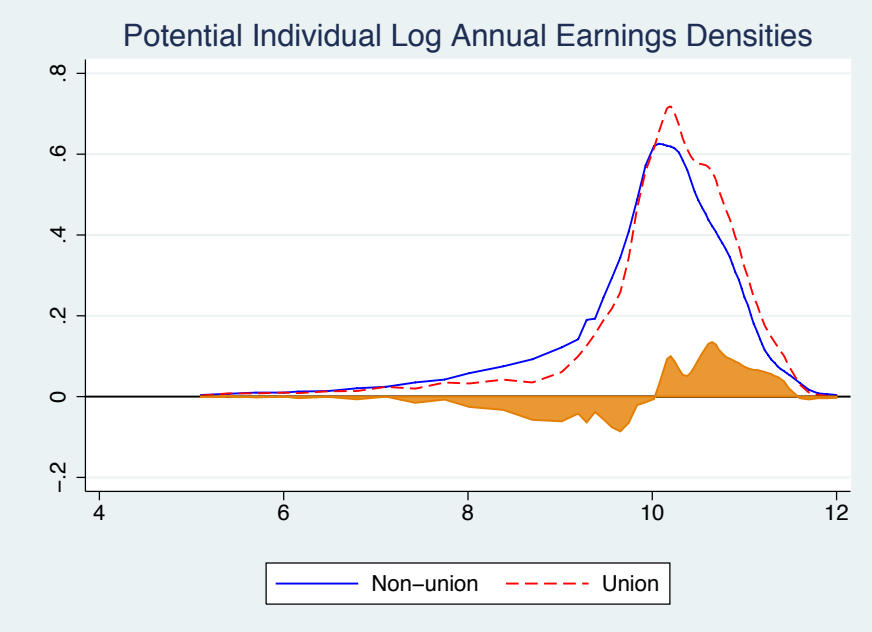

Figure 6: Estimated counterfactual densities of log annual earnings conditional on a close union election.

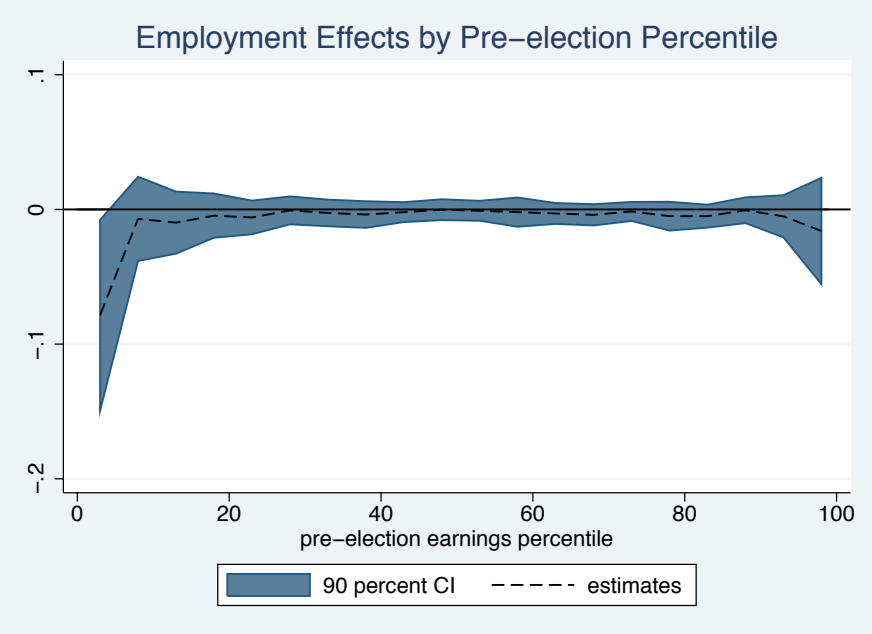

Figure 7: Estimates and 90-percent confidence intervals for the effect of unionization on employees' post-election employment status by percentile of pre-election annual earnings.

on employment status by pre-election earnings percentile. The figure shows around a minus 5 to minus 10 percent effect below the 5th percentile that is marginally significant, but zero effect elsewhere. The results suggest that the lowest-skilled workers were more likely to lose their jobs after unionization, but those who kept their jobs enjoyed higher earnings. Estimates of the effects on worker retention by earnings percentile support this intuition. Figure 8 plots point estimates and confidence intervals for the effect of unionization on whether a worker remained employed at the original plant six quarters after the election, focusing more finely on the lower third of the distribution. The estimates are negative at the bottom of the distribution, reflecting the negative effect on employment in Figure 7, but become positive near the 10th percentile, and decline to near zero farther up the distribution. The effects for the higher percentiles not plotted are all near zero.

These results are consistent with the interpretation that employers respond to higher wages at the very bottom of the distribution by substituting away from the lowest-skilled workers, leading to more turnover (and higher unemployment) there. The positive effect on retention just above the bottom can then be attributed to employers substituting to slightly higher skill levels, a natural prediction if adjacent skill groups are closer substitutes than groups of widely differing skill levels and if wage effects are smaller there. Another interpretation of these results is that unionization

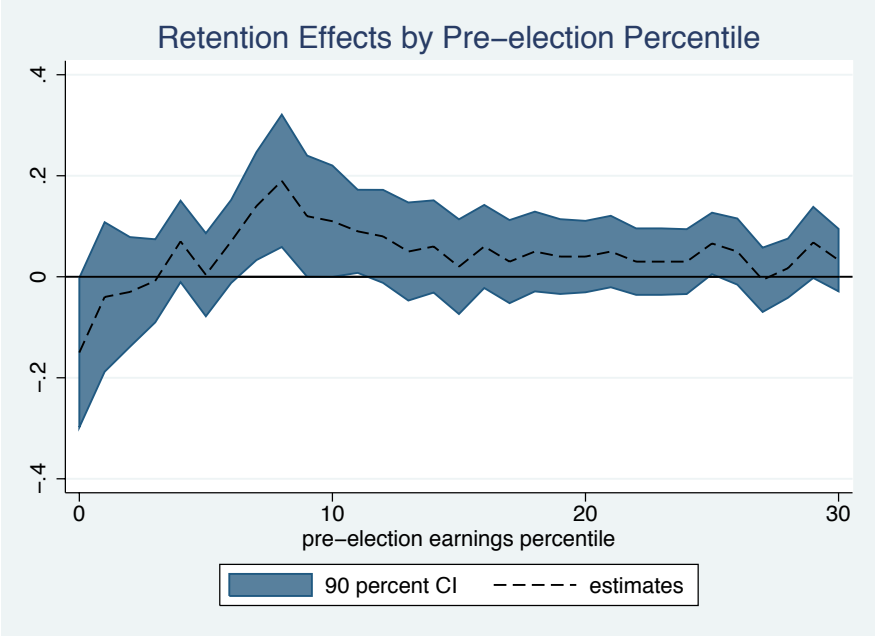

Figure 8: Estimates and 90-percent confidence intervals for the effect of unionization on whether an employee remains employed at the original establishment 6 quarters after a union election.

leads to higher wages for lower-skilled workers, but leads to employment protections only for workers in the bargaining unit, which tends to exclude the very lowest-skilled workers, whom employers are able to lay off in response to higher wages.

5.3 Establishment-level Distributional Effects

The individual-level estimates reported thus far take employees at plants just before a union election as the unit of interest, regardless of where they end up after the election. The next set of results takes election establishments as the unit of interest, and estimates effects on the distribution of earnings actually paid at those establishments. While these estimates do not correspond to causal effects on individuals, they do correspond to the combination of the causal effects on wage-setting within an establishment and on the composition of workers employed at the establishment.

Regression discontinuity estimates of the effect of unionization on earnings paid at election establishments suggest that unionization compresses the distribution of paid earnings to an even greater degree than the individual-level earnings compression seen in the previous section. Figure 9 plots quantile treatment effects estimates from quantile regressions of annual earnings paid at election establishments starting two quarters after a union election. The estimates suggest an

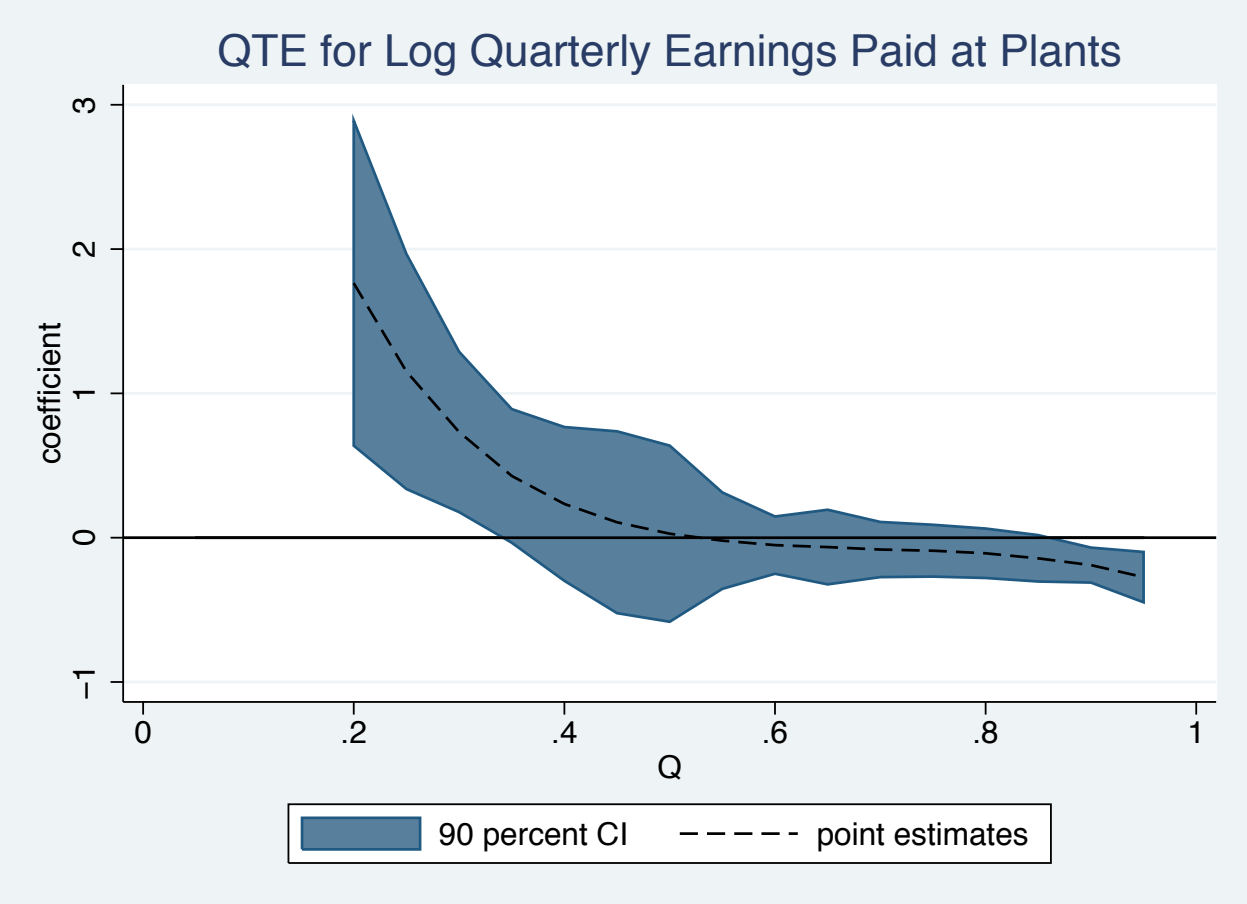

Figure 9: Quantile treatment effects estimates of the effect of unionization on the distribution of earnings paid at an establishment. The outcome variable is annual earnings paid at election establishments beginning two quarters after a union election.

extremely large effect of between 100 and 200 log points at the lower end of the distribution, with much smaller, even negative, effects at the upper end. Figure 10 depicts the same effect in terms of the counterfactual densities of earnings paid at establishments. The movement of mass from the tails to the center of the distribution is especially marked at the lower end, reflecting the large quantile treatment effects in Figure 9.

These results are consistent with the individual-level earnings and employment effects in the previous section. The combined effect of a decrease in employment at the bottom of the distribution and an increase in wages for remaining workers is to effectively cut off the bottom of the distribution of earnings paid at the plant, resulting in the dramatic increase in the lower quantiles of earnings observed in the figure. The negative effect observed at the top of the distribution is consistent with a relative shift in the composition of workers employed at unionizing establishments away from the

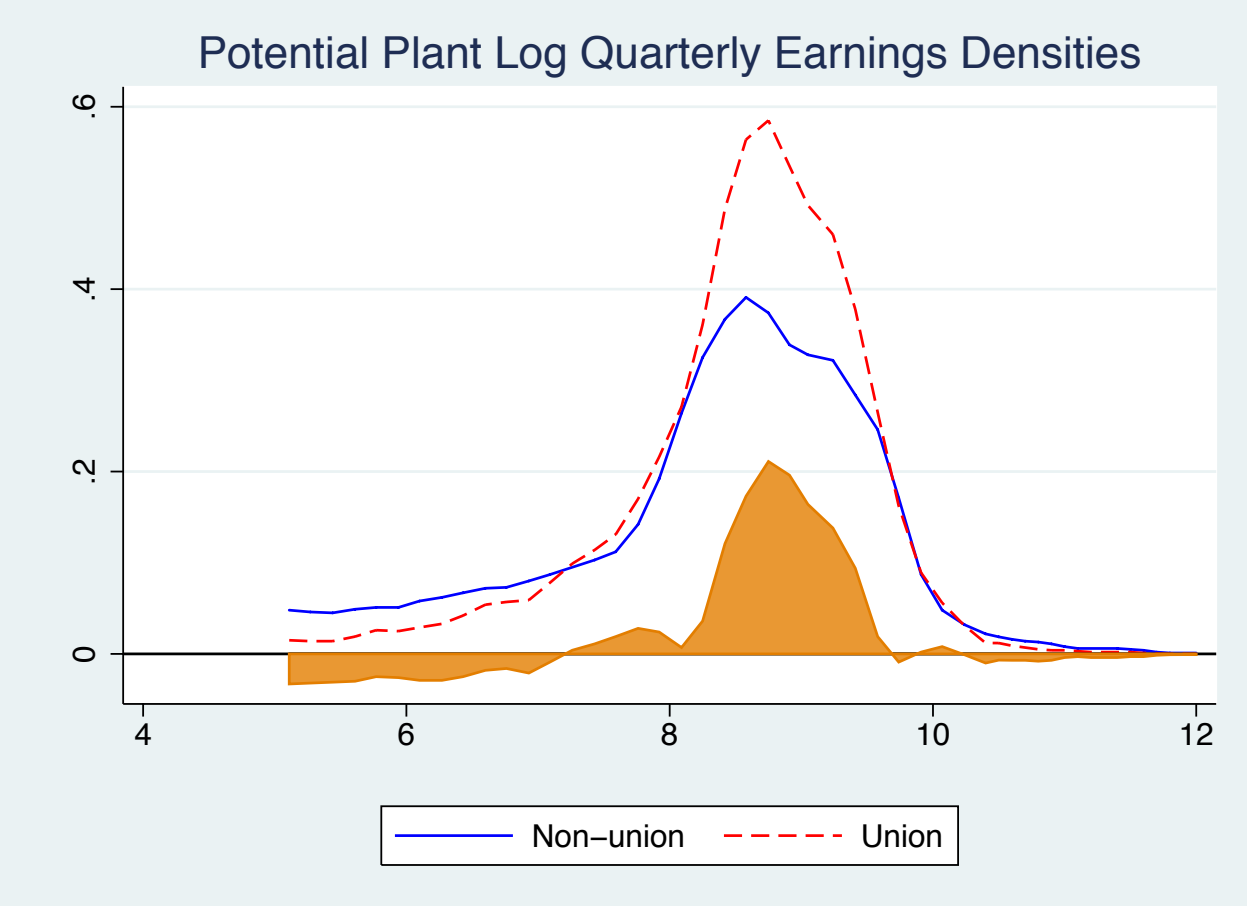

Figure 10: Estimated counterfactual densities of earnings paid at establishments.

highest-skilled workers.

The results in this section provide evidence for the union wage compression found in previous regression-based studies, but against a large average effect of unionization. Thus the principal substantive conclusion of Freeman (1993), Card (1996), DiNardo, Fortin, and Lemieux (1996) that unions reduce dispersion holds up even in a quasi-experimental setting, but the finding in these studies and many others that union earnings are on average significantly higher than non-union earnings appears to be largely a selection effect, consistent with DiNardo and Lee (2004), at least for the typical case of a close union election.[14]

6 Interpretation

The empirical results suggest that unionization has a large positive effect on earnings at the bottom of the distribution, but a declining effect farther up the distribution. The results are consistent with an interpretation that unions impose a wage premium that is large for lower-skilled workers and declines with skill level, and at the same time are able to enforce employment protections for employees within the bargaining unit. This section motivates this interpretation by analyzing the theoretical implications for the distribution of wages of union electoral competition.

The model developed here takes as its starting place that the primary objectives of union leaders are survival and expansion of the organization, and retention of their offices (Atherton, 1973; Ross, 1948; Berkowitz, 1954; Ashenfelter and Johnson, 1969). Accordingly, a union facing a certification election will pursue a wage agreement that maximizes the probability of winning. The theoretical framework in this section is similar to Farber (1978), Booth (1995), Acemoglu, Aghion, and Violante (2002), and Lee and Mas (forthcoming), who also consider the effects of majority-rules politics on union wage policies, but I focus on the implications for the distribution of workers' wages.

The model casts the union election and wage-setting process as a four-stage game. The stages, shown in Figure 11, follow the stylized union certification process described in Section 2. In the first stage, a union petitions to represent the workers at a plant (currently producing in the competitive sector) and proposes a wage schedule, $w(H) = (v + r(H))H$, which gives a worker's wages as a function of his or her human capital, H. The outside (competitive) price of human capital is v,

[14] Positive union gaps may also be due to unions that win by large margins (Lee and Mas, forthcoming).

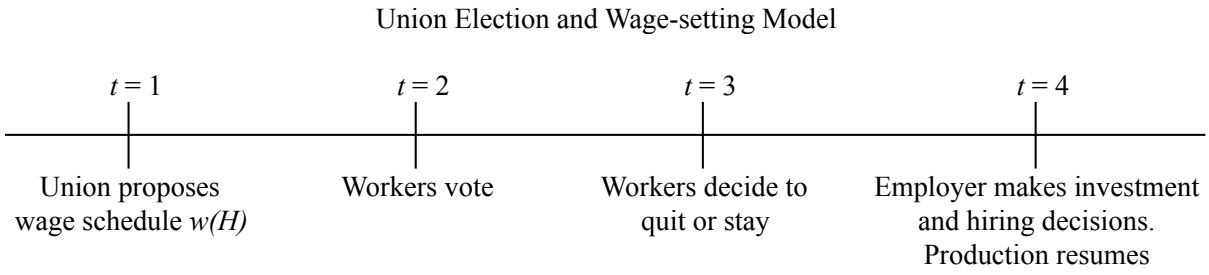

Figure 11: The model

and $r(H)$ denotes the union rent earned by a worker with human capital H. The term "union rent" hightlights that $r(H)$ is the difference between the competitive and union price of skill. In the second stage, workers vote for the union if their union rent, $r(H)$, exceeds their individual cost of union representation, η, provided the wage schedule doesn't cause the plant to shut down. The voting rule can therefore be written $1(\eta \leq r(H))$, subject to the plant not shutting down. The cost of union representation η reflects any pecuniary (e.g., union dues) or nonpecuniary factors affecting workers' preferences for union representation outside of wage differences, and is assumed to be independent of H. In the third stage, workers decide whether to quit or stay, after observing the outcome of the election. Workers stay if their union rent exceeds their cost of unionization, net of an individual-specific switching cost, ε, also independent of H. The worker's decision rule after a union victory is therefore $1(\eta - \varepsilon \leq r(H))$, where the worker stays if the indicator is equal to one. Finally, the employer makes investment and hiring decisions to maximize profits.

The production technology combines H with another factor, K, in fixed proportions with constant returns to scale, so the production function can be written $Y(H) = yH$. Normalizing the H/K ratio to unity, the competitive return to K is $y - v$. A fraction ϕ_K of K is sunk in the production relationship.

Two assumptions in the model setup are important for the results. The first is that workers' voting and quitting decisions depend on more than simply a comparison of wages. These other factors are modeled as union costs η and switching costs ε. The union and switching costs are assumed to reduce the effective union and outside return to human capital, respectively. While the specific functional form is a simplification, the substantive assumption that the "equivalence premium"—the dollar amount by which a union would have to raise a worker's earnings to make him indifferent between a union and no union—is on average increasing in a worker's outside wage

is important. A concrete motivating example is that union dues are commonly collected as a percentage of wages. The consequence of this assumption for the model is that it makes it more efficient for the union to shift resources to attract the votes of lower-skilled workers.

The second key assumption is that employers face short term rigidities in adjusting inputs. The model captures this in a simple way following Caballero and Hammour (1998) by assuming the chosen technology takes fixed ratios of inputs. The assumption is that while technology can adjust over time, it is essentially fixed over the period relevant for an initial collective bargaining agreement. Consequently, employers cannot undo the effects of a union wage schedule by immediately adjusting the production inputs. An alternative way of capturing this that leaves the results substantively unchanged is to allow the employer to fully adjust K, but have the union impose firing costs.

The union's optimal wage schedule is shown in the appendix to maximize the expected vote share, subject to a minimum profit constraint for the employer:

$$\begin{aligned} &\max_{r(h)} \int F_{\eta}\left(r\left(h\right)\right) dF_{H}\left(h\right) \\ &\text{s.t.} \int \left(\phi_{K}\left(y-v\right)-r\left(h\right)\right) hF_{\eta-\varepsilon}\left(r\left(h\right)\right) dF_{H}\left(h\right) \geq 0. \end{aligned}$$

Concrete functional forms illustrate the main implications of the model. Assume H is log-normally distributed and η and ε are exponentially distributed with parameters λ_{η} and λ_{ε}, where the mean switching cost, $1/\lambda_{\varepsilon}$, exceeds the average rent, $\phi_{K}\left(y-v\right)$. In this case the optimal rent schedule is

$$r\left(H\right)=\begin{cases} r^{+}\left(H,\lambda^{*}\right) & , \quad H\leq h_{1}\left(\lambda^{*}\right) \\ -\left(\frac{1}{\lambda_{\varepsilon}}-\phi_{K}\left(y-v\right)\right) & , \quad H>h_{1}\left(\lambda^{*}\right) \end{cases}, \tag{4a}$$

where $r^{+}\left(H,\lambda\right)$ satisfies

$$\exp\left(\lambda_{\eta}r\right)=\frac{\lambda_{\eta}}{\lambda H}+\frac{\lambda_{\varepsilon}}{\lambda_{\eta}+\lambda_{\varepsilon}}\left(1+\left(\phi_{K}\left(y-v\right)-r\right)\lambda_{\eta}\right). \tag{4b}$$

The threshold above which rents are negative is

$$h_{1}\left(\lambda\right)=\frac{1}{\lambda}\frac{\lambda_{\eta}+\lambda_{\varepsilon}}{1-\lambda_{\varepsilon}\phi_{K}\left(y-v\right)}, \tag{4c}$$

and λ^* is the value of λ that satisfies the profit constraint with equality.

The optimal union wage schedule, (4), has two features with stark implications for the union effect on the distribution of earnings. First, the union rent is positive for lower-productivity workers, and negative for higher-productivity workers. Interestingly, the less the union can extract from the employer (i.e., the smaller is ϕ_K), the lower the threshold above which rents are negative. The second feature is that even where rents are positive, they are decreasing in the level of human capital:

$$\left.\frac{dr}{dH}\right|_{H=h\leq h_1} = -\left(\lambda h^2\left(\frac{\lambda_\varepsilon}{\lambda_\eta+\lambda_\varepsilon}+\exp\left(\lambda_\eta r\right)\right)\right)^{-1} < 0.$$

Thus compared to the competitive equilibrium, the union wage schedule compresses the distribution of potential wages and shifts it to the right.

Given the assumption of fixed K, these implications are likely to apply to the short term. In the long term, however, K can adjust, and the employer may be able to terminate lower-productivity workers, who are being paid above marginal product. In the longer term, therefore, the degree to which the union wage schedule inflates the wages at the bottom of the distribution relative to the top will be attenuated.

A numerical example calibrated to workers in the 1998-2000 Current Population Survey (CPS) illustrates the compressing effect of the optimal union wage schedule on the distribution of earnings (see the Theory appendix for the chosen values of the parameters). Figure 12 compares the distribution of potential union and non-union log wages implied by this example. The union distribution is compressed and shifted to the right relative to the non-union distribution.

The model's qualitative prediction of wage compression is consistent with the empirical results in section 5, but the model's main prediction concerns union wage effects by *relative* rank within a firm or industry, rather than by absolute rank as in section 5. A test of the prediction that the union effect is negative for workers above a certain rank within their industry provides a sharper test. Regression discontinuity estimates of the union effect by *within-industry* rank suggests that the union effect for workers of higher rank may be negative. Figure 13 plots estimates and confidence intervals for the effect of union representation on log annual earnings by within-industry percentile of pre-election earnings. The difference between Figure 13 and Figure 5 is that here the ranking is within industry, to more closely approximate the conditions of the model. The same large effect

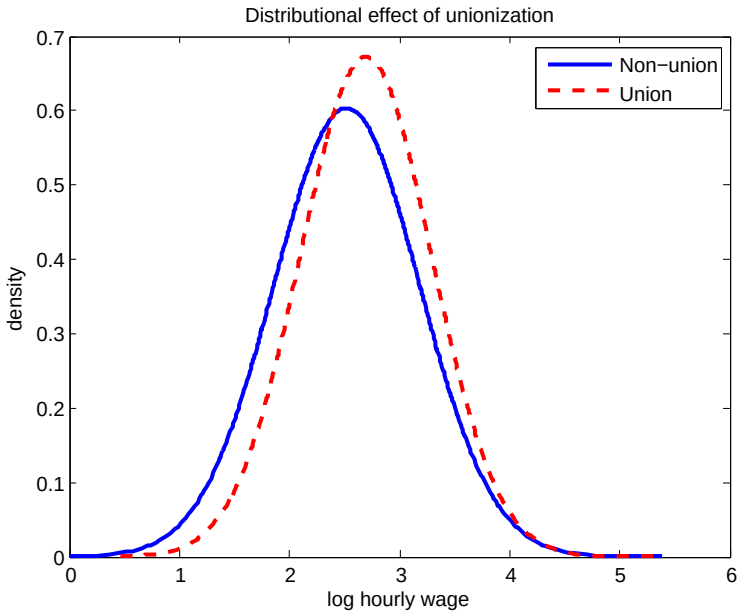

Figure 12: Densities of potential union and non-union log wages implied by the optimal wage schedule from the example in the text.

at the bottom of the distribution appears, with an estimate of around 50 log points at the 10th percentile. Above about the 50th percentile, the point estimates become negative, with an estimate of about minus-10 log points at the 95th percentile, although this is not significantly different from zero. While these results are far from conclusive, they provide suggestive evidence in support of the model's interpretation.

The theoretical discussion shows that unions facing certification elections have an incentive to commit to a wage schedule that favors lower-productivity workers at the expense of higher-productivity workers. In the example, the union wage schedule raises the wages of low-productivity workers and lowers the wages of high-productivity workers, compressing the distribution of wages.

7 Summary and Conclusions

Quasi-experimental estimates based on close union certification elections show unionization substantially compresses the distribution of employee earnings. A union victory raised the tenth percentile of earnings by as much as 30 log points with a much smaller effect in the middle and upper end of the distribution. Estimates of the union effect on employment status on the other hand showed a significant *negative* effect at the lower end of the distribution, with little effect elsewhere, suggesting

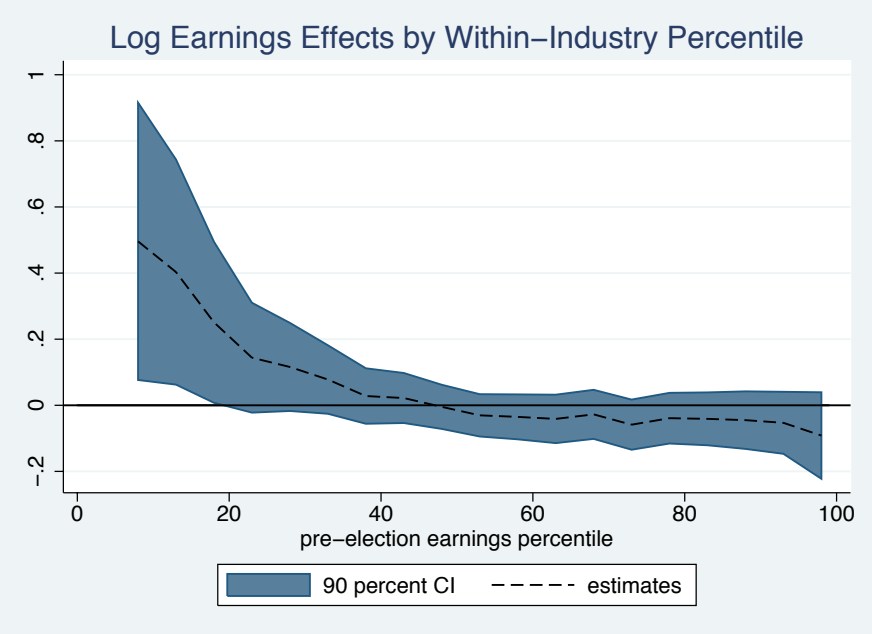

Figure 13: Estimates and 90-percent confidence intervals for the effect of unionization on the employees' log annual earnings by within-industry pre-election earnings percentile.

that employers respond to the increase in wages by substituting away from the very lowest-skilled workers. Union retention effects by earnings percentile support this interpretation, showing a drop in retention at the very bottom, but an increase in retention for workers of just higher skill levels, and little effect elsewhere. Estimates of the effect of unionization on the distribution of earnings actually paid at election plants show that the combination of the individual-level earnings and retention effects leads to a compositional shift that results in even greater compression of earning paid at plants.

The pattern of effects on the distribution of earnings and worker retention is consistent with an interpretation that unions pursue a wage schedule that increases earnings at the bottom of the skill distribution and reduces the return to skill, and increases employment protection for workers within the bargaining unit. Any workers at the bottom of the skill distribution who are excluded from the bargaining unit are then more likely to become unemployed or otherwise leave employment at the plant as employers respond to the union wage schedule by substituting to slightly higher-skilled workers.

The results imply that unions close to the margin of victory unambiguously compress the overall earnings distribution and can therefore can explain part of the increased inequality in the

U.S. income distribution since the 1970s. To compare the earnings compression implied by my estimates with the previous literature, the estimated effect on the variance of log earnings for full-time, full-year workers at plants where workers unionized is about -.10, with essentially zero effect on the mean (see Table 2), which corresponds to a decrease in the overall variance of log earnings of about .047 if the private sector unionization rate had remained at its 1979 level.[15] This reduction amounts to 26.4 percent of the increase in the variance of log earnings reported by Kopczuk, Saez, and Song (2010), a larger compression effect than the previous literature's estimates of around 6 to 21 percent of the overall increase in inequality (DiNardo, Fortin, and Lemieux, 1996; Card, 2001).

In conclusion, private sector unions now cover less than 8 percent of the U.S. workforce (Bureau of Labor Statistics, 2011) and recent quasi-experimental evidence suggests they have little effect on average. Have unions ceased to have meaningful economic impacts? The evidence in this paper suggests not. Unions have signficant distributional effects on both individual earnings and employment, and their decline can account for over one quarter of the rise in inequality over the past 30 years. An open question is whether the causal effects reported here carry over to unions in the public sector, an area where unions have not declined, and on which more research is needed.

Appendix

Construction of the dataset

As described in the text, the dataset used in this paper consists of NLRB certification election results matched to employer-employee wage data from the Census Bureau's LEHD program.

The union certification election records were collected by the NLRB, and in large part maintained by the AFL-CIO. John-Paul Ferguson, Thomas Holmes, and Hank Farber obtained the election records from the NLRB, and made them available for this research. The complete data set covers the period 1963-2009, and contains records from over 250,000 union elections, although

[15] The effect on the overall variance of log earnings was arrived at by computing the within effect for the fraction of workers at establishments that would see an increase in unionization if the unionization rate remained at its 1979 level. Extrapolating Hirsch's (1991) survey data to 2009 suggests the share of employees at unionized plants who are covered by the union is around 32.8 percent, so increasing the unionization rate from 8 percent to 23.3 percent would impact $(23.3\% - 8\%)/32.8\% = 47\%$ of workers. The within effect on overall variance is therefore $-.10 \times 47\% = -.047$. As in DiNardo, Fortin, and Lemieux (1996) and Card (2001), this calculation assumes the union effect has been constant over time. Even if this assumption does not hold, the calculation serves as a way to compare different estimates of union wage compression using different earnings and union coverage concepts.

the main sample used in the analysis covers the years 1990-2007, and includes elections in states included in the LEHD, for a total of 12,894 elections in the matched dataset. The raw data contains results from elections stemming from several different type of petitions, including cases where a union seeks to be certified (RC), an employer seeks an election to remove an existing union (RM), or employees seek to remove a union (RD). I restrict to RC-cases, where a union seeks certification. The dataset contains many duplicate records. In some cases they are true duplicates: one election generated multiple records in the database. In these cases I simply delete the redundant entries. In other cases, multiple entries arise from more than one union being on the ballot. In these cases the relevant union vote share is the largest one; I therefore retain the entry with the largest vote share, and delete the others. Finally, in some cases multiple elections were held at the same establishment because, for example, different groups of workers constituted different bargaining units. Since I can't distinguish between workers in different bargaining units, the relevant vote share is the largest, so again I keep only the entry corresponding to the election where the union received the highest vote share.

The second data component consists of the Employment History Files (EHF) within the LEHD database. As described in the text, the EHF contains employee, employer, and earnings data for each employment relationship that generated at least one dollar of wages. The EHF includes a state employer identification number (SEIN) with each record, and in some cases an identifer for the establishment within the employer, which is important for multi-unit employers. For the cases where there is no establishment identifier, the LEHD provides a Unit-to-Worker (U2W) imputation to assign workers to establishments. The employer name and address of these establishments—obtained from the Business Register's Standard Statistical Establishment List (SSEL)—are then used to link to the union election dataset to determine union coverage status.

The matching process to combine these two data sources is as follows. First, employer name and address information from both the NLRB dataset and the Census Bureau's Business Register (BR) were cleaned and standardized using the SAS Data Quality Server standardization functions. NLRB election records were then matched to BR records by several combinations of state, county, city, employer name, street address, and industry code. The match was performed iteratively in descending order of strictness. The cutoff level of strictness was determined by hand checking matches from each iteration, and stopping once match quality dipped below 95 percent. The match

rate from this procedure ranged from 75 to 80 percent of the election records. The matched BR records were then linked to employers in the LEHD's Employer Characteristics File (ECF) by the Business Register Bridge (BRB) via state, year, county, Employer Identification Number (EIN) and industry code. Finally the work histories (including earnings) of all individuals employed at the matched employers during the quarter of the certification election were drawn from the Personal History File (PHF), using the Unit-2-Worker imputation to complete the match in the case of multi-unit employers.

The main analysis sample includes all workers with positive earnings at election establishments in the quarter just prior to a union election. Some of the analysis focuses on full-time, full-year (FTFY) workers in order to isolate wage effects from hours effects. FTFY status was determined based on pre-election earnings and employment history: workers were deemed FTFY if they were continuously employed at the election establishment for at least four quarters before the election, and had earnings during that period of at least $20,000 (in 2000 dollars). Although crude, this conservative approximation to full-time, full-year status is based on pre-determined wages and so does not affect the validity of the estimation, and aids in interpretation. In addition, due to some extreme outliers in annual earnings, records with annual earnings above the 99th percentile were dropped.

7.1 Equivalence of distributional effects under rank stability

Theorem 1 *Suppose Y_{-1}, Y_0, and Y_1 are continuously distributed and the following rank stability conditions hold almost surely:*

a. $F_{Y_{-1}|R=0}(Y_{-1}) = F_{Y_0|R=0}(Y_0)$

b. $F_{Y_0|R=0}(Y_0) = F_{Y_1|R=0}(Y_1)$.

Then $\beta_Y(\tau) = QTE(\tau)$ for all $\tau \in (0,1)$.

Proof. Define the conditional expectation in (2) in terms of the following limit:

$$\begin{aligned}\beta_Y(\tau) &\equiv \lim_{\varepsilon\to 0} E\left[Y_1 - Y_0 | F_{Y_{-1}|R=0}(Y_{-1}) \in (\tau-\varepsilon, \tau+\varepsilon), R=0\right] \\ &= \lim_{\varepsilon\to 0} E\left[Y_1 | F_{Y_1|R=0}(Y_1) \in (\tau-\varepsilon, \tau+\varepsilon), R=0\right] \\ &\quad -\lim_{\varepsilon\to 0} E\left[Y_0 | F_{Y_0|R=0}(Y_0) \in (\tau-\varepsilon, \tau+\varepsilon), R=0\right],\end{aligned}$$

where the second line uses the rank stability conditions in the theorem. Applying the monotonic transformations $Q_{Y_1|R=0}(\cdot) \equiv F^{-1}_{Y_1|R=0}(\cdot)$ and $Q_{Y_0|R=0}(\cdot) \equiv F^{-1}_{Y_0|R=0}(\cdot)$ to the left and right sides of the conditioning sets, we obtain

$$\begin{aligned}&\lim_{\varepsilon\to 0} E\left[Y_1 | Y_1 \in \left(Q_{Y_1|R=0}(\tau-\varepsilon), Q_{Y_1|R=0}(\tau+\varepsilon)\right), R=0\right] \\ &-\lim_{\varepsilon\to 0} E\left[Y_0 | Y_0 \in \left(Q_{Y_0|R=0}(\tau-\varepsilon), Q_{Y_0|R=0}(\tau+\varepsilon)\right), R=0\right] \\ = \;& Q_{Y_1|R=0}(\tau) - Q_{Y_0|R=0}(\tau),\end{aligned}$$

which is the definition of $QTE(\tau)$ in (1). ■

7.2 Model Solution

The optimal wage schedule is part of a subgame perfect equilibrium: it maximizes the union's probability of winning the certification election, given that workers vote sincerely conditional on the wage schedule, and given that the employer invests and hires to maximize profits conditional on the wage schedule and the outcome of the election. In this setting maximizing the probability of winning is equivalent to maximizing the vote share. To see this, consider an election at a plant with N workers voting. Let $V_i(\theta)$ be the i-th worker's voting rule, as a function of the union's choice parameter, θ. The union solves $\max_\theta \Pr\left(\sum_{i=1}^N V_i(\theta) > N/2\right)$. The number of votes is a Binomial random variable with parameters $(P(\theta), N)$, where $P(\theta)$ is the expected vote share. Denoting the cdf of this binomial random variable by F_B, the union's problem can be rewritten as $\max_\theta 1 - F_B(N/2; (P(\theta), N))$. The first order condition is $\frac{\partial F_B}{\partial P}\frac{\partial P}{\partial \theta} = 0$. Since for the Binomial distribution $\frac{\partial F_B}{\partial P} \neq 0$, this reduces to $\frac{\partial P}{\partial \theta} = 0$, which is the first order condition for maximizing the expected vote share.

Intuitively, for a proposed wage schedule to garner a positive vote share in equilibrium, it must

result in K earning at least $(1-\phi_K)(y-v)$. Otherwise, the employer would simply shut down, and all workers would incur the switching cost to find a job elsewhere, and thus would prefer no union in the first place.[16] Thus, although the union would like to set the wage schedule so as to garner the most votes possible, it must also take into consideration the direct effect of the wage schedule on the firm's profits via the payroll, as well as the indirect effect via the wage schedule's effect on the distribution of workers' human capital employed at the firm. In consequence, any incentive to redistribute rents among workers of different levels of human capital is tempered by the tendency of workers who are losers under the union wage schedule to quit, further tightening the firm's profit constraint. The following proposition formally characterizes the optimal union wage schedule.

Proposition 2 *A subgame perfect Nash equilibrium union wage schedule is* $w^U(H) = (v + r(H, \lambda^*))H$, *where the union rent schedule* $r(H, \lambda)$ *satisfies*

$$\frac{f_\eta(r)}{F_{\eta-\varepsilon}(r)} = \lambda H\left(1 - (\phi_K(y-v) - r)\frac{f_{\eta-\varepsilon}(r)}{F_{\eta-\varepsilon}(r)}\right),$$

and λ^* *satisfies*

$$\int (\phi_K(y-v) - r(h, \lambda^*))\, hF_{\eta-\varepsilon}(r(h, \lambda^*))\, dF_H(h) = 0,$$

and $F_{\eta-\varepsilon}$ *denotes the distribution function of the random variable* $\eta - \varepsilon$, *the unionization cost net of the switching cost.*

Proof. Working backwards, consider the firm's investment and hiring decision given $r(H) = w(H)/H - v$, the outcome of the election, and workers' quitting decisions. First take the case where the union loses. Then the price of H hasn't changed, and the competitive equilibrium remains optimal for the employer. If any workers quit, the employer hires from the pool of applicants (in this case identical in distribution to the population of workers) to replace them and production continues. If no workers quit, no additional hiring or investment takes place and production resumes as before. Now take the case where the union wins. Since in the short run the production technology takes fixed proportions of H and K, if no workers quit, the available actions are hiring additional

[16] A recent paper by Kremer and Olken (2009) highlights unions' incentives to take into account employers' profitability.

H and K, or releasing currently employed H and K. The employer would hire additional H and K only if the return to the additional K is greater than $y - v$ (the purchase price of K). However, the employer must pay at least v in order to hire more H, so the equilibrium cannot involve the employer hiring additional H or investing in more K. Still considering the case where the union wins but no workers quit, the employer releases currently employed H and K if the return to K under the union wage schedule is less than $(1 - \phi_K)(y - v)$. In consequence of constant returns to scale and fixed proportions, if it's optimal to release any H and K, it is optimal to shut down completely. Finally, in the case where the union wins and some workers quit, it is never optimal for the employer to replace the lost workers, for the same reason the employer doesn't hire additional workers. Instead the employer will divest the freed up K (recouping $(1 - \phi_K)(y - v)$ per unit). The employer will either resume production with the remaining H and K, or shut down, again depending on whether the return to the remaining K under the union wage schedule is greater than or less than $(1 - \phi_K)(y - v)$. In summary, the employer's equilibrium strategy is to divest any K freed up by quitting workers, and continue production with what remains if the return to K is at least $(1 - \phi_K)(y - v)$, and shut down otherwise.

Turn now to the worker's decision to stay or leave conditional on $r(H)$ and given the employer's equilibrium strategy. A worker of human capital H chooses to stay if her union rent, $r(H)$, exceeds her cost of unionization net of her switching cost. Thus the equilibrium decision rule is $1(\eta - \varepsilon \leq r(H))$. The density of human capital conditional on staying at the firm is therefore

$$f_{H|\text{stay}}(h) = \frac{F_{\eta-\varepsilon}(r(h)) f_H(h)}{\int F_{\eta-\varepsilon}(r(s)) dF_H(s)},$$

where $F_{\eta-\varepsilon}$ denotes the distribution function of the random variable $\eta - \varepsilon$, the cost of unionization net of the switching cost.

Next, consider the workers' voting choice conditional on $r(H)$ and given the employer's equilibrium strategy. No worker will vote for the union if $r(H)$ is such that the employer shuts down production, since all workers would then incur the switching cost ε, and earn the same wage, v, elsewhere. Conditional on $r(H)$ not resulting in a shut-down, a worker of human capital H votes for the union if her rent under the union schedule exceeds her cost of unionization, η. The equilibrium voting function is therefore $1(\eta \leq r(H))$, again subject to the condition that the plant stays

open under the wage schedule.

Finally, given the workers' and the employer's equilibrium strategies, the union chooses $r(H)$ to maximize the expected vote share, which is equivalent to maximizing the expectation of $1(\eta \leq r(H))$, subject to the firm earning an ex post return of at least $(1-\phi_K)(y-v)$. The problem the union solves can be written:

$$\begin{aligned} &\max_{r(h)} \int F_\eta(r(h))\, dF_H(h) \\ &\text{s.t.} \int (\phi_K(y-v) - r(h))\, hF_{\eta-\varepsilon}(r(h))\, dF_H(h) \geq 0. \end{aligned} \tag{5}$$

This is a straightforward calculus of variations problem of the type treated by, say, Theorem 1 in Gelfand and Fomin (1963, p. 43). The optimal rent schedule therefore satisfies

$$\frac{f_\eta(r)}{F_{\eta-\varepsilon}(r)} = \lambda H\left(1 - (\phi_K(y-v) - r)\frac{f_{\eta-\varepsilon}(r)}{F_{\eta-\varepsilon}(r)}\right),$$

where λ is the Lagrange multiplier on the minimum profit constraint for the firm. This condition can be solved for $r(H,\lambda)$. Substituting this into the profit constraint, (5), and solving for λ gives λ^*, and thus the optimal wage schedule can be written

$$w^U(H) = (r(H,\lambda^*) + v)H.$$

■

Numerical example details

The numerical example in Section 6 is based on the distribution of wages observed in the the sample of full-time, nonunion workers from the 1998-2000 Current Population Survey (CPS). Taking the hourly wage as a measure of human capital, H, I assume the sunk fraction of K is $\phi_K = .3$. Normalizing the competitive return to human capital to be $v = 1$, and assuming labor's share in income is about .7, I set the production function parameter to be $y = 10/7$. Finally, I assume unionization and switching costs both have means of $1.5 \times \phi_K(y-v)$.

The optimal wage schedule in this example distributes rents disproportionately to low-wage

Figure 14: Union wage premium as a function of human capital for parameter values $\phi_K = .3$, $v = 1$, $y = 10/7$, $\lambda_\eta = \lambda_\varepsilon = (1.5\phi_K (y - v))^{-1}$.

workers and substantially compresses the distribution of earnings. Figure 14 plots the union premium (in dollars), $r(H) \times H$, implied by the optimal union wage schedule as a function of human capital. The union premium is positive for lower levels of human capital, and decreasing and eventually negative for higher levels of human capital.

References

Abowd, John M., John Haltiwanger, and Julia Lane (2004): "Integrated Longitudinal Employee-Employer Data for the United States," *American Economic Review Papers and Proceedings*, 94, 224–229.

Abowd, John M., Bryce Stephens, Lars Vilhuber, Fredrik Andersson, Kevin L. McKinney, Marc Roemer, and Simon Woodcock (2009): "The LEHD Infrastructure Files and the Creation of the Quarterly Workforce Indicators," in *Producer Dynamics: New Evidence from Micro Data*, ed. by T. Dunne, J.B. Jensen, and M.J. Roberts, pp. 149–230. University of Chicago Press for the National Bureau of Economic Research, Chicago.

Acemoglu, Daron, Phillipe Aghion, and Gianluca Violante (2002): "Technical Change, Deunionization, and Inequality," in *Carnegie-Rochester Conference Series on Public Policy*.

Ashenfelter, Orley, and George E. Johnson (1969): "Bargaining Theory, Trade Unions, and Industrial Strike Activity," *The American Economic Review*, 59, 35–49.

Atherton, Wallace N. (1973): *Theory of Union Bargaining Goals*. Princeton University Press, Princeton, New Jersey.

Atkinson, Anthony B. (1970): "On the Measurement of Inequality," *Journal of Economic Theory*, 2, 244–263.

Berkowitz, Monroe (1954): "The Economics of Trade Union Organization and Administration," *Industrial and Labor Relations Review*, 7, 575–592.

Blanchflower, David, and Alex Bryson (2003): "What effect do unions have on wages now and would 'What do unions do?' be surprised?," NBER Working Paper No. 9973.

Booth, Alison L. (1995): *The Economics of the Trade Unions*. Cambridge University Press.

Brudney, James J. (2005): "Contractual Approaches to Labor Organizing: Supplanting the Election Paradigm?," in *Labor and Employment Relations Association Series, Proceeding of the 57th Annual Meeting*, pp. 106–113.

Bureau of Labor Statistics (2011): "Union Members Summary," in BLS annual union membership report, 2010.

Caballero, Ricardo J., and Mohamad L. Hammour (1998): "The Macroeconomics of Specificity," *The Journal of Political Economy*, 106, 724–767.

Campbell, Donald Thomas (1969): "Reforms as Experiments," *American Psychologist*, 24, 409–429.

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

——— (2001): "The Effect of Unions on Wage Inequality in the U.S. Labor Market," *Industrial and Labor Relations Review*, 54, 296–315.

Card, David, Thomas Lemieux, and W. Craig Riddell (2004): "Unions and Wage Inequality," *Journal of Labor Research*, 25, 519–559.

Chamberlain, Gary (1994): "Quantile Regression, Censoring and the Structure of Wages," in *Proceedings of the Sixth World Cogress of the Econometrics Society, Barcelona, Spain*, ed. by Christopher A. Sims, and Jean-Jacques Laffont, pp. 179–209. Cambridge University Press, New York.

Chernozhukov, Victor, and Christian Hansen (2005): "An IV Model of Quantile Treatment Effects," *Econometrica*, 73, 245–261.

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

DiNardo, John, and David S. Lee (2004): "Economic Impacts of New Unionization on Private Sector Employers: 1984-2001," *Quarterly Journal of Economics*, 119, 1383–1441.

Farber, Henry S. (1978): "Individual Preferences and Union Wage Determination: The Case of the United Mine Workers," *The Journal of Political Economy*, 86, 923–942.

Ferguson, John-Paul (2008): "The Eyes of the Needles: A Sequential Model of Union Organizing Drives, 1999-2004," *Industrial & Labor Relations Review*, 62, 3–21.

Firpo, Sergio, Nicole M. Fortin, and Thomas Lemieux (2009): "Unconditional Quantile Regressions," *Econometrica*, 77, 953–973.

Frandsen, Brigham R., Markus Frölich, and Blaise Melly (2010): "Quantile Treatment Effects in the Regression Discontinuity Design," Unpublished manuscript, available at http://econ-www.mit.edu/grad/frandsen/research.

Freeman, Richard B. (1993): "How Much Has De-unionisation Contributed to the Rise in Male Earnings Inequality?," in *Uneven Tides: Rising Inequality in America*, ed. by Sheldon Danziger, and Peter Gottschalk, pp. 133–163. Russell Sage Foundation, New York.

Gelfand, I. M., and S. V. Fomin (1963): *Calculus of Variations*. Prentice-Hall, New Jersey, Edited by Richard A. Silverman.

Greenhouse, Steven (2009): "Fierce Lobbying Greets Bill to Help Workers Unionize," New York Times, March 9, 2009, Retrieved from the World Wide Web September 23, 2011 at http://www.nytimes.com/2009/03/11/business/11labor.html?ref=business.

Heckman, James J., Jeffrey Smith, and Nancy Clements (1997): "Making the Most Out of Programme Evaluations and Social Experiments: Accounting for Heterogeneity in Programme Impacts," *The Review of Economic Studies*, 64, 487–535.

Hirsch, Barry T. (1991): *Labor Unions and the Economic Performance of U.S. Firms*. Upjohn Institute for Employment Research, Kalamazoo, Mich.

Holmes, Thomas J. (2006): "Geographic Spillover of Unionism," Federal Reserve Bank of Minneapolis Research Department Staff Report 368.

Imbens, Guido, and Karthik Kalyanaraman (2009): "Optimal Bandwidth Choice for the Regression Discontinuity Estimator," NBER Working Paper No. 14726.

Kopczuk, Wojciech, Emmanuel Saez, and Jae Song (2010): "Earnings Inequality and Mobility in the United States: Evidence from Social Security Data Since 1937," *The Quarterly Journal of Economics*, 125, 91–128.

Kremer, Michael, and Benjamin A. Olken (2009): "A Biological Model of Unions," *American Economic Journal: Applied Economics*, 1, 150–175.

Lane, Julia (2008): "Linking Administrative and Survey Data," Society of Labor Economists 2008 Annual Meeting Presentation.

Lee, David, and Alexandre Mas (forthcoming): "Long-Run Impacts of Unions on Firms: New Evidence from Financial Markets, 1961-1999," *Quarterly Journal of Economics*.

Lee, David S. (2008): "Randomized experiments from non-random selection in U.S. House elections," *Journal of Econometrics*, 142, 675 – 697, The regression discontinuity design: Theory and applications.

Lewis, H. Gregg (1986): *Union Relative Wage Effects A Survey*. University of Chicago Press.

McKinney, Kevin L., and Lars Vilhuber (2008): "LEHD Infrastructure files in the Census RDC - Overview Revision 219," U.S. Census Bureau Technical Paper.

National Labor Relations Board (2010): "Procedures Guide," Available on the world wide web at http://www.nlrb.gov/publications/Procedures_Guide.htm.

——— (2011): "Representation–Case Procedures Proposed Rule," *Federal Register*, 76, 36812–36847.

Ross, Arthur Max (1948): *Trade Union Wage Policy*. University of California Press, Berkeley.

Thistlethwaite, Donald L., and Donald T. Campbell (1960): "Regression-Discontinuity Analysis: An Alternative to the Ex Post Facto Experiment," *Journal of Educational Psychology*, 51, 309–317.

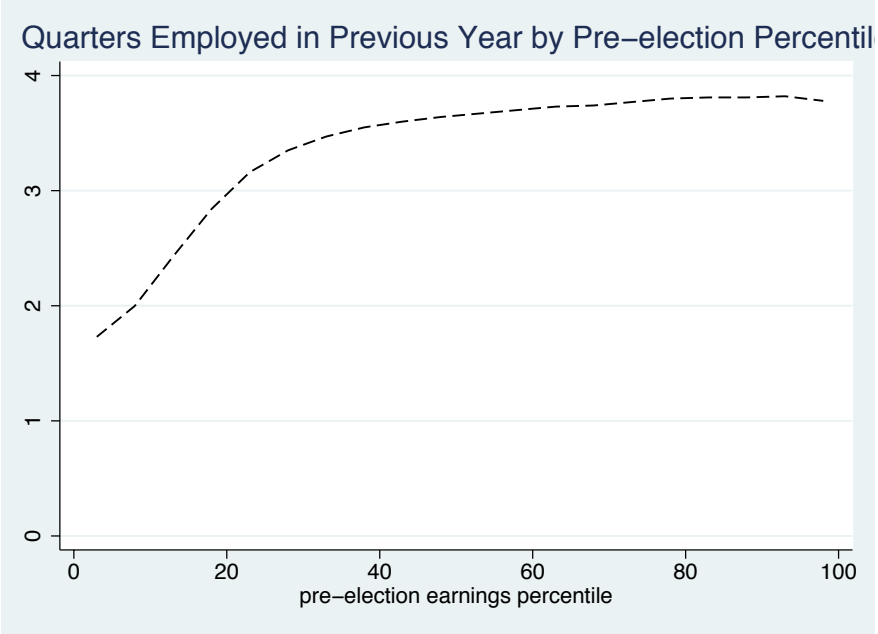

Figure 15: Average number of quarters employed during the year previous to a certification election by pre-election annual earnings percentile. The points are raw fractions by non-overlapping vote share bins of width .05. The lines are predicted values from a piecewise-cubic fit. Data are from the LEHD and NLRB election records.

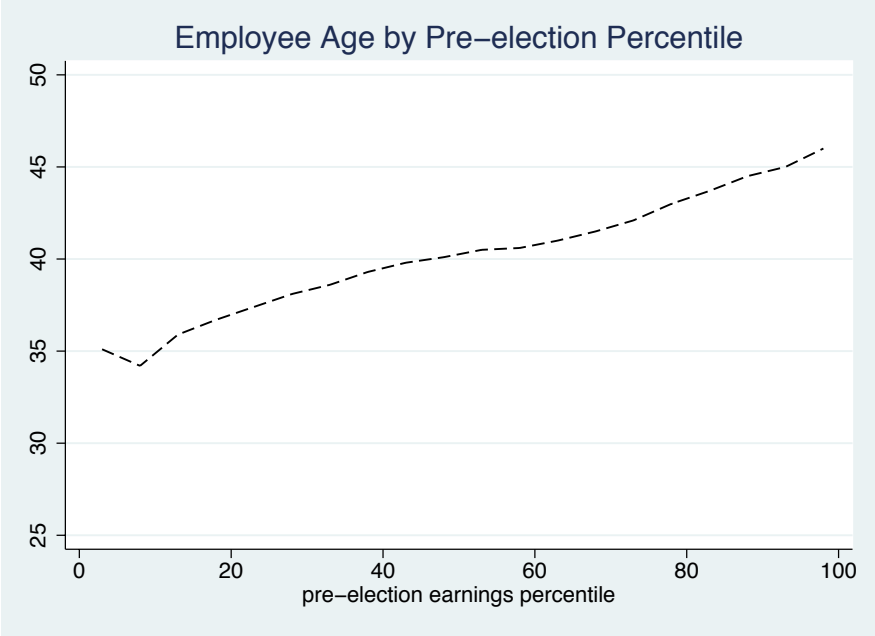

Figure 16: Average age at time of certification election by pre-election annual earnings percentile. The points are raw fractions by non-overlapping vote share bins of width .05. The lines are predicted values from a piecewise-cubic fit. Data are from the LEHD and NLRB election records.

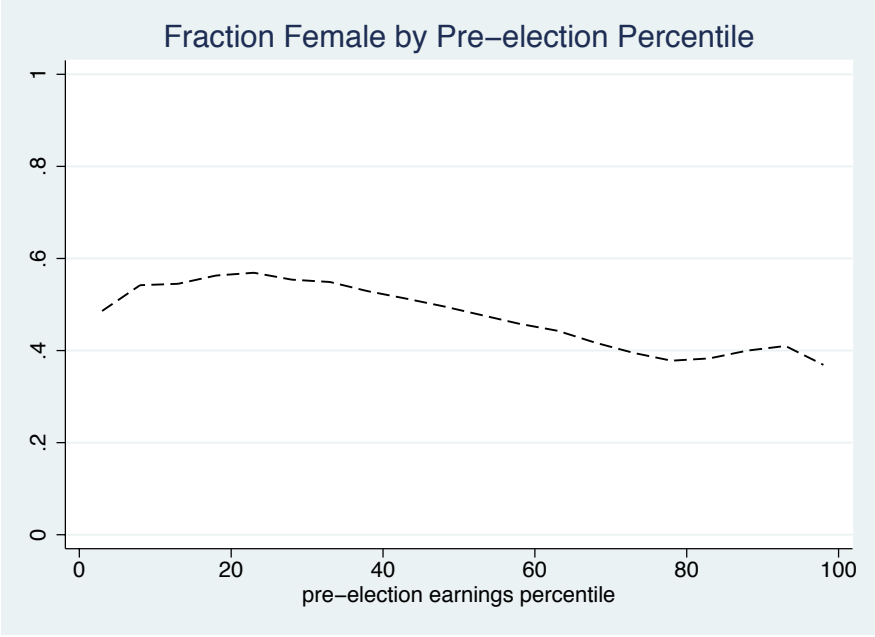

Figure 17: Fraction female by pre-election annual earnings percentile. The points are raw fractions by non-overlapping vote share bins of width .05. The lines are predicted values from a piecewise-cubic fit. Data are from the LEHD and NLRB election records.

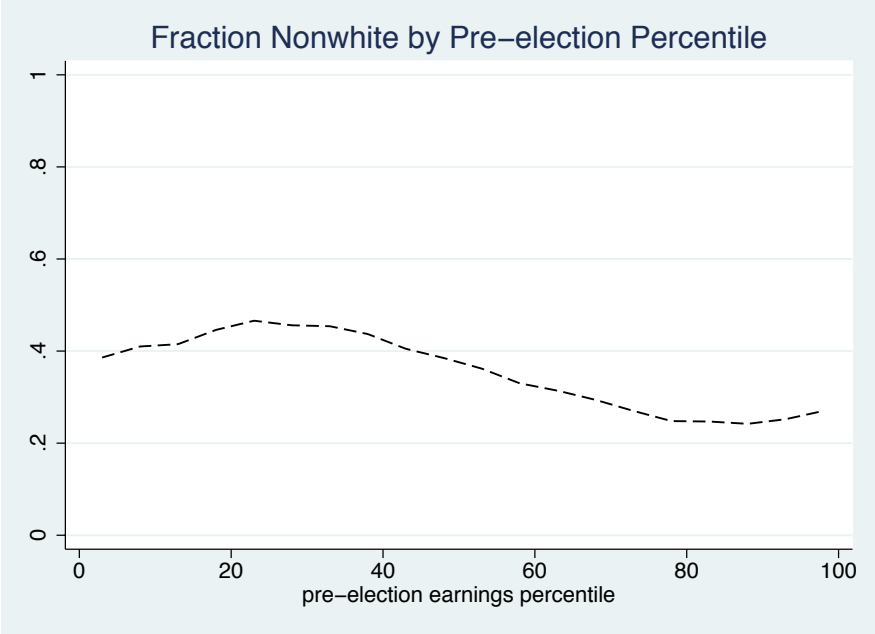

Figure 18: Fraction nonwhite by pre-election annual earnings percentile. The points are raw fractions by non-overlapping vote share bins of width .05. The lines are predicted values from a piecewise-cubic fit. Data are from the LEHD and NLRB election records.

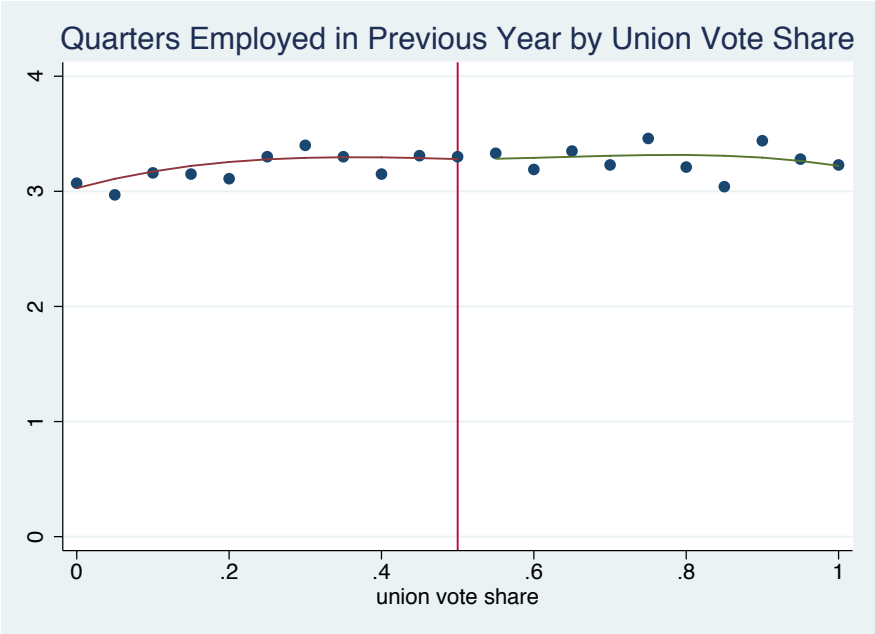

Figure 19: Average number of quarters employed during the year previous to a certification election by union vote share. The points are raw fractions by non-overlapping vote share bins of width .05. The lines are predicted values from a piecewise-cubic fit. Data are from the LEHD and NLRB election records.

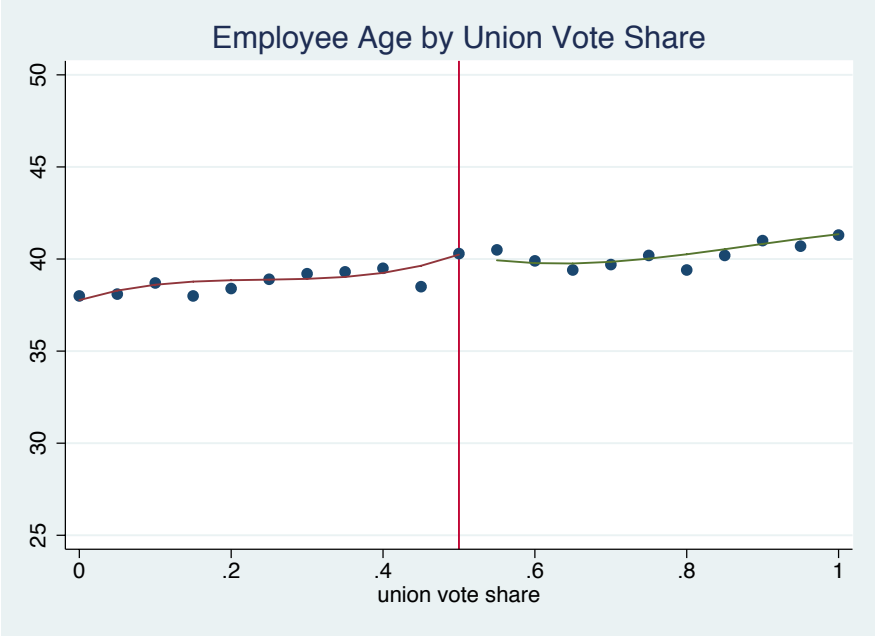

Figure 20: Average age at time of certification election by union vote share. The points are raw fractions by non-overlapping vote share bins of width .05. The lines are predicted values from a piecewise-cubic fit. Data are from the LEHD and NLRB election records.

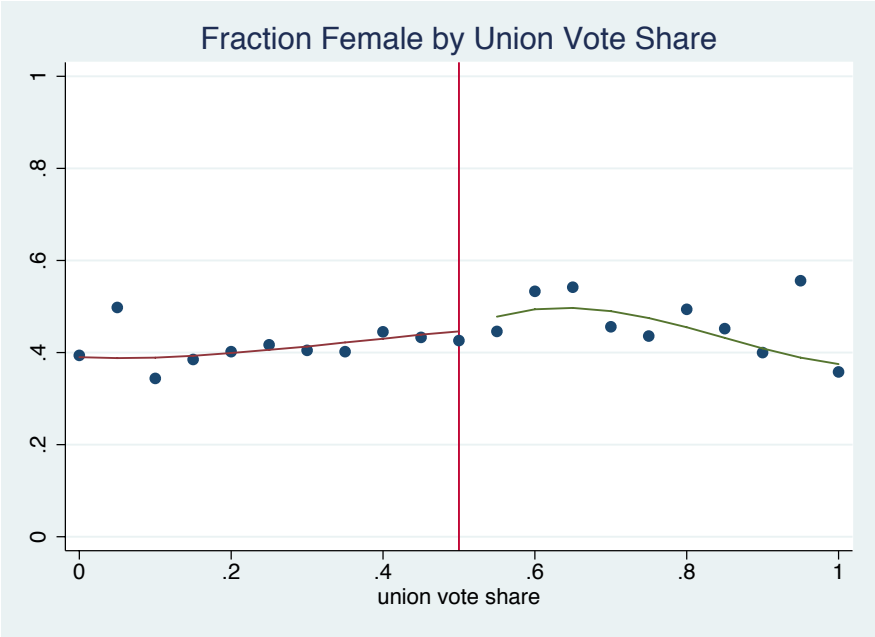

Figure 21: Fraction female by union vote share. The points are raw fractions by non-overlapping vote share bins of width .05. The lines are predicted values from a piecewise-cubic fit. Data are from the LEHD and NLRB election records.

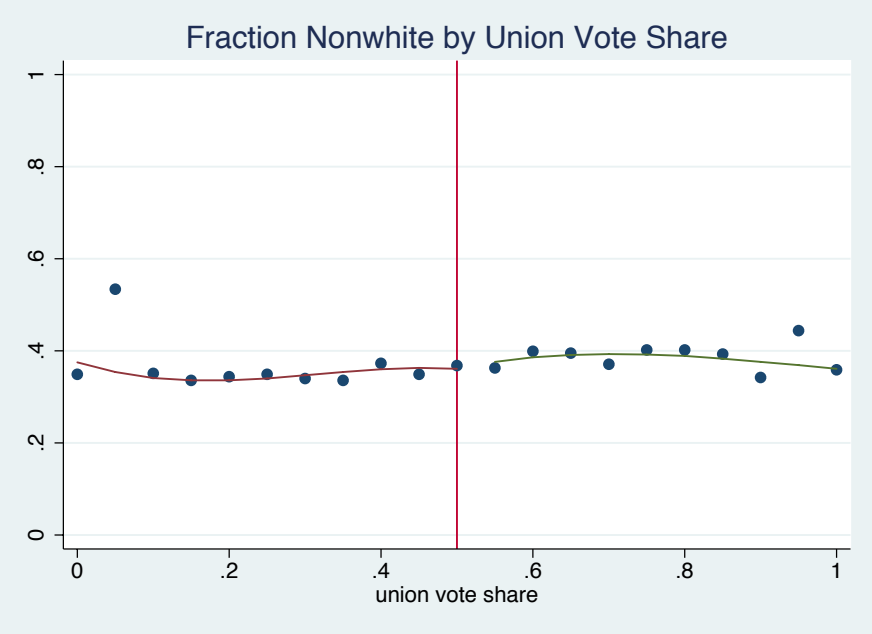

Figure 22: Fraction nonwhite by union vote share. The points are raw fractions by non-overlapping vote share bins of width .05. The lines are predicted values from a piecewise-cubic fit. Data are from the LEHD and NLRB election records.