

WORKING PAPER #547
INDUSTRIAL RELATIONS SECTION
PRINCETON UNIVERSITY
JANUARY 2009

Long-Run Impacts of Unions on Firms: New Evidence from Financial Markets, 1961-1999¹

David S. Lee
Princeton University and NBER

Alexandre Mas
UC Berkeley and NBER

January 2009

Abstract

We estimate the effect of new unionization on firms' equity value over the 1961-1999 period using a newly assembled sample of National Labor Relations Board (NLRB) representation elections matched to stock market data. Event-study estimates show an average union effect on the equity value of the firm equivalent to a cost of at least \$40,500 per unionized worker. At the same time, point estimates from a regression-discontinuity design – comparing the stock market impact of close union election wins to close losses – are considerably smaller and close to zero. We find a negative relationship between the cumulative abnormal returns and the vote share in support of the union, allowing us to reconcile these seemingly contradictory findings. Using the magnitudes from the analysis, we calibrate a structural “median voter” model of endogenous union determination in order to conduct counterfactual policy simulations of policies that would marginally increase the ease of unionization.

¹We thank Jonathan Berk, David Card, John DiNardo, Lawrence Katz, Morris Kleiner, Robert Moffitt, Jesse Rothstein, Eric Verhoogen, Hans-Joachim Voth, and numerous seminar participants for helpful suggestions. Emily Buchsbaum, Mariana Carrera, Elizabeth Debraggio, Briallen Hopper, Sanny Liao, Xiaotong Niu, Zhuan Pei, Andrew Shelton, and Fanyin Zheng provided outstanding research assistance. We gratefully acknowledge research support from the Center for Economic Policy Studies at Princeton University.

“[L]aymen and economists alike tend, in my view, to exaggerate greatly the extent to which labor unions affect the structure and level of wage rates.” – Milton Friedman, 1950¹

“Everyone ‘knows’ that unions raise wages. The questions are how much, under what conditions, and with what effects on the overall performance of the economy.” – Richard Freeman and James Medoff, 1984²

1 Introduction

It is undisputed that employers oppose unions, viewing them as a threat to profitability. An example receiving recent national attention is Wal-Mart’s effort to resist unionization – from its strategic location of stores in areas less favorable to unions to its hard-line stance against organization (Basker, 2007). According to a handbook the retailer distributed to its managers, “Staying union free is a full-time commitment...The commitment to stay union free must exist at all levels of management – from the Chairperson of the “Board” down to the front-line manager...”³ It is easy to find isolated cases that confirm the fears of employers like Wal-Mart. For example, in a March 1999 National Labor Relations Board (NLRB) representation election, workers at National Linen Service (NLS) Corp., a large linen supplier, voted by an over 2 to 1 margin to organize as a local chapter of the Union of Needletrades, Industrial, and Textile Employees (UNITE). The stock market response appeared to punish NLS in a severe, though perhaps not swift, fashion. Figure 1 shows the cumulative return of NLS’ stock for the two years prior to and following the election, as well as the cumulative return of a broad market index over the same period. Before the election, the returns for NLS and the market tracked each other quite closely. But immediately following the election, NLS began to lag. By March 2001, the price of NLS shares had fallen by 25 percent, while the broad market index had increased by 50 percent.

How general is this phenomenon? Is NLS the exception or the rule? Despite an enormous literature documenting numerous aspects of unions and their role in the labor market, the magnitude of an “average” effect of unions on firm performance throughout the economy remains somewhat unclear.

Empirically, there are at least three reasons why measuring these effects is quite challenging. First, large-scale establishment or firm-level micro-data containing the relevant information on the extent of unionization

¹See Friedman (1950).

²See Freeman and Medoff (1984).

³Quoted in Featherstone (2004).

are not readily available. Second, even when such data are available, omitted variables and the endogeneity of unionization at the firm-level makes it difficult to separate causal effects from other unobserved confounding factors. Third, it is difficult to find data that can also be plausibly representative of the population of unionized companies in the United States.

Furthermore, from a theoretical standpoint, it is not obvious to what degree unions should affect firms. One view, articulated by Friedman (1950), is that workers would reject substantially above-market wages, knowing full well that such wages could adversely affect job security. Unions, after taking these considerations into account, would tend to moderate wage demands.⁴ Moreover, firms may respond to a unionization threat by conceding higher wages and better working conditions. Accounting for these forces suggests a reduction in the gap in compensation and working conditions between union and non-union workforces, at least in situations where there is a threat of unionization. The possibility that unions may temper their demands because of electoral pressure may help explain the results of DiNardo and Lee (2004), who found generally small differences in wages, employment, and output between unionized and otherwise comparable non-unionized workplaces in close representation elections.

In this paper, we first assess the extent to which the pattern in Figure 1 is a generalizable phenomenon, measuring an average overall effect of unionization among publicly-traded firms. To do so, we begin with a sample frame that is the universe of all firms with NLRB union representation elections between 1961-1999. Since a large number of unionized workplaces in the U.S. come into existence via a secret-ballot election on the question of representation, this population provides a reasonable representation of newly unionized workplaces and, to the extent they survive, the future stock of unions in the United States.

We begin analyzing the stock market reaction to union victories using event-study methodologies. The most distinctive feature of our data – crucial for our research design – is the long panel (up to 48 months before and after the election) of high-frequency data on stock market returns for each firm. This feature allows us to use the pre-event data to test the adequacy of the benchmarks used to predict the counterfactual returns in the post-event period. The long panel also allows us to examine returns several months beyond the event, so as to capture the long-run expected effects of new unions, without having to rely heavily on the assumption that the stock price immediately and instantaneously adjusts to capture the expected presence of the unions.

Our event-study analysis reveals substantial losses in market value following a union election victory –

⁴It is this line of reasoning that led to Friedman's view that the impact of wages was exaggerated (Friedman, 1950).

about a 10 percent decline, equivalent to about \$40,500 per unionized worker. The evidence supporting this finding is compelling: we find that these firms' average returns are quite close to the predicted returns every month leading up to the election, but at precisely at the time of the election, the actual and counterfactual returns diverge. The results for these firms are robust to a number of different specifications. In the sample of firms where we know that the union is a small fraction of the workforce, we do not find a similar divergence of returns from the benchmark. Notably, while the equity value of firms unionizing begins diverging at the time of the election, the full effect takes approximately 18 months to fully materialize. This slow speed of adjustment is consistent with a number of studies showing that over horizons of 3 to 12 months share prices underreact to bad news (Hong et al., 2000).

The event-study estimate appears to average a great deal of heterogeneity in the effects. We additionally employ a Regression Discontinuity (RD) design, implicitly comparing close union victories to close union losses, and consistent with DiNardo and Lee (2004), we find little evidence of a significant discontinuous relationship between the vote share and market returns. If anything, the RD point estimates show a 4 percent *positive* (though statistically insignificant) effect of union certification (*vis-a-vis* union defeat). The event-study estimates vary systematically by the observed vote share, with the largest negative abnormal returns for cases where the union won the election by a large margin.

We use our estimates to make predictions for the effects of policies that lower the threshold for new unionization, such as the recently proposed Employee Free Choice Act (EFCA). To do so while also incorporating unions' and firms' responses to the new policy requires modeling their behavior and interactions. We choose as our framework a two-party model of electoral competition, where the firm and the union are each seeking to win the sympathies of the "median" voter in an NLRB election. As is standard in this class of models, despite having opposing interests, the two parties may be forced to propose a level of compensation, (accompanied by a risk of job loss), that is closer to the preferences of the median voter.

Within this framework, which is reminiscent of Friedman's view, the RD design estimate of the unionization effect identifies the gap between the union's and firm's proposals for workplaces where the median voter has moderate demands. Depending on how aggressively firms and unions court voters, this gap could be close to zero, even if *on average* – including both small and large electoral victories – unions significantly affect the profitability of firms. Viewed through the lens of this model, the pattern of results imply that voters have a strong desire for higher wages (or better working conditions) in a relatively small number of elections. Overall, our policy simulation exercise suggests that a policy-induced increase in the win rate from 33

to 70 percent would lead to a 4.3 percent decline in market value, averaged across all firms (including firms that unionize under the new policy, as well as those that remain nonunion). For a more dramatic policy that increases the win rate from 33 to nearly 99 percent, the estimate is a decline of about 11 percent averaged across all firms.

The remainder of the paper is organized as follows. Section 2 briefly highlights what is known from the literature and how our study relates to those findings. We provide some institutional details in Section 3 that are relevant to our research design, which we describe along with our data. We present and discuss the empirical results in Section 4. In Section 5 we present a structural model, which we then use to conduct counterfactual policy simulations. Section 6 concludes.

2 Existing Literature and Background

In this section we provide a brief overview of the literature most relevant to our analysis. First, there is an enormous union wage premium literature, discussed and summarized in the landmark works of Freeman and Medoff (1984) and Lewis (1986), with more recent evidence discussed in Blanchflower and Bryson (2007). These studies typically use household-level survey data to compute the wages for workers who are union members, comparing them to “otherwise comparable” non-union members. In some cases, these studies track workers in longitudinal data sets, as they switch from union to non-union status. In their analysis, Freeman and Kleiner (1990) note that these “[e]stimates based on longitudinal data...contrast workers who change union status by moving to or from already organized workplaces rather than contrasting workers in plants that are newly organized with those in plants that remain non-union.” Expanding upon this point, DiNardo and Lee (2004) clearly state that the effect of *unionization* (changing a workplace from non-union to union) is distinct from the effect of moving a *worker* from a non-unionized to a unionized workplace. In particular, a “typical” unionized workplace may differ from a “typical” non-unionized workplace along a number of dimensions (e.g. geography, firm size, industry), which may independently influence wage levels. We therefore view this well-established literature as fundamentally unable to account for the selection of unionism at the firm- or establishment-level, and thus potentially estimating something quite distinct from the causal effect of unionizing a workplace.

Next, there is a literature utilizing firm- or establishment-level data with information on union status. As discussed in Hirsch (2007), a recent study reviewing this evidence, there are a number of important

reasons warranting caution when drawing inferences from the existing research. First, there can be important omitted variables – unobserved determinants of the long-run viability of the firm that could be correlated with the presence of the union. Related, there is a potential endogeneity problem, whereby unions may specifically target a highly profitable firm for organization. Alternatively, it may be that poorly managed, and thus low-performing, firms lead to the demand for worker representation. Examples of studies implicitly relying on the assumption that union status is an exogenous variable include the in-depth analyses of Clark (1984), Hirsch (1991a), and Hirsch (1991b). A second limitation Hirsch (2007) emphasizes is the limited generalizability of many of the studies. For example, the cement industry is examined in Clark (1980a) and Clark (1980b), hospitals and nursing homes in Allen (1986a), the construction industry in Allen (1986b), the trucking industry in Rose (1987), and sawmills in Mitchell and Stone (1992). It is difficult to extrapolate the findings on productivity from these studies to a broader, representative cross-section of firms in the United States. Indeed, our analysis is largely motivated by the belief that it might be easy to find particular incidents and/or companies where unions have imposed large costs on firms. The challenge, however, is determining to what extent isolated examples (such as the one illustrated in Figure 1) generalize to a broader population of interest.

Finally, there are three particular studies that we consider to be most closely related to our analysis: Lalonde et al. (1996), Ruback and Zimmerman (1984), and DiNardo and Lee (2004). We believe our analysis addresses some of the most important limitations of each of these studies.

The main difficulty faced in Lalonde et al. (1996) – which utilizes a “fixed effects” approach with establishment-level panel data from the Longitudinal Research Database (LRD) to examine the impact of a successful union organizing campaign – is one of interpretation. The study shows some differences in employment growth between the eventually successful and failed organizing attempts *prior* to the election event. For example, one sample shows an expanding gap in employment, while another shows a contracting gap. Overall, the estimates and standard errors are consistent with pre-election employment growth differences ranging from -10 to 14 percent. As a result, Lalonde et al. (1996) are careful to note that their examination of pre-election growth rates for many of the outcome variables proved “inconclusive,” and that their “subsequent findings on the effects of unionization may be too large.” Essentially, the main problem is that the data they examine are not rich enough to rigorously test their “difference-in-difference” specification with the pre-event data, and as a result more caution is required in interpreting the post-event patterns.⁵

⁵Another study in the spirit of a “before-after” design is that of Freeman and Kleiner (1990), in which 203 establishments

A similar issue arises in the well-known study of Ruback and Zimmerman (1984), which, like our analysis, examines the stock market reaction to NLRB union certification events.⁶ There, the main estimates of a 3.8 percent drop in stock market valuation is computed within a few months surrounding the unionization event.⁷ Again, the difficulty in interpretation arises from the substantial negative abnormal returns that emerge *well before* the unionization event; specifically, a decline in market value of about 7 percent between the 12th and 7th months preceding unionization. While Ruback and Zimmerman (1984) have no explanation for this significant decline, they argue that it is unlikely to indicate anticipation of the outcome of the election due to its timing.⁸ This pattern raises the question of whether the post-election decline in the stock market valuation reflects unionization or the factors which led to the pre-election trend in the first place.

In our analysis, we address these ambiguities by taking advantage of a very long panel of monthly data on stock returns, using an arguably more disciplined approach to modeling the counterfactual “no union” state. Specifically, we use the data from 24 months prior to the event and just before the event to *test* our specification. If there are significant departures between our predicted returns and the observed returns over the two year period before the event, we consider any estimates obtained from the post-event data to be invalid.⁹ This approach is a direct application of conventional testing of over-identifying restrictions for “difference-in-difference” modeling in labor economics program evaluation.¹⁰ Furthermore, we track abnormal returns over a period of at least 24 months after the unionization. Strictly speaking, perfectly efficient financial markets imply that any changes in valuation caused by the outcome of the election that are known to investors will be fully capitalized into the stock price by the time the outcome of the election is

were surveyed before and after NLRB elections took place and were compared to 161 “control” firms (where elections did not take place). In the study, there was only one period before and one period after, so testing the over-identifying restrictions of the difference-in-difference design was not possible.

⁶There are a number of other studies that examine various aspects of unions through stock market reactions. They typically do not aim to generate effects of unionization (versus the absence of unions), as they use samples of already unionized firms or industries. See Abowd (1989), Becker and Olson (1986), Neumann (1990), DiNardo and Hallock (2002), and Becker (1987). Olson and Becker (1990) is an exception in this regard, as it examines the impact of the passage of the National Labor Relations Act on 75 firms that were at risk of being unionized in the 1930s.

⁷Specifically, the main estimate of -3.84 percent is computed by taking the one-month change associated with the petition date and adding it to the one-month change associated with the date of the actual certification. This can be seen as the summation of the third and fifth rows, which equals the first row of the third column in their Table 2. Their main estimate can also be seen in their Figure 1(c) as the summation of the two downward notches around the petition and certification dates.

⁸Specifically, on p.1145, they note that “[t]he abnormal return for these firms in the 6 months immediately preceding the petition is 0.16 percent. This timing suggests that the pre-petition abnormal returns are not due to unionization. Instead, the results suggest that firms in which unions are successful experienced declines in value prior to the union activity.”

⁹An alternative interpretation of pre-election divergence in the predicted and actual returns is the diffusion of anticipatory information regarding the election outcome. Recognizing this alternative, we allow for non-zero excess returns in a short window prior to the event, but conclude that any significant divergence over a long-period of time prior to the event is evidence of a mis-specified model.

¹⁰For example, see Ashenfelter and Card (1982) and Heckman and Hotz (1989).

revealed. Nevertheless, our approach relies less on the assumption of instantaneous adjustment by examining the patterns of returns for many months following unionization, thus allowing time for the market to adjust.

The final closely related study is the regression discontinuity analysis of union elections, using data from the LRD, in DiNardo and Lee (2004). They exploit the “near-experiment” generated by secret ballot elections, comparing establishments where unions became recognized by a close margin of the vote with workplaces where the union barely lost. The most precise estimates in that study are those on wages: wage increases of 2 percent could be statistically ruled out as far away as seven years after the election.¹¹ There are a number of important limitations to inferring the long-run costs of unions from this evidence. For one, it may take a much longer period of time – perhaps a decade or more – for unions to establish enough support within the workplace to have the required bargaining power to negotiate for substantially higher wages. Secondly, unions impose other costs that are not measured by the LRD, such as the use of seniority rules, work rules, grievance procedures, and other working conditions specified in union contracts. In principal, our approach in this paper of examining the effect of stock market valuation addresses both of these concerns: if the market correctly prices the firm, it should capture the sum of all costs imposed by the union, and effects that might occur many years in the future should be capitalized into the stock market valuation of the firm in the short-run.

A final important limitation is that by estimating a discontinuity in the relationship between wages and the vote share at the 50 percent threshold, the RD analysis can only estimate a weighted average treatment effect, where the weights are proportional to the *ex ante* likelihood an election was predicted to be “close.”¹² That is, among the observed close elections, a disproportionately small number would have had the fundamentals of strong union support. The RD is fundamentally unable to provide a counterfactual for the set of elections where workers voted 90 percent in favor of unionization. By contrast, the present analysis seeks to estimate effects for precisely these “inframarginal cases.” In the analysis we describe below, we present results from both an event-study as well as an RD approach, and provide a framework for interpreting both sets of results.

¹¹Interestingly, the magnitudes are also in line with what was found on wages in Lalonde et al. (1996). Freeman and Kleiner (1990) also find wage effects that are much smaller than those found in cross-sectional worker-level studies.

¹²For a detailed discussion of this interpretation, see Lee (2008).

3 Institutional Background, Data, and Research Design

The National Labor Relations Act provides the legal framework by which most workers in the United States become unionized. Workers who organize into unions through the procedures specified by the NLRA are guaranteed the right to bargain collectively. There are several ways a group of workers may become unionized under the auspices of the NLRA, though it is believed that most new unionization occurs through representation elections (Farber and Western, 2001). There are several steps involved in this process, which are described in detail in DiNardo and Lee (2004). Briefly, when a group of workers decides to organize, they first petition the NLRB to hold a representation election. To be legally granted an election, the petition must be signed by at least 30 percent of the workforce, typically over no longer than a six month period. Once the NLRB determines the appropriate bargaining unit, it holds an election at the work site. The union wins the election with a simple majority of support amongst the workers. Barring objections by the employer, a win means the union is certified as the exclusive bargaining agent for the unit and that the employer is legally required to bargain with the union in good faith.

Our research design and subsequent data collection were motivated by our desire to estimate the average effect of union victories and losses in representation elections on firm profitability, and to attempt to address some of the aforementioned puzzles and challenges in the literature. In collecting the data our goal was to obtain information on the profitability of firms over a long time span, with a panel structure allowing for an event-study design with a long event window. Our sample size needed to be large enough so we could also estimate the cross-sectional relationship between post-event abnormal returns and the union vote share. For these reasons, and because we were also interested in how the union effect evolved over time, we sought to collect information on elections over as many years as possible. Since data on the profits of privately held firms are difficult to come by, we focused on publicly traded firms for which stock market information and other performance measures are available through mandatory disclosure.

3.1 Data Set Assembly

This study primarily uses three sources of data: election results from the NLRB, data from the Center for Research on Security Prices (CRSP), and the CRSP/Compustat Industrial Quarterly Merged Database.

The NLRB began publicly reporting representation election vote tallies in 1961. However, previous studies using NLRB election data typically used records that were already in electronic form (e.g. Farber

and Western, 2001; DiNardo and Lee, 2004; and Holmes, 2006). We use those data for the 1977-1999 period, but augment those with data from 1961-1976 that we digitized for this study.¹³ Data for the 1961-1976 period were hand-entered from hard copies of NLRB monthly election reports. Among other things, the NLRB data set contains the number of voters who voted in favor of the union, the number of voters voting against the union, the number of eligible voters, the name of the company, a two digit industry code, the city and state of the election, and the month that the NLRB closed the election.¹⁴ The CRSP and Compustat data were obtained from Wharton Research Data Services.

The primary objective of the data assembly process was to match companies in the NLRB election files to companies in the CRSP data file. The procedure for matching establishments in NLRB dataset to firms in the CRSP dataset is detailed in the Data Appendix. This matching process is complex because while the NLRB file provides the company name where the election took place, most other identifying information is unknown.^{15,16} However, as explained in the Appendix, we are confident that the match is high quality.

Previous event studies of representation elections use samples of elections with a very large number of eligible voters. Ruback and Zimmerman (1984) and Bronars and Deere (1990) limit their sample to elections with at least 750 eligible voters. Elections of this size are quite rare, thereby resulting in small sample sizes (54 union victories in the main sample of Ruback and Zimmerman, 1984). We believe that the effects of these elections are easier to detect if the number of eligible voters is large *relative* to the size of the firm. However, limiting the sample to large elections is neither necessary nor sufficient to achieve this objective. Because many of these elections take place in very large firms, the ratio of voters to total firm employment is no larger here than for moderately sized elections. While we do not have the exact sample used by Ruback and Zimmerman (1984), we can attempt to replicate it based on their description of the sample selection scheme.^{17,18} Using their sample selection scheme we find that in more than 10 percent of the elections, less

¹³The 1977-1999 period data were obtained from Thomas Holmes' website (http://www.econ.umn.edu/~holmes/data/geo_spill/) and are used in Holmes (2006).

¹⁴For a limited number of years the NLRB data has information on the calendar date of the election and the calendar date the NLRB closed the case.

¹⁵The location of the election is not very useful for matching because the CRSP file only contains the location of company headquarters, which may differ from the location of any establishment undergoing a recognition election.

¹⁶The only additional information that could help us identify a match is the two digit SIC industry code of the establishment. However, the industry of an establishment may differ from the primary industry of the firm. This variable is more useful as a check for the validity of the matches.

¹⁷We contacted Professors Ruback and Zimmerman to request their dataset. As their paper was published more than 20 years ago they understandably could no longer provide it.

¹⁸Using the Ruback and Zimmerman procedure we ended up with almost twice as many elections as they had considered over the same time period. The only information that Ruback and Zimmerman had that we do not is the petition date. They excluded elections where the petition date was unavailable. We therefore infer that this exclusion restriction would have resulted in us dropping 50 percent of the elections in the sample.

than 1 percent of the firm's workforce voted. In our reproduction of their sample, the median percentage of the workforce voting in an election is 5 percent.¹⁹ By contrast, our main analysis limits the sample to elections where at *least* 5 percent of the total workforce voted.²⁰ The median election in our sample consists of 13 percent of the company's workforce voting (mean = 22 percent).²¹ Therefore, our sample selection scheme not only provides us with elections that are relatively salient for a given firm (or, at a minimum, excludes those elections which are clearly not salient), but also yields a substantially larger sample size compared to what we would have obtained using the Ruback and Zimmerman (1984) criterion. Our baseline sample is almost eight times larger than the Ruback and Zimmerman (1984) sample.

We present summary statistics of firm characteristics in Table 1. Columns (1) and (2) correspond to elections where at least 5 percent of the workforce voted (hereafter the " $\geq 5\%$ sample") for UV ("Union Victory") and UL ("Union Loss") firms respectively. Columns (3) and (4) correspond to elections where less than 5 percent of the workforce voted (hereafter the " $< 5\%$ sample") for UV and UL firms respectively. We report the market value of the firm using both the CRSP and the Compustat databases. Because there are a large number of missing observations in the Compustat database, especially before 1970, these measures differ. Companies in the Compustat database have larger market values on average, implying that small firms are underrepresented in the Compustat dataset.

Looking at the first row of Table 1, there are about twice as many elections in the $< 5\%$ sample than in the $\geq 5\%$ sample, and in both samples there are about twice as many firms where the union lost than where the union won. Not surprisingly, firms in the $\geq 5\%$ sample tend to be substantially smaller than firms in the $< 5\%$ sample. This inference can be made by comparing a variety of measures, including employment (4,530 vs. 73,223 employees) and market value (\$338 million vs. \$5.9 billion in 1998 dollars, using the more broadly available CRSP measure). However, the $\geq 5\%$ sample corresponds to bigger elections, with an average of 453 workers voting as compared to an average of 291 in the $< 5\%$ sample.

In addition to the mean and standard deviation, for variables derived from the Compustat database we report in braces the average percentile rank of that variable relative to all other firms in the Compustat database for the year and quarter of election. The average percentile rank is convenient for assessing how the

¹⁹Huth and MacDonald (1990) conduct an event-study of decertification elections. Their sample selection scheme involves all decertification elections involving at least 250 workers between June 1977 and May 1987. They also do not condition on there being a sufficiently high fraction of a firm's workers involved in the election. Our (inexact) reproduction of their sample has a median fraction of the workplace voting of 2 percent, with approximately 30 percent of elections in the sample involving less than 1 percent of the company's workforce.

²⁰Total employment in the year of the election is from the Compustat annual files.

²¹We do not use elections where employment information is missing.

firms in our sample compare to companies in the Compustat universe, and is advantageous as a statistic that is "robust" to outliers. From the percentile rankings it can be seen that firms in the $< 5\%$ sample tend to be around the 75th percentile in the size distribution of all Compustat companies, whereas firms in the $\geq 5\%$ sample are, on average, in the 35th percentile. In both samples, firms tend to be fairly representative with respect to profit margins, return on assets, Tobin's Q, and the dividend ratio. At the time of the election, UL and UV firms appear to be similar in most measures, including employment, market value, profit margin, profit per employee, Tobin's average Q, and industry composition.

Table 1 also shows the delisting rate for companies. We report the fraction of companies delisted in the two years before or after the election. UV firms are slightly more likely to delist than UL firms (10 versus 8 percent delisting rates respectively).²² While this difference is not large, we will consider several approaches to address this issue, as well as the presence of missing returns more generally. These approaches involve imputing missing returns, estimating all models excluding periods with missing returns, or limiting the sample to firms that have no missing returns in the event window. Simply excluding missing values has the disadvantage that some of the changes in cumulative returns over time may reflect firms that are entering or dropping out of the sample. Using a balanced panel has the advantage that we can be sure that any differences over time are not caused by compositional differences. However, a balanced panel does involve discarding a large number of elections and implies that inclusion into the sample may depend on the realization of the dependent variable. We will demonstrate that the results are not sensitive to the approach employed.

3.2 The Event-Study Method

Our objective is to assess the impact of union elections on the stock market value of firms. Ideally, we would like to compare the firm's stock returns to the returns the firm would have experienced in the absence of a union organizing event. The event-study method provides a framework for estimating this counterfactual return.

As is standard in financial economics literature, we define the abnormal return as the difference between a stock's actual return and the expected return given market conditions. For the company corresponding to

²²We define delisting as any company with a non-missing delisting return in the CRSP dataset.

union representation election i , in month t , the abnormal return is:

$$AR_{it} \equiv r_{it} - E[r_{it}|X_t],$$

where r_{it} is the actual return and $E[r_{it}|X_t]$ is the predicted return. For this study, r_{it} is the CRSP monthly holding-period return including distributions, which is constructed using prices that are adjusted for splits and distributions.²³

For convenience, we express time in terms of months relative to the event:

$$AR_{i\tau} \equiv r_{i\tau} - E[r_{i\tau}|X_\tau],$$

where $AR_{i\tau}$ is the abnormal return of the security corresponding to election i in the τ 'th month relative to the event.

Because returns of companies with unionization events may vary systematically before the elections, perhaps due to anticipation of the event, and because the market may not react instantaneously, we are interested in the cumulative abnormal return (CAR) in a window surrounding the election. The CAR corresponding to event i between months T_1 and T_2 relative to the event is:

$$CAR(T_1, T_2)_i \equiv \sum_{\tau=T_1}^{T_2} AR_{i\tau}.$$

The statistic of interest is the average (across N firms in the sample) cumulative abnormal return:

$$ACAR(T_1, T_2) \equiv \frac{1}{N} \sum_{i=1}^N CAR(T_1, T_2)_i.$$

We will present the average cumulative abnormal return for the set of union victory (UV) and union loss (UL) firms beginning two years prior to the election. Our decision to use such a long event window is in part the consequence of having information on the month that the NLRB closed the case, rather than the exact calendar date. By considering a very long pre-event window we can verify that any difference in the cumulative return of the UL and UV firms and any counterfactual (or “benchmark”) portfolio is not simply a

²³When stocks are delisted we use CRSP delisting returns. We replace missing returns with the predicted return ($E[r_{it}|X_t]$) in order to mitigate survivorship bias, though the results are not sensitive to how missing values are treated. Specifically, the results are not sensitive to simply ignoring missing values, nor to only selecting companies with no missing returns in the entire event-period.

continuation of differential pre-event trends. The long panel also allows us to examine returns in the months beyond the event, so as to capture the long-run expected costs to the firm without having to rely on the assumption that the stock price immediately and instantaneously adjusts to the presence of the union.

A critical decision in event-studies is how to model $E[r_{it}|X_t]$. In traditional short-run event-studies the counterfactual is often estimated from a market-model fit to historical data (as described, for example, in Campbell et al., 1997). In this approach, denoting R_{mt} as the return of a broad market index in month t , one uses historical data to estimate:

$$E[r_{it}|X_t] = a_i + b_i R_{mt} \quad (1)$$

This approach is theoretically attractive because the Capital Asset Pricing Model (CAPM) predicts that market beta is sufficient to describe cross-sectional expected returns. While this choice of benchmark is theoretically justified, a voluminous literature has discredited this idea (see Fama, 1998 for a review), leading to the use of additional explanatory factors for the expected return. For example, it is common practice to include the company's size and the book-to-market equity ratio (BE/ME) (Fama and French, 1993; Carhart, 1997) in these market models.

As pointed out in the literature though, there are a number of difficulties with estimating counterfactual returns using out-of-sample data in long-run event-studies. The approach requires that the estimated parameters remain time-invariant, an assumption that is known to not hold (Harvey, 1989). Additionally, estimation of the market-model parameters in the pre-event period must be done over an interval that is free of unusual pre-event returns, perhaps owing to the event. A solution to this problem is to estimate the market-model parameters using data from a long time (perhaps several years) prior to the event. But doing so exacerbates the first problem – that the market-model parameters may have changed – and leads to a new one: stocks that were not listed during the estimation window will be excluded from the analysis. As a result of these, and other concerns, the traditional methodologies developed for short-run studies are rarely used for long-run examinations.

A common approach for computing abnormal returns in long-run event-studies involves the use of reference or “benchmark” portfolios matched on a firm's characteristics (see Barber and John D. Lyons, 1997; Lyon et al., 1999; and Brav, 2000). The advantages of this approach are that the benchmark can be constructed in-sample and that it allows for shocks occurring by chance that affect firms with similar character-

istics. We employ this approach, matching every firm in our sample to a portfolio of firms in the same size-decile.²⁴ As a probe for robustness we have also used the CRSP equally-weighted NYSE/AMEX/NASDAQ index as a benchmark, comparing firms both in the same size decile and in the same one-digit SIC industry.²⁵

A second commonly used approach in long-run event-studies is the calendar time portfolio (CTP) approach developed by Jaffe (1974) and Mandelker (1974) and advocated by Fama (1998). For each calendar month we compute the return of an equally-weighted portfolio of companies that unionized in the last T months, where T is either 18 or 24 in our study. The return of this “unionization portfolio” is denoted R_{ut} , where u indicates that the portfolio consists of companies where workers voted for unionization and t denotes the calendar month. The unionization portfolio is rolling, because companies with new unionization events are added in any given month, while firms without a unionization event within the last T months are dropped. The Fama-French three factor model (Fama and French, 1993) is used to compute the abnormal return of this portfolio:

$$R_{ut} - R_{ft} = \alpha_u + b_u(R_{Mt} - R_{ft}) + s_uSMB_t + h_uHML_t + \varepsilon_{ut}, \quad (2)$$

where R_{ft} is the one-month treasury bill rate, R_{Mt} is the monthly return on a value-weight market portfolio of NYSE, AMEX, and NASDAQ stocks, SMB is the difference in the returns on portfolios of small and big stocks (below or above the NYSE median), and HML is the difference in the returns of portfolios of high- and low-BE/ME stocks.²⁶ In practice, Equation 2 is estimated by weighting the number of equities in the R_{Mt} portfolio at time t , as suggested by Fama (1998). Assuming that the broad-market return and the Fama-French factors adequately describe average returns, the parameter of interest, α_u , can be interpreted as the average abnormal return associated with holding this simulated portfolio.

The CTP methodology has been used in many long-run event-studies, for example Loughran and Ritter (1995), Brav and Gompers (1997), Mitchell and Stafford (2000), and Greenstone et al. (2006). This approach is thought by some, including Fama (1998) and Mitchell and Stafford (2000), to have better statistical properties than leading alternatives. For example, firms clustered in event-time can lead to over-stated

²⁴CRSP produces indices for such purposes. Specifically, every year CRSP allocates companies into one of ten size deciles, based on market-value. The value-weighted average return of securities in these deciles are then calculated on a monthly basis. CRSP also produces a cross-walk that allows one to link each security to the appropriate size decile.

²⁵We cannot match on the book-to-market equity ratio, as many studies do, because this variable is unavailable for a large number of companies in our sample, especially in the earlier periods.

²⁶The three factors, R_{Mt} , SMB , and HML , were taken from Kenneth French’s web page (http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html). The web page contains additional information on the construction of these series.

test statistics in the matched-portfolio approach described above.²⁷ Since the CTP methodology uses a time-series of portfolio returns, cross-correlations of firm abnormal returns are incorporated in the portfolio variance. Additionally, this approach allows for classical statistical inference because the distribution of the estimator is well-approximated by the normal distribution (Mitchell and Stafford, 2000). A disadvantage to this approach is that the market-model parameters of the portfolio are assumed constant. But, because the model is estimated over a long time-period (1961-1999) and because the firms in the portfolio are changing, that assumption is unrealistic.

A complication arising in both methodologies is how one defines the “event.” The appropriate event is the date on which most of the information on the probability of future unionization is incorporated. For much of the sample (1961-1976) we only observe the month that the NLRB closed the case. While we have a well-defined event, it is not the only relevant event and it may not be the most important one. Alternative, potentially important events are the petition and election dates. Using post-1977 data, where both the election and case closure calendar dates are available, we find that the median time between the election and NLRB case closure is ten days. In some cases, typically when one of the parties issues a challenge, this gap can be considerably longer. In 5 percent of the elections it took at least six months for the NLRB to close the case. While we do not have data on when the petition was submitted to the employer, it is known from Roomkin and Block (1981) that elections usually occur very soon after the petition. In their sample, 42 percent of elections occurred within one month of petition and 83 percent within two months. Therefore, we do not believe that using the month the NLRB closed the election presents serious problems for estimation if most of the new information is revealed at or after the petition date. To assess whether gradual diffusion of news led to abnormal returns prior to the closing date it is useful to examine a long pre-event window. We believe, however, that it will be difficult to empirically distinguish the market’s anticipation of unionization from an inadequate comparison portfolio.

The event-study method can inform us on how the equity value of firms responds to certification elections. We can also estimate event-study models for elections with varying degrees of union support in order to explore heterogeneity in the effect size. A more complete investigation of heterogeneity in the impact of certification elections on stock market performance involves estimating the post-event cumulative abnormal return for every election and relating these to the vote share in a flexible way. We conduct this analysis to

²⁷Though, it should be noted, we will allow for such correlations in computing standard errors by clustering on election and calendar month, using the formula from Cameron et al. (2006).

examine the heterogeneity in the stock market reaction to election outcomes and to determine whether there is a discontinuous relationship between cumulative abnormal returns and the vote share at the 50 percent threshold.

4 Empirical Results

4.1 Event-Study Estimates

In Figure 2 we plot the average cumulative return of union victory firms against the average cumulative return of the size-matched reference portfolios over the same time period.²⁸ The figure reveals that both UV firms and the corresponding reference portfolios have almost identical trends in returns prior to the union victory. However, near the time of the election there is a pronounced downward break in the returns of UV firms relative to the benchmark, persisting for approximately a year and a half. The average cumulative abnormal return implied by this divergence is approximately -10 percent.

In order to assess the magnitudes and statistical significance of the effect implied by Figure 2, in Figure 3, Panel A we plot $ACAR(-24, \tau)$, for $\tau = -24$ through $\tau = 24$, with 95 percent point-wise confidence intervals. In Panel B we plot $ACAR(0, \tau)$, for $\tau = 0$ through $\tau = 24$. This second panel is relevant for assessing the overall effect size and for determining statistical significance. The figures show that the downward shift in abnormal returns emerging soon after NLRB case closure is statistically significant. We can reject the null hypothesis that the average abnormal returns are equal to zero five months after the event at a 5 percent level of significance. We interpret Figures 2 and 3 as providing evidence that union election wins correspond to large negative abnormal returns.

Figure 4 contains the plot of the average cumulative return for union loss firms against the average cumulative return of the size-matched reference portfolios. As with the UV firms, the reference portfolios closely track the progression of UL firms prior to the election, but unlike UV firms, the returns of UL firms do not diverge from the benchmark after NLRB case closure. If anything, there is a moderate increase in the cumulative return of UL firms relative to the benchmark, though in Figure 5, which presents the difference in these series with confidence bands, we see this increase is not statistically significant at conventional levels.

We have conducted a variety of analyses to determine whether the patterns seen in Figure 2 and Figure

²⁸For convenience, we will often refer to the event month as the “election month,” though it should be understood that we actually only know when the NLRB closed the case.

4 are robust. These analyses include: not imputing missing returns (Appendix Figures 1 and 7); using a balanced panel (Appendix Figures 2 and 8); excluding elections where cumulative abnormal returns following case closure are less than or equal to the 5th percentile or greater than or equal to the 95th percentile of all post-event cumulative returns (Appendix Figures 3 and 9); using a four year pre-event window (Appendix Figures 4 and 10); using an industry \times size matched-reference portfolio (Appendix Figures 5 and 11); and using the CRSP equally-weighted market index as the reference portfolio (Appendix Figures 6 and 12). In all cases the overall pattern of cumulative returns look very similar to those seen in Figures 2 and 4.²⁹

Table 2, Panel A presents average cumulative abnormal returns following union victories. The first column corresponds to the use of the size-matched benchmark. Column (2) corresponds to the industry \times size-matched benchmark. Column (3) corresponds to the CRSP equally-weighted NYSE/AMEX/NASDAQ index benchmark. In the first row of Panel A we report ACAR(0,24) for each of the three benchmarks. The estimated post-election average cumulative abnormal returns range from -9 to -10 percent and are significant at the 1 percent level. To gauge magnitudes, we calculate that a 10 percent negative return corresponds to approximately \$20 million in lost market value (in 1998 dollars), or \$40,522 per worker eligible to vote. This appears to be a plausible value. Assuming that the entire profit effect is coming about from an increase in wages, the annual income of workers prior to unionization is \$25,000 (in 1998 dollars), and a 6 percent discount rate, the magnitudes are equivalent to a 10 percent union wage premium.^{30,31} Of course, we are unable to say whether the loss in equity value reflects increases in compensation, benefits, or inefficiencies.

In the second row of Table 2 we report ACAR(-24,-4), the average cumulative abnormal return prior to case closure, excluding the three months immediately preceding the event. ACAR(-24,-4) is statistically indistinguishable from zero in all three specifications. The lack of significant abnormal returns prior to the election indicates that the market did not anticipate these events, on average, and also suggests that all three benchmarks do a reasonable job of predicting average returns of the portfolio of UV firms. In the third row we compute ACAR(0,24) after adjusting abnormal returns in the post-event period for the equity-specific trends in abnormal returns in the months before the election. Specifically, before computing ACAR(0,24) we subtract the average abnormal return for the equity in months -24 through -7 relative to case closure from the post-event abnormal return. The point estimates are very close to the unadjusted version. But, not

²⁹A possible exception is Appendix Figure 10, which shows that UL firms experienced a period of positive abnormal returns three years before the election.

³⁰In 1980 (the mid-point of our sample frame) the average non union wage was \$12.43 in 1998 dollars (Hirsch and Macpherson, 2008)(Hirsch and MacPherson 2008), translating to approximately \$25,000 in annual income

³¹This premium is a lower-bound because approximately 25% of union election victories do not lead to contracts (Cooke, 1985).

surprisingly, they are somewhat less precise. Table 2, Panel B reports the same set of estimates for union loss firms. Consistent with what we observe in Figure 5, the cumulative abnormal returns are close to zero and statistically insignificant.

One possible concern is that elections are endogenous to the performance of firms. However, we find little evidence that this is the case. The firms in our sample track their benchmarks quite closely prior to the election, so it does not appear to be the case that the election is a result of the firms under- or over-performing the benchmark. There is also no indication that the firm's performance in the two years prior to the election the union fares in the election. This can be seen in a number of ways. For example, looking at Figure 2, winners and losers are not trending differentially prior to the election. To test this hypothesis more directly we have regressed the union vote share in the election of the cumulative abnormal return from -24 to -4 and found no significant relationship between the two variables.³² If workers are deciding on the performance of the firm, they are basing their decision on forecasts of future performance rather than past performance. While we cannot rule out this possibility, it is not obvious how workers could forecast future share prices of the firm, and why it would be optimal for them to ignore past performance. Moreover, it is not clear why it would be optimal to unionize when the firm is projected to perform poorly.

Our sample selection scheme was partly predicated on choosing elections where a sizable fraction of the firm's workforce was voting: in practice we used a 5 percent cutoff. As a falsification exercise we examine elections where a small fraction of the firm's total workforce voted. The idea is that we should not see effects in firms where only a very small share of the employees voted. In Table 3 we examine whether cumulative abnormal returns following an election become more pronounced when a larger share of the firm's workforce is participating. Specifically, using the full sample of elections we relate $ACAR(0,24)_i$, where i denotes an election, to the share of the firm's total workforce in the bargaining unit. As seen in Column (1) of Panel A, when the union wins the election and the fraction of the firm's workforce in the bargaining unit is essentially zero, the firm experiences a small and positive abnormal return. As we would expect, as the share of the firm involved in the election increases, the resulting effect on the abnormal return becomes more pronounced. Each percentage point increase in the share of the firm's employees voting in the election is associated with a third of a percentage point decline in the post-event cumulative abnormal return, a relationship which is statistically significant at the 5 percent level. In Column (2) we use adjusted- $ACAR(0,24)$ and the relation-

³²Specifically, we estimate a coefficient of -0.006 with a standard error of 0.09. This estimate is not sensitive to the pre-event window over which the CAR is calculated.

ship continues to hold. Panel B presents these estimates for the union loss sample. The negative relationship in the post-event cumulative abnormal return and the share of the workforce voting is not present. In fact, there is a positive relationship, which is what we would expect if union losses resulted in positive abnormal returns.

Finally, Table 4 presents the estimates from the calendar time event-study methodology. The portfolio of stocks consisting of all firms with a unionization win in the previous 24 months has a precisely estimated alpha of -0.005 (t-ratio=-3.6). In the second row we consider a hypothetical portfolio of firms that are purchased two years prior to case closure and are sold four months prior to case closure (-24 to -4 months relative to closure). This portfolio corresponds to a small and statistically insignificant alpha. Likewise, we do not observe an economically or statistically significant alpha for portfolios of firms recently experiencing union losses (Table 4, Panel B), nor for portfolios consisting of firms with small elections relative to the size of the company (Table 4, Panels C and D). These results give us confidence in our finding: negative alphas are only present when the union wins, and even then, only when the electorate is a large fraction of the firm's total workforce. Moreover, the results are robust to the use of two standard methodologies for long-run event studies.

4.2 Discussion of the Results and Additional Analyses

Speed of Adjustment

Perhaps a surprising feature of Figure 3 is that, while the efficient market hypothesis would predict the entire unionization effect should be fully realized by the time of the election, we instead see an effect which grows over a longer period, with an abnormal return beginning around the time of election and persisting for approximately 15 months. Ours is not the first study showing that markets under-react to seemingly important events. Systematic under-reactions have been reported in response to IPOs and SEOs (Loughran and Ritter, 1995), mergers (Asquith, 1983; Mitchell and Stafford, 2000), stock splits (Ikenberry et al., 1996), share repurchases (Mitchell and Stafford, 2000), exchange listings (Dharan and Ikenberry, 1995), dividend initiations (Michaely et al., 1995), spin-offs (Cusatis et al., 1993), earnings announcements (Ball and Brown, 1968), and predictable changes in demographics (Dellavigna and Pollet, forthcoming). While Fama (1998) questions the robustness of some of these findings, even he acknowledges that the short-term continuation of returns documented by Jegadeesh and Titman (1993) is “an open puzzle,” and that the slow post-earnings

announcement drift “has survived robustness checks, including extension to more recent data.”³³

Hong et al. (2000) show that stock price momentum is largely confined to smaller stocks, so perhaps it is not surprising that there is a gradual-adjustment following an union election win as our sample selection ensures that the firms in our sample are relatively small. Hong et al. (2000) theorize that stock prices exhibit momentum because information, especially negative information, diffuses gradually to investors. To examine whether the slow diffusion of information explains the speed of adjustment of stock prices in our study we follow their example and compare firms with and without analyst coverage.

According to I/B/E/S International analyst data, only 50 percent of the firms in our sample had analyst coverage at the time of the election, meaning that these elections may not have been widely publicized or followed.³⁴ In Figure 6 we compare average cumulative abnormal returns for companies that did and did not have analyst coverage at the time of the election. Companies with analyst coverage appear to have experienced negative abnormal returns earlier than those without analyst coverage. But even these experienced a relatively slow-reaction to the event on average, suggesting that the lack of analyst coverage is not the complete story.³⁵

While we do not have an explanation for why the adjustment happens slowly in this case, it is true that by definition these large unionization events are typically singular for a given firm. Investors may simply not have known how to process this new information, and instead reacted primarily to news on fundamentals. In Section 4.3 we will see that union wins are associated with a trend break in the growth rate of these companies, as measured by assets, shareholder equity, and profits – an effect that may not have been immediately obvious to investors, but which may have become apparent over time.

Evolution of the unionization effect over time

Next we turn to the evolution of the effect over time. The DiNardo and Lee (2004) sample includes elections beginning in 1984. It is possible that unions did not affect firm performance in this latter period, while in earlier years the effects may have been more pronounced. In Figure 7 we compare the average cumulative abnormal return of UV firms for elections occurring in the 1961-1983 period to those occurring

³³Quoted in Fama (1998).

³⁴The 50 percent figure is derived from I/B/E/S International analyst data for years 1976-1999.

³⁵We are aware that companies not appearing in I/B/E/S may still have analyst coverage. This kind of misclassification tends to reduce the measured difference in excess returns between these two groups of firms, if in fact there are actual differences. It is unlikely that this measurement problem will affect the relatively slow speed of adjustment for companies covered by analysts, as these are presumably measured correctly, meaning that our basic conclusion—that analyst-covered companies exhibit a relatively slow speed of adjustment—still holds.

in the 1984-1999 period. The figure indicates that the average effect of a union certification win on firm performance has remained fairly stable over time. Therefore, we do not believe that the lack of an estimated unionization effect in DiNardo and Lee (2004) is due to their sample frame.³⁶

4.3 Compustat Analysis

The results presented up to this point suggest that union victories are associated with negative abnormal returns. We complement this analysis with an additional investigation of accounting variables. Using quarterly data from Compustat, we examine whether shareholder equity, assets, total liabilities/total assets (a measure of leverage), plant, property and equipment, sales, the dividend ratio, Tobin's average Q, profit margins, and returns on assets are affected by the outcome of representation elections. We compute the average value of these variables (logged when appropriate) over the twelve quarters before and after the event date, comparing UV and UL firms.³⁷ As before, we assess whether these series were trending differentially prior to the event and whether their trend breaks around the time of the event. While we regard this analysis as informative, we are cognizant that the coarseness of these data means that it is more challenging to assess the specifications, as compared to the event-studies, because the pre-event window are shorter.

Unfortunately, the early part of the sample period is unusable in the Compustat analysis because many of these variables were not reported until the late 1960s, and not universally until the early 1970s. Moreover, the fraction of missing observations is substantially higher in the Compustat dataset than in the CRSP dataset. As a result, for this analysis we will only consider elections over the 1973-1999 period. In order to mitigate composition bias to due to unbalanced panels we de-mean the variables, but do not drop elections with missing values.

In the nine panels of Figure 8 we plot averages of the de-meaned variables over event-time, in each case comparing elections where the union won to those where the union lost. The figures show that the time pattern of variables proxying for "size" are consistent with the pattern in equity value. UV firms display a downward break in trend in total assets (Panel A), shareholder equity (Panel B), and sales (Panel

³⁶We have also compared the effects for states with and without right-to-work laws. Conditional on a union winning its election, the stock-market effects of unionization tend to be more pronounced in states with right-to-work laws than those without. This finding does not mean that states with right-to-work laws are more favorable to unions – these firms differ in other dimensions, and it is also probable these laws (or the business climate more generally) affect the likelihood that a union organizes, as well as the likelihood that a union wins an election. The result does call into question though, whether right-to-work laws fundamentally weaken unions because of a potential free-riding problem. This finding lends qualified support to the conclusion of Farber (1984) and Moore and Newman (1985) that right-to-work laws are primarily symbolic, reflecting a taste against union representation rather than having any real effect.

³⁷All variables in 1998 dollars, when appropriate.

C) near or just before certification. The reduction in asset growth is, in large part, due to reduced growth in plant, property, and equipment (Panel D).³⁸ The smaller sample sizes mean these series are not as well-behaved as those for equity values, though they have a similar pattern. We see little effect of union wins on the measure of leverage, defined as long-term debt divided by total assets (Panel E). This last finding can be viewed as circumstantial evidence that companies are not using leverage strategically to influence bargaining negotiations, at least in this sample.³⁹

The marked reduction in the growth rate of assets is notable because if unionization increases the price of labor, there should be substitution from labor to capital (though, as seen in Panel F, Tobin's average Q appears stable). The fact that assets are actually declining implies that the "scale" effect from reduced reinvestment dominates the possible substitution effect. The time pattern of these variables also sheds light on the seemingly slow reaction of investors to unionization events that we see in Figure 3. The pattern of abnormal returns mirrors the time-pattern we observe in shareholder equity, assets, sales, and pre-tax income. The evidence is consistent with the stock market pricing the effect of unionization only after changes in these variables become known.

While the reduced relative size of the UV firms is associated with lower pretax income (Panel G), variables that proxy for operating performance, for example return on assets and profit margins, appear stable.⁴⁰ At first blush, the finding that companies that undergo unionization experience lower growth rates but stable returns on assets and profit margins may seem puzzling. But if firms only select projects that are sufficiently profitable and unionization reduces the number of these high net present value (NPV) projects, then it is possible for the company's growth rate to decline in spite of experiencing no change in its operating performance.

In Table 5 we present difference-in-difference estimates for the effect of a union victory relative to a union loss on each of the six aforementioned variables. The sample consists of election \times event-time observations. We regress each of the (non-demeaned) variables on election fixed-effects, an indicator for whether the NLRB closed the election on or after the given quarter ("post"), and the interaction of "post" with an indicator for whether the union won the election ("post \times union win"). The point estimates suggest

³⁸We have also examined the corresponding figures using a balanced panel. The overall patterns are the same as when using the unbalanced sample, but because we lose so many elections the confidence intervals are substantially wider.

³⁹Bronars and Deere (1991) show there is a positive association between financial leverage and unionization in the cross-section. Matsa (2006) provides evidence that firm measures of leverage were affected by state-level changes in right-to-work laws.

⁴⁰The profit margin in UV firms appears to decline a bit relative to UL firms, but not until about seven quarters after the election (Panel I).

that assets, shareholder equity, and sales fall by approximately 10 percent in UV firms after the election, relative to UL firms. Pre-tax profits of UV firms are approximately 17 percent lower in the post-election period relative to the pre-election period (relative to UL firms). These statistically significant estimates are consistent with the 10 to 14 percent negative abnormal returns we observe in equities.

4.4 Heterogeneous Impacts of Unionization

In view of the findings summarized in the preceding discussion, a natural question comes to mind: how can these large effects be consistent with the substantially smaller ones found in DiNardo and Lee (2004)? This section aims at providing a partial answer to this question.

While DiNardo and Lee (2004) identify the “unionization effect” by focusing on an implicit comparison of winning and losing establishments among close elections, we can learn how unions affect firms by examining the heterogeneity in the effects of unionization at all points in the vote share distribution. This analysis is possible because of the long-panel structure we have at our disposal.

We begin by relating the security-level cumulative abnormal return in the two years following the election to the union vote share. Specifically, we are interested in the shape of $E[\text{CAR}(0,24)_i|v_i]$, where v_i denotes the union vote share in election i . We graphically plot this function by: (1) averaging $\text{CAR}(0,24)_i$ over 20 equally-spaced vote share bins⁴¹ and (2) plotting the predicted values from the model $E[\text{CAR}(0,24)_i|v_i]=p(v_i)+\beta 1(v_i > 0.5)$, where $p(\cdot)$ denotes a sixth-order polynomial and $1(v_i > 0.5)$ is an indicator function for whether the union vote share in a given election exceeded 50 percent. Figure 10 presents estimates of $E[\text{CAR}(0,24)_i|v_i]$ using both of these approaches. (For reference, Figure 9 shows the histogram of the union vote share variable.)

Figure 10 shows clear evidence that the effect of a certification election is heterogeneous, and that it depends on the union vote share. As in the Dinardo and Lee study, there is no discernible discontinuity in the $E[\text{CAR}(0,24)_i|v_i]$ at the 50 percent union vote share threshold. In fact, the estimated discontinuity is somewhat perverse: firms with close union wins experience elevated post-election cumulative returns vis-a-vis firms with close union losses. On the other hand, union victories with higher union vote shares correspond to negative excess returns, and the negative impact of a union election win appears to become markedly more pronounced when the union has a higher vote share. A greater than 60 percent union vote share is associated with negative cumulative abnormal returns of 20 to 30 percent.

⁴¹See DiNardo and Lee (2004) for a description of construction of these 20 equally spaced bins.

Firms with union losses also exhibit a downward sloping relationship between abnormal returns and vote share. Much of the decline appears to occur at the largest vote shares, but there is also greater variability in the predicted cumulative abnormal returns due to small sample sizes. Close union losses are associated with marginally-significant negative abnormal returns, though as we will show, these declines can be explained by a small amount of pre-election trending in the abnormal returns.

We now turn to several robustness checks. In Figure 11 we overlay the predicted CAR in months 0 through 24 (shown in Figure 10) with the predicted CAR computed over event-months -24 to -4. The figure shows that the gradient in CAR by vote share, seen for months 0 to 24, is not present for months -24 through -4. This plot reassures us that the negative CAR observed for higher union vote shares is not a continuation of a pre-event trend.

In order to address the issue of pre-event trends more completely, we consider an additional analysis where we adjust abnormal returns in the post-event period for possible pre-event trends. Specifically, we calculate the cumulative abnormal return in the post-event period deviated from the average abnormal return in the pre-event period (from months -24 to -7 relative to case closure):

$$\text{adjusted-AR}_{i\tau} \equiv AR_{i\tau} - \frac{1}{18} \sum_{\tau=-24}^{-7} AR_{i\tau}$$

We then calculate:

$$\text{adjusted-CAR}(0,24)_i \equiv \sum_{\tau=0}^{24} \text{adjusted-AR}_{i\tau}.$$

Figure 12 plots the predicted adjusted-CAR with 95 percent pointwise confidence intervals.⁴² The figure shows virtually the same pattern of heterogeneity seen in the earlier figures, though with wider confidence intervals. The main difference between the pattern in this figure and Figure 10 is that there is weaker evidence here of a negative CAR among firms with close union losses.

In Table 6 we conduct formal statistical inference. Using the same sample of 1,436 elections used to construct Figure 10, in Column (1) we regress $\text{CAR}(0,24)$ on a dummy for whether the union won the election. Consistent with earlier analyses, we find that union wins are associated with cumulative abnormal returns that are 12.1 percentage points lower than firms with union losses (t-ratio = -3.5). In Column (2)

⁴²As before, the abnormal returns from the pre-event period are calculated using an estimation window that ends 29 months prior to the closing month. The abnormal returns from the post-event period are calculated using an estimation window that ends 5 months prior to the closing month.

we add the union vote share as a covariate.⁴³ The introduction of this variable alone is enough to change the sign on the coefficient of the union win dummy, resulting in a union effect of 0.048 (t-ratio = 0.89). Adding higher-order polynomial terms in the vote share (Column 3) only makes the estimated union win coefficient more positive; the “regression discontinuity” estimate of a union win is 8 percentage points, but is statistically indistinguishable from 0. In Column (4) we examine whether the negative gradient between CAR and the vote share differs among elections where the union won and lost. Specifically, we regress CAR(0,24) on a union win indicator, the vote share, and the vote share interacted with the win indicator. The interaction term is statistically insignificant in all specifications.

In Columns (5)-(8) we estimate the same set of models using CAR(-24,-4) as the dependent variable. None of the patterns observed when using CAR(0,24) as the dependent variable are evident here. In Columns (9)-(12) we re-estimate these models using adjusted-CAR(0,24) as the dependent variable. The point estimates in this set of specifications are very close to the ones obtained using CAR(0,24), but are less precise, with standard errors approximately 50 percent larger than those in Columns (1)-(4).

5 Interpretation and Policy Implications

In this section, we investigate what our empirical results could imply about the potential effect of a policy that makes it easier for workers to unionize. An example of such a policy shift can be seen in the Employee Free Choice Act, recently proposed legislation that is meant to amend the National Labor Relations Act. Specifically, one of the provisions of the legislation would allow employees to authorize a union via “card check”, a showing that the majority of the workers signed cards to authorize a union, without having to win certification via a secret-ballot election process. It is widely believed that the legislation, supported by the AFL-CIO, would make it much easier for workers to unionize, if it were to become law.

In essence, we view such a policy change as a *ceteris paribus* marginal increase in the probability of unionization. One way to conceive of such an exogenous increase, would be to consider the thought experiment of lowering the necessary vote share threshold for certification. After all, the card check process is not unlike the petitioning that constitutes the first step in the NLRB election process.

As a thought experiment, consider lowering the threshold from 50 percent to say, 45 percent. One conjecture is that such a policy change would only effect those firms with vote shares between 45 and 50

⁴³Vote share is grouped into one of 20 equally spaced bins, ranging from 0 to 1. We transform this variable in order to avoid the “integer” problem described in DiNardo and Lee (2004).

percent, and that the effect could be approximated by the RD estimate. The shortcoming of this conjecture is that it assumes that unions, firms, and workers do not respond to the increased ease of unionization. As we noted in the introduction, Friedman (1950) suggested that unions might be forced to moderate promises to raise wages when seeking the support of their workers. In a representation election, this might mean moderating wage expectations to increase their chance of winning. With these forces at work, an exogenous increase in the probability of a union victory could very well lead unions to be more aggressive, resulting in increased negative impacts on profitability – not just for those firms near the 50 percent threshold, but also for those where the union won by a wider margin. Exogenously easing the unionization process might also affect the outcome for firms that eventually do not unionize, through union threat.

Thus, in order to make quantitative predictions regarding the impacts of making unionization easier – predictions that both use the magnitudes we estimate, and allow for behavioral responses to a change in policy – it is necessary to adopt assumptions about the behavior of unions and firms and how profitability is affected by changes in the probability of unionization. We consider a “median voter”-type model of endogenous union determination. The basic idea of the model is that in anticipation of the representation election, the firm and the union each propose an outcome (e.g. a wage level), and voters, recognizing that wages can be both too low or “too high” (if it poses too large a risk of job loss), vote on the two choices in the election. Both the union and the firm face similar trade-offs: the union (firm) would benefit from higher (lower) wages, but proposing those wages loses votes among those workers who have more moderate preferences.

We present a parsimonious parameterization for the model, and then calibrate it by choosing parameters such that the model produces both an equilibrium vote share distribution and event-study estimates that most closely match that which we observe in the data (shown in Figure 12).

This calibrated model yields a distribution of voter preferences, and also allows us to simulate the effects of lowering the vote share threshold, a policy which exogenously increases the probability of unionization. We also assess the model’s predictions for the impact on equity value of two sub-populations, a marginal group (the firms that are not currently unionized, but would be in the new regime), and two inframarginal groups (firms that are either already unionized or not unionized, and whose status does not shift after the policy change).

5.1 Endogenous Voting Model

There are surely an unlimited number of distinct ways to model the interaction between unions, workers, and firms in an election context. Arguably, an obvious starting point is to adopt a “textbook” model of electoral competition.⁴⁴ Indeed, median voter-type models have previously been considered in the theoretical literature on unions (see Atherton, 1973; Farber, 1978; and Booth, 1995).⁴⁵

We assume there are three optimizing entities involved in a representation election, the workers, the union, and the management.⁴⁶

Workers: Each worker is assumed to maximize their own individual utility, and faces the decision to either vote for or against union recognition. In doing so, each forward-looking worker compares the anticipated outcome if the union wins to the expected outcome if the union loses. For example, the main issue could be wages, where the anticipated wage level is higher if the union prevails in the election than if it fails. Workers will not always vote for higher wages, because it may also carry a higher risk of job loss as the firm must respond to those higher wages. So for each worker, there is an “ideal wage” or a “bliss point”.

It is most natural to discuss workers’ (and unions) preferences over wages, benefits and other working conditions. But as long as improved (inferior) compensation and conditions lead to lower (higher) profits for the firm, we can equivalently consider workers’ and unions preferences over profit levels, by applying an appropriate monotonic transformation from wages, for example, to profits. In the discussion below, we use this equivalent formulation, focusing our attention on “profit levels” (strictly speaking, the change in stock market value of the firm).

Thus, the actions of the workers are summarized by the probability of the firm winning the election

$$P(\pi_M, \pi_U)$$

where π_M is the resulting anticipated profit level if the firm wins, and π_U is the anticipated level if the union wins. $\frac{\partial P}{\partial \pi_M}$ and $\frac{\partial P}{\partial \pi_U}$ are both negative: as the outcome under a firm victory becomes more “extreme” and more profitable to the firm, fewer workers find that outcome attractive, lowering the chance of an electoral victory

⁴⁴See Persson and Tabellini (2000) for a guide to models of this sort.

⁴⁵Interestingly, this model has many parallels to the model of final offer arbitration developed in Farber (1980). The two bargaining parties face the same trade-offs as the union and firm do here, and the role of arbitrator is played by the median voter in this context.

⁴⁶The setup is similar to Ashenfelter and Johnson (1969) who also consider management, workers, and unions as separate maximizing entities.

for the firm. Conversely, if the anticipated profit level is more moderate, the “middle” of the electorate gravitates towards voting for the firm. The same is true for the union: the firm has a lower chance of winning if the outcome under a union victory π_U (which will always be less than π_M) is higher (and hence more moderate).

Note that we assume a probabilistic voting model (e.g. workers, firm, and union cannot perfectly predict the outcome) as is common in many electoral competition models. It will be clear that without some uncertainty, there can be no equilibrium where $\pi_M \neq \pi_U$. Thus, introducing some uncertainty as to the exact location of the median voter expands the range of possible equilibria.

Management: The firm influences the anticipated result of a firm electoral victory. Essentially, they propose a profit level π_M in order to maximize expected profits

$$\pi_M \cdot P(\pi_M, \pi_U) + \pi_U \cdot (1 - P(\pi_M, \pi_U))$$

taking the union’s proposal as given. The management faces a clear trade-off: higher profits are desired, but proposing an outcome that leads to higher profits raises the chance that the workers will vote to unionize, which would lead to lower profits.

In this sense the model captures the possibility of “union threat”, where the presence of unions can compel firms to offer above-market wages, even if the workers ultimately do not unionize.

Union: The union faces a similar problem with similar trade-offs. It controls anticipated outcome π_U under a union victory. Essentially, they make a proposal π_U to maximize the objective function

$$\bar{U} \cdot P(\pi_M, \pi_U) + U(\pi_U) (1 - P(\pi_M, \pi_U))$$

taking π_M as given. \bar{U} is the level of utility the union obtains if it loses the election, and $U(\pi_U)$, which is decreasing in π_U , is obtained if the union prevails. We assume that for all the feasible π_U , $\bar{U} \leq U(\pi_U)$, so that the union would never prefer to lose the election. Again, the union benefits from a lower-profit outcome *if it prevails in the election*. But it must also take into account that the further away their proposal is from the median worker, the more likely the less desirable outcome π_M will occur.

Equilibrium: We consider the Nash Equilibrium, which is characterized by the first order conditions

$$\begin{aligned} P(\pi_M, \pi_U) + \frac{\partial P(\pi_M, \pi_U)}{\partial \pi_M} (\pi_M - \pi_U) &= 0 \\ \frac{\partial U(\pi_U)}{\partial \pi_U} (1 - P(\pi_M, \pi_U)) + \frac{\partial P(\pi_M, \pi_U)}{\partial \pi_U} (\bar{U} - U(\pi_U)) &= 0 \end{aligned} \quad (3)$$

The solution to this system yields equilibrium proposals for π_U and π_M as well as the equilibrium probability $P(\pi_M, \pi_U)$.

Finally, we introduce two elements of heterogeneity to make it possible for the model to generate a relationship between the vote share and the observed profit level. First, we allow for heterogeneity across workplaces in the preferences of the workers (i.e. the median voter): heterogeneity in $P(\pi_M, \pi_U)$. Second, we allow for heterogeneity in preferences among workers *within* each workplace. It is possible to include this kind of heterogeneity without affecting the specification of $P(\pi_M, \pi_U)$ and hence the equilibrium π_M and π_U , but without some heterogeneity, realized vote shares could only equal 1 or 0.⁴⁷

5.2 Parameterization and Estimation

Our policy extrapolation exercise requires us to parameterize the model. We choose the following functional forms.

1. In bargaining over wages, profits are bounded. We let $\bar{\pi}$ be the maximum feasible profits, given the constraints of the market. For example, $\bar{\pi}$ could be the profit level if the post-election wage equaled the competitive market wage. If firms are price takers in the labor market, then any wage below that level would mean that they could not hire any workers and would be forced to shut down.
2. We let $U(\pi_U) = -\left(\frac{\pi_U}{c} - 1\right)^2$, which is representative of the entire class of concave quadratic functions in π_U .⁴⁸ c is the union's "ideal" profit level. We also set $\bar{U} = U(\bar{\pi})$ so that the union gains exactly nothing if it wins the election but achieves a wage level no different than the market competitive wage.
3. Voters ideal profit levels ("bliss points") are uniformly distributed over the interval $[\mu - \varepsilon - \sigma, \mu - \varepsilon]$, where μ varies across workplaces and σ quantifies the degree of heterogeneity of voter preferences

⁴⁷If all voters had the same ideal profit level as the median, then either all workers will vote for or against the union.

⁴⁸A quadratic function has three parameters, but the expected utility is invariant to affine transformations, so that it is innocuous to rescale and shift the function so that the peak of the function equals zero, and that the function equals -1 when $\pi_U = 0$. This is therefore a one-parameter function.

within the workplace. ε is a stochastic component, uniformly distributed on $[0, \lambda_\varepsilon]$, reflecting the uncertainty that both union and firm face regarding the exact location of the workers. If individual workers' utility over π are symmetric around their bliss point, this implies that the vote share for the union will be $VS = \frac{1}{\sigma} \left\{ \frac{\pi_M + \pi_U}{2} - (\mu - \varepsilon - \sigma) \right\}$, and that $P(\pi_M, \pi_U) = \Pr[VS < \frac{1}{2}] = \frac{1}{\lambda_\varepsilon} \left[\mu - \frac{\sigma}{2} - \left(\frac{\pi_M + \pi_U}{2} \right) \right]$.⁴⁹ This specification satisfies the above assumption that as the firm or union raises its proposal, the probability of a firm victory declines.

4. μ is distributed across workplaces, such that $-\mu$ follows an exponential, $F(x; \lambda_\mu, \bar{\mu}) = 1 - \exp(-\lambda_\mu \cdot (x - \bar{\mu}))$ for $x - \bar{\mu} \geq 0$, and 0 otherwise. The distribution of μ thus has a long left tail, and a maximum at $\bar{\mu}$.

To summarize, the model contains 6 parameters in total, $\bar{\mu}, \lambda_\mu, \lambda_\varepsilon, \sigma, \bar{\pi}, c$. $\bar{\mu}, \lambda_\mu$ characterize how worker preferences are approximately distributed across workplaces, λ_ε quantifies the degree of uncertainty of the precise location of the voters' preferences, and σ quantifies heterogeneity in workers' preferences within a firm. $\bar{\pi}$ represents the limit on how low the firms' wages can be, and c is the union's "ideal" profit level.

These six parameters are sufficient for generating a joint distribution of π^{obs} (an event-study estimate of the impact of the union on the firm) and the vote share in favor of the union, the two variables that we observe in the data. Specifically, a μ is drawn from the distribution given by the parameters $\bar{\mu}, \lambda_\mu$. Conditional on this value of μ , and the remaining 4 parameters $(\lambda_\varepsilon, \sigma, \bar{\pi}, c)$, the firm and the union make optimal proposals according to the marginal conditions in 3. Subsequent to these optimal choices π_M^* and π_U^* , an ε is drawn and the VS is determined as above, and the observed profit level is given by $\pi^{obs} = \pi_U^* \cdot 1[VS > .5] + \pi_M^* \cdot 1[VS < .5]$.

At the same time, the model has a minimal number of parameters. There is one parameter for the union's objective function (c), one for the firm ($\bar{\pi}$), and two parameters for the distribution of worker preferences across firms ($\bar{\mu}, \lambda_\mu$). Without allowing for λ_ε , there would be no uncertainty in the precise location of voters' preferences, which would imply that the firm's and union's proposals could never be different in equilibrium. Finally, without σ , a vote share would never be anything except 0 or 1.

⁴⁹Additionally, VS and $P(\pi_M, \pi_U)$ must be between 0 and 1.

To calibrate this model, we choose parameters that most closely generate 1) the pattern of event-study estimates in Figure 12, and 2) the distribution of vote shares. Specifically, we minimize the quadratic form

$$f(\theta)' \hat{V}^{-1} f(\theta), \text{ where } f(\theta) = \begin{pmatrix} \hat{\alpha}_1 - E[\pi^{obs} | VS > .5] \\ \hat{\alpha}_2 - E[\pi^{obs} | VS < .5] \\ \hat{\alpha}_3 - \lim_{\Delta \rightarrow 0^+} E[\pi^{obs} | VS = .5 + \Delta] \\ \hat{\alpha}_4 - \lim_{\Delta \rightarrow 0^+} E[\pi^{obs} | VS = .5 - \Delta] \\ \hat{\alpha}_5 - E[VS] \\ \hat{\alpha}_6 - E[VS^2] \end{pmatrix}$$

and θ is the vector of parameters from the model, the expectations are the moments predicted by the model given θ , and the $\hat{\alpha}$ s are the corresponding observed moments. $\hat{\alpha}_1$ is the event-study estimate for all union victories, $\hat{\alpha}_2$ is the event-study estimate for all union losses, $\hat{\alpha}_3$ and $\hat{\alpha}_4$ are the event study estimates close to, and on either side of, the 50 percent union vote share threshold, and $\hat{\alpha}_5$ and $\hat{\alpha}_6$ are the first and second moments of the vote share.⁵⁰ π^{obs} is the change in market value predicted by the model and VS is the predicted vote share for the union. \hat{V} is the estimated variance-covariance matrix of these 6 estimators.

Although our model is parsimonious and we chose simple functional forms (e.g. uniform distributions for ε , the distribution of voter preferences, and quadratic utility), it leads to somewhat complicated (and not particularly illuminating) analytic expressions for the theoretical moments in $f(\theta)$. Therefore, we estimate the parameters via Monte Carlo simulation. For each set of parameter values, we take 50000 Monte Carlo draws of μ and ε , and for each of those draws compute π^{obs} and VS as described above. We then use that simulated data to compute the theoretical moments in the same way the observed moments are calculated.

Before reporting the results, we provide some intuition as to how various parameters would affect the theoretical moments. First, as the distribution of μ (given by the parameters $\bar{\mu}$ and λ_μ) shifts in the negative direction, one can expect π^{obs} to become more negative, as both firm and union proposals respond to the location of μ . Second, $\bar{\pi}$ is essentially an upper bound to the union and firm proposals, so decreases will generally lead to lower π^{obs} as well. Third, a very small σ implies that workers within a firm have very similar preferences, and therefore will vote similarly, implying that the only observed vote shares would be close to either 0 or 1. If σ is very large, then vote shares would be clustered around an intermediate value.

⁵⁰Specifically, $\hat{\alpha}_3$ and $\hat{\alpha}_4$ are the values of the regression prediction on either side of the 50 percent threshold, from a regression of π^{obs} on a quartic in the vote share and a dummy variable for the vote share being greater than 50 percent.

Fourth, a very small λ_ε , which would represent very little uncertainty in the distribution of voters, would lead union and firm offers to converge towards each other. If the proposals are virtually identical, then we would expect no discontinuity in the event-study estimate with respect to the vote share.

Finally, we recognize that the observed data on π^{obs} and VS does not reveal the magnitude of c in any obvious way. For example, given the first order condition in 3 and quadratic utility function, a less negative c would raise the marginal gain to the union of lowering an offer, but at the same time it would increase the potential penalty of losing the election; this suggests an ambiguous impact of c on π^{obs} .⁵¹ For this reason, we investigate the extent to which our qualitative results are sensitive to the value of c by estimating the remaining 5 parameters, conditional on varying values of c .

5.3 Results and Policy Extrapolation

We estimate the model by minimizing the quadratic form described above. In doing so, we discovered that the objective function was virtually flat with respect to the parameter c , and that the estimated five parameters were not sensitive to the magnitude of c . For example, estimating the full six-parameter model gave estimates of $c = -17.20$, $\bar{\pi} = 0.042$, $\bar{\mu} = 0.339$, $\lambda_\mu = 7.80$, $\sigma = 0.311$, $\lambda_\varepsilon = 0.101$, whereas estimating the remaining 5 parameters conditional on fixing the value of c at -2.29 yielded $\bar{\pi} = 0.043$, $\bar{\mu} = 0.343$, $\lambda_\mu = 7.90$, $\sigma = 0.314$, $\lambda_\varepsilon = 0.102$. We concluded that c was not well-identified, and therefore we report the results from fixing c at -2.29 .⁵²

To illustrate the fit of the model, we generated simulated data according to the estimated parameters. Figure 13 shows a histogram of simulated equilibrium vote shares. Overall, the distribution shares a similar shape to the actual distribution of vote shares in Figure 9. As expected – since the estimation procedure only used the first two moments – there are some notable discrepancies. First, the simulated data yields a ratio of union losses to victories is about 2 to 1, compared to the actual ratio of about 2.5 to 1. Second, the simulated data produced no observations with vote shares above 83 percent, whereas Figure 9 shows a small number of cases in that upper tail.

The fit of the model can also be seen in Figure 14, which provides the predicted change in market value,

⁵¹Adding to the ambiguity of how c might affect the equilibrium offers is the fact that c is a lower bound on both union and firm offers.

⁵² -2.29 seems to be the lower bound on the change in market value (relative to the broad market index): historically, over an 18 month period, the most the broad market index has ever increased has been 129 percent. Since an individual firm's stock price cannot lose more than 100 percent of its initial value, we take as the most negative excess returns to be -229 percent. Fixing c to be half of that value (-1.15) has almost no effect on the magnitude of the remaining parameters.

as a function of the observed vote share, using the simulated data. The figure gives the same overall shape as that in Figure 12, with the union effects flat and near zero to the left of the 50 percent vote share threshold, and a negative slope to the right of the threshold. In our judgment, while this five-parameter model certainly does not capture every feature of the observed data, it does seem to provide a reasonable approximation.

Importantly, our modest “calibration” exercise of this electoral competition framework suggests that unions are responding to workers’ preferences. Using the simulated data, the regression of the union offers on the expected median position $(\mu - \frac{\sigma}{2} - \frac{\lambda \varepsilon}{2})$ yields a coefficient of 0.734. Furthermore, our model suggests that firm and union offers are generally more “moderate” than the positions of the median voter. Figure 14 plots the average realized position of the median $(\mu - \frac{\sigma}{2} - \varepsilon)$ worker by the realized vote share using the simulated data. It shows that when the union loses, the median worker’s ideal profit is higher than the firm’s offer, while the worker’s ideal level is more negative than what the union offers, when the union prevails in the election.

The simulation results also provide insights into worker preferences. The simulations imply that the distribution of worker preferences for higher compensation is highly skewed left. The 50th percentile of the median voter distribution (across firms) has a median voter with a preference for the change in equity value of *positive* 4.6%. At the same time, 25 percent of median voters have preferences that are more negative than -4.3%, while 5% have preferences for the change in equity value of less than -24%. This distribution suggests that the taste for large compensation packages amongst workers considering unionization is present in only a small number of establishments. We can think of these preferences as related to the establishment’s elasticity of labor demand in the sense that workers are willing to accept larger compensation packages when demand is more inelastic and when their jobs are not at risk. From this perspective these simulations imply that the great majority of establishments undergoing elections have fairly elastic labor demand.

For the policy simulation we hold all of the parameters fixed at their estimated values, and c at -2.29, and then vary the threshold for a union election win. Changing the threshold alters the probability $P(\pi_M, \pi_U)$ in our model. For example, if the threshold is 25% of affirmative votes required to unionize, then the probability of a firm victory becomes $P(\pi_M, \pi_U) = \Pr[VS < \frac{1}{4}] = \frac{1}{\lambda \varepsilon} [\mu - \frac{3}{4}\sigma - (\frac{\pi_M + \pi_U}{2})]$. We then conduct a Monte Carlo simulation with 50,000 draws to compute the equilibrium union and management offers, and the union vote share distribution. We view this analysis as applicable to proposals that would make it easier for unions to organize workers, such as EFCA. Under the EFCA scenario there would no longer be elections, but it is arguably still true that we can view workers as deciding between two options (sign card or not), which is not

unlike an election with a low union vote threshold. As in our model, firms and unions would try to influence that decision. Our policy simulation yields predicted effects of making it marginally easier to unionize (as indexed by the fraction that would win certification).

In Table 7 we present the results from the policy simulations. The columns represent different scenarios according to the union vote share threshold for certification. The population is split into five mutually exclusive groups, represented by the rows. Each sub-group will either be all unionized or not, depending on the scenario, as indicated by the labels “YES” and “NO”. In the “Proportion” column, it is seen that as the threshold decreases, more and more elections result in a union win. For example, a 25 percent threshold corresponds to a $(33+37+15=)$ 85 percent union win rate while a 50 percent threshold corresponds to a 33 union win rate.

The first row shows the predicted average percentage change in market value for the entire population for the different scenarios. Note that here we include both sets of establishments (winners and losers) because the model allows for management to vary their offers in response to a union threat and because lowering the threshold changes the composition of establishment in each category, as more establishments are unionizing. We find that a more than doubling of the union win rate (from 33 to 70 percent, as we move from the 50 percent to the 33 percent vote share threshold), leads to an overall decrease in equity value of about 4 percent. If the union vote share threshold were lowered to 10 percent, it would increase the proportion unionized to 0.99, and the simulation predicts a further 6.6 percent decline (from -0.058 to -0.124) in equity value.

Our fully specified model allows us to examine the main sources of these changes. We are able to examine the changes for the sub-groups defined in the second through sixth rows of Table 7. We point to three general patterns. First, as we lower the vote share threshold, the market value change of the group of firms that would continue to lose under the new scenario remains fairly stable. Indeed the group in the second row (“Inframarginal Loss”) experiences no change in market value. This pattern is consistent with management not being highly responsive to increased union threat as a result of the policy change. By contrast, we do see important changes in equity value among “inframarginal” unions who are already victorious with the higher threshold. This can be seen most clearly in the “Inframarginal win” row of Table 7, where the union effect drops from -0.117 to -0.153, moving from the 50 percent to the 33 percent threshold. This negative equity effect falls to -0.205 when the threshold falls to 10 percent, which according to the simulation would mean nearly the entire population would be unionized. Finally, we observe that when a marginal group shifts from the union losing the election to it winning the election there is a significant reduction in the market value of

the firm. We note that each time this occurs, the change in the marginal group (ranging from -0.08 to -0.10) is reasonably approximated by the the estimated RD estimate using the simulated data (also is in the -0.08 to -0.10 range), which is shown in the last row of Table 7. Thus, one reason to be cautious about the simulated overall effect is that it is to some extent being driven by the simulated RD estimate, which is somewhat larger than the point estimates we obtain from the actual data.

There are other reasons for caution in making these policy predictions, particularly because of our choice of model. For example, we are not modeling which establishments hold union representation elections in the first place. It is possible that lowering the threshold for unionization will change the composition of which establishments hold elections. We speculate that the marginal firms induced to hold an election by the policy change would be ones where wage demands are relatively weak in the first place, since one could argue that the cost of holding an election outweighed the potential benefits to the union.

A perhaps more fundamental concern is that our conclusions are made through the lens of a model of electoral competition. But it is possible that workers are not voting on compensation packages, and hence that unions and management are not acting strategically to influence the vote. In this case, we might expect to see the observed relationship between the vote share and the change in the market value because unions require widespread support in order to be effective, for example to impose a credible strike threat. Distinguishing this model from the one we propose would involve examining the employment changes following representation elections. One prediction of a model of electoral competition is that there should be limited employment effects from new unionization, something that we view as plausible given the results in DiNardo and Lee (2002, 2004). Exploring this further would be a fruitful avenue for future research.

We are not aware of any other attempt to estimate the impact of policies that ease unionization. Thus, in spite of the above caveats, we believe that our modeling and simulation exercise, which is disciplined by the magnitudes we find in our event-study and RD analyses, provides a useful benchmark for policy predictions.

6 Conclusion

The economic effects of unions on the labor market and the economy have been a longstanding area of interest for economists. The literature has considered the impact of unions on wages, their potential role as monopolies, their role in work stoppages, their effect on the aggregate economy, as well as the question of how they can even exist and survive in a competitive labor market. In order to even partially address many

of these questions, we must first understand how unions affect firms.

We began by asking whether the case of National Linen Services was the rule or the exception. In one respect, it is the rule. We have shown that among publicly traded firms where the workforce attempting to organize is not too small, new unionization is associated with a reduction in the firm's market value, in a way that parallels the experience of NLS. Like the NLS case, the stock market reaction to union victories is somewhat slow, as has been found in a number of other event-study contexts. This finding is robust to the use of a variety of specifications and to the use of several different methodologies. The negative effects of unionization on the equity value of firms appears fairly stable over time, showing no major differences before or after 1984. An examination of accounting variables of both sets of firms reveals that union wins are associated with relatively lower growth, though there is little evidence to suggest that these firms experienced lower return on assets or profit margins as compared to firms with union losses. The evidence is therefore consistent with the claim that unionization reduces the number of sufficiently positive NPV projects available to a firm.

In another respect, however, the case of NLS is a clear exception. By two years after the union victory, NLS stock had earned negative 75 percent abnormal returns. By contrast, for our sample we estimate abnormal returns of about negative 10 percent, and our sample is somewhat representative of publicly-traded firms at risk of unionization. Based on the market capitalization of these firms, this 10 percent equity loss translates to a total loss of about \$40,500 (in 1998 dollars) per voter. Since this amount represents a combination of a transfer to workers as well as lost profit due to inefficiencies caused by the union, one can view this magnitude as an upper bound on the redistributive effect or the efficiency effect.⁵³ For example, if the true average union wage effect is 8 percent and if our back-of-the-envelope calculation (that a \$40,500 loss would translate to a pure transfer equivalent to a 10 percent wage premium) is correct, then this would imply a 2 percent loss in terms of efficiency due to unions.

The large difference in magnitude between the case of NLS and the estimated average effects serves to highlight the importance of heterogeneous effects, which we carefully document in our analysis. Using a different sample from DiNardo and Lee (2004), we also find RD estimates that imply unionization is largely ineffective for firms where there is more moderate support for unions, at least to the extent that unions do not affect a firm's equity value. This finding can be reconciled with the findings from the event-study analysis

⁵³Treating this magnitude as an upper bound requires assuming that unions can only impose efficiency costs, and cannot lead to increases in profitability (after netting out compensation costs).

through the negative gradient in abnormal returns in relation to the union vote share.

Finally, we consider a voting model of endogenous union determination, and calibrate it with the magnitudes we find in our empirical analysis to make a first-cut prediction on the likely impact of policies that increase likelihood of unionization. Policy simulations show that easing the threshold necessary to gain recognition would not lead to union threat effects (firms losing value by having to respond to the threat of unionization), but would cause unions to use this increased voter slack to be more aggressive. While the RD estimates reasonably approximate effects for small policy changes, the approximation leads to an increasingly larger understatement of the effects of larger policy shifts. Our exercise suggests that a policy-induced doubling of unionization would lead to a 4.3 percent decrease in the equity value of all firms at risk of unionization.

References

- Abowd, John M.**, “The Effect of Wage Bargains on the Stock Market Value of the Firm,” *American Economic Review*, September 1989, 79 (4), 774–800.
- Allen, Steven G.**, “The Effect of Unionism on Productivity in Privately and Publicly Owned Hospitals and Nursing Homes,” *Journal of Labor Research*, 1986, 7, 59–68.
- , “Unionization and Productivity in Office Building and School Construction,” *Industrial and Labor Relations Review*, 1986, 39, 187–201.
- Ashenfelter, Orley and David Card**, “Time Series Representations of Economic Variables and Alternative Models of the Labour Market,” *Review of Economic Studies*, 1982, 49 (5), 761–81.
- **and George Johnson**, “Trade Unions, Bargaining Theory, and Industrial Strike Activity,” *American Economic Review*, May 1969, 59, 35–49.
- Asquith, Paul**, “Merger bids, uncertainty, and stockholder returns,” *Journal of Financial Economics*, April 1983, 11, 51–83.
- Atherton, Wallace N.**, *The Theory of Union Bargaining Goals*, Princeton, NJ: Princeton University Press, 1973.
- Ball, Ray and Philip Brown**, “An Empirical Evaluation of Accounting Income Numbers,” *Journal of Accounting Research*, Autumn 1968, 6 (2), 159–178.
- Barber, Brad M. and William John D. Lyons**, “Detecting long-run abnormal stock returns: The empirical power and specification of test statistics,” *Journal of Financial Economics*. Volume 43, Issue 3, March 1997, Pages 341-372, March 1997, 43 (3), 341–372.
- Basker, Emek**, “The Causes and Consequences of Wal-Mart’s Growth,” *Journal of Economic Perspectives*, 2007, Summer, 177–198.
- Becker, Brian E.**, “Concession Bargaining: The Impact on Shareholders’ Equity,” *Industrial and Labor Relations Review*, 1987, 40, 268–279.
- **and Craig A. Olson**, “The Impact of Strikes on Shareholder Equity,” *Industrial and Labor Relations Review*, April 1986, 39 (3), 425–438.
- Blanchflower, David G. and Alex Bryson**, “What Effect Do Unions Have on Wages Now and Would Freeman and Medoff Be Surprised?,” in James T. Bennett and Bruce E. Kaufman, eds., *What Do Unions Do? A Twenty-Yr. Perspective*, Transaction Publishers, 2007.
- Booth, Alison Lee**, *The Economics of the Trade Unions*, Cambridge: Cambridge University Press, 1995.
- Brav, Alon**, “Inference in Long-Horizon Event Studies: A Bayesian Approach with Application to Initial Public Offerings,” *The Journal of Finance*, October 2000, 55 (5), 1979–2016.
- **and Paul A. Gompers**, “Myth or Reality? The Long-Run Underperformance of Initial Public Offerings: Evidence from Venture and Nonventure Capital-Backed Companies,” *Journal of Finance*, 1997, 52 (5), 5.
- Bronars, Stephen G. and Donald R. Deere**, “Union Representation Elections and Firm Profitability,” *Industrial Relations*, Winter 1990, 29, 15–37.

- and — , “The Threat of Unionization, the Use of Debt, and the Preservation of Shareholder Wealth,” *Quarterly Journal of Economics*, 1991, 106, 231–54.
- Cameron, Colin A., Jonah B. Gelbach, and Douglas L. Miller**, “Robust Inference with Multi-Way Clustering,” Technical Working Paper T327, National Bureau of Economic Research September 2006.
- Campbell, John Y., Andrew W. Lo, and A. Craig MacKinlay**, *The Econometrics of Financial Markets*, Princeton, NJ: Princeton University Press, 1997.
- Carhart, Mark M.**, “On Persistence in Mutual Fund Performance,” *Journal of Finance*, 1997, 52, 57–82.
- Clark, Kim B.**, “The Impact of Unionization on Productivity: A Case Study,” *Industrial and Labor Relations Review*, 1980, 33, 451–69.
- , “Unionization and Productivity: Micro-Econometric Evidence,” *Quarterly Journal of Economics*, 1980, 95, 613–39.
- , “Unionization and Firm Performance: The Impact on Profits, Growth, and Productivity,” *American Economic Review*, December 1984, 74 (5), 893–919.
- Cooke, William N.**, “The Failure to Negotiate First Contracts: Determinants and Policy Implications,” *Industrial and Labor Relations Review*, January 1985, 38 (2), 163–178.
- Cusatis, Patrick J., James A. Miles, and J. Randall Woolridge**, “Restructuring through spinoffs: The stock market evidence,” *Journal of Financial Economics*, June 1993, 33 (3), 293–311.
- Dellavigna, Stefano and Joshua Pollet**, “Demographics and Industry Returns,” *American Economic Review*, forthcoming.
- Dharan, Bala G. and David L. Ikenberry**, “The Long-Run Negative Drift of Post-Listing Stock Returns,” *The Journal of Finance*, December 1995, 50 (5), 1547–1574.
- DiNardo, John and David S. Lee**, “The Impact of Unionization on Establishment Closure: A Regression Discontinuity Analysis of Representation Elections,” Working Paper 8993, National Bureau of Economic Research 2002.
- and — , “Economic Impacts of New Unionization on Private Sector Employers: 1984-2001,” *Quarterly Journal of Economics*, November 2004, 119 (4), 1383–1441.
- and **Kevin F. Hallock**, “When Unions “Mattered”: The Impact of Strikes on Financial Markets, 1925-1937,” *Industrial and Labor Relations Review*, 2002, 55.
- Fama, Eugene**, “Market efficiency, long-term returns, and behavioral finance,” *Journal of Financial Economics*, September 1998, 49 (3), 283–306.
- Fama, Eugene F. and Kenneth R. French**, “Common Risk Factors in the Returns on Stocks and Bonds,” *Journal of Financial Economics*, 1993, 33, 3–56.
- Farber, Henry**, “Individual Preferences and Union Wage Determination: The Case of the United Mine Workers,” *Journal of Political Economy*, 1978, 86, 923–942.
- Farber, Henry S.**, “An Analysis of Final-Offer Arbitration,” *Journal of Conflict Resolution*, 1980, 24, 683–705.

- , “Right-to-Work Laws and the Extent of Unionization,” *Journal of Labor Economics*, July 1984, 2 (3), 319–352.
- **and Bruce Western**, “Accounting for the Decline of Unions in the Private Sector, 1973-1998,” *Journal of Labor Research*, 2001, 22, 459–486.
- Featherstone, Liza**, “Will Labor Take the Wal-Mart Challenge?,” *The Nation*, 2004, June.
- Freeman, Richard B. and James L. Medoff**, *What Do Unions Do?*, New York: Basic Books, 1984.
- **and Morris M. Kleiner**, “The Impact of New Unionization on Wages and Working Conditions,” *Journal of Labor Economics*, January 1990, 8 (1), S8–S25.
- Friedman, Milton**, “Some Comments on the Significance of Labor Unions for Economic Policy,” in David McCord Wright, ed., *The Impact of the Union: Eight Economic Theorists Evaluate the Labor Union Movement*, New York: Harcourt, Brace and Company, 1950. Institute on the Structure of the Labor Market, American University, Washington D.C.
- Greenstone, Michael, Paul Oyer, and Annette Vissing-Jorgensen**, “Mandated Disclosure, Stock Returns, and the 1964 Securities Acts Amendments,” *Quarterly Journal of Economics*, May 2006, 121 (2), 399–460.
- Harvey, Campbell**, “Time-varying conditional covariances in tests of asset pricing models,” *Journal of Financial Economics*, 1989, 24 (2), 289–317.
- Heckman, James and Joseph Hotz**, “Alternative Methods for Evaluating the Impact of Training Programs,” *Journal of the American Statistical Association*, 1989.
- Hirsch, Barry T.**, *Labor Unions and the Economic Performance of U.S. Firms*, Upjohn Institute for Employment Research, 1991.
- , “Union Coverage and Profitability among U.S. Firms,” *Review of Economics and Statistics*, 1991, 73, 69–77.
- , “What Do Unions Do for Economic Performance?,” in James T. Bennett and Bruce E. Kaufman, eds., *What Do Unions Do? A Twenty-Yr. Perspective*, Transaction Publishers, 2007.
- **and David A. Macpherson**, *Union Membership and Earnings Data Book*, Bureau of National Affairs, Inc., 2008.
- Holmes, Thomas**, “Geographic Spillover of Unionism,” April 2006.
- Hong, Harrison, Terrence Lim, and Jeremy C. Stein**, “Bad News Travels Slowly: Size, Analyst Coverage, and the Profitability of Momentum Strategies,” *Journal of Finance*, 2000, 55 (1), 2650295.
- Huth, William L. and Don N. MacDonald**, “Equity Market Response to union Decertification Petitions and Elections,” *Journal of Labor Research*, Spring 1990, 11 (2), 193–201.
- Ikenberry, David L., Graeme Rankine, and Earl K. Stice**, “What Do Stock Splits Really Signal?,” *The Journal of Financial and Quantitative Analysis*, September 1996, 31 (3), 357–375.
- Jaffe, Jeffrey**, “Special Information and Insider Trading,” *The Journal of Business*, July 1974, 47 (3), 410–428.

- Jegadeesh, Narasimhan and Sheridan Titman**, “Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency,” *The Journal of Finance*, March 1993, 48 (1), 65–91.
- Lalonde, Robert J., G. Marschke, and Kenneth Troske**, “Using Longitudinal Data on Establishments to Analyze the Effects of Union Organizing Campaigns in the United States,” *Annales D’Economie et de Statistique*, 1996, 41/42, 155–185.
- Lee, David S.**, “Randomized Experiments from Non-random Selection in U.S. House Elections,” *Journal of Econometrics*, 2008, 142, 675–697.
- Lewis, H. G.**, *Union Relative Wage Effects: A Survey*, Chicago: University of Chicago Press, 1986.
- Loughran, Tim and Jay R. Ritter**, “The New Issues Puzzles,” *The Journal of Finance*, March 1995, 50 (1), 23–51.
- Lyon, John D., Brad M. Barber, and Chih-Ling Tsai**, “Improved Methods for Tests of Long-Run Abnormal Stock Returns,” *The Journal of Finance*, February 1999, 54 (1), 165–201.
- Mandelker, Gershon**, “Risk and return: The case of merging firms,” *Journal of Financial Economics*, December 1974, 1 (4), 303–335.
- Matsa, David**, “Capital Structure as a Strategic Variable: Evidence from Collective Bargaining,” September 2006.
- Michaely, Roni, Richard H. Thaler, and Kent L. Womack**, “Price Reactions to Dividend Initiations and Omissions: Overreaction or Drift?,” *The Journal of Finance*, June 1995, 50 (2), 573–608.
- Mitchell, Mark L. and Erik Stafford**, “Managerial Decisions and Long-Term Stock Price Performance,” *The Journal of Business*, July 2000, 73 (3), 287–329.
- Mitchell, Merwin W. and Joe A. Stone**, “Union Effects on Productivity: Evidence from Western Sawmills,” *Industrial and Labor Relations Review*, 1992, 46, 135–45.
- Moore, William J. and Robert J. Newman**, “The Effects of Right-to-Work Laws: A Review of the Literature,” *Industrial and Labor Relations Review*, July 1985, 38 (4), 571–585.
- Neumann, George R.**, “The Predictability of Strikes: Evidence from the Stock Market,” *Industrial and Labor Relations Review*, July 1990, 33 (4), 535–535.
- Olson, Craig A. and Brian E. Becker**, “The Effects of the NRLA on Stockholder Wealth in the 1930s,” *Industrial and Labor Relations Review*, October 1990, 44 (1), 116–129.
- Persson, Torsten and Guido Tabellini**, *Political Economics: Explaining Economic Policy*, MIT Press, 2000.
- Roomkin, Myron and Richard Block**, “Case Processing Time and the Outcome of Representation Elections: Some Empirical Evidence,” *University of Illinois Law Review*, 1981, 1, 75–97.
- Rose, Nancy**, “Labor Rent Sharing and Regulation: Evidence from the Trucking Industry,” *The Journal of Political Economy*, December 1987, 95 (6), 1146–1178.
- Ruback, Richard and Martin Zimmerman**, “Unionization and Profitability: Evidence from the Capital Market,” *Journal of Political Economy*, 1984, 92, 1134–1157.

Data Appendix

This Appendix describes how we match establishments in the NLRB data to firms in the CRSP database. When matching we looked for similarities in the name listed in the NLRB election file to names that were *ever* present in the CRSP files. To this end, we created two data sets: one containing the company names in the NLRB election file and the other containing every company name that has ever appeared in the CRSP database.⁵⁴ This second data set will be hereafter referred to as the “master names file.” In addition to the company names, the master names file also contains a unique company id, the “PERMNO”, which allows for further matching to the CRSP and Compustat databases.

There are 195,889 certification elections in the NLRB data set that could potentially be matched to companies in the master names file. Because the matching process is tedious, and must almost entirely be done manually, we excluded any election with less than 100 voters. This resulted in 24,709 firms in the certification election file that potentially matched firms in the master list of CRSP company names.⁵⁵ These elections are comprised of 61 percent of all workers eligible to vote in NLRB certification elections. Using this smaller subset, firms in the election file were compared to firms in the master CRSP file using the matching algorithm employed by DiNardo and Lee (2004), which makes use of the SAS SPEDIS function. The algorithm matches company names in the NLRB file to company names in the master names file based on a so-called “spelling distance,” which considers those comparisons with a spelling distance above a pre-determined threshold as candidate matches.⁵⁶ The algorithm may match a company in the election file to more than one company name in the CRSP file. In these cases we selected the lowest spelling distance as the candidate match. If there was a tie in spelling distance between two candidate comparisons, we selected one match at random.

Because we matched firms on names only, manual inspection of the matches revealed that our automated procedure resulted in many matches that were obviously incorrect. Therefore, research assistants reviewed

⁵⁴Many companies have multiple names.

⁵⁵Because a firm can have multiple elections, this number includes multiple cases of the same firms. There are 18,344 unique firm spellings, though there are fewer unique firm names because of misspellings and abbreviations.

⁵⁶We refer the reader to DiNardo and Lee (2004) for further details on this algorithm. That study relied heavily on the establishment’s street address, which is unavailable here. Therefore, the spelling distance threshold was quite specific to that application. As a first pass, we modified the program to match only on firm name, and discovered that in this application, that same threshold led to “too many” matches. As we describe below, we therefore augmented the process with a manual review.

every match and dropped those where they judged the two firm names as different companies.^{57,58} We then collected all of the unmatched companies in the election file, from the initial set of 24,709, and attempted to locate each one in Dun and Bradstreet's Million Dollar Directory and the Lexis/Nexis' Directory of Corporate Affiliations for the year of election. This step identified subsidiaries of publicly traded parent companies, and allowed us to spot companies that were dropped erroneously in the previous step.

We ultimately matched 7,693 elections from the NLRB election file to companies in the CRSP master file. In 1,579 cases, the firm in the CRSP file was not publicly traded at the time of the election. After excluding the private firms, our final sample contained 6,114 elections, consisting of 20 percent of all workers eligible to vote in NLRB elections.

In order to determine whether the matches appeared reasonable, we compared the reported two-digit SIC industry code and the state of the establishment from the election file to the corresponding variables in the CRSP and Compustat files, for industry and state respectively. Because companies are diversified, the main SIC code for a company in the CRSP database need not be the same as the SIC code for a particular establishment in the NLRB election file. Similarly, an establishment may not be located in the same state as the company's headquarters. However, the comparisons are reassuring: the two digit SIC codes in the two data sets are the same for 50 percent of the matches, while 40 percent of the matches show the same state. For reference, if we randomly pair companies from the final NLRB data set to companies in the master names file that were never matched to the NLRB data through our procedure, the corresponding match rate is 5 percent for industry and 4 percent for state.

⁵⁷For example, the algorithm determined that any company in the election file with the word "American" as part of its name was a sufficiently good match for the company "American Enterprises" in the CRSP file, if a better match did not exist. Therefore, a disparate set of companies like "American Laundry," "American Envelope," and "Pan American Screws" were all matched to "American Enterprise." All of these matches were dropped by our research assistants.

⁵⁸Because there was an element of judgment, these exclusions were recorded in a log file for replication purposes.

Figure 1: Cumulative stock market returns surrounding National Linen Service's 1999 representation election

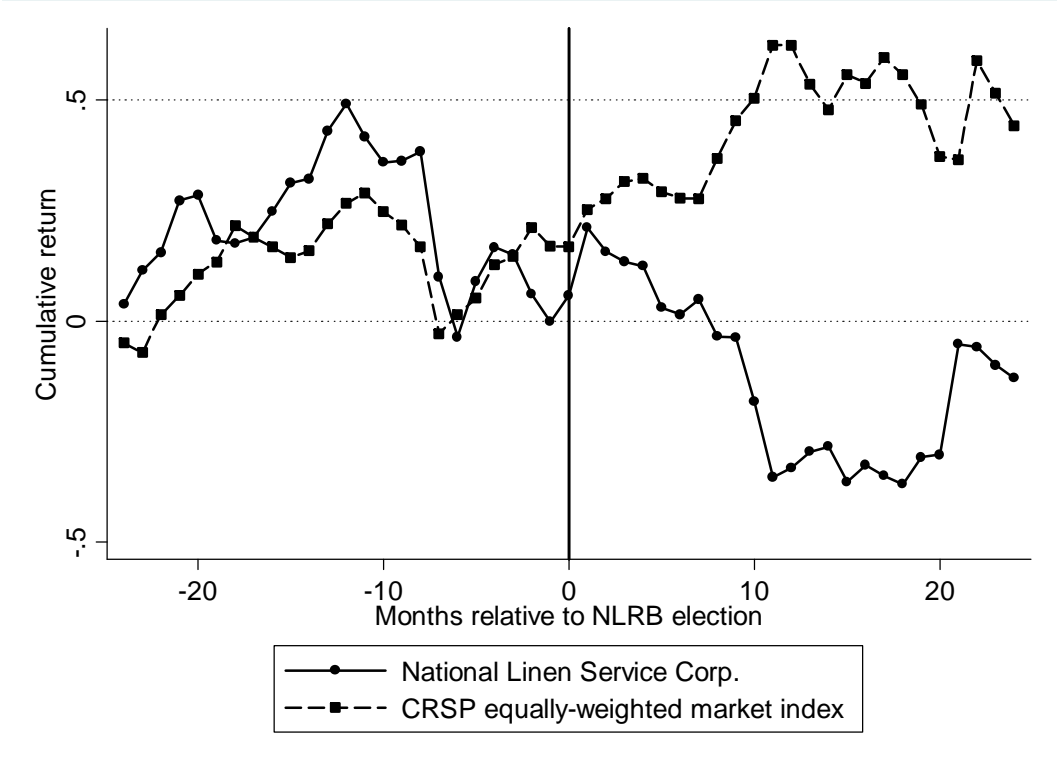
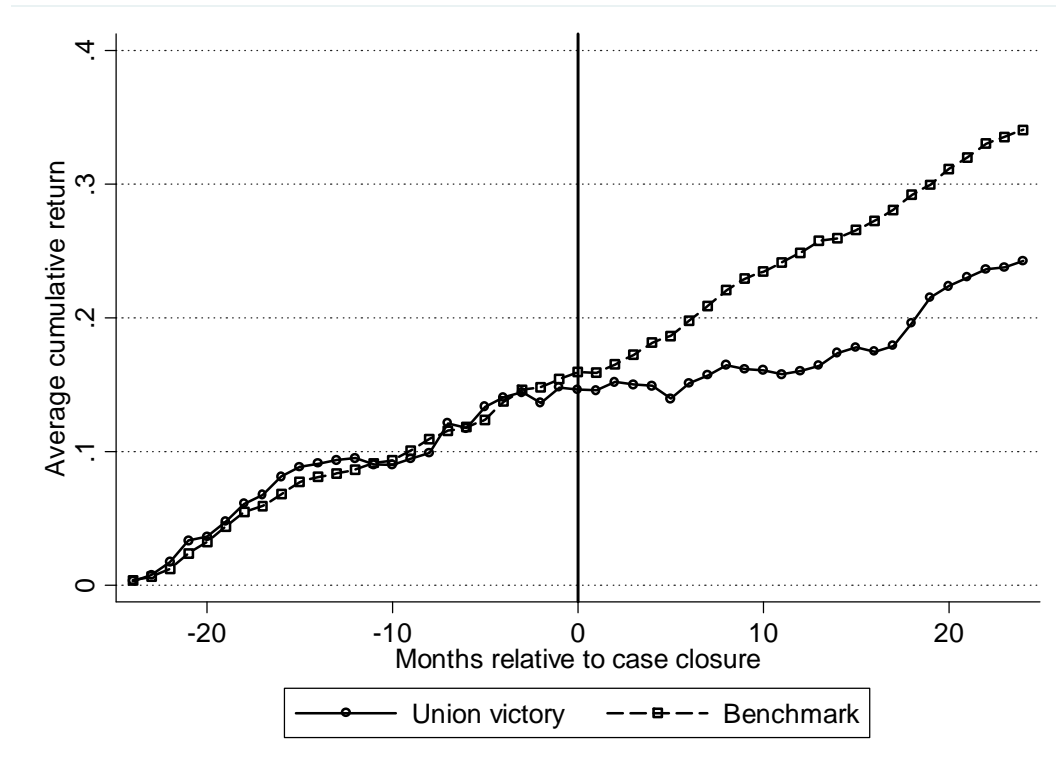


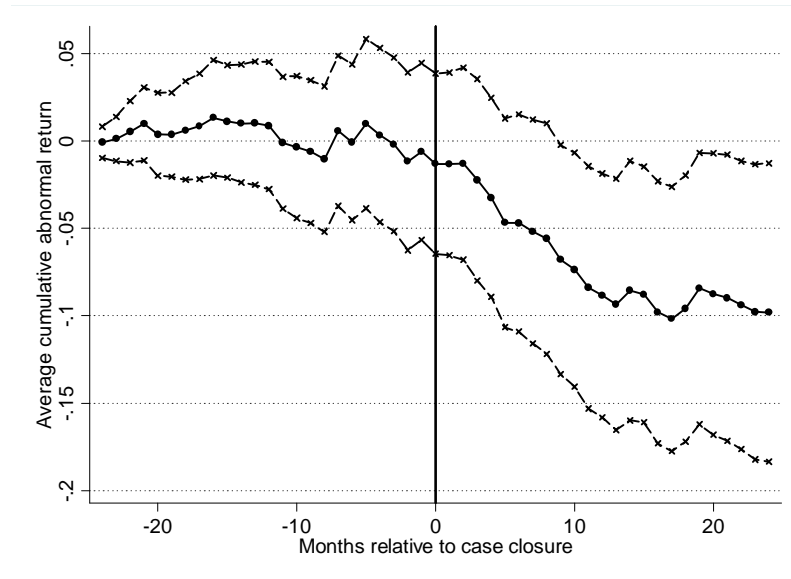
Figure 2: Average cumulative returns of union victory firms and of the size-matched reference portfolio, by month relative to NLRB case closure



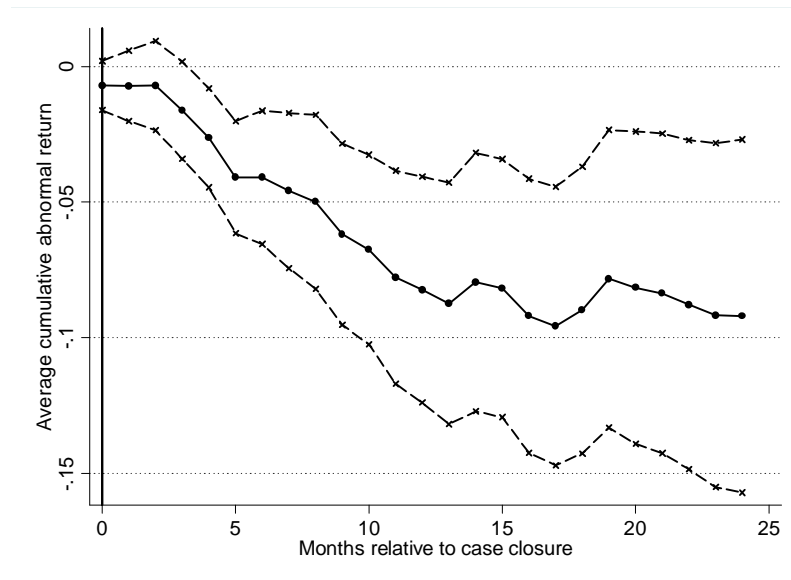
Note: Union victory firms consist of publicly traded companies holding representation elections where at least 5% of the company's workforce voted, and where the union won. Each point is the average cumulative return up to the month relative to case closure, beginning 24 months prior to case closure. Each firm in the sample is associated with a benchmark portfolio matched on size. The benchmark series corresponds to the average cumulative return of these size-matched reference portfolios. Returns are expressed net of the risk-free rate.

Figure 3: Average cumulative abnormal return of union victory firms, by month relative to NLRB case closure

Panel A: Beginning 24 months prior to case closure

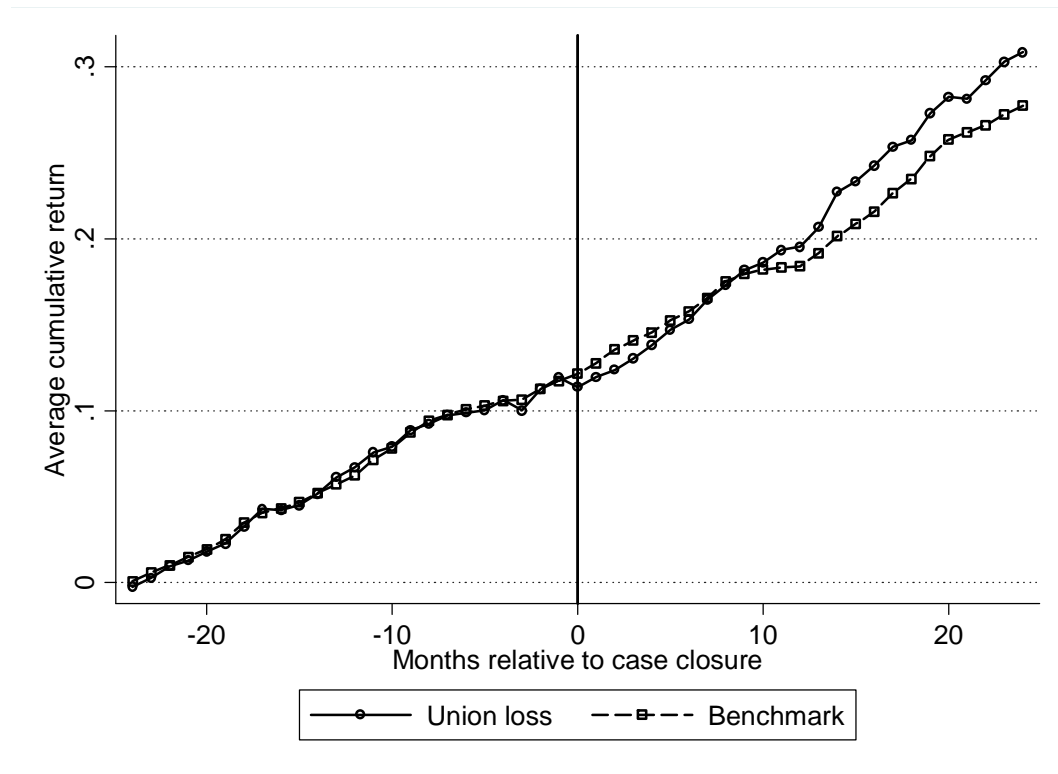


Panel B: Beginning month of case closure



Notes: Both panels show the difference in the average cumulative return of union victory firms and the size-matched reference portfolio, as shown in Figure 2. Panel A corresponds to the average cumulative abnormal return computed beginning 24 months prior to case closure. Panel B corresponds to the average cumulative abnormal return computed beginning in the month of case closure. The dashed lines represent the 95% confidence intervals, which are computed using standard errors clustered on elections and calendar months. We use the formula in Cameron, Gelbach, and Miller (2006) to compute standard errors with multi-way clustering.

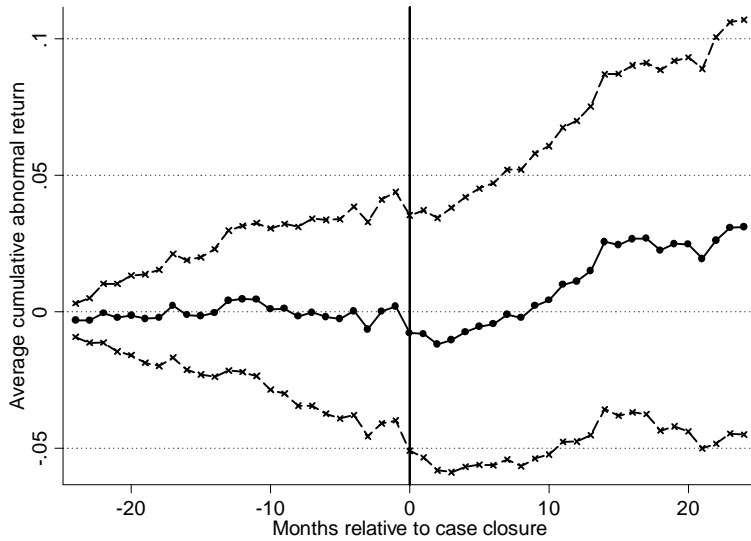
Figure 4: Average cumulative returns of union loss firms and of the size-matched reference portfolio, by month relative to NLRB case closure



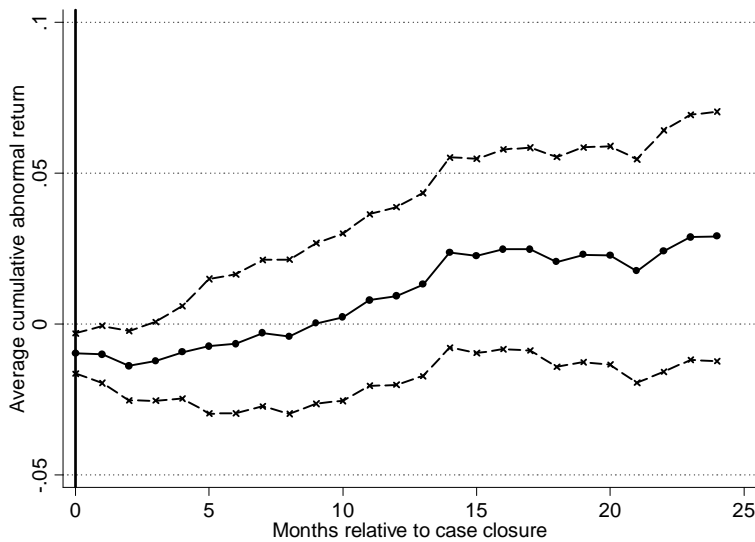
Note: Union loss firms consist of publicly traded companies holding representation elections where at least 5% of the company's workforce voted, and where the union lost. Each point is the average cumulative return up to the month relative to case closure, beginning 24 months prior to case closure. Each firm in the sample is associated with a benchmark portfolio matched on size. The benchmark series corresponds to the average cumulative return of these size-matched reference portfolios. Returns are expressed net of the risk-free rate.

Figure 5: Average cumulative abnormal returns of union loss firms, by month relative to case closure

Panel A: Beginning 24 months prior to case closure

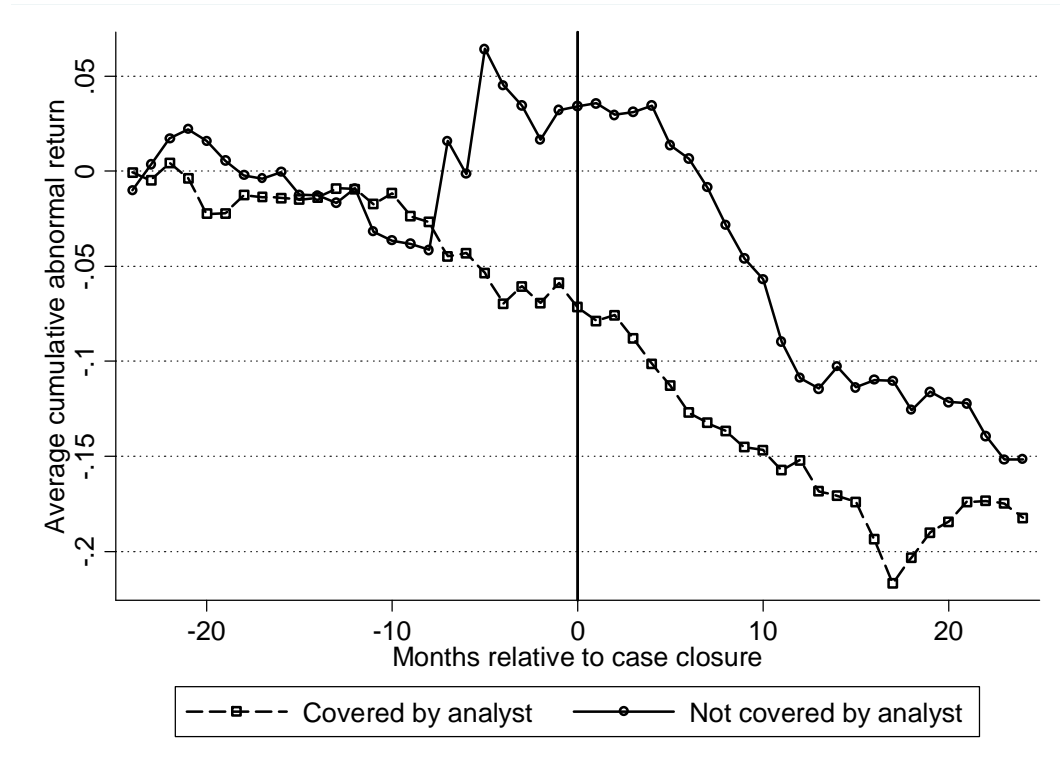


Panel B: Beginning month of case closure



Notes: Both panels show the difference in the average cumulative return of the union loss portfolio and the size-matched reference portfolio, shown in Figure 4. Panel A corresponds to the average cumulative abnormal return computed beginning 24 months prior to case closure. Panel B corresponds to the average cumulative abnormal return computed beginning at the month of case closure. The dashed lines represent the 95% confidence intervals, which are computed using standard errors clustered on elections and calendar months. We use the formula in Cameron, Gelbach, and Miller (2006) to compute standard errors with multi-way clustering.

Figure 6: Average cumulative abnormal return, by analyst coverage

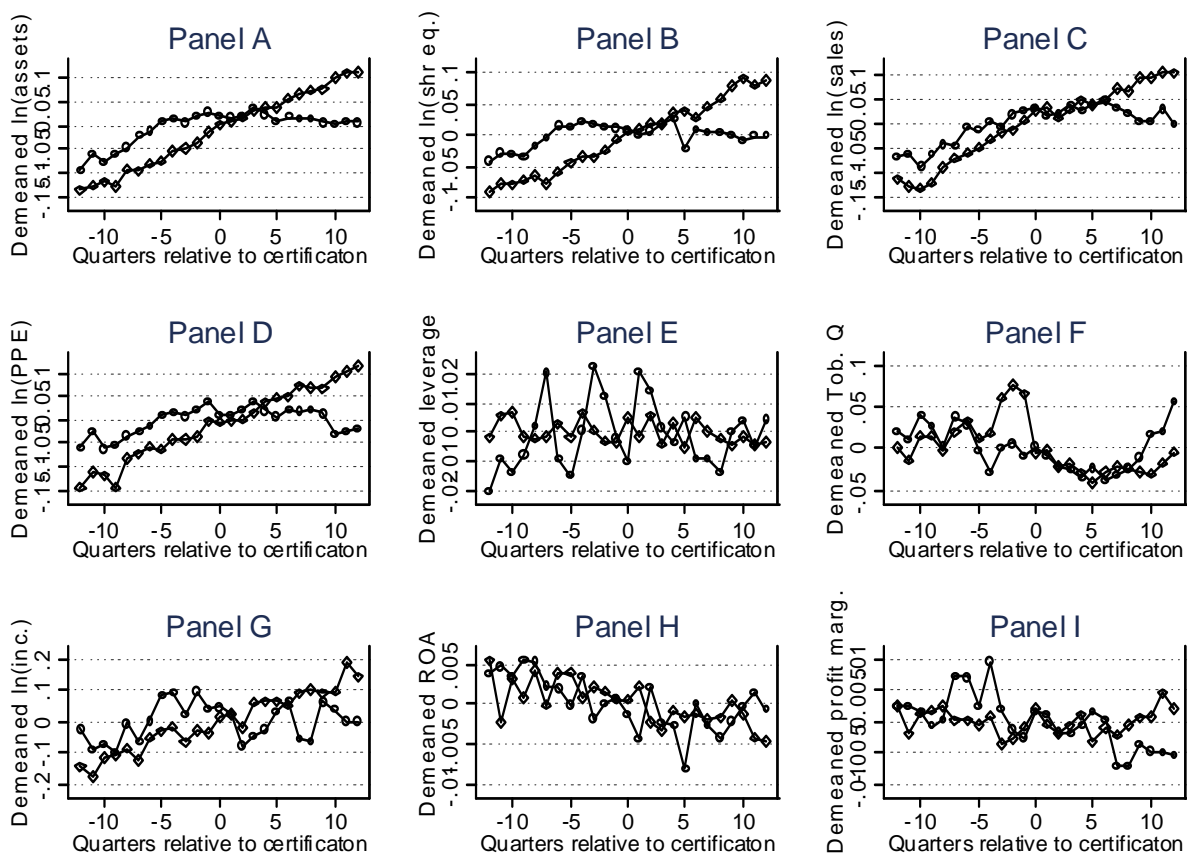


Note: A company is considered to have analyst coverage if it appears in the I/B/E/S dataset in the year of the election. The sample is limited to elections occurring in years where I/B/E/S data were available, between 1976 and 1999

Figure 7: Average cumulative abnormal return, by time period of election



Figure 8: Compustat variables; Union victory/loss comparisons



Notes: The sample consists of publicly traded companies with elections taking place between 1973-1999 where at least 5% of the workforce voted. Lines with circles correspond to union victory companies. Lines with diamonds correspond to union loss companies. All variables are drawn from the Compustat quarterly database. Each variable is demeaned, where the mean is taken within each election panel.

Figure 9: Histogram of the union vote share

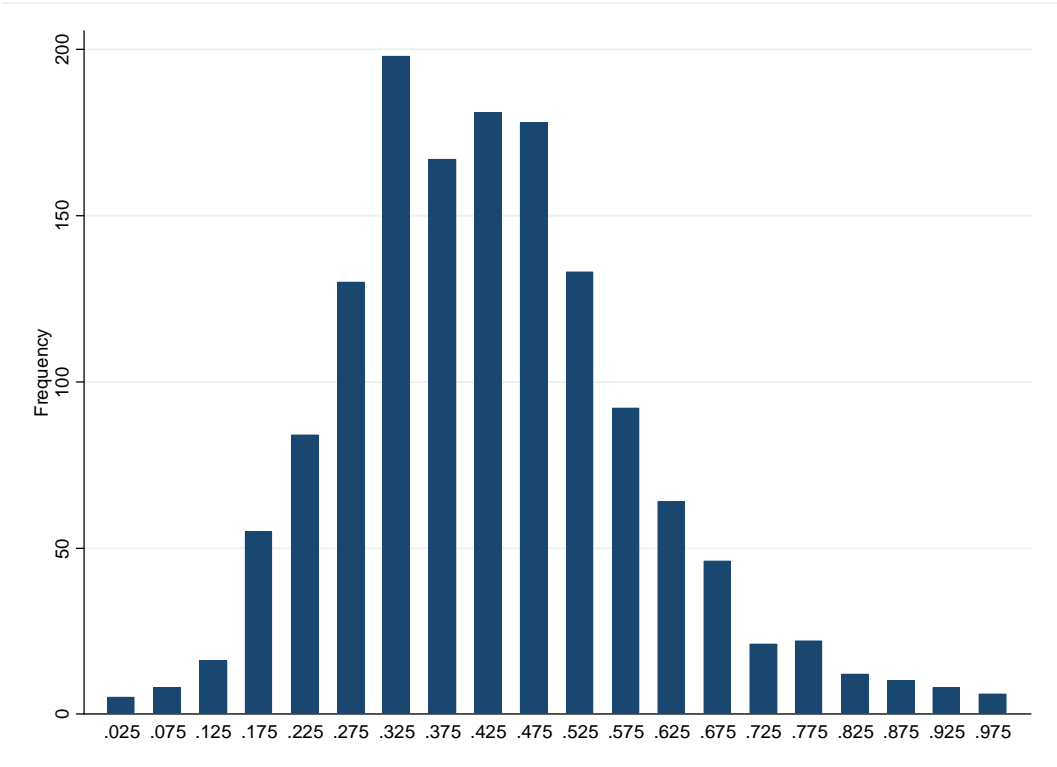
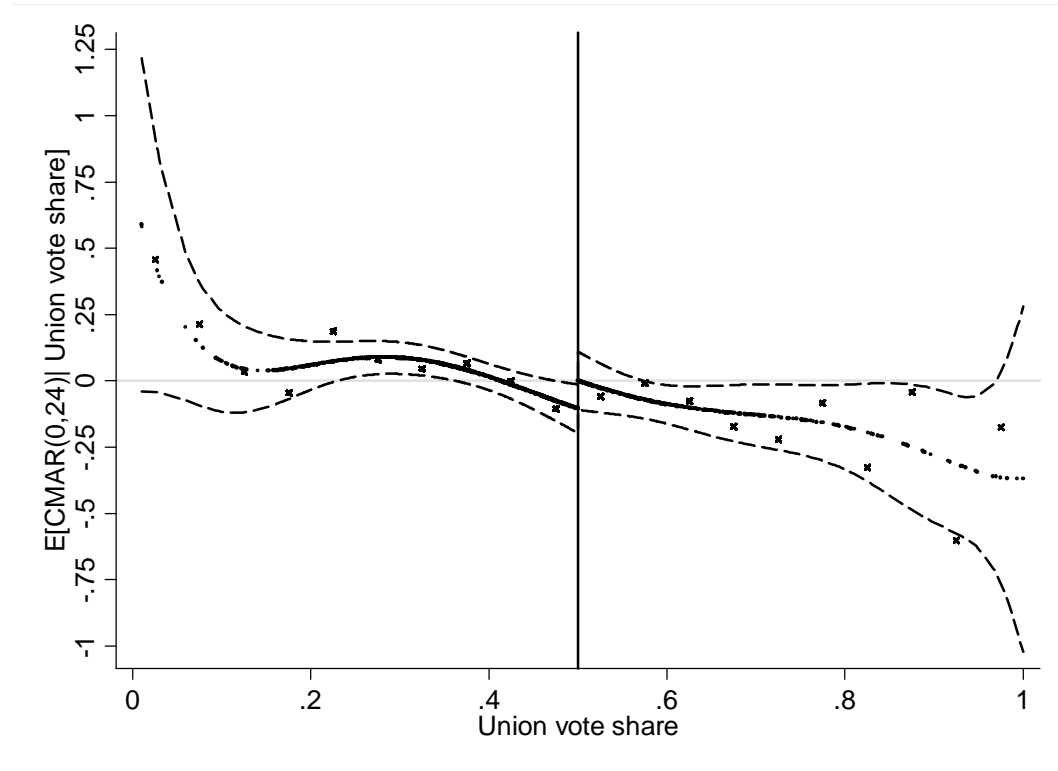
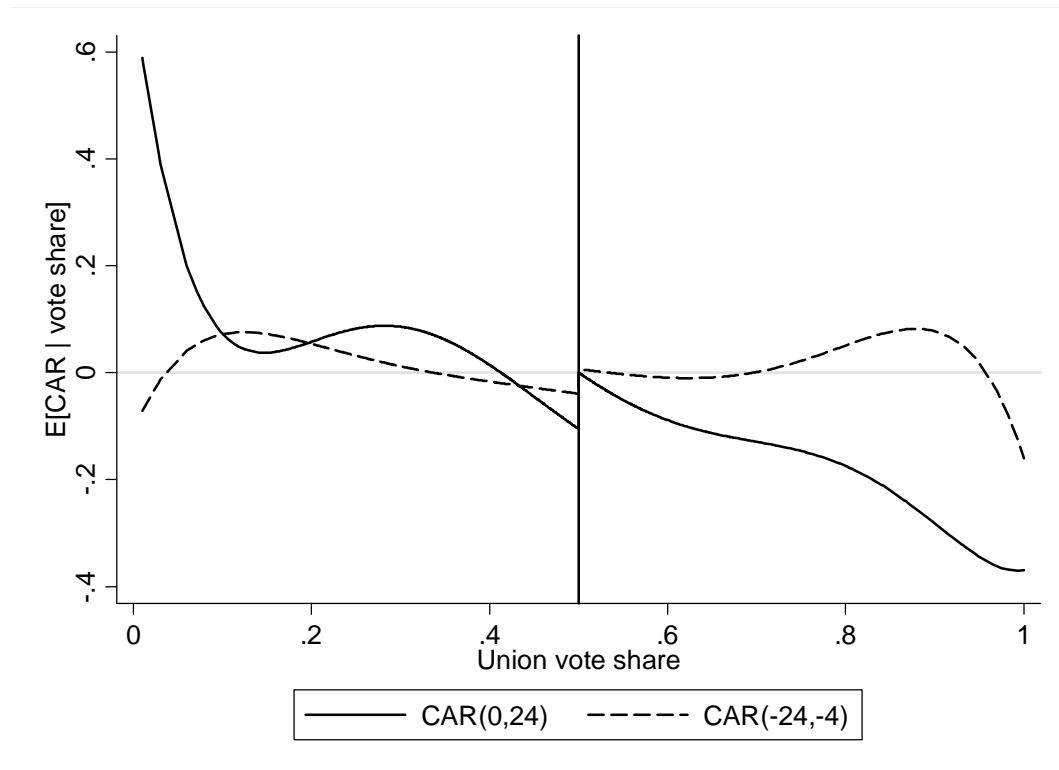


Figure 10: Cumulative abnormal returns in the two years after NLRB closes election, by relation to vote share



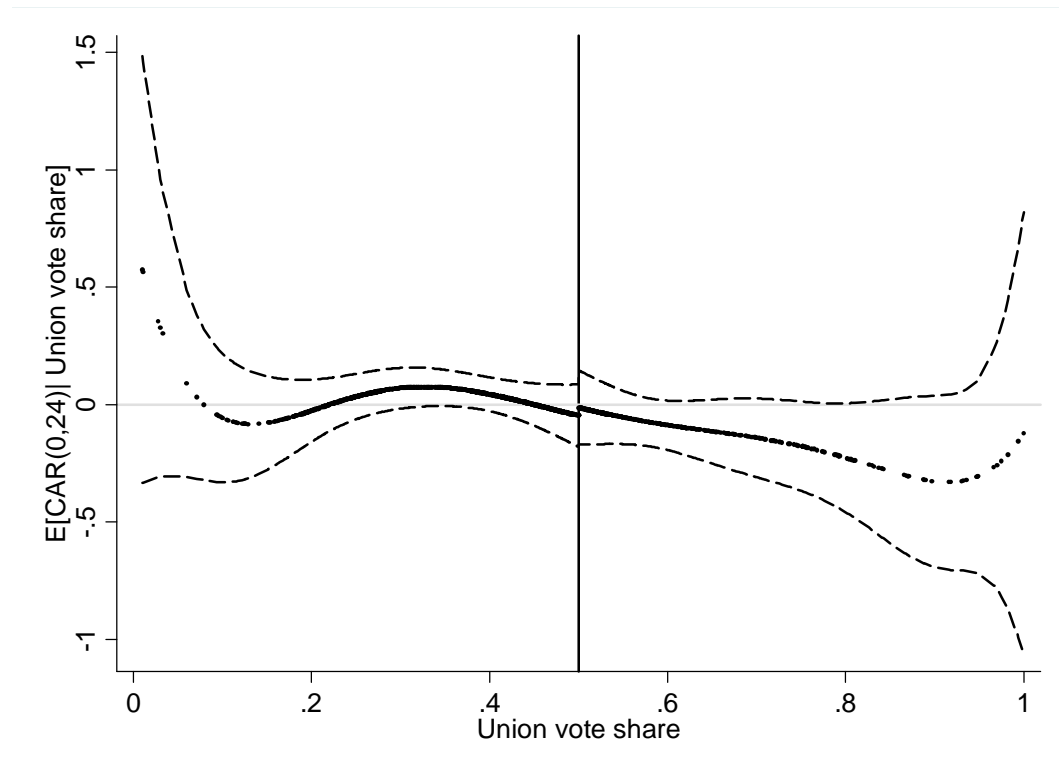
Note: Abnormal returns are the simple difference in the security's return and the size-matched benchmark portfolio in the same month. Cumulative abnormal returns are the sum of the abnormal returns over a two year period beginning in the month of case closure. Predicted values are calculated using a sixth-order polynomial, and an indicator for whether the union won. Dashed lines are the 95% confidence interval. Dots are the average cumulative excess return in 20 equally spaced bins. See Section 4.4 for further details on the construction of this figure.

Figure 11: Cumulative abnormal returns in the pre- and post-event periods, by relation to vote share



Notes: Predicted values are calculated using a sixth-order polynomial and an indicator for whether the union won. The solid line corresponds to the predicted cumulative excess return in the two years following case closure, conditional on union vote share. The dashed line corresponds to the predicted cumulative abnormal return calculated starting 24 months prior to the election through four months prior to case closure, conditional on union vote share. See Section 4.4 for further details on the construction of this figure.

Figure 12: Cumulative adjusted abnormal returns in the two years after NLRB closes election, by relation to vote share



Notes: Adjusted-cumulative abnormal returns are cumulative abnormal returns that have been adjusted for security-specific pre-election trend in abnormal returns. See Section 4.4 for details on the construction of this variable. Predicted values are calculated using a sixth-order polynomial and an indicator for whether the union won. Dashed lines are the 95% confidence intervals.

Figure 13: Histogram of simulated equilibrium vote shares

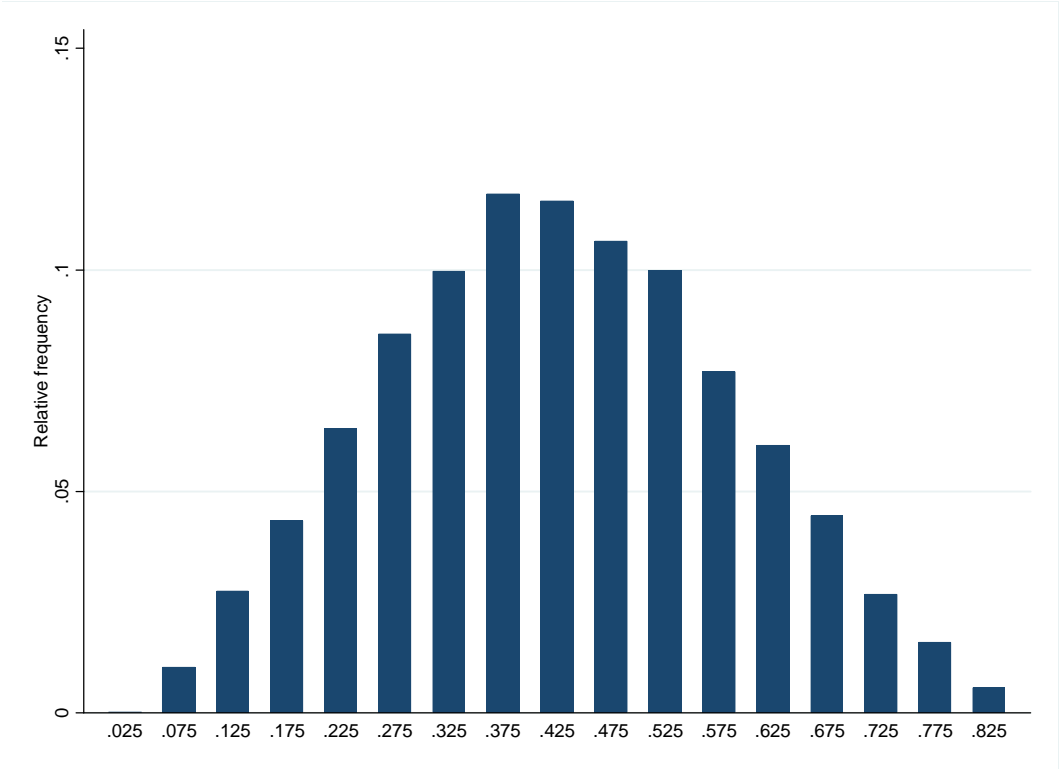


Figure 14: Predicted change in market value, as a function of the observed vote share, using the simulated data

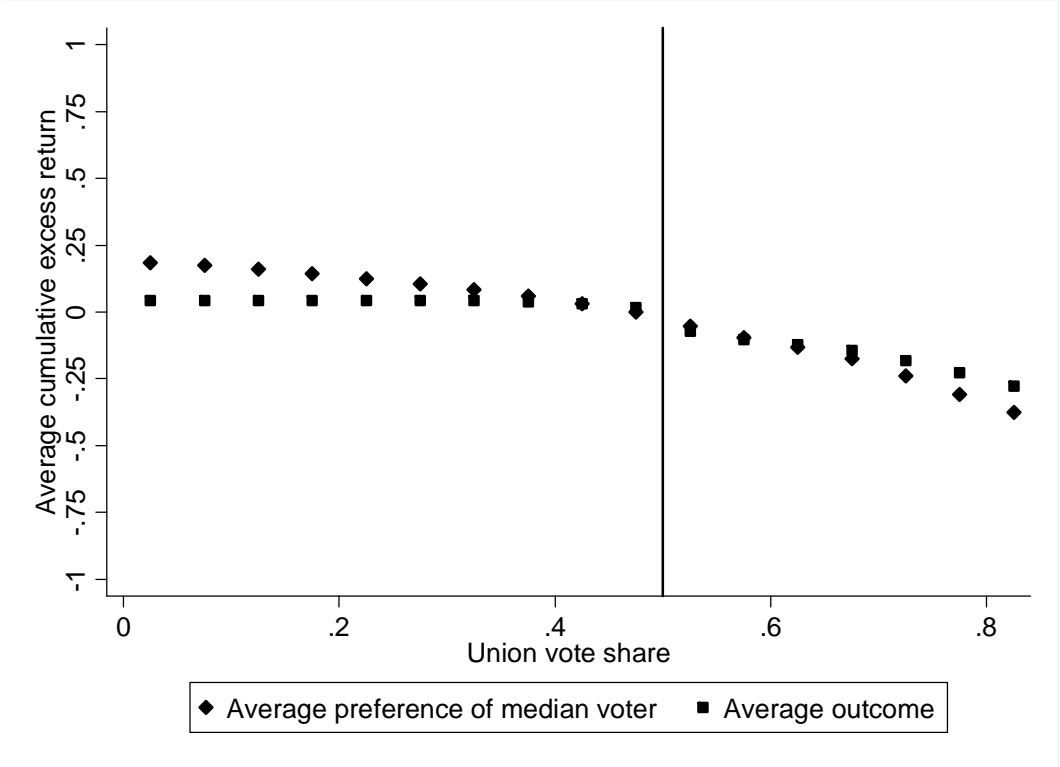


Table 1: Summary Statistics

	At least 5% of workforce voting		Less than 5% of workforce voting	
	Union wins	Union loses	Union wins	Union loses
Number of elections	414	1022	1163	2682
Vote share for union	0.62 [0.11]	0.35 [0.10]	0.64 [0.13]	0.35 [0.10]
Number of voters	449.1 [534.9]	454.2 [558.5]	276.5 [263.4]	297.6 [301.6]
Number eligible	496.0 [649.3]	494.0 [638.9]	286.4 [286.1]	317.9 [330.4]
Fraction of employees voting	0.21 [0.21]	0.23 [0.21]	0.01 [0.01]	0.01 [0.01]
Year of election	1975.2 [9.17]	1976.9 [9.11]	1974.9 [9.24]	1976.6 [9.42]
Fraction in Manufacturing	0.78	0.75	0.79	0.81
Number of employees	3813.3 [5377.5]	3430.8 [5195.4]	68468.6 [134336.5]	75284.6 [123610]
Market Value (CRSP)	353.8 [880.3]	330.9 [783.8]	4734.1 [10,547]	6350 [13,660]
Market Value (Compustat)	308.7 [614.9] {0.34}	329.80 [799.0] {0.33}	6334.1 [13372.0] {0.76}	7580.9 [16,343.1] {0.78}
Shareholder equity	242.6 [433.0] {0.34}	233.2 [497.7] {0.31}	4991.7 [13859.3] {0.77}	4479.8 [9432.4] {0.77}
Total Assets	588.4 [1243.3] {0.37}	683.8 [1876.5] {0.31}	13974.4 [36396.5] {0.78}	14164.9 [33308.0] {0.79}
Total Liabilities/Total Assets	0.060 [0.118] {0.58}	0.068 [0.162] {0.44}	0.062 [0.112] {0.55}	0.071 [0.183] {0.60}
Pretax income	15.11 [46.97] {0.35}	9.76 [41.9] {0.36}	249.3 [731.7] {0.74}	276.3 [731.1] {0.74}

Table 1 (cont.)

	At least 5% of workforce voting		Less than 5% of workforce voting	
	Union victory (UV firms)	Union loss (UL firms)	Union victory (UV firms)	Union loss (UL firms)
Sales	160.7 [238.7] {0.33}	144.2 [225.1] {0.31}	2693.5 [5306.3] {0.80}	3041.2 [5534.1] {0.80}
Tobin's Q	1.17 [0.658] {0.44}	1.30 [0.694] {0.50}	1.29 [0.642] {0.48}	1.31 [0.625] {0.56}
Profit margin	0.069 [0.119] {0.44}	0.060 [0.167] {0.50}	0.084 [0.073] {0.46}	0.084 [0.074] {0.52}
Income/Employees	0.004 [0.023] {0.41}	0.003 [0.008] {0.49}	0.004 [0.006] {0.48}	0.004 [0.007] {0.51}
Return on Assets	0.013 [0.051] {0.48}	0.022 [0.037] {0.53}	0.026 [0.023] {0.47}	0.027 [0.027] {0.25}
Dividend Ratio	0.633 [3.42] {0.44}	0.259 [1.100] {0.50}	1.15 [6.99] {0.58}	0.941 [11.02] {0.59}
Fraction of stocks delisted	0.10	0.08	0.049	0.028

Notes: Summary statistics are based on the NLRB election, Compustat, and CRSP data. Standard deviations are in brackets. For Compustat variables, the average percentile rank, relative to all Compustat companies in the year and quarter of the election, are in braces. Market value, shareholder equity, total assets, pretax income, and sales are in millions of dollars. Summary statistics for market value are derived from both the CRSP and Compustat databases. These measures differ because there are more missing values in the Compustat database. Fraction of stocks delisted is computed as the fraction of stocks with a non-missing delisting return in a two year window surrounding the NLRB case closure month. Profit margin = pre-tax income/sales. Dividend ratio = dividends/pre-tax income.

Table 2: Estimates of post-election cumulative abnormal returns

	(1)	(2)	(3)
	Size-matched benchmark	Size x industry- matched benchmark	Broad-market benchmark
Panel A: Union Victory			
ACAR(0,24)	-0.092 (0.028)	-0.096 (0.028)	-0.103 (0.029)
ACAR (-24,-4)	-0.010 (0.020)	-0.009 (0.020)	-0.010 (0.020)
Adjusted- ACAR (0,24)	-0.100 (0.042)	-0.103 (0.042)	-0.111 (0.042)
Panel B: Union Loss			
ACAR (0,24)	0.029 (0.021)	0.020 (0.020)	0.016 (0.022)
ACAR (-24,-4)	0.034 (0.022)	0.004 (0.014)	-0.009 (0.014)
Adjusted- ACAR (0,24)	0.029 (0.030)	0.016 (0.032)	0.028 (0.031)

Notes: ACAR(X,Y) denotes the average cumulative abnormal return from month X to month Y relative to the NLRB case closure month. There are 414 elections in the sample in Panel A, and 1022 elections in Panel B. See Section 3.2 for details on the construction of the benchmark portfolios and estimation.

Table 3: Relating post-event cumulative abnormal returns to the share of the workforce in the bargaining unit

	(1) ACAR(0,24)	(2) Adjusted- ACAR (0,24)
Panel A: Union win		
Constant	0.03 (0.01)	-0.01 (0.02)
Share of workforce in bargaining unit	-0.31 (0.08)	-0.26 (0.13)
Observations	1577	1577
Panel B: Union loss		
Constant	0.03 (0.01)	0.01 (0.01)
Share of workforce in bargaining unit	0.06 (0.05)	0.17 (0.08)
Observations	3704	3704

Note: Sample includes all NLRB elections that we matched to publicly traded firms. See note to Figure 2 for details on how ACAR(0,24) and Adjusted-ACAR (0,24) were constructed.

Table 4: Fama-French calendar time Portfolio estimates

Panel A: Union Win Portfolio ($\geq 5\%$ sample)

Event-window:	Alpha	MKTRF	HML	SMB
(0,24)	-0.0051 (0.0014)	0.909 (0.035)	0.421 (0.054)	1.12 (0.048)
(-24,-4)	-0.0015 (0.0015)	0.996 (0.038)	0.487 (0.062)	1.14 (0.054)

Panel B: Union Loss Portfolio ($\geq 5\%$ sample)

Event-window:	Alpha	MKTRF	HML	SMB
(0,24)	-0.0001 (0.0017)	1.04 (0.031)	0.469 (0.048)	1.01 (0.043)
(-24,-4)	-0.0005 (0.0011)	0.970 (0.020)	0.264 (0.040)	1.04 (0.035)

Panel C: Union Win Portfolio ($< 5\%$ sample)

Event-window:	Alpha	MKTRF	HML	SMB
(0,24)	0.0010 (0.0014)	1.10 (0.037)	0.395 (0.055)	0.222 (0.048)
(-24,-4)	-0.0009 (0.001)	1.10 (0.026)	0.283 (0.042)	0.373 (0.037)

Panel D: Union Loss Portfolio ($< 5\%$ sample)

Event-window:	Alpha	MKTRF	HML	SMB
(0,24)	-0.0015 (0.0007)	1.14 (0.023)	0.509 (0.035)	0.212 (0.030)
(-24,-4)	-0.0009 (0.0008)	1.10 (0.017)	0.220 (0.031)	0.335 (0.027)

Note: The “ $\geq 5\%$ sample” consists of elections where at least 5% of the firm’s workforce voted. The “ $< 5\%$ sample” corresponds to elections where less than 5% of the firm’s workforce voted. MKTRF is the monthly return of the CRSP value-weighted NYSE/AMEX/NASDAQ broad market index, SMB is the monthly return on the zero investment portfolio for the common size factor in stock returns, and HML is the monthly return on the zero investment portfolio for the common book-to-market equity factor in stock returns. The unit of observation is the calendar month. Observations are weighted by the number of firms in the event-window.

Table 5: Compustat Analysis

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	ln(Assets)	ln(Shareholder equity)	ln(PPE)	ln(Sales)	ln(pretax income)	Dividend Ratio	Profit margin	ROA	Tobin's Q	Liabilities/ Assets
post	0.150 (0.021)	0.106 (0.020)	0.137 (0.028)	0.132 (0.019)	0.168 (0.031)	-0.197 (0.118)	0.0001 (0.002)	-0.004 (0.001)	-0.054 (0.026)	-0.001 (0.005)
post x union win	-0.110 (0.037)	-0.098 (0.035)	-0.113 (0.048)	-0.077 (0.034)	-0.168 (0.062)	0.045 (0.263)	-0.005 (0.003)	-0.001 (0.002)	0.031 (0.038)	0.003 (0.008)
Observations	14,319	16,220	14,223	17,028	14,042	6,127	14,585	13,960	14,035	5,791
R-squared	0.97	0.95	0.96	0.94	0.75	0.084	0.64	0.32	0.66	0.28

Notes: Variables are derived from Compustat data; 1973-1999. Each column corresponds to a different model estimated using OLS. Standard errors clustered on election are in parentheses. Observations are event quarter \times firm cells. The dependent variables are demeaned, where the mean is taken over all non-missing observations in an election panel. Sample sizes vary due to the presence of missing values. PPE stands for plant, property, and equipment. ROA stands for return on assets.

Table 6: Cumulative abnormal returns and vote share

	CAR(0,24):				CAR(-24,-4):				Adjusted-CAR(0,24):			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Constant	0.029 (0.021)	-0.065 (0.030)	-0.075 (0.039)	-0.064 (0.035)	0.0003 (0.016)	-0.018 (0.025)	-0.029 (0.033)	-0.035 (0.030)	0.029 (0.030)	-0.024 (0.045)	-0.029 (0.058)	-0.002 (0.053)
Union won	-0.121 (0.035)	0.048 (0.054)	0.080 (0.066)	0.049 (0.053)	0.003 (0.028)	0.037 (0.046)	0.021 (0.057)	0.027 (0.047)	-0.129 (0.051)	-0.032 (0.081)	0.035 (0.101)	-0.019 (0.082)
Union won ×vote share				-0.016 (0.321)				0.332 (0.255)				-0.439 (0.469)
vote share		-0.616 (0.160)		-0.610 (0.207)		-0.123 (0.126)		-0.235 (0.162)		-0.353 (0.233)		-0.205 (0.301)
p(vote share)			X				X				X	
Observations	1436	1436	1436	1436	1436	1436	1436	1436	1436	1436	1436	1436

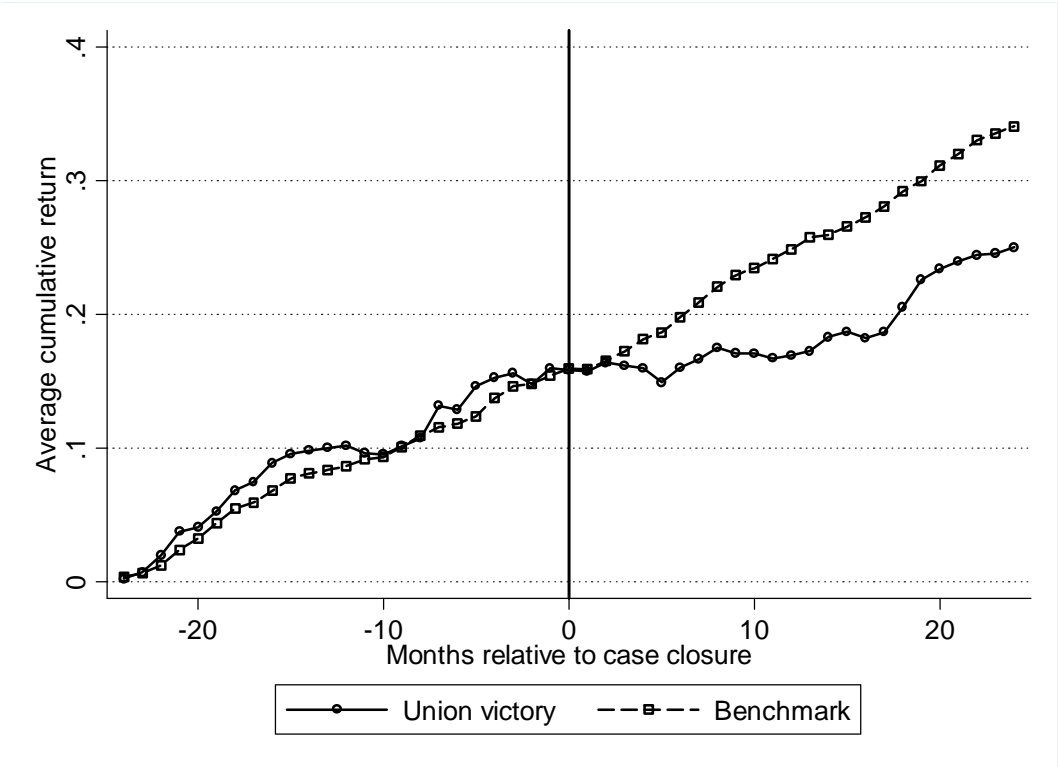
Note: Robust standard errors are in parentheses. The sample consists of all elections where at least 5% of the workforce voted. The variable “vote share” denotes the union vote share, minus 0.5. Following Dinardo and Lee (2004), the vote share is aggregated to 20 discrete bins. The dependent variable is the cumulative abnormal return from months 0 to 24 relative to case closure (columns 1-4), the cumulative abnormal return from -24 through -4 months relative to case closure (columns 5-8), and the adjusted-cumulative abnormal return from 0 to 24 (columns 9-12),. See Section 4.4 for details on the construction of these variables. The term p(vote share) denotes a fourth-order polynomial in the union vote share.

Table 7: Simulated Market Value Changes Under Different Vote Threshold Scenarios

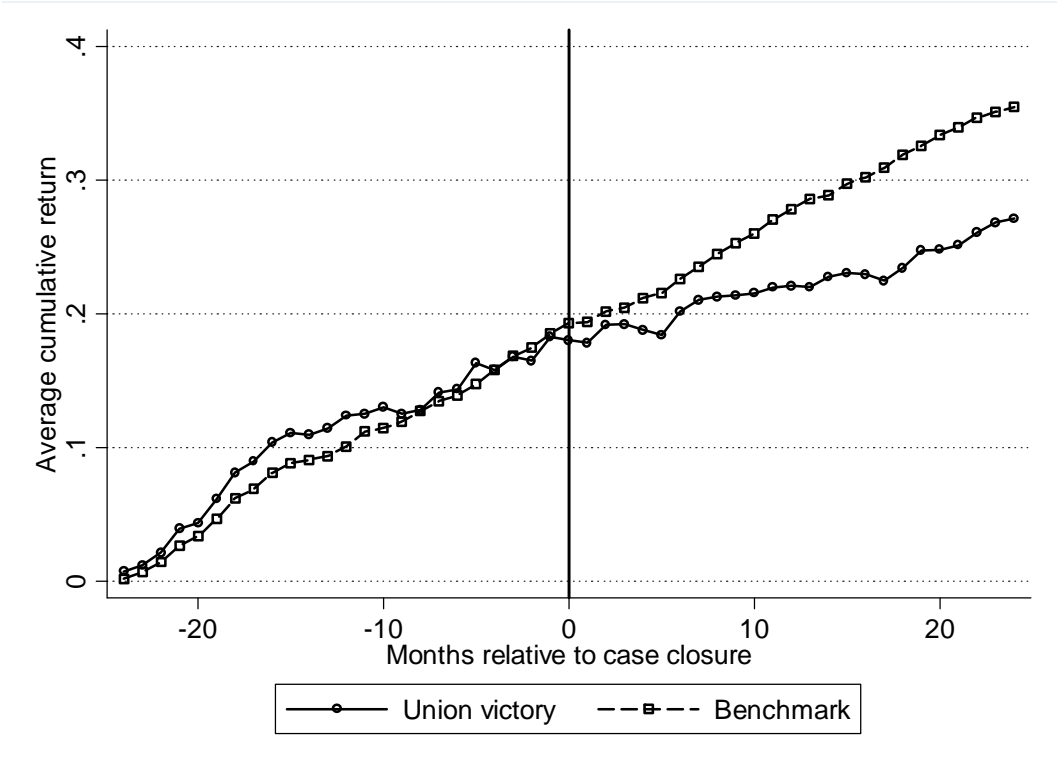
Group	Proportion	Threshold=0.50		Threshold=0.33		Threshold=0.25		Threshold=0.10	
		Effect	Win?	Effect	Win?	Effect	Win?	Effect	Win?
Overall	1.00	-0.015		-0.058		-0.083		-0.124	
Inframarginal Union Loss	0.01	0.042	NO	0.042	NO	0.042	NO	0.042	NO
Marginal Group 1	0.13	0.042	NO	0.042	NO	0.042	NO	-0.062	YES
Marginal Group 2	0.15	0.042	NO	0.040	NO	-0.041	YES	-0.074	YES
Marginal Group 3	0.37	0.031	NO	-0.050	YES	-0.069	YES	-0.098	YES
Inframarginal union win	0.33	-0.117	YES	-0.153	YES	-0.171	YES	-0.205	YES
Simulated Discontinuity		-0.080		-0.081		-0.082		-0.108	

Note: Each column represents a different scenario for necessary union vote share necessary for certification. The population is split into five groups (represented by rows). "Inframarginal loss" denotes firms that would not be unionized under any scenario. "Inframarginal win" denotes firms that would be unionized under all scenarios. "Marginal Groups" denote firms in which unions would lose under one or more scenarios, but would win with a lower threshold (as indicated under the sub-column "Win?"). e.g. Marginal Group 3 comprises of firms where the union vote is marginally below the 50 percent when the threshold is 0.50; they would become unionized in any of the other scenarios. "Simulated Discontinuity" is the RD estimate - via a 4th order polynomial regression -- using the simulated data.

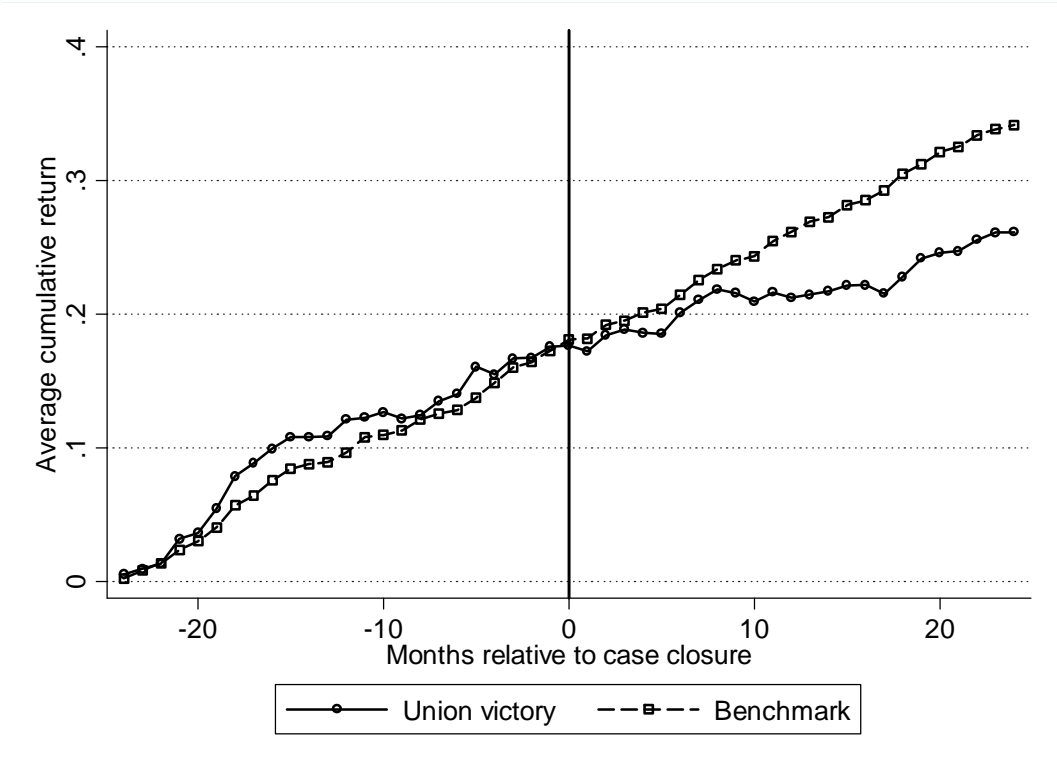
Appendix Figure 1: Average cumulative returns of union victory firms and of the sized-matched benchmark; non-imputed data



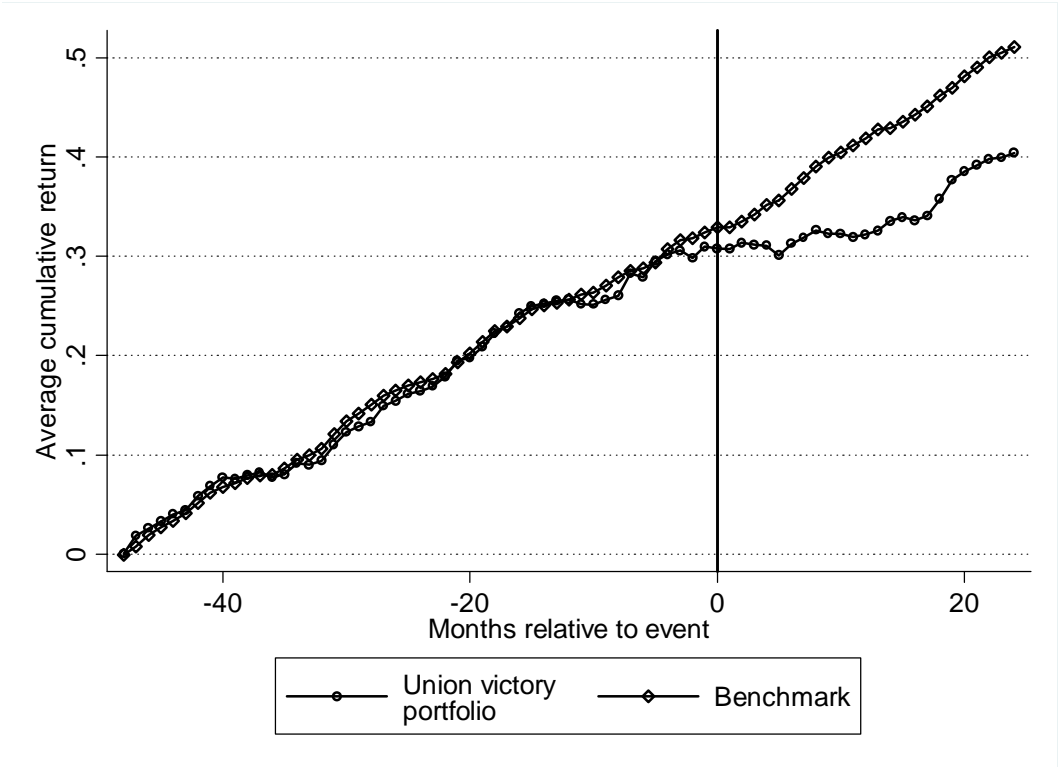
Appendix Figure 2: Average cumulative returns of union victory firms and of the sized-matched benchmark; Balanced panel



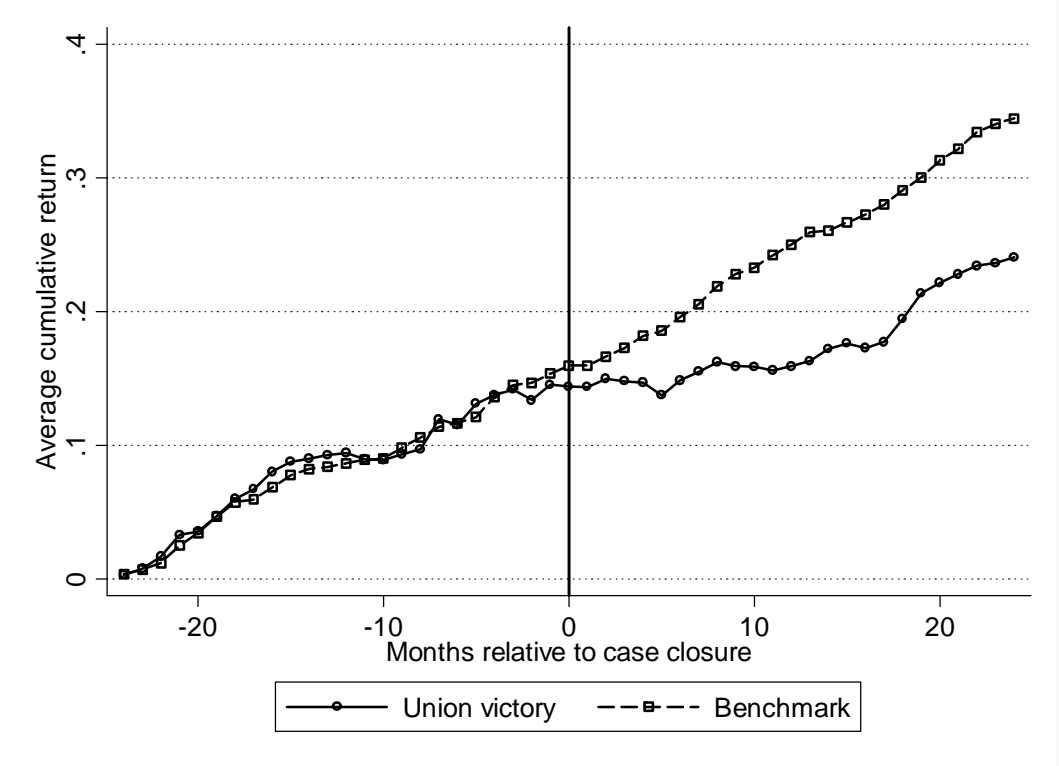
Appendix Figure 3: Average cumulative returns of union victory firms and of the sized-matched benchmark; eliminate 5% most positive and 5% most negative post-event abnormal return elections



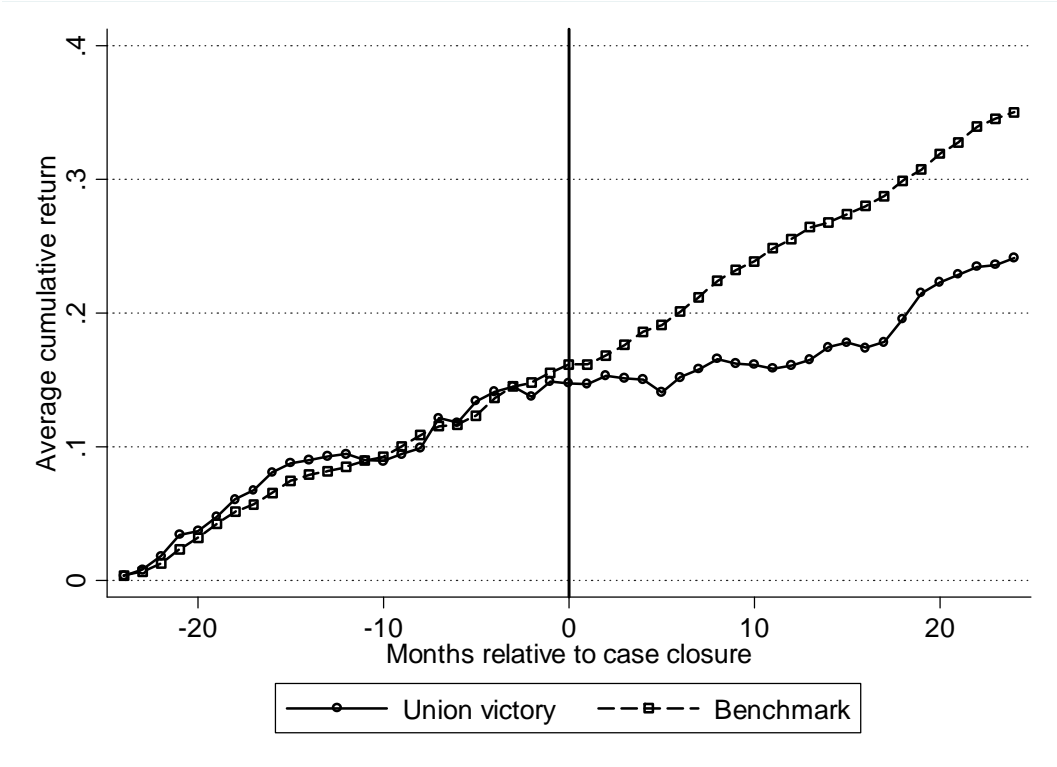
Appendix Figure 4: Average cumulative returns of union victory firms and of the sized-matched benchmark; Four year pre-event window



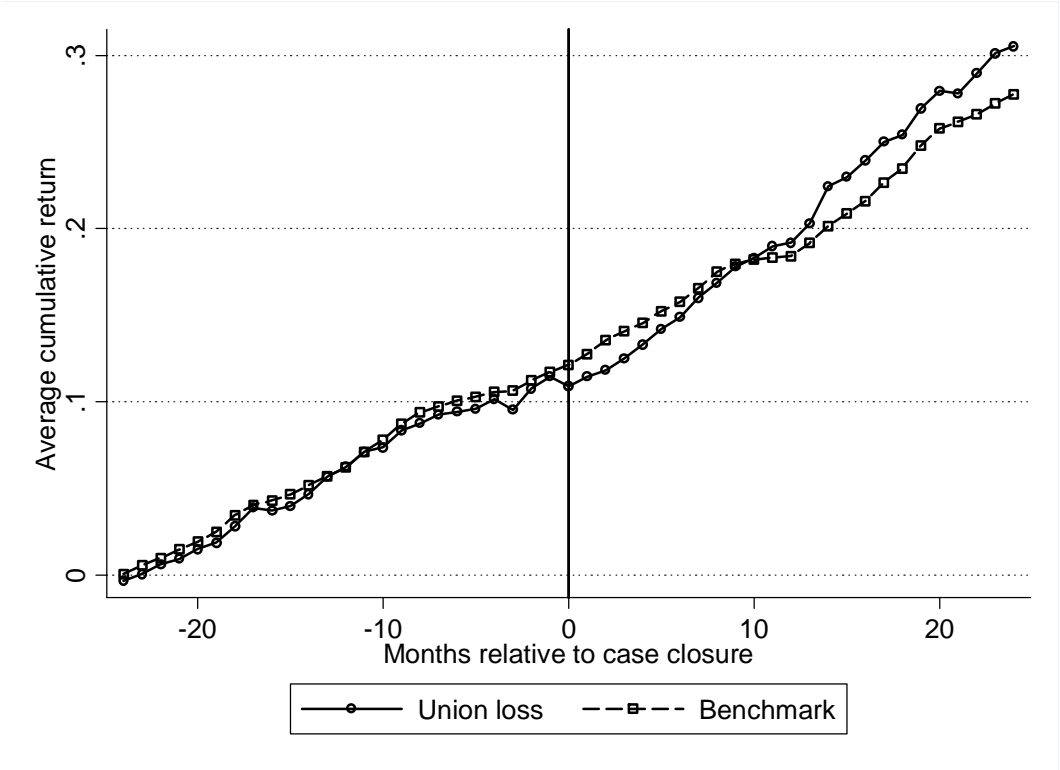
Appendix Figure 5: Average cumulative returns of union victory firms and of the industry×sized-matched benchmark



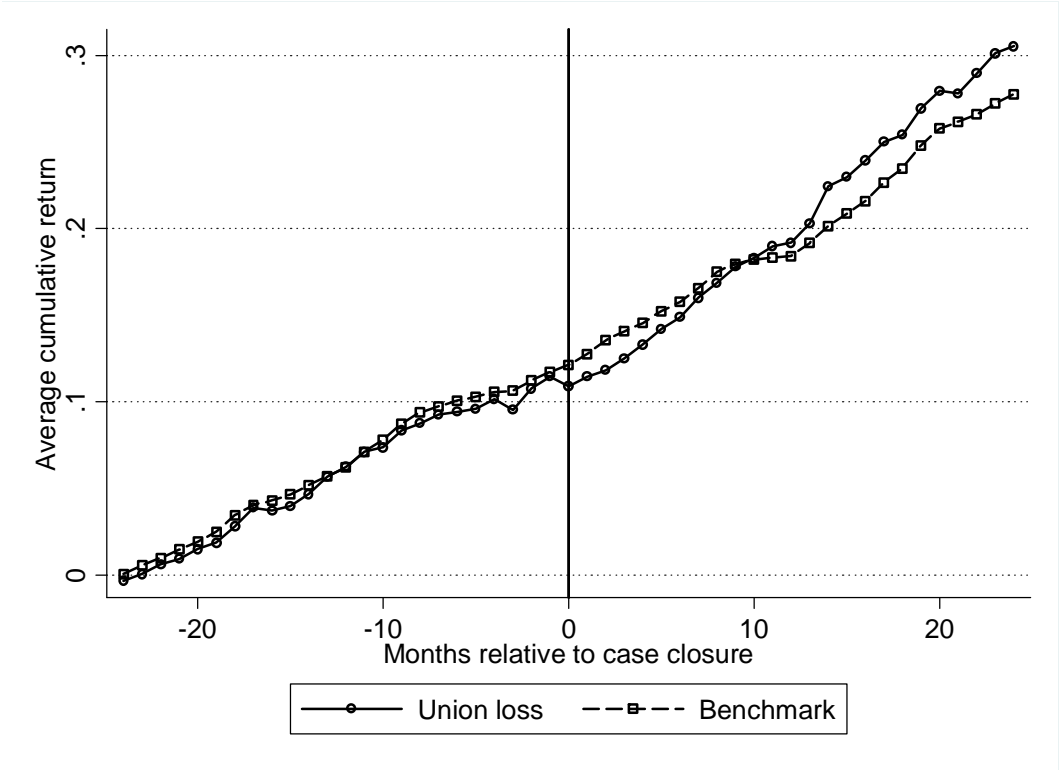
Appendix Figure 6: Average cumulative returns of union victory firms and of the CRSP equally-weighted index benchmark



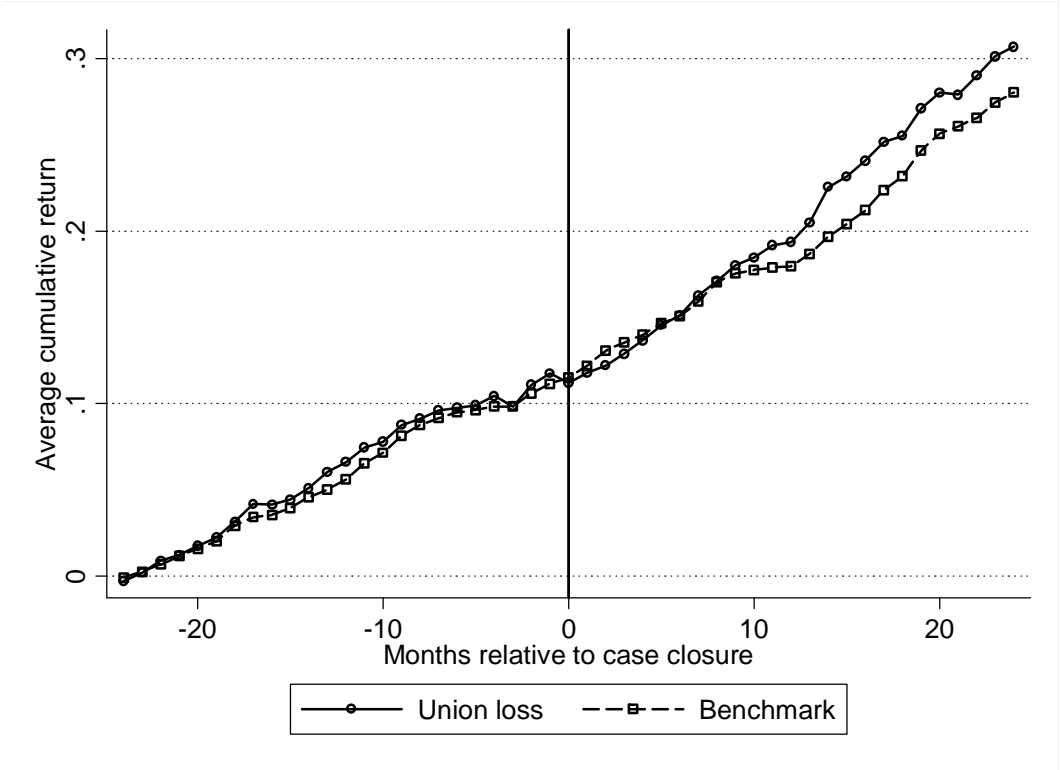
Appendix Figure 7: Average cumulative returns of union loss firms and of the sized-matched benchmark; non-imputed data



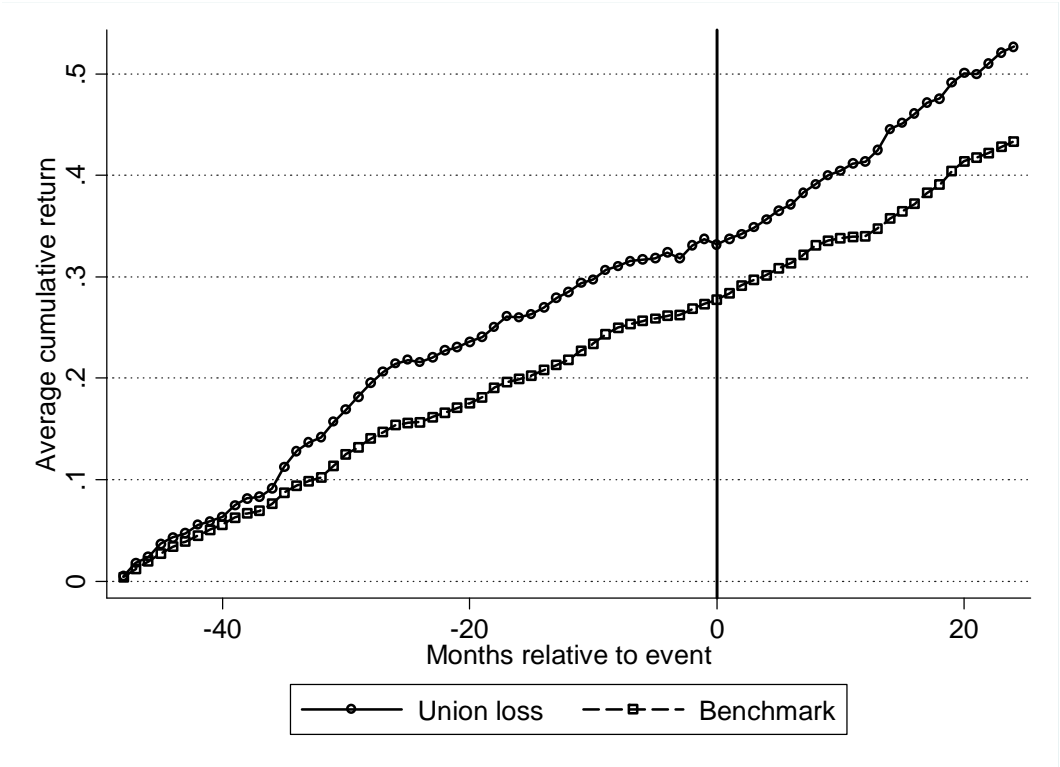
Appendix Figure 8: Average cumulative returns of union loss firms and of the sized-matched benchmark; Balanced panel



Appendix Figure 9: Average cumulative returns of union loss firms and of the sized-matched benchmark; eliminate 5% most positive and 5% most negative post-event abnormal return elections



Appendix Figure 10: Average cumulative returns of union loss firms and of the sized-matched benchmark; Four year pre-event window



Appendix Figure 11: Average cumulative returns of union loss firms and of the industry×sized-matched benchmark



Appendix Figure 12: Average cumulative returns of union loss firms and of the CRSP equally-weighted index benchmark

