PRELIMINARY
DO NOT CIRCULATE
COMMENTS WELCOME

Long-Run Impacts of Unions on Firms: New Evidence from Financial Markets, 1961-1999[1]

David S. Lee
Princeton University and NBER

Alexandre Mas
UC Berkeley and NBER

January 2008

Abstract

We estimate the effect of new unionization on the equity value of firms over the 1961-1999 period using a newly available sample of National Labor Relations Board (NLRB) representation elections matched to stock market data. As in Dinardo and Lee (2004), the point estimates from a regression-discontinuity design – where we compare the stock market impact of close union election wins to close losses – imply that unions are not costly to firms. However, event-study estimates show that new unionization is associated with at least -10% excess returns, equivalent to $40,500 per worker unionized. We find a negative gradient in the cumulative excess returns by union vote share, thus allowing us to reconcile these two seemingly contradictory findings. The lack of a discontinuity in excess returns about the 50% union vote share threshold is a result of firms experiencing increasingly worse financial performance when workers are more supportive of unionization. When viewed through the lens of a "median voter" model of endogenous union determination, the evidence is consistent with firms that are unresponsive to the threat of new unionization, and unions that moderate their demands in order to gain electoral advantage.

[1]We thank Jonathan Berk, David Card, Jesse Rothstein, Eric Verhoogen, Hans-Joachim Voth, and seminar participants at Stanford for helpful suggestions. Emily Buchsbaum, Mariana Carrera, Briallen Hopper, Sanny Liao, and Andrew Shelton provided outstanding research assistance.

"[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, 1950[1]

"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, 1984[2]

1 Introduction

It is uncontroversial that employers oppose unions, viewing them as a threat to profitability. An example that has received recent national attention is Wal-Mart's efforts in resisting 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 that 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 also easy to find isolated cases that confirm the fears of employers like Wal-Mart. As an 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 a more than 2 to 1 margin to organize as a local chapter of Union of Needletrades, Industrial, and Textile Employees (UNITE). The stock market appears to have punished NLS a severe, though perhaps not swift, fashion. Figure 1 shows the cumulative return of NLS' stock beginning two years prior to the election through two years after, as well as the cumulative return of a broad market index over the same period. Prior to the election, NLS and the market returns 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%, while the broad market index had increased by 50%.

How general is this phenomenon? Is National Linen Service Corp. the exception or the rule? Despite an enormous literature that has documented numerous aspects of unions and their role in the labor market, the magnitude of the average effect of unions on firm performance throughout the

[1] Quoted from Friedman (1950).
[2] Quoted from Freeman and Medoff (1984).
[3] Quoted in Featherstone (2004).

economy remains somewhat unclear.

Empirically, there are at least three reasons why measuring these effects has proven to be exceptionally challenging. First, large-scale establishment or firm-level micro-data with the relevant information on the extent of unionization is not readily available. Second, even when such data are available, omitted variables and the endogeneity of unionization at the level of the firm makes it difficult to separate causal effects from other unobserved confounding factors. Finally, it is even more difficult to meet these two challenges with data that can plausibly be representative of the population of unionized companies in the United States.

Theoretically, it is not even obvious to what degree unions should affect firms. One view, articulated by Milton Friedman, is that workers would reject substantially above-market wages, knowing full well that such wages could have adverse effects on job security.[4] Unions, taking these considerations into account, would then therefore tend to moderate wage demands. Moreover, firms may respond to a unionization threat by conceding higher wages and better working conditions. These forces would tend to reduce the gap in compensation and working conditions between union and non-union workforces, at least in those situations where there is a threat of unionization. These ideas may help explain the results of Dinardo and Lee (2004), who find generally small differences in wages, employment, and output between unionized and otherwise comparable non-unionized workplaces.

In this paper we formalize these concepts with a simple model of endogenous unionization. This framework offers predictions for the various kinds of union impacts that can exist when firms can alter working conditions to prevent the union from successfully organizing, and when unions can elect to moderate their demands in order to attract more support. Specifically, we consider a two-party model of electoral competition, where the firm and the union are seeking to win the sympathies of the "median" voter in an NLRB election. As is standard in this class of models, the parties, although having opposing interests, may be forced to propose a level of compensation, and the accompanying risk of job loss, that is closer to that preferred by the median voter. The model implies that the regression-discontinuity design estimate of the unionization effect identifies the gap between the union and firm's proposal when the median voter is indifferent between the two. Depending on how aggressively firms and unions court voters, this gap could be close to zero, even if unions affect the

[4]Which led Friedman to say the above quote.

profitability of firms on average.

In the empirical analysis, we first assess the extent to which the pattern in Figure 1 is a more general phenomenon among publicly-traded firms, so as to measure an average overall effect of unionization. To do so, we use a sample frame that is the universe of all firms at which NLRB union representation elections took place between 1961-1999. Since most unionized workplaces in the U.S. come into existence via a secret-ballot election on the question of representation, this sample frame can provide a reasonable representation of the stock of unions today and particularly into the future.

The most distinctive feature of our data – crucial for our research design – is its long panel, up to 48 months before and after the election, of high-frequency data on stock market returns for these firms. We use a standard event-study approach to analyzing the stock market reaction to union victories, but we are able to use the pre-event data to test the adequacy of the market models and specifications 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 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, equivalent to $40,500 per worker unionized. The evidence is compelling in the sense that we find that these firms' average returns are quite close to the predicted returns for every month leading up to the election. Precisely at the time of the election, the actual and counterfactual returns diverge. The results are robust to a number of different specifications, and notably, 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.

We then employ a regression discontinuity design to where we compare close union victories to close union losses. As in Dinardo and Lee (2004), we find no evidence of a significant discontinuity in market returns with respect to the vote share for the union. The RD point estimates show, if anything, 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, and even some evidence of positive abnormal returns for cases where the union lost by a large margin.

Viewed through the lens of the theoretical framework, the pattern of the effects is consistent

with firms that do not (or cannot) compete for voters by changing working conditions. The evidence is also consistent with unions that modify their positions in order to gain electoral advantage, in a way that is reminiscent of Friedman's view. The estimates imply that voters have strong desire for higher wages in only a relatively small number of elections. However, because the firm is largely unresponsive to voters, these elections are associated with a large pro-union vote share, and a substantial reduction in equity value of the firm. The results imply that whatever salary increase is enacted following an NLRB election, it should rarely be larger than the amount desired by the median voter, regardless of who wins the election. Therefore, new unionization should lead to relatively limited employment effects, even when the firm experiences a considerable reduction in value.

The paper is organized as follows. Section 2 briefly highlights what is known from the literature, and how our study relates to that literature. Our theoretical framework for precisely defining the different kinds of union effects is presented in Section 3. We then provide some institutional details in Section 4 that are relevant to our research design, which is also described, along with our data. We present and discuss the results in Section 5. Section 6 concludes.

2 Existing Literature and Background

In this section we provide a brief overview of the literature that is most related 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). Most of these studies 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, this involves following workers in longitudinal datasets, 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 nonunion." Following this point, DiNardo and Lee (2004) make clear that the effect of *unionization* (changing a workplace from non-union to unionized) is distinct from the effect of moving a *worker* from a unionized workplace to a non-unionized workplace. In particular,

a "typical" unionized workplace may differ from a "typical" non-unionized workplace along a number of different dimensions (e.g. geography, firm size, industry) which themselves may independently influence wage levels. We therefore view this well-established literature as being fundamentally unable to account for the selection of unionism at the firm- or establishment-level, and therefore potentially estimating something quite distinct from the causal effect of unionization.

Next, there is a literature that does utilize firm- or establishment-level data containing information on union status. As pointed out in Hirsch (2007), a recent review of this evidence, there are a number of important reasons why caution is warranted in 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 to this is a potential endogeneity problem, whereby unions may specifically target a highly profitable firm for organization. Alternatively, it may be poorly managed, and hence low-performing firms, that lead to the demand for worker representation. Examples of studies that implicitly rely 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 that 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), hospital and nursing homes in Allen (1986a), construction industry in Allen (1986b), and sawmills in Mitchell and Stone (1992). It is difficult to extrapolate from these studies on productivity to a broader, representative cross-section of firms in the U.S. Indeed, our analysis is largely motivated by the notion that it might not be too difficult to find particular incidents and or companies, or industries where unions have imposed large costs on firms, but the question is to what extent isolated examples such as that illustrated in Figure 1 generalizes to a broader population of interest.

Finally, there are three particular studies that we consider to be the most related to our analysis – that of 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) – who utilize a "fixed effects" approach to examine the impact of a successful union organizing campaign with establishment level panel data from the Longitudinal Research Database (LRD) – 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 percent 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]

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 the 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, a decline in market value of about 7 percent between the 12th and 7th months preceding unionization. Ruback and Zimmerman (1984) argue that the timing of this substantial decline is unlikely to indicate anticipation of the outcome of the election, but have no explanation for this important decline.[8] This therefore raises the question as to whether to attribute the further decline in the stock market valuation to unionization or instead to whatever factors that might have led to the pre-election trend.

In our analysis, we address these ambiguities by taking advantage of a very long panel of monthly

[5]Another study in the spirit of a "before-after" design is that of Freeman and Kleiner (1990), in which 203 establishments were surveyed before and after NLRB elections took place, and were compared to a 161 "control" firms, where elections did not take place. There, there was only one period before and one period after, and so testing the over-identifying restrictions of the difference-in-difference design was not possible.

[6]There are a number of other studies that examine various aspects of unions as seen through stock market reactions. They generally 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.

[7]Specifically, their main estimate of -3.84 is computed by taking the 1-month change associated with the petition date and adding it to the 1-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.

[8]Specifically, 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 prepetition 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."

data on stock returns, using an arguably more disciplined approach to modeling the counterfactual "no union" state. Specifically, using standard stock market event-study approaches, we use stock price data strictly earlier than the 24 months prior to the election date in order to estimate the relation between the overall market index and the price of our "portfolio" of firms facing organizing drives. We then use the data between 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, we then consider to be invalid any estimates we obtain from the post-event data.[9] This can be recognized as directly parallel to the conventional testing of over-identifying restrictions for "difference-in-difference" modeling in program evaluation research in labor economics.[10] Furthermore, we track abnormal returns throughout a long time period (at least 24 months) after the unionization. Strictly speaking, perfectly efficient financial markets predict that any changes in valuation caused by the outcome of the election will be instantaneously capitalized into the stock price, the moment the outcome of the election is known. Nevertheless, our approach is to rely less on the assumption of instantaneous adjustment, and thus allowing some time for the market to adjust, by examining the patterns of returns for many months following unionization.

The final study to which the present study relates is the regression discontinuity analysis of union elections, using the data from the LRD, as in DiNardo and Lee (2004). There, they exploit the "near-experiment" generated by secret ballot elections, comparing establishments where unions became recognized by a close margin of the vote with those 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 even 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 time – perhaps 10 more more years – for unions to establish support within the workplace to have enough bargaining power to negotiate for substantially higher wages. Second, there are potentially other costs imposed by unions that are not measured in the LRD, such as the use of seniority rules, work rules, grievance procedures, and other working conditions specified in union

[9] An 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 nonetheless take a more conservative approach and simply conclude from any significant divergence as evidence of a mis-specified model.

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

[11] Interestingly, 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 that found in cross-sectional worker-level studies.

contracts. Our approach of examining the effect of stock market valuation can in principle address both of these concerns: if the market correctly prices the firm, it should capture the entire effect of the presence of the union, and "real" effects that might occur many years into the future should be capitalized into the stock market valuation of the firm in the short-run.

A final important limitation is that the RD analysis, by estimating a discontinuity in the relation between wages and the vote share at the 50 percent threshold, can only estimated a weighted average treatment effect, where the weights are proportional to the *ex ante* likelihood that an election was to be "close". That is, among the observed close elections, a disproportionately small number of them would have had the fundamentals of strong union support, for example. The RD is fundamentally unable to provide a counterfactual for the set of elections in which the 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 all of the results.

3 Conceptual Framework

Here we summarize a simple conceptual framework that we believe is a reasonable way to describe the interaction between the management, union, and the workers at a firm in the events leading up to potential unionization.[12] We use this framework to precisely define the different kinds of union effects, and the various counterfactuals needed for our empirical analysis. Specifically, our goal is to develop a model that allows us to translate the observed relationships in the data into the parameters governing the behavior and composition of unions, firms, and worker during representation elections.

We consider a variant of "median voter model" of endogenous union determination. While the median voter-type model has previously been considered in the theoretical literature on unions (see Atherton (1973), Farber (1978) and Booth (1995)), we recognize there are likely to be a virtually unlimited number of different ways to model union elections, bargaining, and union threat. Our goal, therefore, is to describe a theoretical framework that has as its backbone a few elements that we believe any sensible model of the process would include.

The basic idea of the model is that the firm and the union each submit a proposal for vote by

[12] Some of these ideas are discussed, but not fully or formally developed, in DiNardo and Lee (2002, 2004).

the workers in a secret ballot election. For purposes of exposition, suppose the key issue is the wage. Both the union and firm try to persuade the workers to vote for or against union representation by proposing different wage levels. The voters will select a proposal that maximize their utility, taking into account that too high a wage increases the likelihood of job loss as a result of the firm moving up the labor demand curve. We formalize the model below.

3.1 Setup

This model has three agents who are involved in a recognition election: workers, the firm, and the union.

- **Workers.** Workers are considered to maximize their own individual utility, which could be, for simplicity, lifetime wealth. In this environment, each worker will face the decision to either vote for or against union recognition. In doing so, the forward-looking workers each compares the anticipated outcome if the union wins to the expected outcome if the union loses. Reasons to vote for the union might include higher wages, benefits, better working conditions; reasons to vote against might include the potential responses of the firms (e.g. lower employment). We assume that each worker has an "ideal" bargaining unit-level unit wage increase, with the median preference in the establishment denoted μ. To simplify the problem, we assume that positive and negative deviations of the enacted wage from this ideal wage enter symmetrically into workers loss functions. We assume that there is heterogeneity in μ across establishments, perhaps depending on the firm's elasticity of labor demand.

- **Union.** The unions choose a proposed pay raise, denoted y_U, which is the percent increase from the pre-election wage. As we only consider situations whereby the union proposes a larger salary increase than the firm, if the union raises its offer, it is a more extreme offer, and all those who already were in favor of the lower y_U will not change their vote, while some on the margin will be swayed from switching from the union to the management. We model the union's offer as: $y_U = \alpha + \beta\mu$.[13,14] The parameters α and β can be thought of as an equilibrium

[13] More precisely, the union offers $y_U = \max(\alpha + \beta\mu, 0)$.

[14] We recognize that, strictly speaking, this model is deterministic in the sense that vote share and election outcomes are perfectly predictable given the parameters of the model. This raises the question of why unions and firms would choose a proposal that is known to produce a loss in an election. A natural extension to the model is to include an unpredictable component to the vote share. We anticipate that many of the qualitative predictions of the model will withstand this modification. Exploring this issue more fully is a natural next step in our ongoing research.

outcome of the union's optimization problem. An example of such a problem would be if the union were to choose its proposal strategically, so as to make the best response given the firm's offer (which may also be chosen strategically), taking into account the trade-off between the probability of winning the election and a higher wage, which presumably benefits the union.[15] The parameter α describes the extent to which the union moderates its offer when the electorate has a low tolerance for higher wages. The parameter β controls the responsiveness of the union to the preferences of the workforce. For example, if $\beta = 0$, the union always offers α to the bargaining-unit. If $\beta = 1$, the union offers an αpremium relative to the median voter's ideal wage.

- **Management.** The firm proposes a salary increase, denoted $0 \leq y_M < y_U$. Raising its offer – and hence proposing something more moderate – will influence more workers to vote against the union. However, the firm's objective function penalizes higher wages. We parametrize the firm's offer as $y_M = \gamma + \delta\mu$.[16] The parameter δ affects the degree to which the firm responds to the union threat. If γ and δ both equal zero, the firm offer the prevailing wage, regardless of the preferences of the median voters. These parameter values are consistent with firms complying with the National Labor Relations Act (NLRA), which prohibits employers from increasing wages or benefits in response to a petition.[17]

3.2 Various Union Effects

The sequence of events are as follows: an election is announced, union and management propose pay raises, the election occurs and, depending on the true distribution of the worker preferences, either the firm or union wins and its offer is implemented. The impact of an election on the change in profits is $\Delta\pi = \Delta\pi(y_M)$ if the firm prevails, and $\Delta\pi = \Delta\pi(y_U)$ if the union prevails, where $\pi()$ is a decreasing function satisfying $\pi(0) = 0$ and $\pi(x) \leq 0 \ \forall \ x \geq 0$. The union wins the election whenever $|y_U - \mu| < |y_M - \mu| \Rightarrow \overline{y} \equiv (y_U + y_M)/2 < \mu$. The union wins the election (receives a greater than 50% vote share) whenever $\mu > \widetilde{\mu} \equiv \frac{\alpha+\gamma}{2-(\beta+\delta)}$. Note that, offers fixed, the union vote share is becoming larger as $\mu - \widetilde{\mu}$ becomes more positive.

[15] Such a game is isomorphic to the model of final-offer arbitration in Farber (1980).
[16] More precisely, the firm offers $y_M = \min(\max(\gamma + \delta\mu, 0), y_U)$.
[17] See LRM Packaging, Inc., 308 NLRB 829, 834 (1992)

We consider the thought experiment of repeated elections, with values of μ drawn from a distribution with pdf $h()$. We wish to explore how the parameters of this model translate into $E[\Delta\pi|v]$, a function that we will be estimating in this paper. Denoting the union vote share as v, we are interested in the following properties of $E[\Delta\pi|v]$:

$$B_1 = E[\Delta\pi|v \leq 0.5] = \int_{\mu \leq \widetilde{\mu}} \Delta\pi(\gamma + \delta\mu) \frac{h(\mu)}{\int_{\widehat{\mu} \leq \widetilde{\mu}} h(\mu) d\widehat{\mu}} d\mu,$$

$$B_2 = E[\Delta\pi|v > 0.5] = \int_{\mu > \widetilde{\mu}} \Delta\pi(\alpha + \beta\mu) \frac{h(\mu)}{\int_{\widehat{\mu} > \widetilde{\mu}} h(\mu) d\widehat{\mu}} d\mu,$$

$$B_3 = lim_{\Delta \to 0^+} E[\Delta\pi|v = 0.5 + \Delta] - lim_{\Delta \to 0^+} E[\Delta\pi|v = 0.5 - \Delta] = \Delta\pi(\alpha + \beta\widetilde{\mu}) - \Delta\pi(\gamma + \delta\widetilde{\mu}),$$

$$B_4 = lim_{\Delta \to 0^+} E[\Delta\pi|v = 0.5 - \Delta] = \Delta\pi(\gamma + \delta\widetilde{\mu}).$$

B_1 and B_2 are the average effect of a union loss and win, respectively. B_3 is the discontinuity in $E[\Delta\pi|v]$ at the 50% union vote share threshold. B_4 is the effect of a close union loss on profits.

In Figure 2 we graph the proposed union and firm offers, and their average, as functions of μ. The discontinuity–B_3–corresponds to the point $\widetilde{\mu}$, where the average of y_U and y_M intersects the 45 degree line. The size of the discontinuity can be inferred from the distance between the y_U and y_M curves at that point. Specifically, B_3 is more negative in that difference. The union wins when $\mu > \widetilde{\mu}$. Therefore, B_1 corresponds to the area under the y_M curve, to the left of $\widetilde{\mu}$. In particular, B_1 is the area under the y_M weighted by the conditional density of μ. Similarly, B_2 corresponds to the area under y_U, to the right of $\widetilde{\mu}$.

To illustrate how knowledge of B_1 through B_4 helps us infer the behavior governing unions and firms, as well as the underlying preferences of workers, we explore the implications of the model in several cases. We first consider the behavior of firms. From the above formulation, it can be seen that $B_1 = B_4 = 0$ implies that the firm does not respond to the threat of an election ($\gamma = 0$), nor to the preferences of voters ($\delta = 0$). $B_1 = B_4 < 0$ implies that the firm responds to the election by offering a pay raise ($\gamma > 0$), but not to the preferences of the median voter ($\delta = 0$), while

$B_4 < B_1 < 0$ implies that the firm responds to the election and to the preference of the median voter. Therefore, by estimating B_1 and B_4 we can test whether and how firms respond to the threat of unionization as a result of the election.

If the parameters α and β are such that $y_U > \mu$ over most values of $\mu > \widetilde{\mu}$, then the union tends to "overshoot" the preference of the median voter in the majority of elections. Such behavior is consistent with the union placing considerable weight in obtaining higher wages, at a cost of losing support of the workers. By contrast, if $y_U \leq \mu$ generally holds, the union is acting relatively conservatively, which would be consistent with the idea that the union reduces its demands in order to gain electoral support. What we infer about the union's behavior from realizations of B_2 and B_3 depends in part on whether the firm responds to threat. For example, if the firm abides by the NLRA, and does not increase wages in response to the election, then $B_2 = B_3$ implies that $\beta = 0$, in which case $B_2 = B_3 = \triangle\pi(\alpha)$. This scenario corresponds to the situation where the union offers a fixed wage increase that is independent of the preferences of voters. The degree to which the union is acting conservatively in such a scenario depends on α. $B_2 < B_3 < 0$ implies that unions are responsive to voter preferences, so that $\beta > 0$.[18] However, in order to infer to what extent the union is moderating its offers, we need to know both α and β, or to what extent, and over what range of values, is y_U greater or less than μ.

We will revisit this model in the empirical section of the paper. After discussing our empirical analysis, where we estimate B_1 through B_4, we discuss how we would interpret the findings that result from our research design in light of this framework.

4 Institutional Background, Data and Research Design

The National Labor Relations Act provides the legal framework through which most workers in the United States become unionized. Workers that organize into unions through the procedures specified by the NLRA are guaranteed the right to bargain collectively. There are several ways that 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,

[18] More generally, B_2 will be more negative than B_3 provided that $\partial\overline{y}/\partial\mu < 1$.

when a group of workers decides to organize into a union, they first petition the NLRB to hold a representation election. To be legally granted an election by the NLRB, the petition must be signed by at least 30% of the workforce, typically over no longer than a six month period. Once the NLRB determines the appropriate bargaining unit, the NLRB holds an election at the worksite. The union wins the election with a simple majority of support amongst the workers. Barring objections by the employer, a win means that 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.

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

4.1 The event study method

Our objective is to assess the impact of union certification elections on firm stock prices. Ideally, we would like to compare the firm's stock returns to those that the firm would have experienced in the absence of the event. The event study method provides a framework for predicting counterfactual returns based on aggregate market conditions. The approach involves comparing the returns of a set of firms around an event to the returns of a comparison portfolio, and other factors.

The return for security i in time t is defined as:

$$r_{it} = \frac{p_{it} - p_{it-1} + d_{it}}{p_{it-1}},$$

where p_{it} is the price of the portfolio i in period t, and d_{it} denotes dividends.[19] We group stocks into equally-weighted "union victory" and "union loss" portfolios defined over event-time.[20] For ease of exposition we will denote firms were the union won in an election as "union victory" (UV) firms, and those were the union lost as "union loss" (UL) firms. The return for the portfolio k in event-time τ is defined as:

$$r_{k\tau} = \frac{1}{n}\sum_{i=1}^{n} r_{i\tau},$$

where $r_{i\tau}$denotes the return of security i in the portfolio k, τ months relative to the event.

The excess return in a given month is the difference between the actual return and the predicted return:

$$\mathrm{ER}_{k\tau} = \ r_{k\tau} - E[r_{k\tau}|X_\tau].$$

If certification elections lead financial markets to revalue the firms, we would expect to observe non-zero average excess returns in an interval around the event date for the portfolio.

We compute the cumulative average excess return over "windows" of time around the event. These cumulative returns allow us to track the evolution of the performance of these firms over time. The cumulative excess return over a 24 month period, beginning at the event is computed as:

$$\mathrm{CER}(0,24)_k = \sum_{\tau=0}^{24} \mathrm{ER}_{k\tau},$$

where τ is the month relative to the event and $\mathrm{CER}(0,24)_k$ denotes the cumulative excess return from 0 through 24 months after the event.

Note that causal interpretation of the cumulative average excess returns requires that our estimate of $E[r_{k\tau}|X_\tau]$ be unbiased. Otherwise we will detect excess returns when in fact prices are simply reacting normally to aggregate market movements.

There are many ways to estimate $E[r_{k\tau}|X_t]$. The counterfactual is often modeled as the predicted returns given the observed realization of returns from a broad market index. For example, denoting

[19] The return is adjusted for splits. When appropriate, we include the CRSP delisting return.

[20] We find that estimating market-model parameters using monthly security-level data leads to imprecise and unstable estimates, in the sense that there is a relatively weak correlation between estimated market-model parameters estimated over consecutive intervals (for example, two years). However, the results are not sensitive to estimating the market-models using individual data.

$R_{m\tau}$ as the average return of a broad market index across firms in event date τ, one could estimate:

$$E[r_{k\tau}|X_\tau] = a + bR_{m\tau} \tag{1}$$

using historical data (Campbell et al., 1997). Alternatively, one could estimate a model with additional factors, such as the Fama and French three factor model with a momentum term:

$$E[r_{k\tau}|X_\tau] = a + b_1 R_{m\tau} + b_2 F_{2\tau} + b_3 F_{3\tau} + b_4 F_{4\tau}, \tag{2}$$

where $F_{2\tau}$, $F_{3\tau}$, $F_{4\tau}$ are size, book-to-market, and momentum factors, again averaged over event-time ((Fama and French, 1993), (Carhart, 1997)).[21]

While in several specifications we will be benchmarking returns using these commonly used market models, our analysis will emphasize transparent summaries of the data. We will begin by presenting the cumulative average excess return for the set of union victory and union loss forms, beginning two years prior to the election closing event, and tracking them for two years after. We will overlay each of these series on the average cumulative return of a broad market index over the same time period. By comparing the slopes of the average cumulative returns prior to the event relative to the market, it is possible to assess the effective beta of the UV and UL portfolios. By considering a very long pre-event window we can verify that any difference in the cumulative return of the UL and UV portfolios and the broad market are not simply a continuation of differential pre-event trends. The long panel also allows us to examine returns several 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 capture changes in the resulting the long-run cost implications of unions.

A complication that arises is how one defines the "event". The appropriate event is the date on which most 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 election. While we have a well-defined event, it is not the only relevant event, and it may not be the most important one. Alternative events that are potentially important are the petition and election dates. Using

[21] See Kenneth French's web page for a description of these series (http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html).

post-1977 data, where both the election and closing calendar dates are available, we find that the median time between the election and the NLRB closing is 10 days. In some case, typically when one of the parties issues a challenge, this gap can be considerably longer. In five 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 1 month of petition, and 83 percent within 2 months (Ruback and Zimmerman, 1984). Therefore, we do not believe that using month that the NLRB closed the election will present serious problems for estimation if most of the new information is being revealed at or after the petition date. To assess whether gradual diffusion of news led to excess returns prior to the closing date it is useful to examine a long pre-event window. However, empirically, we believe that it will be difficult to distinguish the market's anticipation of unionization from a mis-specified market model.

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 election on stock market performance involves estimating the post-event cumulative excess return for every election and relating these to the vote share in a flexible way. We will also conduct this analysis in order to examine the heterogeneity in the stock market reaction to election outcomes, and to determine whether there is a discontinuous relationship between cumulative excess returns and the vote share at the 50% threshold.

4.2 Dataset Assembly

This study 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 that used NLRB election data have 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.[22] Data for the 1961-1976 period were hand-entered from hard copies of NLRB monthly

[22]The 1977-1999 period data were obtained from Thomas Holmes' website

election reports.[23] Among other things, the NLRB dataset contains the number of voters that 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.[24] 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. This matching process is complex because while we know the company name where the election took place in the NLRB file, we know little else.[25,26] Therefore, when matching we are looking for similarities in the name listed in NLRB election file to names that were *ever* present in the CRSP files. To this end, we created two datasets, 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. This second dataset will be hereafter referred to as the "master name file". The master names file contains every name of every company that has ever appeared in CRSP.[27] The master name 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 dataset that could potentially be matched to companies in the master names file. Because the matching process is tedious, and must be largely done by hand, we excluded any election with less than 100 voters from consideration. This exclusion results in 24,709 firms in the certification election file to be potentially matched to firms in the master list of CRSP company names.[28,29]. These elections compromise 61% of all workers eligible to vote in NLRB certification elections. Using this smaller subset of the elections, firms in the election file were compared to firms in the master CRSP file. The algorithm matches company names in the NLRB

(http://www.econ.umn.edu/~holmes/data/geo_spill/) and are used in Holmes (2006).

[23]The records were digitized by a data-entry service. We took a random sample of records and estimated that they were 99.9% accurate.

[24]For a limited number of years the NLRB data has information of the calendar date of the election and the calendar date the NLRB closed the case.

[25]Location of the election is not very useful because the CRSP file only contains the location of company headquarters, which may differ from the location of any establishment undergoing a recognition election.

[26]The only additional information that could help us in 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 will be more useful as a check for the validity of the matches.

[27]Many companies have multiple names.

[28]Because 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 misspelled names and abbreviations.

[29]In future iterations of this paper we plan on taking a random sample of companies with less than 100 voters in order to construct a more representative sample.

file to company names in the master names file based on "spelling distance." We used the same matching algorithm employed by DiNardo and Lee (2004), which makes use of the SAS SPEDIS function. This algorithm considers comparisons with a spelling distance above a pre-determined threshold as candidate matches.[30] 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 firm on names only, manual inspection of the matches revealed that our automated procedure resulted in many matches that were obviously incorrect. Therefore, we had research assistants review every match, dropping those where the two firms appeared, in their judgment, to be a different company.[31,32] We then collected all of the unmatched companies in the election file, from the initial set of 24,709, and looked each of them up in Dun and Bradstreet's Million Dollar Directory and the Lexis/Nexis' Directory of Corporate Affiliations for the year of election. This step was taken in order to identify subsidiaries of publicly traded parent companies, and as a way to add back in companies that were dropped erroneously in the previous step.

We ultimately matched 7693 elections from the NLRB election file to companies in the CRSP master file. In 1579 cases, the firm in the CRSP file was not publicly traded at the time the election took place. Therefore, our final sample contains 6114 elections. This sample consists of 20% of all workers eligible to vote in elections.[33]

In order to determine whether the matches appear reasonable, we compare the reported two-digit SIC industry code and 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

[30] We refer the reader to DiNardo and Lee (2004) for further details on this algorithm.

[31] For example, the algorithm determined that companies in the election file that had the word "American" as part of their name as a sufficiently good match for the company "American Enterprises" in the CRSP file, if there was no better match. 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.

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

[33] Further inspection of the matched dataset reveals at least a handful of matches that are clearly incorrect. However, because we had already begun data analysis we thought it would be inappropriate to conduct further sample selection at that stage. We will review the matches at a later date.

may not be located in the same state as the company's headquarters. However, the comparisons are reassuring, in 50% of the matches the two digit SIC codes in the two datasets are the same, while in 40% of the matches the state of the establishment are the same. For reference, if we randomly match companies originating in the NLRB file to names from the master names file, from the subset of companies that are included in the final dataset, the corresponding match rate is 5% for industry and 4% for state.

Previous event studies of certification election use samples of very large elections. Ruback and Zimmerman (1984) and Bronars and Deere (1990) limit their sample to election with at least 750 eligible voters. Elections of this size are quite rare, resulting in small sample sizes (54 union victories in the main sample of Ruback and Zimmerman (1984)). We believe that the effects of certification 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. 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.[34,35] Using their sample selection scheme we find that in more than ten percent of the elections, less than 1% of the firm's workforce voted. In our reproduction of their sample, the median fraction of the workforce voting in an election is 5%.[36] By contrast, in our main specifications we only consider elections for which at *least* 5% of the total workforce voted.[37] In our sample the median election consists of 13% of the company's workforce voting (mean = 22%).[38] 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 as compared to what we would have obtained using the Ruback and Zimmerman criterion. Our baseline sample is almost eight

[34]We contacted Professor's Ruback and Zimmerman to request their dataset. As their paper was published more than 20 years ago they understandably could no longer provide it.

[35]Using 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% of the elections in the sample.

[36]Huth 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%, with approximately 30% of elections in the sample involving less than 1% of the company's workforce.

[37]Total employment in the year of the election is from the Compustat annual files.

[38]We do not use elections where employment information is missing.

times larger than the Ruback and Zimmerman sample.

We present summary statistics of firm characteristics in Table 1. Column (1) corresponds to all elections. Columns (2) and (3) correspond to elections where at least 5% of the workforce voted (hereafter the "$\geq 5\%$ sample") for UV and UL firms respectively. Columns (4) and (5) correspond to elections where less than 5% 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 value 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 (3621 vs.74,117 employees) and market value ($200 million vs. $2.2 billion, using the more broadly available CRSP measure). However, the $\geq 5\%$ corresponds to bigger elections, with an average of 450 workers voting as compared to an average of 292 in the $< 5\%$ sample.

In addition to the mean and standard deviation, for variables derived from the Compustat database we report in the braces the average percentile rank of that variable relative to all other firms in the Compustat database in the year and quarter of election. The average percentile rank is convenient for assessing how the firms in our sample compare to companies in the Compustat universe, and also has the advantage that it is 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 within two years of the election, conditional on being listed 2 years prior to the election. UV firms

are a somewhat more likely to delist than UL firms in the event-window (10% versus 7%). We will consider several approaches to address this potential problem, as well as the presence of missing returns in general. 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 coming about from compositional differences. However, a balanced panel does involve throwing away 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 considered.

5 Empirical Results

In Figure 3 we plot the average cumulative return of the portfolio of union victory firms against the average cumulative return of the broad market index over the same time period. Specifically, for each election we compute the cumulative sum of that security's returns beginning 24 months prior to the election through 24 months after the election.[39] For each election we also take the cumulative sum of the broad market index's returns over the same 25 months. We then average these cumulative returns over each month relative to certification (-24 to +24) for companies with union victories and for the market. In construction of this figure we impute missing returns with market returns in the same month, though we will also consider several alternative specifications where we do not impute returns. We use the CRSP equally weighted market index including dividends to proxy for the market.

The figure reveals that prior to the certification event both the UV and the market portfolios have almost identical trends, suggesting that the union victory portfolio has a beta of close to 1. The union victory portfolio displays a noticeable downward break in the average cumulative return trend around the month that the NLRB recognizes the union. There is no similar discernible break in trend in the broad market index. The union victory portfolio than gradually experiences lower returns

[39] For 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.

than the market until about 15 months after certification, when it reverts back to the market trend. We interpret the figure as providing evidence that union election certification wins are associated with large and negative excess returns.

In order to assess magnitudes and statistical significance of the excess returns presented in Figure 3, in Figure 4 we plot the market-adjusted cumulative return with a 95% confidence band. The market-adjusted return is computed as the difference in the average cumulative return of the UV portfolio and that of the market index. In Panel A we plot this return for the entire 49 month window, while in Panel B we compute this return beginning at the certification month. This second panel is the relevant one for assessing the overall effect size, and for assessing statistical significance. The figures show that the downward shift in the market-adjusted return that emerges soon after the election is closed by the NLRB is statistically significant. We can reject with 95% confidence that the average excess returns are equal to zero after 4 month.

We have conducted a variety of analyses to probe whether the pattern seen in Figure 3 is robust. These include plotting the average cumulative return of the union victory portfolio against the market without imputing missing data (Appendix Figure 1), using a balanced panel (Appendix Figure 2), excluding elections where following certification the cumulative abnormal returns are in the 5th or 95th percentile among all elections (Appendix Figure 3), and plotting the average cumulative return of the UV portfolio and of the market over an eight year window, four years prior to the election closing month, and four years after (Appendix Figure 4). In all cases the overall pattern of cumulative returns look very similar to that of Figure 3.[40]

In Figure 5 we plot the average cumulative excess return of the UV portfolio using a market-model benchmark. As is standard in event-studies, we estimate these excess returns in two steps. We first estimate the parameters α and β in equation (1). We use an estimation window which is 60 months long, ending 25 months prior to the event. Then, beginning 24 months prior to the event we compute and plot the cumulative prediction errors $\text{CER}(0,T)$, for $T = 0 - 24$. The pre-event data allows us to test the adequacy of the model used to predict the counterfactual returns in the post-event period. As before, we use the CRSP equally-weighted index that includes dividends to proxy

[40]Perhaps the only exception is Appendix Figure 4 in which the cumulative returns of the UV portfolio appears to converge towards the market sometime after two year. We can only speculate as to why this convergence is happening. Possible reasons include changes in composition as increasingly more firms are delisted, changed beta of the portfolio, or a very long-run effect of certification wins that leads to some improvements in firm performance.

for the market. The resulting pattern in excess returns looks very similar to the market-adjusted version in Figure 4. Average cumulative excess returns are close to 0 prior to the election, and are negative after. The total post-event average cumulative excess return is approximately -10%.

In Figure 6 we present average cumulative excess returns using the four-factor Fama-French model described in equation (2) as the benchmark.[41] This method yields the same pattern of excess returns, and overall magnitude of effects, that we observe with the market-adjusted and one factor models.

A surprising feature of Figures 4-7 is that, while the efficient market hypothesis would predict that the entire certification effect should be fully articulated by the time of the election month, we instead see a sustained effect, with a divergence in excess returns between the UV portfolio and the market index that begins around the time of election and persists for approximately 15 months. There are several reasons for why the efficient market prediction may not hold. It may be that while these events were known, investors or analysts did not know how to revalue the company with this information. This explanation is consistent with he fact that large unionization events are relatively infrequent. It may also be the case that this news was slow to diffuse to investors. As the companies in our sample are relatively small, only 50% of them had analyst coverage at the time of the election (over the 1983-1999 period).[42] It is therefore possible that these elections were not widely publicized or known.[43,44] Both explanations would imply that negative excess returns be accompanied by inferior outcomes of performance measures that were readily available and interpretable by investors, for example assets or profits. In Section 6.1 we will see that, in fact, the growth rate of these companies breaks trend at the same time, and for about the same number of months, as equity value.

We now turn to the evolution of the unionization effect over time. The DiNardo and Lee (2004) sample includes elections beginning in 1984. It is possible that unions no longer affected firm performance in this latter period, while in earlier years the effects may have been more pronounced.

[41] We use the same estimation and event window as with the basic market model.

[42] The 50% figure is derived from I/B/E/S International analyst data for years 1983-1999, the maximum possible time span.

[43] We have attempted to compare the pattern of excess returns for those companies with and without analyst coverage. However, because data on analyst coverage is only available starting in 1983, these tests are too low-powered to make any meaningful inferences.

[44] "Inattention" to predictable events in financial markets is not unprecedented, for example, see Dellavigna and Pollet (forthcoming).

In Figure 7 we compare the average cumulative market adjusted return of UV firms for election occurring in the 1961-1983 period (Panel A) to those occurring in the 1984-1999 period (Panel B). The figure indicates that the average effect of a union certification win on firm performance has remained fairly stable over time.

Table 2 presents regression estimates. In the first row of Panel A we report CER(0,24) for each of the three benchmark models. Columns (1)-(3) correspond to the market-adjusted, one factor, and four-factor models respectively. The estimated post-event cumulative excess returns range from -9% to -13%, and all are significant at the 1% level.[45] In order to capture possible stock market reaction from the petition, which typically occurs two to three months prior to the election, in the second row we calculate the cumulative return starting four months prior to the election closing. The estimates are slightly more negative than when starting on the recognition month though, as could be seen from the graphs, most of the effect emerges after the NLRB recognizes the union. In the third and fourth rows we report the cumulative excess returns over two pre-event spans. These returns are statistically indistinguishable from 0. The lack of significant excess returns prior to the election indicates that the market did not anticipate these events, on average.

To gauge magnitudes we calculate that a 10% negative return corresponds to approximately $20 million in lost market value (in 1998 dollars), or $40,522 per worker who is eligible to vote. This appears to be a plausible value. For example, if the average income of workers prior to certification is $35,000, then our magnitudes imply approximately a 9% union wage premium.

In Figure 8 we present average cumulative returns for the union loss portfolio against the market, as in Figure 3. As with the UV elections, this portfolio tracks the market quite well prior to the election, but unlike the UV portfolio it does not appear to break trend after the election. If anything, there is a gradual increase in the cumulative return of the UL portfolio after the election, though in Figure 8, which presents the difference in these series with confidence bands, we can see that this difference is not significant.

Panel B of Table 2 reports from cumulative excess returns for the two years following the election closing event for the union loss portfolio. Consistent with what we observe in Figure 4, the cumulative excess returns are close to zero and statistically insignificant in all cases.

Our sample selection scheme was in part predicated on selecting elections for which a sizable

[45] Standard errors are derived from the formula in Patell (1976).

fraction of the firm's workforce was voting. In practice we used a 5% cutoff. As an additional falsification check we examine whether elections with low fraction of workforce voting experience similar patterns of abnormal returns. In Figure 9 we plot the average cumulative return of these low percentage elections, for both firms where the union won, and for firms where the union lost. We overlay these series on the average cumulative return of the market index. For reasons that are unknown to us, both sets of firms are under-performing the market over the entire event-window. However, neither the UV or the UL portfolio exhibit a break in trend relative to market, or to each other around certification.

5.1 Compustat analysis

The results presented up to this point suggest that union victories are associated with negative excess returns. We complement this analysis with an additional investigation using variables from the balance sheets and income statements of these firms. 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 margin, and return on assets are affected by the outcome of representation elections. We compute the average value of these variables (logged when appropriate) over the 12 quarters before and after the event date, comparing UV and UL firms.[46] As before, we are looking to see whether these series were trending differentially prior to the event, and whether their trend breaks around the time of the event.

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 1960's, and not universally until the early 1970's. Moreover, as seen in Table 1, the fraction of observations that are missing is substantially higher in the Compustat dataset than in the CRSP dataset. For these reasons, 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 demean the variables, but do not drop elections with missing values.

In nine panels of Figure 11 we plot averages of these de-meaned variables over event-time, in each case comparing elections where the union won to those where the union lost. Figure 12 plots

[46] All variables in dollar units are expressed in 1998 dollars.

the difference in these series with a 95% confidence band.[47]. The figures show that the time pattern of variables that proxy for size are consistent with the pattern in equity value. UV firms display a downward break in trend near, or just before certification, in total assets (Panel A), shareholder equity (Panel B), and sales (Panel C). The reduction in asset growth is coming about in large part to reduced growth of plant, property, and equipment (Panel D). Because of small sample sizes, these series are not as well-behaved as those for equity values, but they look very similar.[48] We see little effect of union wins on a 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.[49]

The marked reduction in the growth rate of assets is notable because if unionization is increasing the price of labor, there should be substitution from labor to capital (though, as seen in Panel F, Tobin's average Q appears stable). That assets are actually declining implies that the "scale" effect from reduced reinvestment dominates this 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 excess returns mirrors the time-pattern we observe in shareholder equity, assets, sales and pre-tax income. The evidence is consistent with stock market pricing in 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), return on assets appears stable (Panel H). The profit margin in UV firms appears to decline a bit relative to UL firms, but not until about 7 quarters after the election (Panel I). At first blush, the finding that companies that undergo unionization experience lower growth rates but stable return on assets and profit margins may seem puzzling. However, if firms only select projects that are sufficiently profitable, and unionization reduces the number of these high net present value (NPV) projects, it is then possible for the company's growth rate to decline in spite of experiencing no change in its operating performance.

In Table 3 we present difference-in-difference estimates for the effect of a union victory relative

[47] Standard errors are clustered on the election id.

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

[49] Bronars and Deere (1991) shows that there is a positive association of 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.

to a union loss on each of the six aforementioned variables. The sample consists of election × 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 × union win."). The point estimates suggest that assets, shareholder equity, and sales fall by approximately 10% in UV firms after the election, relative to UL firms. Profits of UV firms are approximately 17% lower in the post-election period relative to the pre-election period, relative to UL firms. These estimates, which are in the realm of statistical significance, are consistent with the 10-14% negative excess returns we observe in equities.

5.2 Heterogeneity

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 sections aims at providing a partial answer to this question.

Whereas DiNardo and Lee (2004) identify the "unionization effect" by focusing on close elections, the discussion in Section 3 suggests that we can learn about 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 that we have at our disposal.

We begin by relating the cumulative market-adjusted return, computed at the security-level over the two years following the month that the NLRB closes the election, to the union vote share. For election i, we denote this cumulative market-adjusted return as $\text{CMAR}(0,24)_i$. We will estimate $E[\text{CMAR}(0,24)_i|v_i]$, where v_i denotes the union vote share in election i.

A more complete investigation into this relationship involves computing the cumulative excess return for every firm in the sample using a market model.[50]. To this end, we estimate the alpha and beta for every firm using the specification:

$$r_{it} = \alpha_i + \beta_i R_t.$$

We estimate this model using up to five years of historical returns ending five months prior to the

[50]The results are not sensitive to accumulating excess returns over shorter or longer windows.

election.[51] We do not estimate this model for stocks that had less than 10 months of returns available in the estimation window.

Because we are using monthly data, we have found that the estimated parameters from the market-model on individual stocks are quite imprecise. This imprecision in turn leads to imprecision in the cumulative excess returns. In a handful of cases the estimated betas are very large (for example, greater than 10 in absolute value) resulting in extreme outliers. The reason for this imprecision is that by estimating a different market model for every security, our model is over-parametrized. Because we are primarily interested in how the cumulative excess returns relates to vote share, we can mitigate this problem by assuming that the parameters from the market model evolve smoothly over the vote share, thus allowing us to impose some additional structure. For a given value of the vote share we predict $\alpha$and β using:

$$\text{pred-}\beta_i = \widehat{E[\beta_i|v_i]} = c + d_1 v_i + d_2 v_i^2 + d_3 v_i^3 + d_4 v_i^4,$$

$$\text{pred-}\alpha_i = \widehat{E[\alpha_i|v_i]} = c + e_1 v_i + e_2 v_i^2 + e_3 v_i^3 + e_4 v_i^4.$$

We estimate these models using the estimated α_i and β_i for each of the individual equities. In practice we compute pred-α_i and pred-β_i even when there was not sufficient data to estimate α_i and β_i for a given stock, though the results are not sensitive to excluding these cases. Excess returns are then computed as:

$$ER_{it} = r_{it} - (\text{pred-}\alpha_i + \text{pred-}\beta_i R_t).$$

As before, the cumulative excess returns are calculated as:

$$\text{CER}(0,T)_i = \sum_{\tau=0}^{T} \text{ER}_{i\tau}.$$

We are interested in the shape of $E[\text{CMAR}(0,24)_i|v_i]$ and $E[\text{CER}(0,24)_i|v_i]$. We graphically plot these functions by: (1) averaging $\text{CMAR}(0,24)_i$ and $\text{CER}(0,24)_i$ over 20 equally-spaced vote share bins,[52] and (2) plotting the predicted values from the model $E[y_i|v_i] = p(v_i) + \beta 1(v_i > 0.5)$, where $p(\cdot)$ denotes a sixth-order polynomial, $1(v_i > 0.5)$ is an indicator function for whether the

[51] Not all stocks have 60 months of returns available in the estimation window.

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

union vote share in a given election exceeds 50%, and $y_i$is either $\text{CMAR}(0,24)_i$ or $\text{CER}(0,24)_i$, depending on the specification.

Figure 14 presents estimates of $E[\text{CMAR}(0,24)_i|v_i]$ using both of these approaches. (For reference, Figure 13 shows the histogram of the union vote share variable.) Figure 14 shows clear evidence that the effect of a certification election is heterogeneous, and that it depends on the union vote share. The negative impact of a union election win appears to become markedly more pronounced when the union has a higher vote share. Firms with union losses also exhibit a downward sloping relationship between excess returns and vote share.[53] As in the Dinardo and Lee study, there is no discernible discontinuity in the CMAR at the 50% 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. While this difference is not statistically significant, it is the case that the negative cumulative market-adjusted returns experienced by firms with close union losses are significant.

In Figure 15 we plot the estimated function $E[\text{CER}(0,24)_i|v_i]$. The predicted values resemble those obtained using market-adjusted returns. There is no visible discontinuity in excess returns at the 50% threshold and, as before, $E[\text{CER}(0,24)_i|v_i]$ is negatively sloped. Greater than 60% union vote share is associated with 20-30 percent negative cumulative excess returns. The most noticeable difference between the estimated $E[\text{CMAR}(0,24)_i|v_i]$ and $E[\text{CER}(0,24)_i|v_i]$ functions is that the latter appears to be a bit flatter for union vote shares lower than 50%.

We now turn to several checks of validity. For these analyses we will focus on CER. In Figure 16 we overlay the predicted CER in months 0 through 24 (shown in Figure 15) with the predicted CER for months -24 through -4.[54] The figure shows that the gradient in CER by vote share, seen for months 0 to 24, is not present for months -24 through -4. This plot reassures us that the negative CER observed for higher 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 excess returns in the post-event period for possible pre-event trends. Specifically, we calculate the excess return in the post-event period deviated from the average excess return in

[53] At extreme vote shares there is more variability in the predicted cumulative market-adjusted returns due to small sample sizes.

[54] CER(-24,-1) is calculated just as CER(0,24) except that the market model uses an estimation window that ends 29 months prior to the election, as opposed to 5 as was the case for the post-event excess returns. We use a different estimation window in order to avoid having it overlap with event window.

the pre-event period:

$$\text{adjusted-}ER_{it} = ER_{it} - \frac{1}{18}\sum_{\tau=-24}^{-6} ER_{i\tau}$$

We then calculate:

$$\text{adjusted-CER}(0,24)_i = \sum_{\tau=0}^{24} \text{adjusted-}ER_{i\tau}.$$

Figure 17 plots the predicted adjusted-CER with a 95% confidence band.[55] The Figure shows virtually the same pattern of heterogeneity seen in the earlier figures, though with a wider confidence band.

In Table 4 we conduct formal statistical inference. Using the same sample of 1347 elections used to construct Figures 14-17, in column (1) we regress CMAR(0,24) on a dummy for whether the union won the election. Consistent with earlier analyses, we find that union wins experience cumulative excess returns that are 12.5% percentage points lower than firms with union losses (t-ratio = -3.44). In column (2) we add the union vote share as a covariate.[56] The introduction of this variable alone attenuates the coefficient on the union win dummy to 0.024 (t-ratio = 0.06). Adding higher-order polynomial terms in the vote share (column 3) only increases the estimated union win coefficient. The "regression discontinuity" estimate of a union win is 4.8%, which is statistically indistinguishable from 0. In column (4) we examine whether the negative gradient between CMAR and the vote share differs among elections where the union won and lost. Specifically, we regress CMAR(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 columns (5)-(8) we estimate the same models using CER(0,24) as the dependent variable. The point estimates are very similar to those found using CMAR(0,24) and, qualitatively, the conclusions do not change. As already indicated, the main difference in results is on the slope on vote share for losing elections. In column (8) it can be seen that there is a more pronounced negative relationship between CER(0,24) and vote share amongst UV firms than UL firms, though this difference is only at the margins of significance.

For completeness, in columns (9)-(12) we estimate the same set of models using CER(-24,-4) as

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

[56] Vote 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).

the dependent variable. None of the patterns observed when using CER(0,24) as the dependent variable are evident here. In columns (13)-(16) we re-estimate these models using adjusted-CER(0,24) as the dependent variable. The point estimates are very close to the ones using CER(0,24), but are estimated less precisely, with standard errors approximately 50% larger than those in columns (5)-(8).

5.3 Discussion

In order to facilitate interpretation of these results, we revisit the theoretical framework outlined above. Using the notation from the model, the point estimates imply that B_1 and B_3 are small, while B_2 is relatively large. As discussed in Section 3, $B_1 \approx 0$ implies that the firm is largely unresponsive to the election, or to the preferences of the median voter. This result is consistent with firms that are complying with NLRA rules that prohibit firms from changing working conditions in order to gain votes in an NLRB election.

The limited responsiveness on the part of the firm helps us pin down the parameters governing the behavior of the union. Figure 18 graphically shows the circumstances under which there could be a large average effect of unionization, while at the same time there be a limited difference in the effects of close union wins relative to close union losses. Because the firm does not adjust its offer in relation to worker preferences, it is only possible to have a close election when workers have low tolerance for high wages, and when the union's proposal is small. Specifically, as can be seen, a small discontinuity ($B_3 \approx 0$) implies a small α. Therefore, when workers have limited tolerance for higher wages, unions tends to make very modest proposals. Also evident from Figure 18 is that we must select a small β so as to make the discontinuity small. Provided that there are some elections with sufficiently large μ (corresponding to large y_U), in doing so we can preserve a large average unionization effect.

A small β implies that unions sluggishly adjust their proposals to accommodate the underlying preferences of the workers. In particular, unions will make larger, but increasingly conservative offers as the median worker's preferred wage increases. Such behavior suggests that as μ becomes larger, the enacted wage is below the preferred wage for an increasingly larger share of workers. However, as Figure 18 demonstrates, the pattern of results is consistent with unions generally offering proposals that are lower than the preference of the median voter over most values of μ. Such behavior is

consistent with unions that place a high premium on winning elections, even at a cost of lower wages.

The results also have implications for the relationship between unionization and outcomes that we have not considered in this paper. If we think of the preference of the median voter as related to the establishment's elasticity of labor demand–in that preferences for higher compensation correspond to establishments with relatively inelastic labor demand–these results suggest that new unionization should lead to relatively limited employment effects, even when the firm experiences considerable equity loss from new unionization. This is because, through selection, unionization leads to large profit impacts only when labor demand is inelastic. We should not expect to see plant closure following unionization, as the median voter would necessarily be adversely impacted by that outcome.

In light of the model, the histogram of the union vote share presented in Figure 13 provides additional information about the underlying distribution of preferences. The histogram shows that there are more union losses than wins. Moreover, the vote share distribution is bell shaped, so that large union wins or defeats are relatively infrequent. If we think about the union vote share distribution as an endogenous outcome of the model described in Section 3, it must be the case that most elections correspond to small μ–below $\widetilde{\mu}$ in Figure 18. Furthermore, since there are not many elections where the union wins overwhelmingly, there can only be a relatively small number of elections with very large μ. If we again relate preferences to the elasticity of labor demand, the model therefore suggests that most establishments undergoing elections have relatively elastic labor demand. As already suggested, it is the relatively few cases where demand is inelastic that lead to the large average effects of new unionization on equity value.

Of course, there are important caveats to be made. First, while the estimated RD estimate is close to zero (and positive), it is measured with some imprecision. The RD estimates do not allow us to rule out cumulative excess returns between -8% and +18%. However, it is the case that the negative effects of unionization that we observe on average are largely driven by elections with more extreme vote shares. A second caveat is that the model outlined above is not the only one that can explain the pattern in the results. For example, it may be the case that unions need to have widespread support among the workers in order to be effective, perhaps because it is the only way that a strike threat would be credible.

6 Conclusion

The economic effects of unions on the labor market and the economy has been a longstanding area of interest for economists with debates regarding the impact of unions on wages, their potential role as monopolies, and their role in work stoppages, as well as the question of why they can even exist and survive in a competitive labor market. Moreover, a complete understanding of the economic implications of declining unionization is not possible without precisely knowing 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 new unionization is associated with a reduction in the firm's market value totaling approximately $40,500 per worker eligible to vote. This finding is robust to the use of a variety of specifications, and to the use of several different commonly used market benchmarks. 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 the balance sheets and income statements of both sets of firms reveals that unionization 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 UL firms. The evidence is therefore consistent with unionization reducing the number of (sufficiently) positive NPV projects available to a firm.

However, when viewed in isolation, the NLS case misses what turns out to be important heterogeneity in the stock market reaction to recognition elections. Using a different sample than Dinardo and Lee, we also find RD estimates that imply that unions are largely ineffective, at least to the extent that they do not affect a firm's equity value. This finding can be reconciled with the findings from the event-study analysis through the negative gradient in excess returns in relation to the union vote share. There is smoothness in the predicted excess returns about the 50% vote share threshold because more overwhelming union victories are associated with worse financial performance relative to more limited union wins. These results are consistent with firms that do not respond to the threat of an election, the existence of a limited number of elections where workers demand substantially higher wages, and unions that place a high premium on winning elections.

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.

Atherton, Wallace N., *The Theory of Union Bargaining Goals*, Princeton, NJ: Princeton University Press, 1973.

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.

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.

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.

Dellavigna, Stefano and Joshua Pollet, "Demographics and Industry Returns," *American Economic Review*, forthcoming.

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 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.

_ 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.

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.

Holmes, Thomas, "Geographic Spillover of Unionism," April 2006.

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.

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.

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

Matsa, David, "Capital Structure as a Stretegic Variable: Evidence from Collective Bargaining," September 2006.

Mitchell, Merwin W. and Joe A. Stone, "Union Effects on Productivity: Evidence from Western Sawmills," *Industrial and Labor Relations Review*, 1992, *46*, 135–45.

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.

Patell, James M., "Corporate Forecasts of Earnings Per Share and Stock Price Behavior: Empirical Tests," *Journal of Accounting Research*, 1976, *14*, 246–276.

Roomkin, Myron and Richard Block, "Case Processing Time and the Outcome of Representation Elections: Some Empirical Evidence," *Univeristy of Illinois Law Review*, 1981, *1*, 75–97.

Ruback, Richard and Martin Zimmerman, "Unionization and Profitability: Evidence from the Capital Market," *Journal of Political Economy*, 1984, *92*, 1134–1157.

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	404	965	1153	2681
Vote share for union	0.621	0.348	0.640	0.346
	[0.108]	[0.098]	[0.128]	[0.101]
Number of voters	446.9	454.315	278.146	297.602
	[537.100]	[567.966]	[265.554]	[301.021]
Number eligible	492.5	495.798	286.114	317.454
	[649.683]	[656.433]	[286.914]	[329.610]
Fraction of employees voting	0.199	0.223	0.013	0.013
	[0.206]	[0.203]	[0.012]	[0.012]
Year of election	1975.4	1977.048	1974.812	1976.542
	[9.191]	[9.067]	[9.135]	[9.384]
Fraction in Manufacturing	0.525	0.528	0.487	0.501
	[0.500]	[0.499]	[0.500]	[0.500]
Fraction in Construction	0.252	0.230	0.310	0.312
	[0.435]	[0.421]	[0.463]	[0.463]
Number of employees	3885.5	3510.0	70261	75775.740
	[5413.8]	[5355.3]	[136508]	[123899.100]
Market Value (CRSP)	197.0	201.476	2249.3	2169.632
	[434.0]	[518.9]	[6248.941]	[6342.550]
Market Value (Compustat)	308.717	329.805	6334.1	7581
	[614.9]	[799.0]	[13372]	[16343]
	{0.336}	{0.329}	{0.760}	{0.780}
Shareholder equity	242.6	233.2	4991.712	4480
	[433.0]	[497.721]	[13859.310]	[9432]
	{0.344}	{0.314}	{0.774}	{0.774}
Total Assets	588.4	683.8	13974	14165
	[1243.327]	[1876.515]	[36396]	[33308]
	{0.373}	{0.313}	{0.779}	{0.788}
Total Liabilities/Total Assets	0.060	0.068	0.062	0.071
	[0.118]	[0.162]	[0.112]	[0.183]
	{0.535}	{0.421}	{0.562}	{0.595}
Pretax income	15.114	9.762	249.3	276.292
	[46.971]	[41.908]	[731.7]	[731.061]
	{0.351}	{0.356}	{0.738}	{0.741}

Table 1 (cont.)

	At least 5% of workforce voting		Less than 5% of workforce voting	
	Union wins	Union loses	Union wins	Union loses
Sales	160.734	144.2	2693.538	3041
	[238.7]	[225.088]	[5306.252]	[5534]
	{0.329}	{0.308}	{0.800}	{0.804}
Tobin's Q	1.167	1.297	1.287	1.309
	[0.658]	[0.694]	[0.642]	[0.625]
	{0.444}	{0.503}	{0.477}	{0.557}
Profit margin	0.069	0.060	0.084	0.084
	[0.119]	[0.167]	[0.073]	[0.074]
	{0.435}	{0.497}	{0.464}	{0.518}
Income/Employees	0.004	0.003	0.004	0.004
	[0.023]	[0.008]	[0.006]	[0.007]
	{0.414}	{0.489}	{0.485}	{0.510}
Return on Assets	0.013	0.022	0.026	0.027
	[0.051]	[0.037]	[0.023]	[0.027]
	{0.476}	{0.530}	{0.470}	{0.510}
Dividend Ratio	0.633	0.259	1.147	0.941
	[3.415]	[1.100]	[6.993]	[11.023]
	{0.424}	{0.491}	{0.582}	{0.576}
Fraction of stocks delisted	0.10	0.07	0.04	0.02
	[0.30]	[0.30]	[0.19]	[0.15]

Notes: Summary statistics are based on NLRB election, Compustat, and CRSP data. Standard deviations in brackets. For Compustat variables, 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. Profit margin = pre-tax income/sales. Dividend ratio = dividends/pre-tax income.

Table 2: Estimates of post- and pre-election cumulative excess returns

	(1) Market-Adjusted	(2) Basic Market-Model	(3) Fama-French Four Factor Model
Panel A: Union Wins			
CER(0,24)	-0.092	-0.108	-0.134
	(0.028)	(0.031)	(0.025)
CER(-4,24)	-0.106	-0.131	-0.159
	(0.030)	(0.034)	(0.034)
CER(0,15)	-0.085	-0.097	-0.115
	(0.023)	(0.023)	(0.023)
CER(-4,15)	-0.100	-0.12	-0.14
	(0.025)	(0.027)	(0.027)
CER(-24,-1)	-0.001	0.001	-0.02
	(0.025)	(0.030)	(0.03)
CER(-24,-8)	-0.007	-0.008	-0.024
	(0.021)	(0.024)	(0.024)
Panel B: Union Loses			
CER(0,24)	0.029	-0.032	-0.026
	(0.022)	(0.032)	(0.025)
CER(-4,24)	0.031	-0.04	-0.035
	(0.023)	(0.035)	(0.027)
CER(0,15)	0.023	-0.012	-0.008
	(0.017)	(0.024)	(0.02)
CER(-4,15)	0.025	-0.02	-0.017
	(0.018)	(0.027)	(0.023)
CER(-24,-1)	0.007	-0.042	-0.041
	(0.017)	(0.031)	(0.025)
CER(-24,-8)	0.005	-0.033	-0.033
	(0.014)	(0.025)	(0.021)

Notes: Each estimate corresponds to a different estimated model. There are 404 elections in the sample reported in Panel A and 965 elections in Panel B. CER(X,Y) denotes the average cumulative excess return from month X to month Y relative to the NLRB election closing date. Column (1) corresponds to market-adjusted returns. These returns are computed as the simple difference between portfolio and market returns. Column (2) corresponds to excess returns computed using the basic market model as the benchmark (equation 1 in the text). Column (3) corresponds to excess returns computed using a Fama-French four factor model (equation 2 in the text). Both market models use a 60 month estimation window that ends 5 months prior to the event.

Table 3: Compustat Analysis

	(1) ln(Assets)	(2) ln(Shareholder equity)	(3) ln(PPE)	(4) ln(Sales)	(5) ln(pretax income)	(6) Dividend Ratio	(7) Profit margin	(8) ROA	(9) Tobin's Q	(10) Liabilities/ Assets
post	0.151	0.014	0.137	0.132	0.170	-0.184	0.000	-0.004	-0.054	-0.001
	(0.021)	(0.006)	(0.028)	(0.019)	(0.032)	(0.117)	(0.002)	(0.001)	(0.026)	(0.005)
post	-0.106	-0.016	-0.011	-0.075	-0.168	0.033	-0.006	-0.001	0.029	0.003
× union win	(0.037)	(0.008)	(0.048)	(0.034)	(0.062)	(0.261)	(0.003)	(0.002)	(0.038)	(0.008)
Observations	14235	16309	14150	16903	13925	6110	14465	13878	13956	5776
R-squared	0.966	0.988	0.956	0.936	0.745	0.075	0.639	0.325	0.657	0.282

Notes: Notes: Variables are derived from Compustat data; 1973-1999. Each column corresponds to a different model estimated using OLS. Standard errors that are clustered on election are in parentheses. Observations are event quarter × 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 4: Cumulative excess returns and vote share

	CMAR(0,24)				CER(0,24)				CER(-24,-4)				Adjusted-CER(0,24)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
Constant	0.033	-0.049	-0.053	-0.049	0.011	-0.050	-0.042	-0.008	-0.021	-0.014	-0.016	-0.014	0.039	-0.013	-0.008	0.027
	(0.022)	(0.031)	(0.039)	(0.035)	(0.022)	(0.030)	(0.038)	(0.035)	(0.017)	(0.028)	(0.036)	(0.033)	(0.034)	(0.053)	(0.069)	(0.061)
Union won	-0.125	0.024	0.048	0.024	-0.125	-0.016	0.045	0.291	0.016	0.004	-0.008	0.000	-0.149	-0.055	0.005	0.166
	(0.036)	(0.055)	(0.068)	(0.055)	(0.036)	(0.055)	(0.068)	(0.188)	(0.031)	(0.050)	(0.061)	(0.157)	(0.056)	(0.091)	(0.110)	(0.288)
Union won				-0.002				-0.579				0.006				-0.417
× vote share				(0.343)				(0.344)				(0.287)				(0.537)
vote share		-0.538		-0.537		-0.395		-0.201		0.045		0.043		-0.339		-0.199
		(0.168)		(0.214)		(0.169)		(0.216)		(0.139)		(0.176)		(0.267)		(0.345)
p(vote share)			X				X				X				X	
Observations	1347	1347	1347	1347	1347	1347	1347	1347	1316	1316	1316	1316	1347	1347	1347	1347

Notes: Notes: Robust standard errors in parentheses. The variable "vote share" stands for the union vote share minus 0.5. The dependent variable is the cumulative market-adjusted return from 0 to 24 (columns 1-4), cumulative excess return from 0 through 24 (columns 5-8), cumulative excess return from -24 to -4 (columns 9-12), and cumulative excess return adjusted for pre-event trends (columns 13-16). See Section 6.2 for details on the construction of these variables. p(vote share) is a fourth-order polynomial in the union vote share.

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

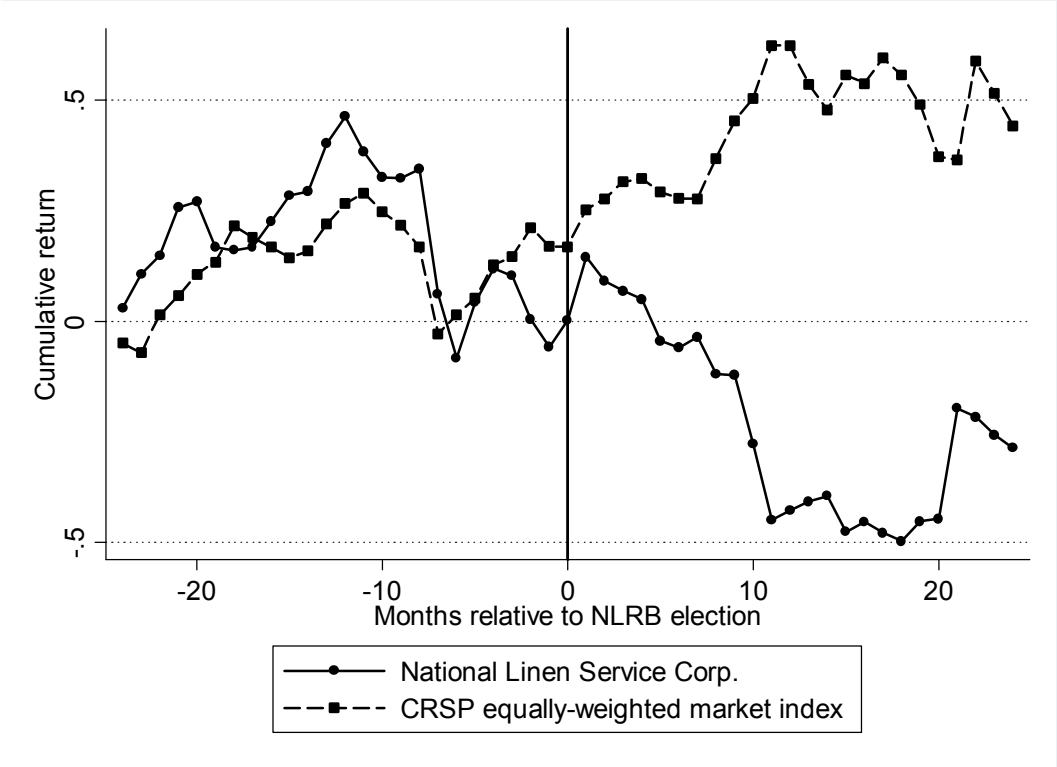

Figure 2: Union and firm offers in relation to median voter preferences

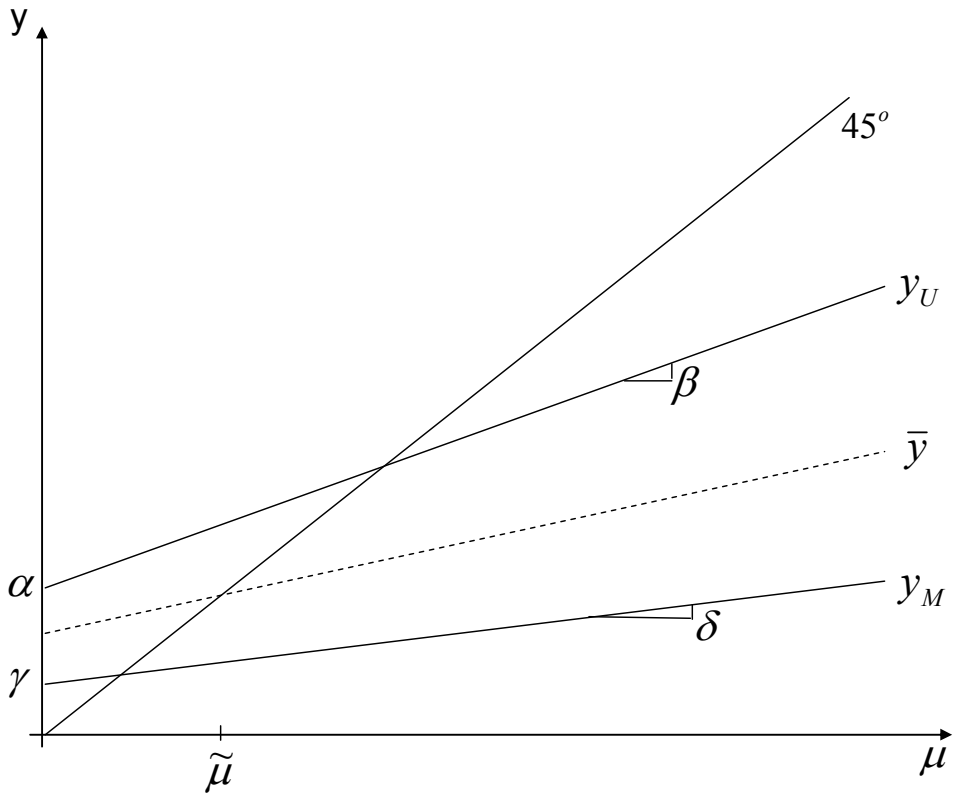

Figure 3: Cumulative returns of the union victory and market portfolios

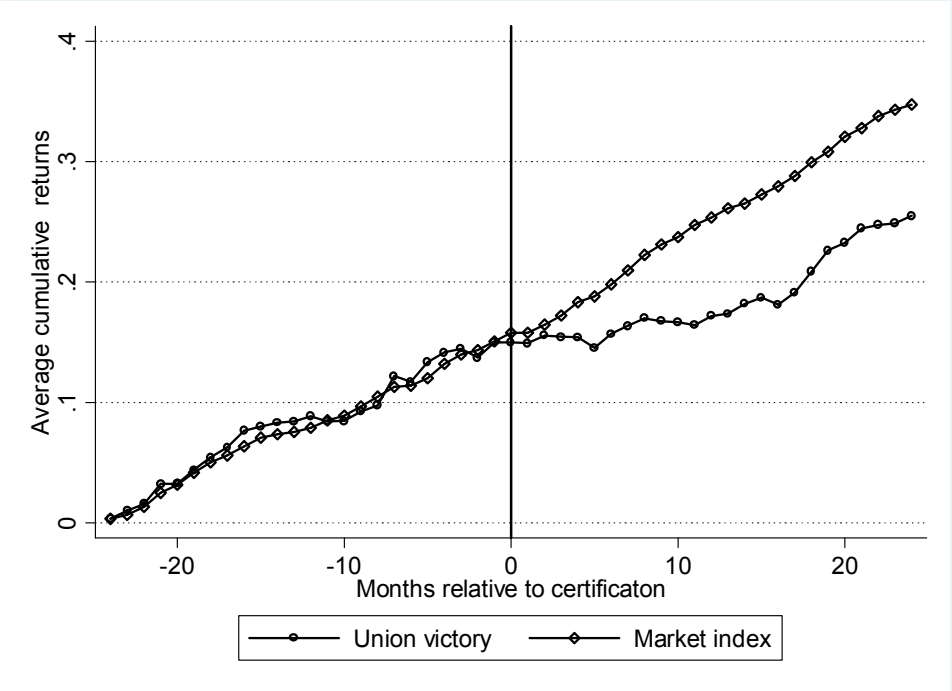

Note: The "union victory" portfolio consists of publicly traded companies holding elections where at least 5% of the workforce voted and the union won. When the company's stock market return is missing it is imputed with the stock market return. Returns are net of the risk-free rate.

Figure 4: Average cumulative market-adjusted returns.

Panel A: Beginning 24 months prior to certification

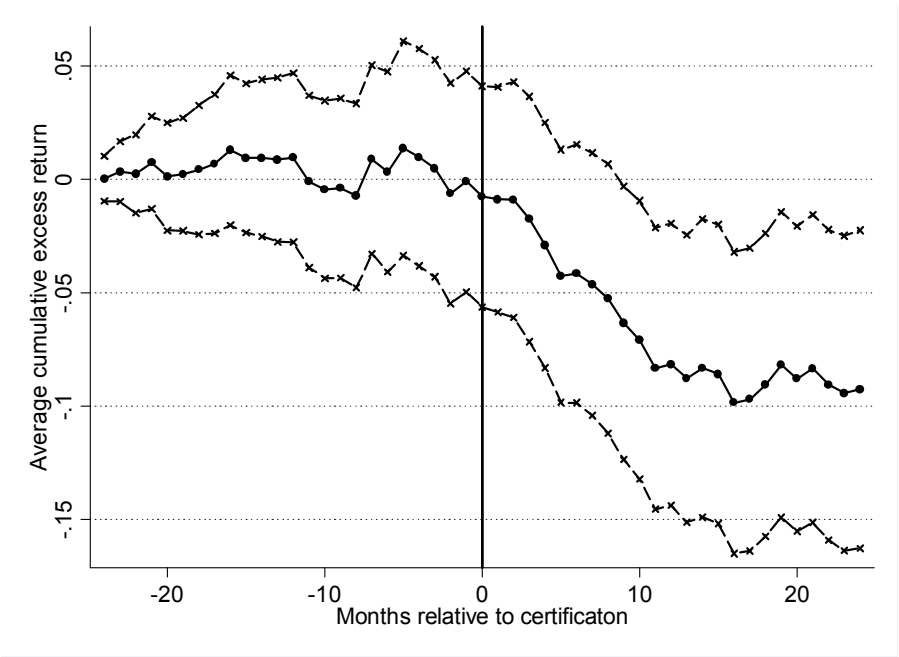

Panel B: Beginning at certification

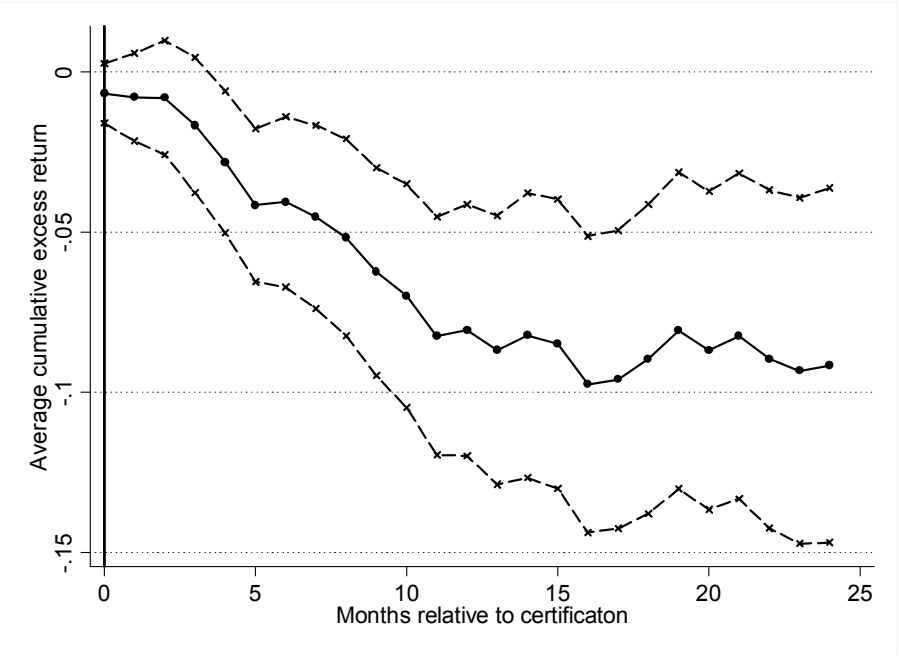

Notes: Both panels show the difference in the average cumulative excess returns between the union victory portfolio and the broad market index, as shown in Figure 3. The sample is the same one used in Figure 3. Panel A corresponds to the average cumulative excess return computed beginning 24 months prior to the NLRB closing date. Panel B corresponds to the average cumulative excess return computed beginning on the NLRB closing date. The dashed line is the 95% confidence interval, which is computed using OLS standard errors clustering on elections.

Figure 5: Average cumulative excess returns of the union victory portfolio; Market-model method.

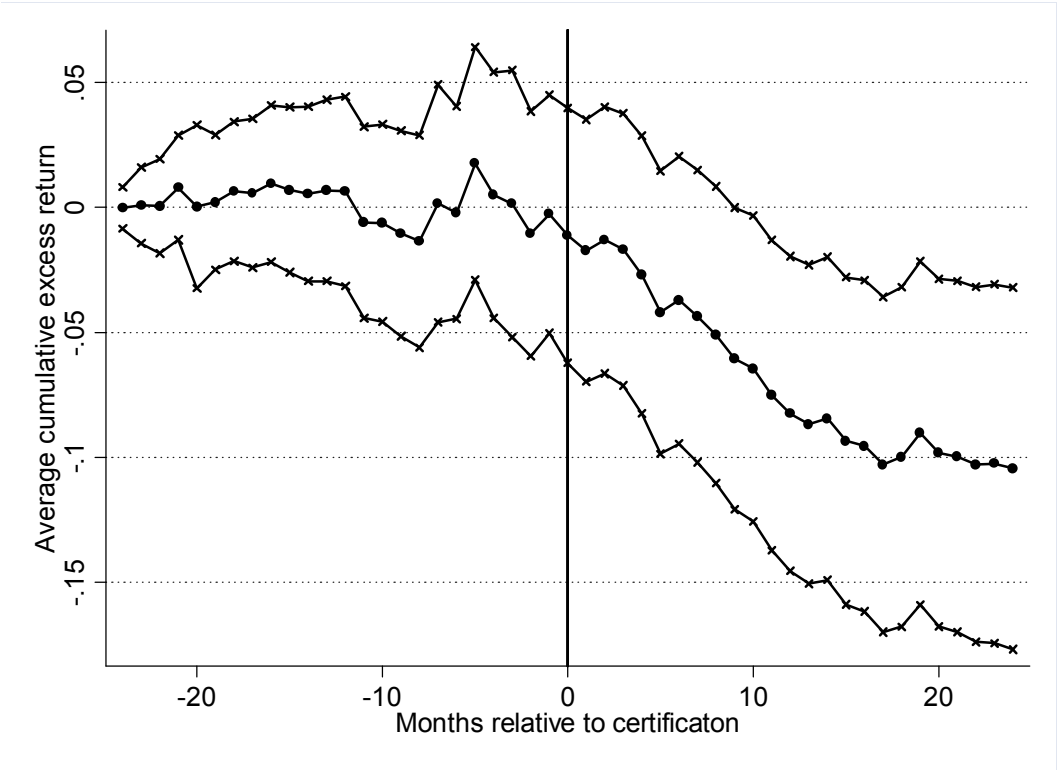

Notes: Excess returns are calculated using the model specified in equation (1) in the text.

Figure 6: Average cumulative excess returns of the union victory portfolio; Fama-French three factor model with a momentum term.

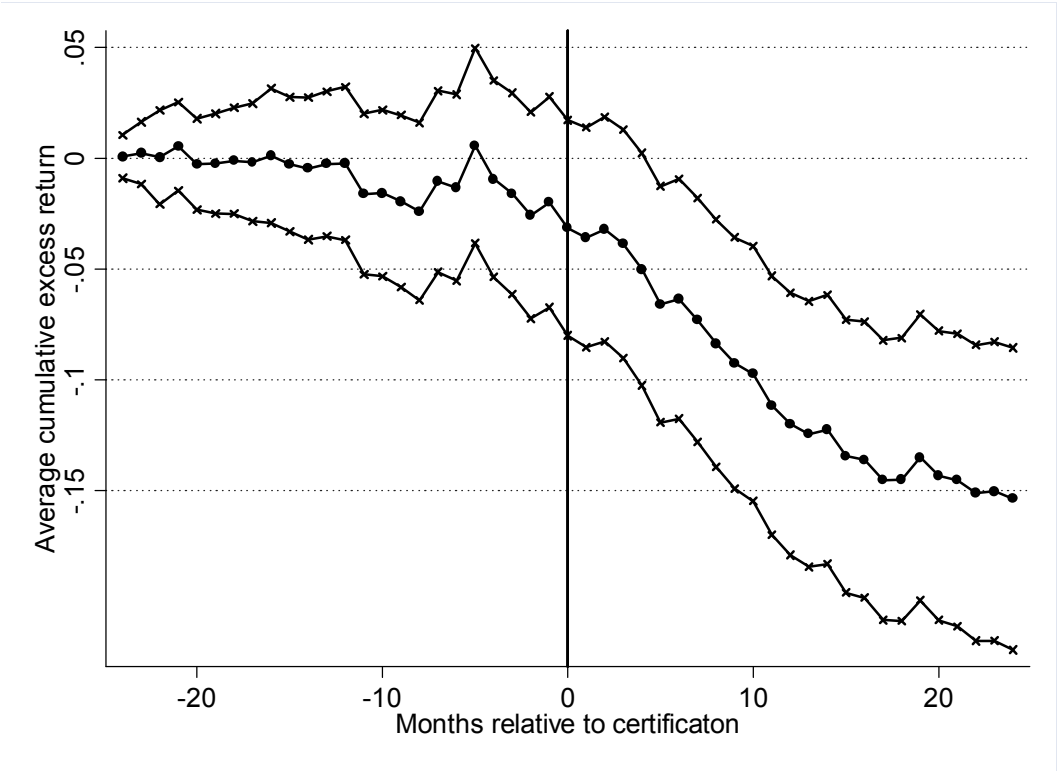

Notes: Excess returns are calculated using the model specified in equation (2) in the text.

Figure 7: Average cumulative excess returns stratified by year

Panel A: 1961-1983

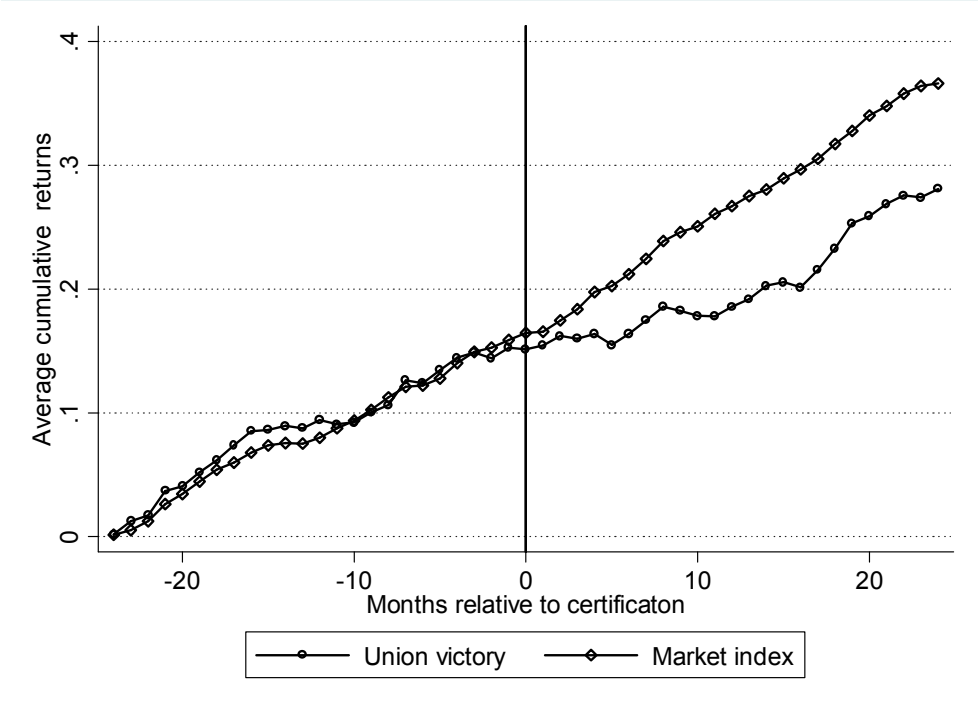

Panel B: 1984-1999

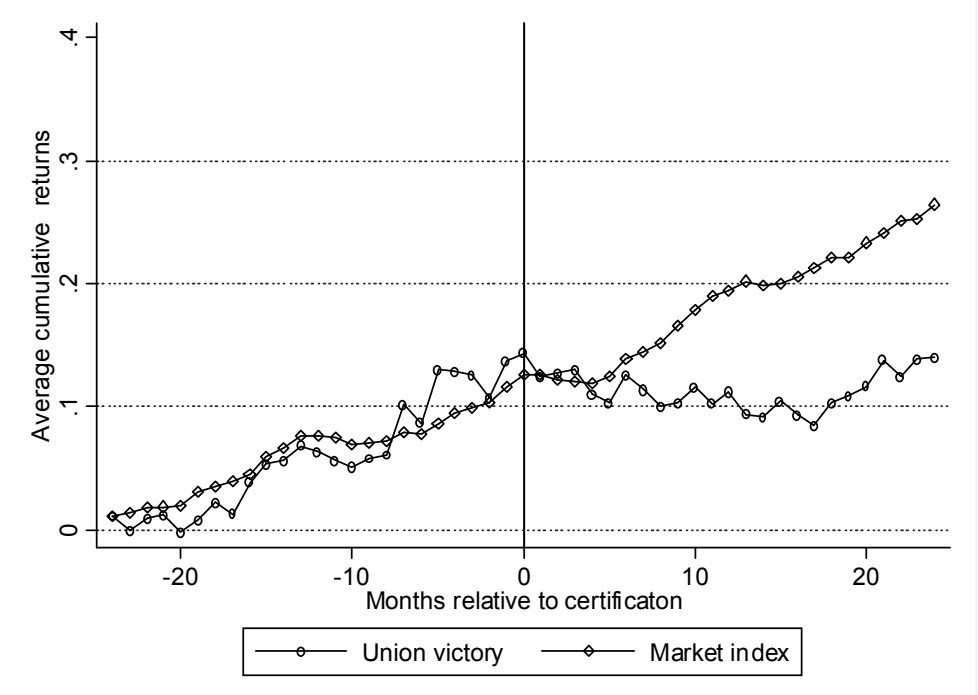

Figure 8: Cumulative excess returns of the union loss and market portfolios

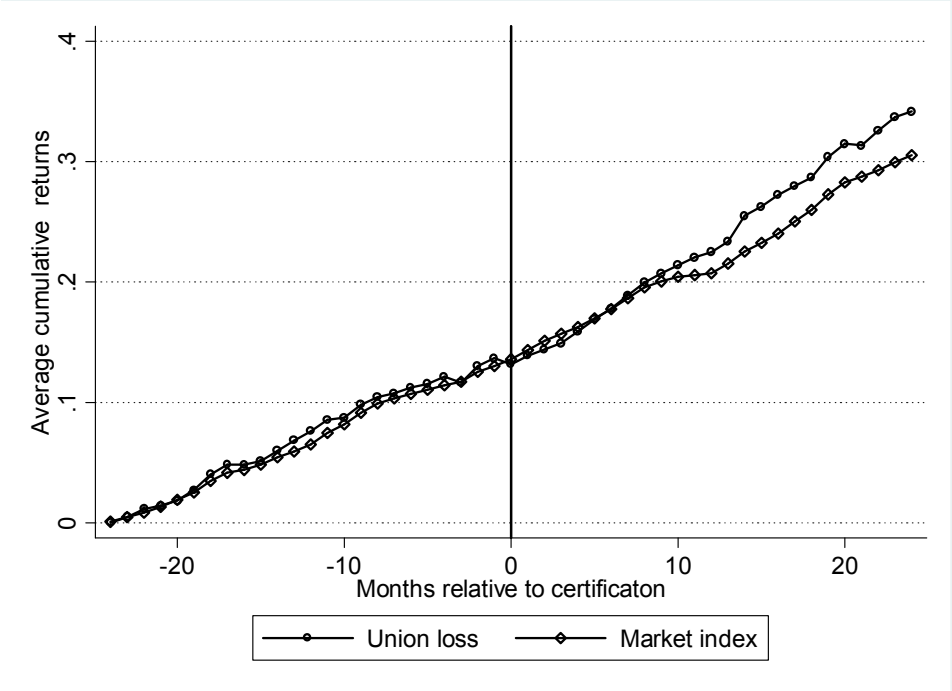

Note: The "union loss" portfolio consists of publicly traded companies holding elections where at least 5% of the workforce voted and the union lost. When the company's stock market return is missing it is imputed with the stock market return. Returns are net of the risk-free rate.

Figure 9: Difference in average cumulative return of the union loss portfolio and the market index

Panel A: Beginning 24 months prior to certification

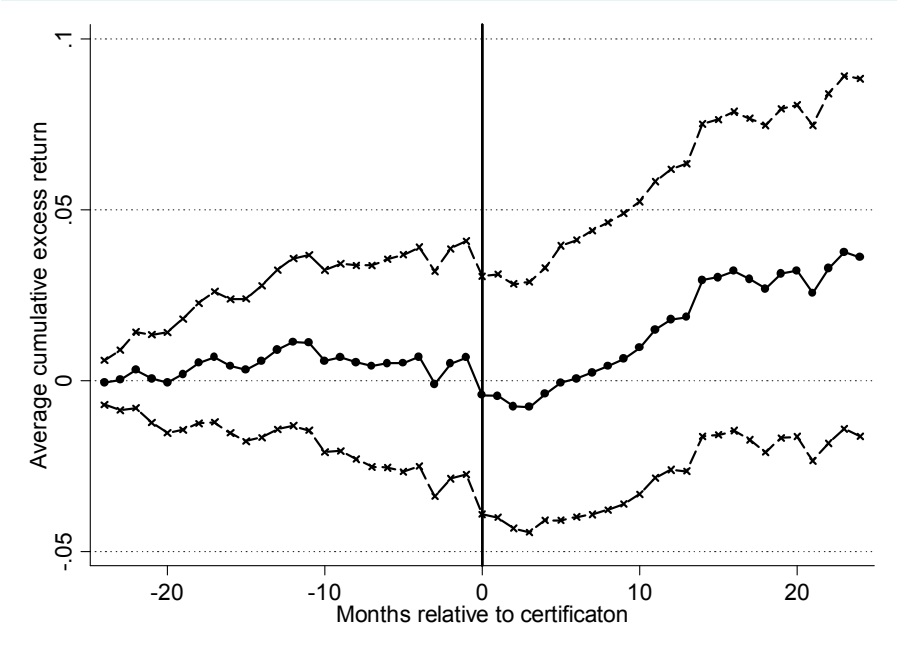

Panel B: Beginning certification month

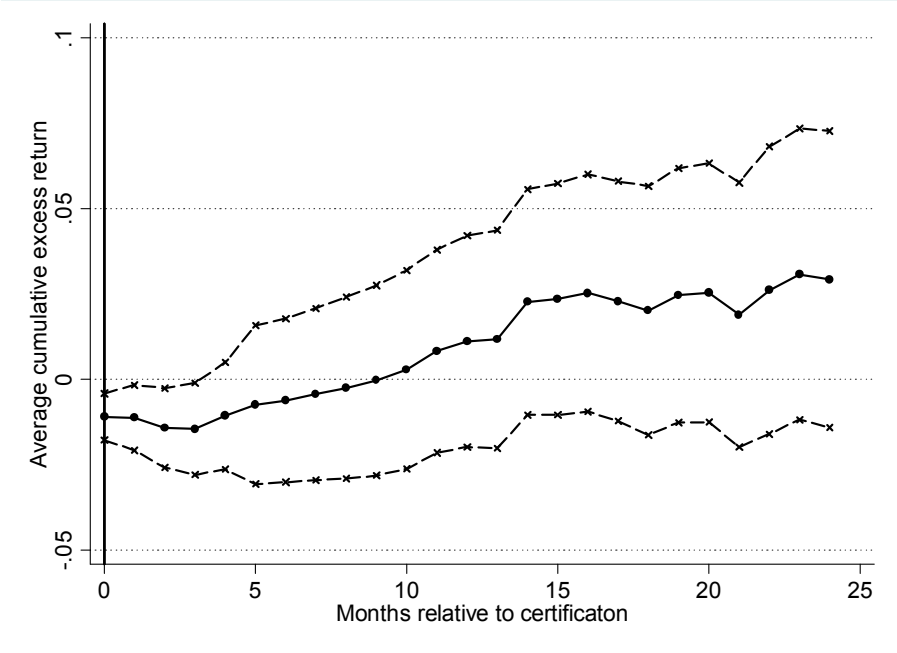

Notes: Both panels show the difference in the average cumulative excess returns between the union loss portfolio and the broad market index, as shown in Figure 9. The sample is the same one used in Figure 8. Panel A corresponds to the average cumulative excess return computed beginning 24 months prior to the election. Panel B corresponds to the average cumulative excess return computed beginning the month of the election. The dash line is the 95% confidence interval, which is computed using OLS standard errors clustering on elections.

Figure 10: Average cumulative returns of the union victory and union loss portfolios, and the market index; Elections with less than 5% of the workforce voting

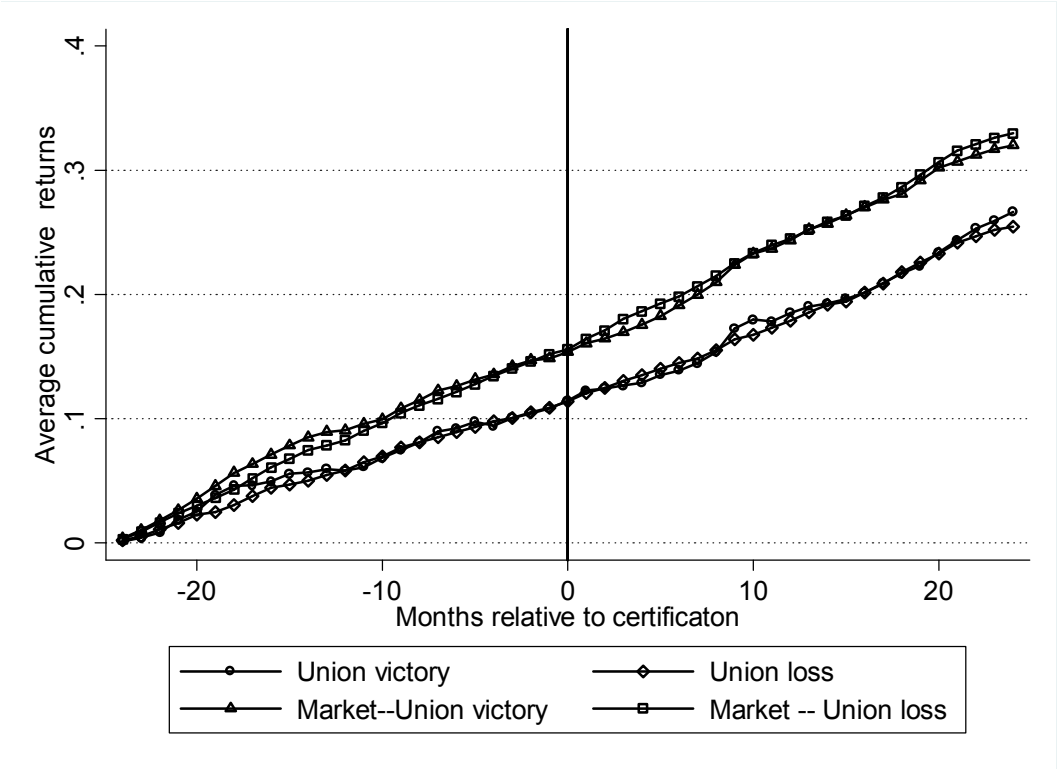

Note: The "union loss" portfolio consists of publicly traded companies holding elections where the union lost. "union victory" portfolio consists of publicly traded companies holding elections where the union won. Sample consists of publicly traded companies where less than 5% of the workforce voted. When the company's stock market return is missing it is imputed with the stock market return. Returns are net of the risk-free rate.

Figure 11: Compustat variables; Union win/loss comparisons

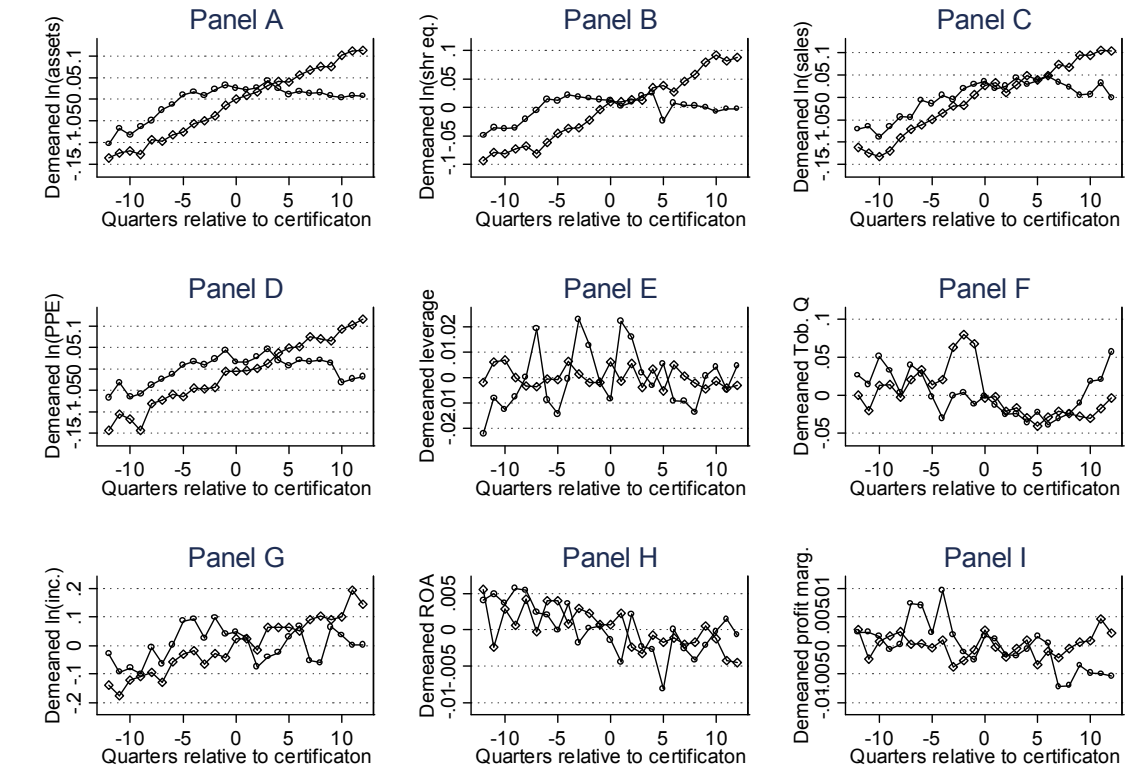

Notes: Sample consists of publicly traded companies with elections taking place between 1973-1999 where at least 5% of the workforce voted. Lines with circles correspond to union victory companies. Lines with diamonds correspond to union loss companies. All variables are drawn from the Compustat quarterly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

Figure 12: Compustat variables; Union win/loss differences

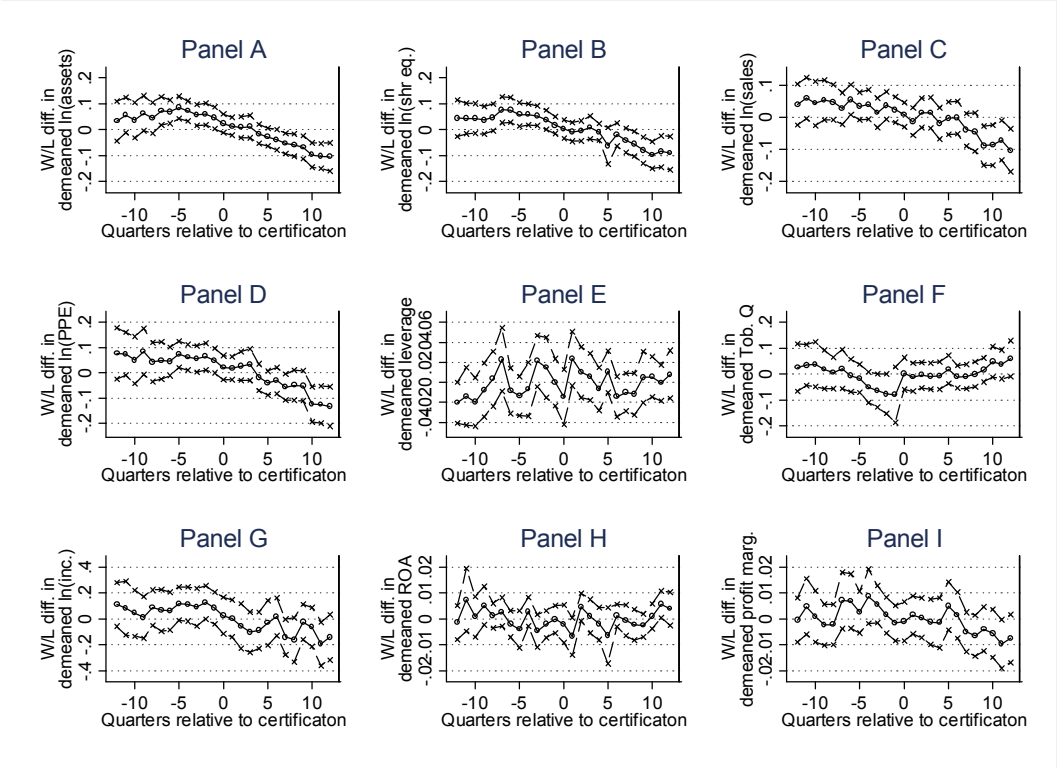

Note: See notes to Figure 11. Dashed lines are 95% confidence bands computed using OLS standard errors clustered on election.

Figure 13: Histogram of the union vote share

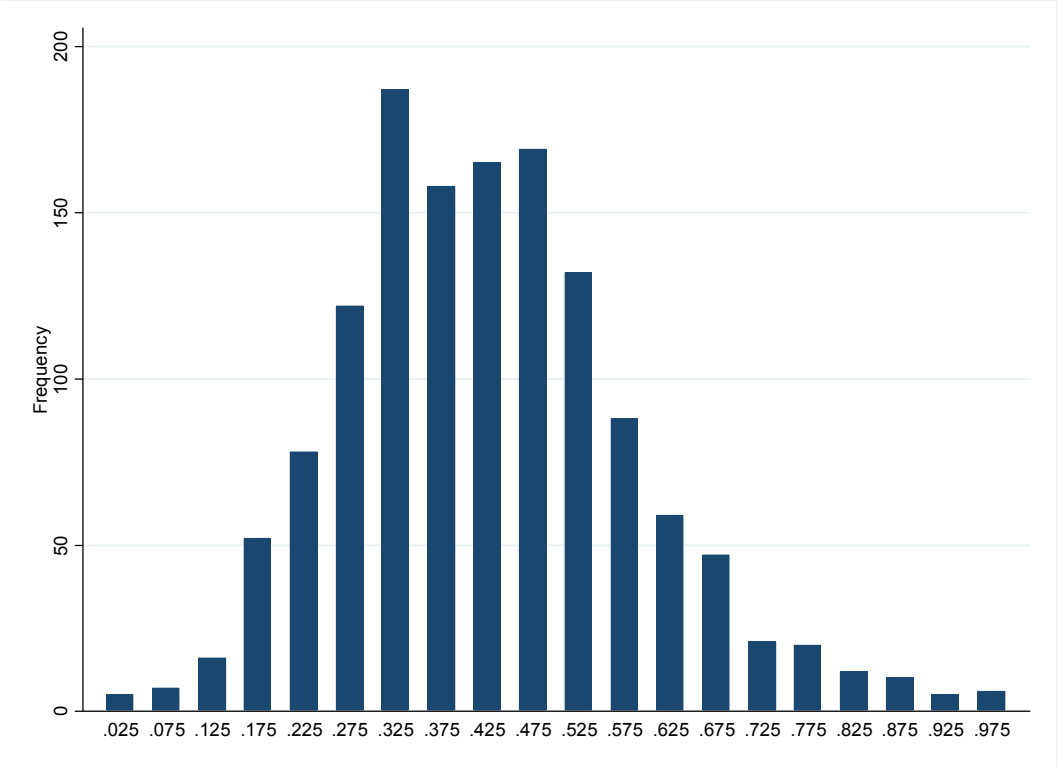

Figure 14: Cumulative market-adjusted returns in the two years after NLRB closes election, by relation to vote share

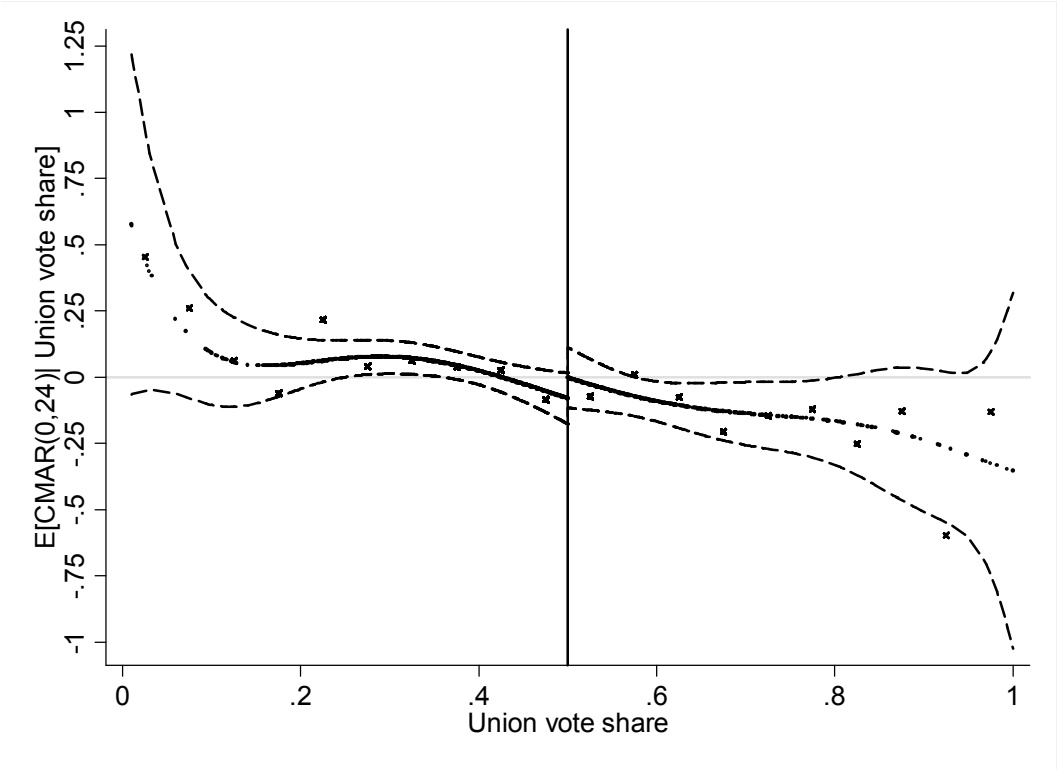

Note: Market-adjusted returns are the simple difference in the security's return and the CRSP equally weighted market index (including dividends) in the same month. Predicted values calculated using a sixth-order polynomial, and an indicator for whether the union won. Dashed lines are the 95% confidence interval. Dots are the average cumulative excess return in 20 equally spaced bins. See text for further details on the construction of this figure.

Figure 15: Cumulative excess returns in the two years after NLRB closes election, by relation to vote share

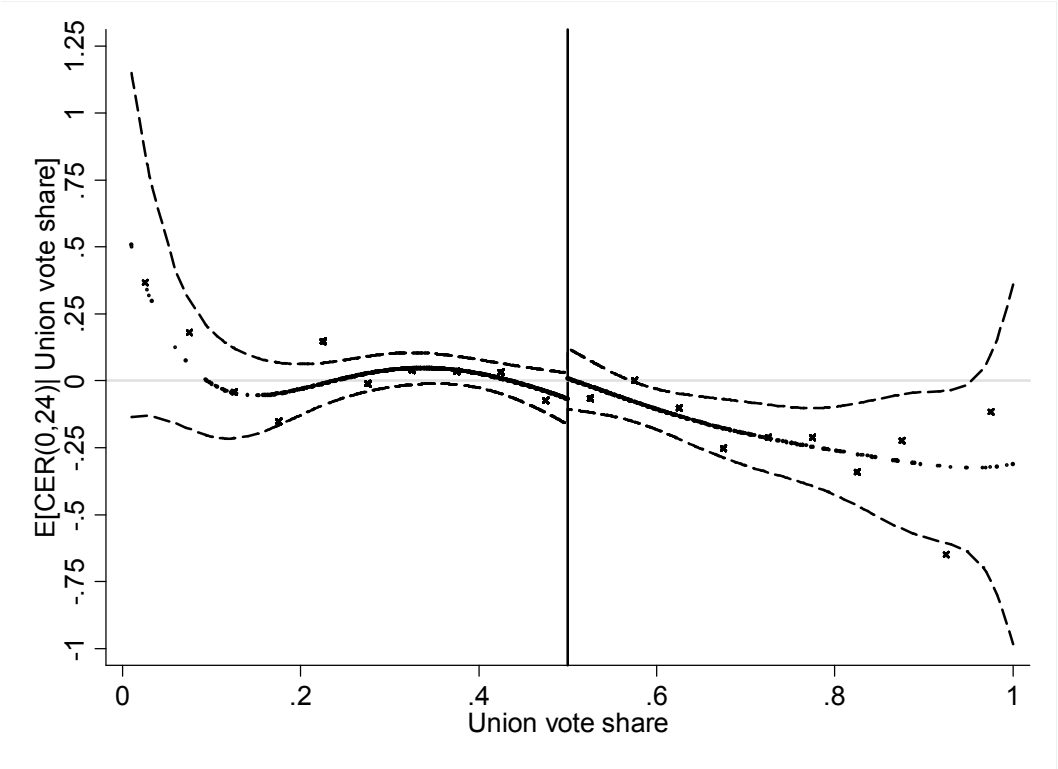

Notes: Predicted values calculated using a sixth-order polynomial, and an indicator for whether the union won. Dashed lines are the 95% confidence interval. Dots are the average cumulative excess return in 20 equally spaced bins. See text for further details on the construction of this figure.

Figure 16: Cumulative excess returns in the pre- and post-event periods, by relation to vote share

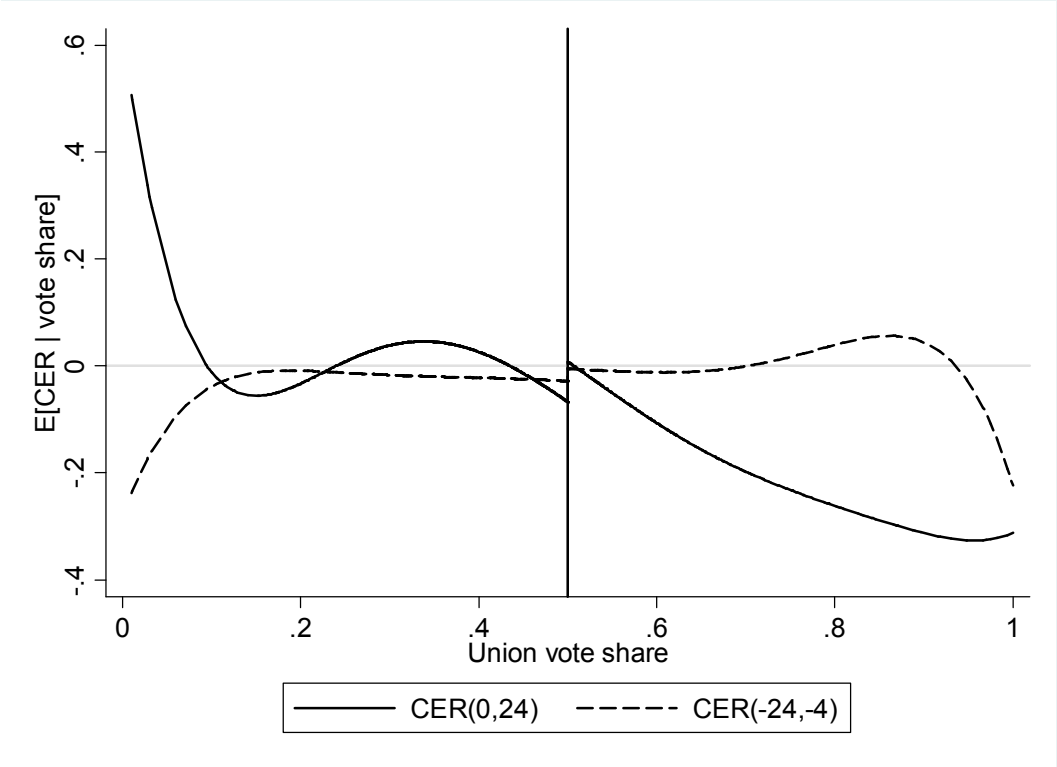

Notes: Predicted values calculated using a sixth-order polynomial and an indicator for whether the union won. Solid line corresponds to the predicted cumulative excess return in the two years following the election, conditional on union vote share. Dashed line corresponds to the predicted cumulative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

Figure 17: Cumulative adjusted excess returns in the two years after NLRB closes election, by relation to vote share

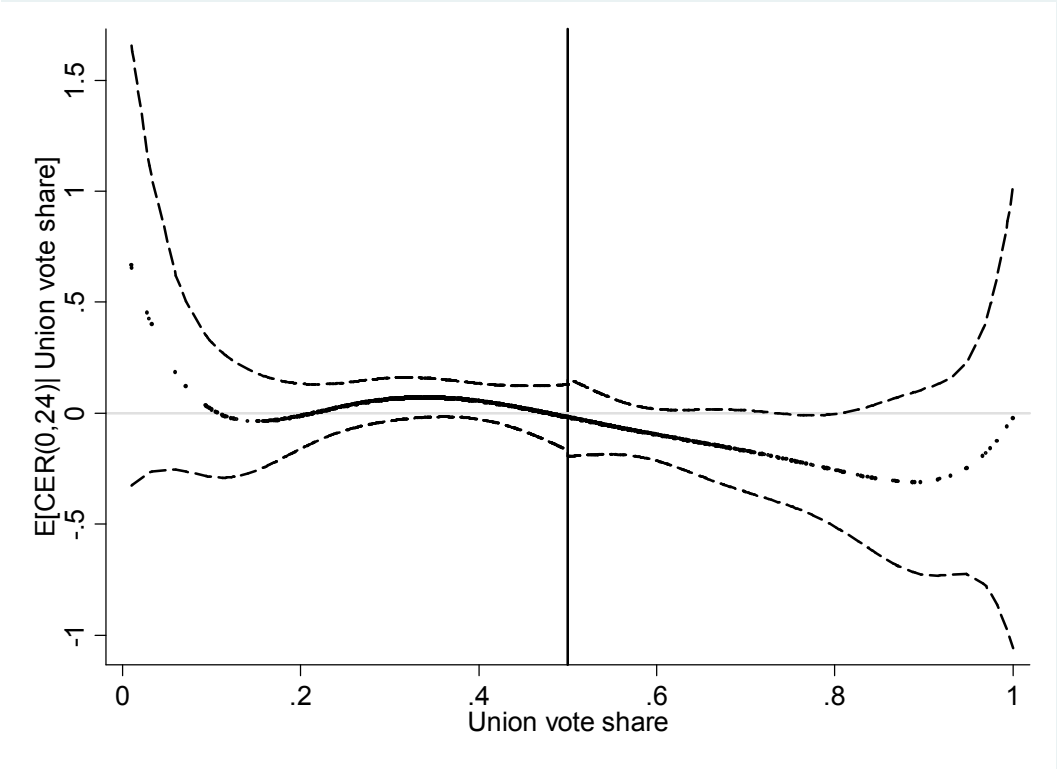

Notes: See text for details on the construction of the adjusted excess returns. Predicted values calculated using a sixth-order polynomial and an indicator for whether the union won. Dashed lines are the 95% confidence interval.

Figure 18: Reconciling the model with estimates from the research design

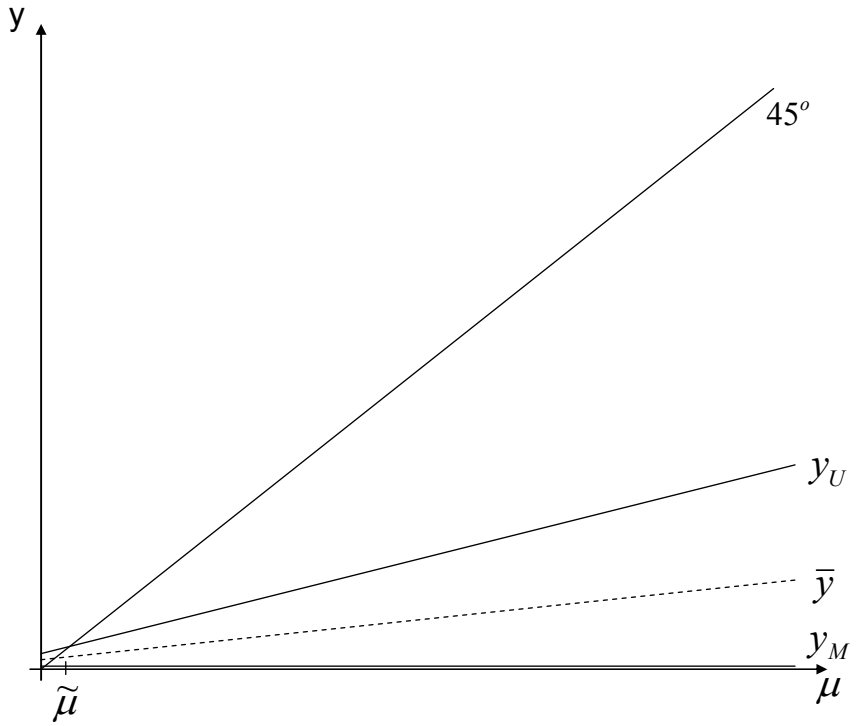

Appendix Figure 1: Union victory versus market; Average cumulative returns; non-imputed

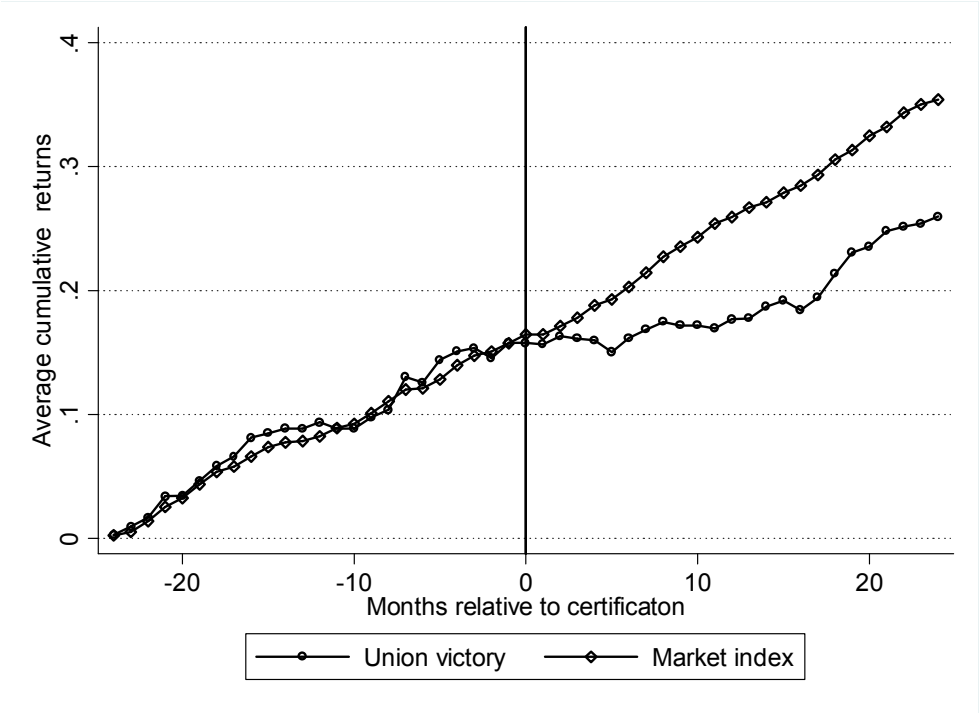

Note: All elections.

Appendix Figure 2: Union victory (with >50% vote share) versus market; Average cumulative returns; Balanced panel

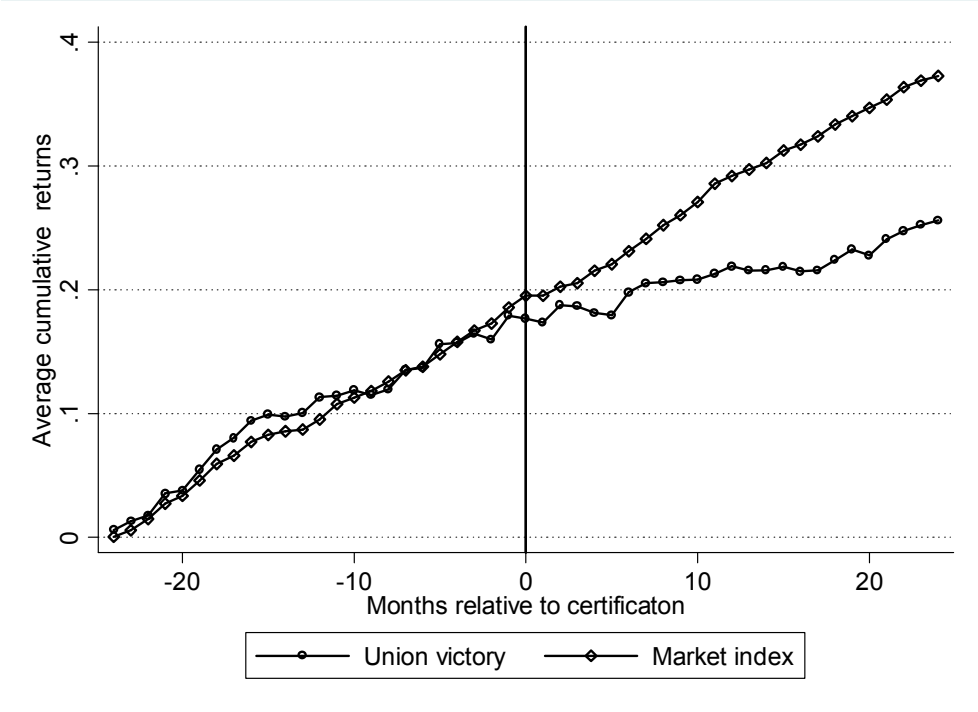

Note: Sample includes only elections with non-missing values

Appendix Figure 3: Union victory versus market; Average cumulative returns; Robustness probe; eliminate 5% most positive and 5% most negative post-event excess return elections;

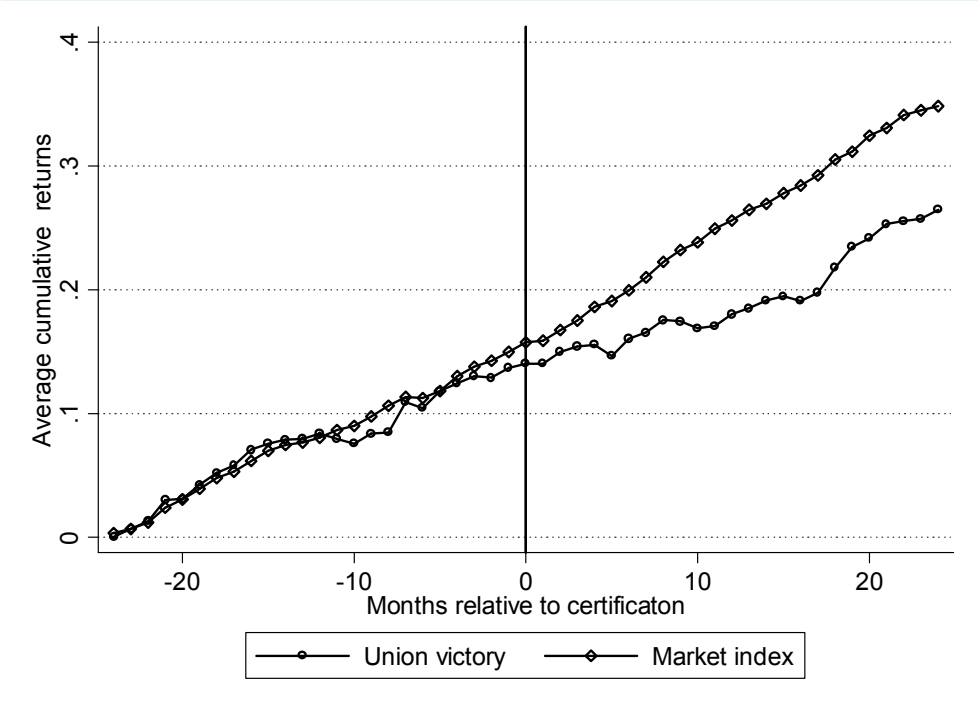

Appendix Figure 4: Union victory (with >50% vote share) versus market; Average cumulative returns; Eight year window