

EPI BRIEFING PAPER

ECONOMIC POLICY INSTITUTE • MARCH 20, 2009 • BRIEFING PAPER #230

STILL OPEN FOR BUSINESS

Unionization Has No Causal Effect on Firm Closures

BY JOHN DINARDO

*"The worst crime against working people is a company which fails to operate at a profit."
—Samuel Gompers, founder, American Federation of Labor*

The most prominent fear employers voice regarding unionization is that it will drive them out of business. But is that fear well-founded? This brief summarizes recent research showing that unionization simply does not cause firm failure: firms that become unionized are no more likely to fail than comparable firms that remain nonunion. This finding may surprise some readers. Because unions clearly do aim to give workers a larger share of the benefits of economic growth, the possibility does exist that if they succeed in transferring income to workers and away from profits, then a firm's solvency could conceivably be threatened.

However, it is also possible that even unions that successfully redistribute income from profits to wages can coexist with firms that remain viable over the long run. Only the most simple-minded and unrealistic economic models argue that there is an inexorable link between *any* such redistribution and a firm's death.

Recent research has been able to make convincing claims about the *causal* impact of unions winning recognition through an election certified by the National Labor Relations Board (NLRB) on the subsequent survival of the newly unionized firms. This research provides evidence that this causal effect of union recognition is zero and has been zero since at least the 1960s, which is how far back we can go with the available data. In short, the biggest fear voiced by employer groups regarding unionization—that it will inevitably drive them out of business—has no evidentiary basis.

Background

Students in introductory economics are often taught a simple supply and demand analysis of union impacts on firms. This analysis makes a prediction that surprises many readers: any union success in raising wages in a firm

TABLE OF CONTENTS

Background	1
The regression discontinuity evidence	3
Conclusion	7
Appendix A: What About Wages?.....	8
Appendix B: Event study estimates on firm equity values from Lee and Mas (2008).....	9

www.epi.org

above the going market wage-rate will inexorably drive that firm out of business. This brief examines some of the arguments and the evidence for this view.

Theory

The basic story that has been retold to millions of undergraduates boils down to this: In “perfectly competitive” markets, firms are passive “price-takers” (meaning that they take the market wage for workers of a given quality as pre-determined) and demand fewer workers as this market wage rises. Further, in this story there is an unlimited supply of workers available to any given firm at the pre-determined “market wage.”

In this story, unions work like a cartel that refuses to provide labor for wages less than some level that is set *higher* than this market wage. If successful, unions raise industry-wide wages and firms passively react to this increase in costs by hiring fewer workers. Depending on what is going on in other labor markets, these newly unemployed workers may remain unemployed or look for work elsewhere, depressing wages and eroding job security in nonunionized sectors.

The upshot, in this scenario, is that the union creates a “wage floor” for unionized workers (Mincer 1981) that distorts an otherwise ideally functioning labor market. The seeming logic of this is so compelling that many accounts do not go beyond it. Indeed, in some textbook accounts (see Mankiw (1997) for example) this is the only mention of labor unions.

Part of what makes such an analysis appealing is, in fact, its core weakness—its simplicity. In this abstraction of a “perfectly competitive” labor market, firms have absolutely no discretion to set wages. Workers are competing with each other for jobs, firms are competing with each other to hire workers, and these workers can both costlessly change jobs and have full information about all alternative employment possibilities.

Manning (2003) has observed the many unappealing or even absurd predictions that stem from such labor market models. First, consider what might happen if a given firm decided to lower its wage by a single penny: every last worker would quit and instantly find equivalent work at another firm paying precisely one penny more than their old job. Second, consider what might happen if

a single firm tried to offer their workers’ wages higher than the going market-wage: unit costs would rise, and the firm would lose all business to competitors (who, by definition, sell an identical product using identical technology) and it would inevitably be driven out of business.

If we abandon the single assumption (among the many embedded in this model) that workers can freely change jobs anytime a better offer is available, then it is possible that the firm could pay a range of wages while still attracting workers and remaining viable against its competitors. More importantly, this means that any conclusion about the impact of unionization on firm survival cannot be arrived at by mere “deductive logic” from a theoretical model with “reasonable assumptions”—fundamentally, the question is an empirical one.

The previous empirical research: Freeman and Kleiner (1999)

Because the necessary information about firms and unionization is difficult to obtain, there is surprisingly little research directly examining whether there is a link between unionization and firm survival. An important exception is Freeman and Kleiner (1999). Their study uses several different types of comparisons and data sets to investigate whether there is a causal relationship between unionization and firm survival.

They first use financial data to compare unionized and nonunionized firms and business lines in industries with a significant union presence. Since unionized and nonunionized firms vary in ways other than their unionization status, they use statistical controls to try to ensure that any estimated difference between the two types of firms represents a *causal* effect of unionization and not merely a correlation that represents the influence of omitted factors. Their statistical model also allows the effect of unionization to be different for different types of firms. Their best estimate of the average causal effect of unions on firms is zero—unions have no net effect on a firm’s closure.

Next, they compare union *workers* to nonunion workers who are otherwise similar in terms education, work experience, etc. Their chief finding is that unionized workers are no more likely than nonunion workers to be displaced due to the permanent closure of an establishment.

Finally, a key complementary finding of Freeman and Kleiner (1999) is that unionization does negatively affect a firm's profitability. Despite this *profitability* effect, it remains the case that unionization has *no significant effect* on a firm's survival. Given these findings, and an enormous literature finding that unions raise workers' wages, they conclude that unions redistribute income away from firm owners and managers to workers *without* committing Gomper's "worst crime against working people."

The regression discontinuity evidence

Freeman and Kleiner (1999) note that there is an "ideal way" to test whether unionization affects the survival of firms (which they are unable to conduct)—randomly assign union status to otherwise identical business units

and observe their survival over time—that is similar to the way the efficacy of new drugs are often tested.


The rest of this brief discusses research by DiNardo and Lee (2002, 2004) and Lee and Mas (2008). Their research takes up Freeman and Kleiner's suggestion by exploiting a particularity of American labor law that creates conditions very close to this "ideal way" to assess whether unionization *causes* firms to fail or whether the relationship is mere "correlation."

NLRB elections as economic experiments

By international standards, the process by which U.S. workers gain the right to bargain collectively is quite unusual. In many developed countries, collective bargaining is done on an industry-wide or even nationwide basis. In the United States, however, the right of

FIGURE A

Most elections are close NLRB election results from 2 data sets



SOURCE: DiNardo and Lee (2004).

private-sector workers (not previously unionized) to bargain collectively must frequently be secured through a process that results in workers voting in an election decided by a simple majority. If 50% plus one worker vote in favor of the union, the workers win the right to bargain collectively. With 50% or less, the workers do not win the right.

Figure A provides a histogram of union vote shares in NLRB elections. Significantly, most elections are close, clustering near the 50% mark. This confirms that these close elections will provide enough data for the empirical investigation.

This institutional set-up of NLRB elections, particularly the simple majority-win rule, affords a rare opportunity in empirical economic research to examine “quasi-experimental” data. The claim to having experi-

mental data is valid because firms where the union wins 50.01% of the vote are likely to be very similar to firms where the union wins only 49.99% of the vote—*except* for the fact that in the “50+1 firms” the workers will be unionized.


This last point is key. Near-winners should look the same as near-losers in every dimension except the union one. If so, then the near-losers will be a good “control group” for firms where the workers have just won the right to bargain collectively.

More importantly, this is not merely an assertion or assumption—it can be tested. Specifically, one can compare the “pre-determined” characteristics of firms and workers in the treatment and control groups to see if they are the same. This is exactly analogous to the procedure in a randomized controlled trial where the age, sex, and

FIGURE B

Union vote share and probability of union recognition

The “jump” at 50% means that a majority win leads to recognition



SOURCE: DiNardo and Lee (2004).

other baseline characteristics of the treatment and control group are examined post-randomization to ensure that they are the same. If they are the same, then randomization worked. If not, then the comparison is suspect.

No “jump” means no effect

After confirming that randomization has worked, DiNardo and Lee (2002, 2004) conduct the experiment by graphing a simple display of the data. Along the horizontal axis is the fraction of votes won by the union, with this fraction rising from left to right along the axis.

At each possible value of this “vote share” the probability of subsequent firm survival is averaged and then plotted. The key to assessing the causal impact is to look for a discontinuity, or, a “jump.”¹ Given enough data, firm survival where the union won (say) 3% of the vote should

look close to the survival rate where the union won 4% of the vote; and these rates at 4% should be similar to rates where the union won 5% and so on.

This pattern should continue until one reaches the 50% vote share. Points to the right of a 50% vote share represent firms where workers have collective bargaining rights, whereas those just to the left of 50% represent firms where workers do not have such rights.


Figure B presents just such a plot where the outcome is (and should be) clear: is the union recognized by the NLRB? Indeed, things are smooth (continuous) until the 50% vote share. When workers get less than 50% of the vote they almost never win recognition, but after 50% workers essentially universally secure recognition.

Of course, this simply suggests that the narrowest goal of a NLRB-sponsored election—that workers’ right to be

FIGURE C

Union vote share and probability of firm survival

No “jump” at 50% means no effect on firm survival



SOURCE: DiNardo and Lee (2004).

recognized as a union to bargain collectively—is accomplished when the majority of workers vote for the union.

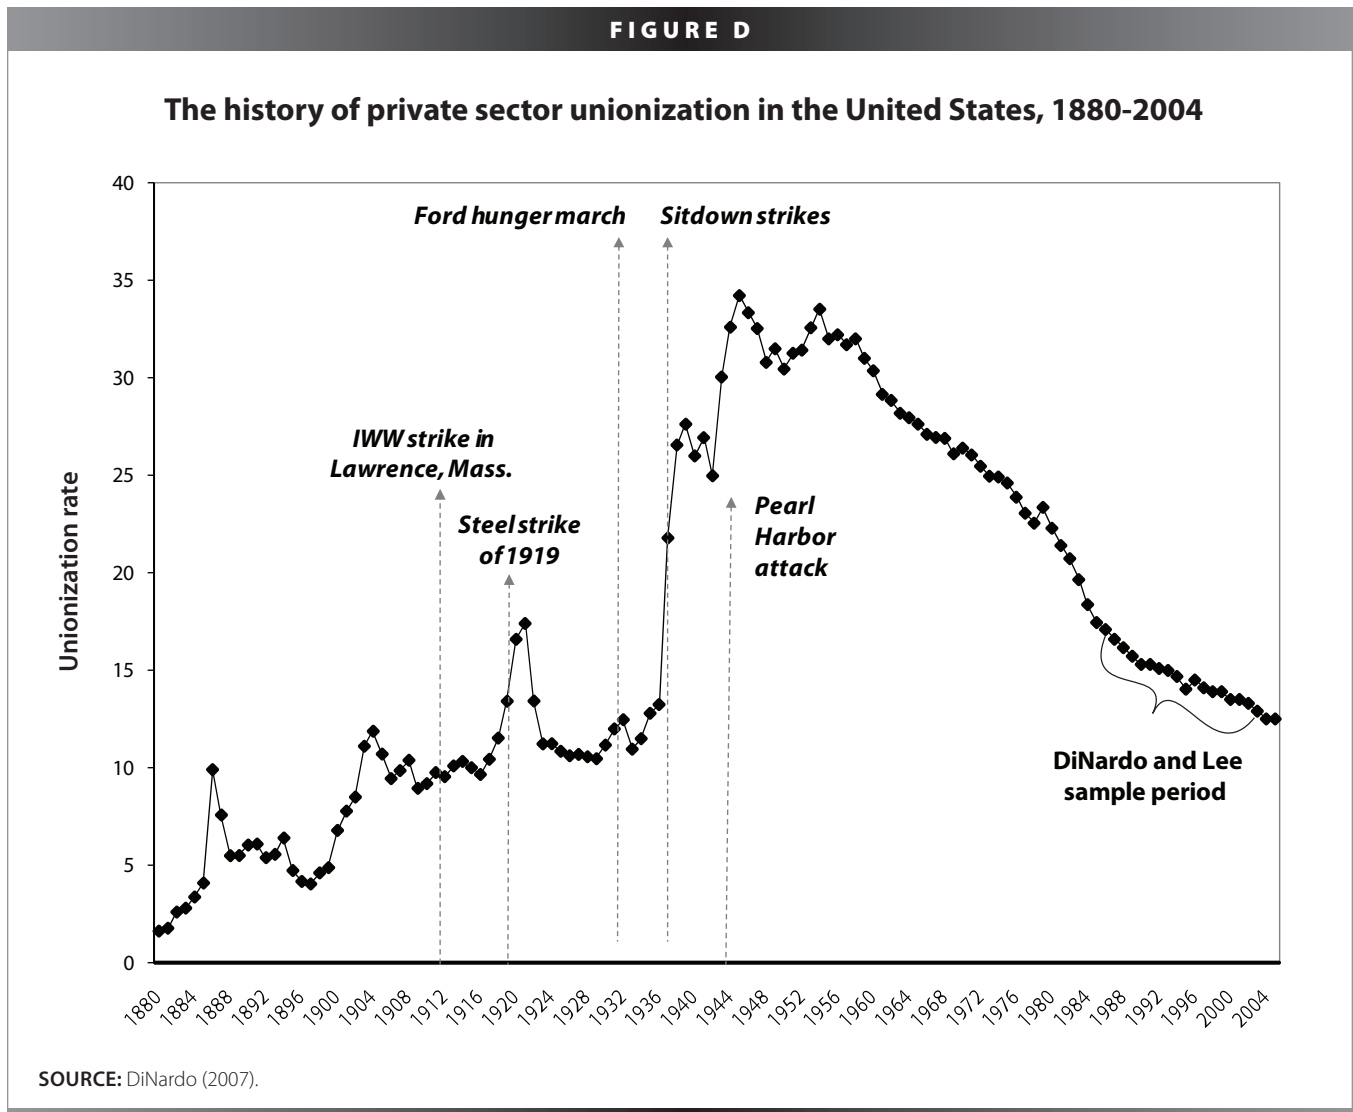
However, the basic idea is the same when assessing the impact of unionization on firm survival following an election: if there is any effect of union recognition on firm survival one should see a “jump” just at the point where the vote share exceeds 50%.

Figure C confirms the central conclusion of this brief—there is no “jump” and the right of workers to bargain collectively has no causal effect on firm survival. Firms just to the right of 50% vote share have failure rates that are almost identical to firms just to the left. This means, therefore, that firms that become unionized are no more likely to fail than firms that remain nonunion.

Has the causal effect always been zero?

One possible caveat should be discussed—the sample period studied by DiNardo and Lee (2002, 2004). **Figure D** gives a highly terse “history” of private sector unionization in the United States. As is evident from the figure, the period analyzed by DiNardo and Lee (2002) and DiNardo and Lee (2004) is on the downward sloping portion of the graph, where union coverage is declining. Private sector unionization rates have been declining for a long time, and the beginning of our sample period comes just after the end of the famous strike by the air traffic controller union (PATCO) which ended in defeat for the union.

Nonetheless, in very recent work, Lee and Mas (2008) extend the analysis in DiNardo and Lee back to the year 1961. Their conclusions from the regression discontinuity



design evidence are exactly the same as for the later period analyzed in DiNardo and Lee (as well as the evidence in Freeman and Kleiner (1999)).²

Therefore, our results demonstrating that firms that become unionized are not more likely to fail than firms that remained nonunion are not an artifact of that particular time period.

Conclusion

John L. Lewis, president of the United Mine Workers of America (UMWA), who challenged Gompers for leadership of the AFL in 1921 made the following remarks after his union signed its first industry-wide contract:

“...[with this contract] the industry can apply itself —both management and labor, to the prob-

lem of producing coal in quantity...at the lowest cost possible by modern techniques. The Mine Workers stand for the investors in the industry and for a return on capital. They stand for the public to have coal at the lowest possible price consistent with the Mine Workers having a decent life...”

There is a growing body of economic evidence that suggests Lewis’ remarks have more than a grain of truth in them, at least as a depiction of U.S. unionism in the post-World War II era. The obvious fact that unions have no stake in driving employers out of business is reinforced by this evidence. It seems clear that American employers as a group need not fear firm insolvency as a result of granting workers rights to collective bargaining.

Appendix A: What about wages?

A substantial body of economics research has established that unions raise wages. The union “wage premium,” defined as how much union status raises an individual workers’ wages holding other influences constant, is consistently estimated as falling between 5% and 40%. Other literature has established that union representation is associated with better non-wage benefits as well. The finding of a wage premium for unionized workers is robust across a wide range of researchers and methodological investigations, including:

DiNardo et al. (1996): constructs control groups based on demographic characteristics that are associated with the same “risk of being unionized.”

Ashenfelter (1978): constructs control groups based on industry, race, and worker type (i.e., craftsmen, operatives, laborers, etc.).

Freeman (1984): compares wage rates for the same individual at different points in time. At one point in time the worker is in a unionized job; at a different point in time the worker is in a nonunionized job.

Lemieux (1998): compares wage rates for the same individual who holds jobs, one of which is unionized, the other of which is not.

Krashinsky (2004): compares wage rates of *identical twins*, one who is unionized and one who is not.

Card (1992): constructs control groups based on observable characteristics that tend to receive the same wage in the nonunion sector as well as exploiting the longitudinal aspect of the data to control for differences in permanent but unobservable characteristics of individuals (i.e., in econometric parlance, person-specific fixed effects).

The DiNardo and Lee (2002, 2004) experiments looked at many outcomes following union recognition besides just firm survival. One of their findings is a puzzle that seems, at first glance, to contradict the clear consensus of the economic literature on unions and wages noted above.

The seeming puzzle is their failure to detect causal effects of union recognition upon establishment wages; closer examination, however, reveals that it is eminently possible to reconcile the DiNardo and Lee (2002, 2004) findings with the larger economics literature that finds unions are unambiguously good for workers’ wages.

This reconciliation is spelled out in great detail in DiNardo (2007). Essentially, unions increase the *pool of good jobs* available to American workers by changing the contours of wage-setting across the economy in ways both more subtle and more profound than just bestowing a “wage premium” on each individual worker who happens to be in one. Instead, over the course of decades the main impact of unions on wages was to transform large groups of firms into “good firms” that pay high wages for all employees, unionized or not. However, these good pay practices of firms did not spring out of thin air—they were surely shaped predominantly by the spread of unions.

In econometric parlance, there exist “firm specific” effects that are correlated with unionism. Essentially, the earlier literature on unions and wages was unable to control for the effect of *specific firms* on workers’ wages, and what showed up in statistical analysis as the effect of *unions* on wages may actually have been picking up the effect of working for a “good” (i.e., high-wage) employer. Put briefly, what looked to the econometrician like the effect of “unions” may have actually been largely a function of “good firms.”

Of course, this begs a couple of questions: (1) what “makes good firms,” and, related, (2) why are “good firms” so tightly correlated with unionization? Can this correlation really be nothing but coincidence? Almost surely not, rather good firms are *created* by the spread of unions.

If all of this is true, the oft-heard lament about declines in the prevalence of “good union jobs” is not merely a nostalgic memory and unions have indeed expanded the pool of well-paying jobs available to American workers and increased wages. However, Figure D made clear that the NLRB process stopped being an effective mechanism slightly more than a decade after it began with the Wagner Act, which made unions “legal.” Despite this, the beneficial effect on firms’ wage practices has apparently persisted for more than 50 years.

In closing, it should be noted that those opposed to workers' unionization have made their own verdict on this issue clear: they fully believe that unions raise wages, as evidenced by their fierce opposition to current unionization drives and any legislative change that would make unionization easier.³

Appendix B: Event study estimates on firm equity values from Lee and Mas (2008)

As noted above, Lee and Mas (2008) extend the DiNardo and Lee (2002, 2004) results reported above. They examine evidence from a longer time-period (1969-99) and look at the effect of unionization upon the *equity value* of firms. Lee and Mas (2008) replicate the DiNardo and Lee (2002, 2004) results of no causal effect of unionization when they use the regression discontinuity design. This allows them to rule out, for example, that the results in Lee and DiNardo are specific to the post-1984 period they investigate.

When they turn to another research design, however, they do find a negative effect of unionization on firms' subsequent equity values. The first thing to note about this finding is that it is actually quite consistent with much of the extant literature on unionization, as it may just provide evidence that a union victory results in a modest wage increase for workers at the possible expense of economic rents transferred from CEOs, managers, and stockholders. As noted above, especially in the Freeman and Kleiner (1999) findings, union-led redistribution from profits to wages does not necessarily imply firms will inevitably go out of business.

The additional research designs examined by Lee and Mas (2008) are "event studies" of unionization. The basic idea is simple and the method is widely employed in the empirical finance literature. A firm which experiences an "event" (in this case, a successful union election) is compared to a group of "similar firms" before and after the "event." Lee and Mas go to great lengths to construct appropriate control groups and can verify that before the election (which they take as the beginning of the event) the "to-be-unionized" firm and the comparison firms have similar time series of equity prices, etc. The differences that emerge in the period after the "event" are taken to be an assessment of the investor's reaction to the "event."

One advantage of an event study approach is that, unlike the regression discontinuity (RD) design, with an event researchers are able, in principle, to estimate an effect for *each* firm in their sample. By contrast, the RD method delivers only a specific sort of "average" effect across all the firms in their sample. When Lee and Mas apply the RD method to their financial data, their estimate of the average effect of unionization on firms is still essentially zero. With this advantage, however, there is also a significant disadvantage to an event study method in this context even if we agree that Lee and Mas are able to estimate successfully the time path of financial variables that would have occurred absent an NLRB election. In particular, event study estimates are ultimately about investors' *reactions to information*. Drawing an inference about the effect of *unionization requires* that investors' reactions (even if one is willing to maintain that they are being rational, well-informed, etc.) are *exclusively* about unionization and not any other information revealed by the event.

An alternative explanation for their results would be that the occurrence of unionization—in particular a high vote share for the union—may be a sign of managerial incompetence. It is possible that better run workplaces are less likely to face the prospect of a unionized workforce, and a decisive vote in favor of the union may indicate mismanagement of labor relations. In this case, the "event study" will confound changes in the "economic value" of the firm induced by unionization with changes due to investors' reassessment of their view about the efficacy of management. Indeed, it is a veritable cliché in the "union avoidance" industry that a first step in avoiding unionization is management attention to the needs of the workers.

Although it is not unique to Lee and Mas' carefully conducted analysis, investor reaction is *slow*—Lee and Mas use two-year windows after the "event" to measure union effects. For some in finance, such a slow diffusion of information mitigates against the validity of the estimate as a measure of the "true value of the firm" (although it may be a valid measure of investor "reactions"). In many variants of the "efficient markets hypothesis" the price of a firm's equity is supposed to reflect all information available at all points in time; such a slow reaction suggests that this is unlikely to be correct (although as Lee and

Mas point out, this may only mean it may take time for investor's to process the information—for example, half the firms in their sample do not have analyst coverage).

Further, when Lee and Mas turn to a more “speculative” analysis that involves “describing” difference in the event study estimates and attempting to estimate a “calibrated” model that matches some of the key characteristics of the actual data, they turn up some more evidence that is consistent with (if not dispositive for) the “signaling” interpretation of their results described above. The essence of the Lee and Mas calibration is to *correlate* the event study estimates with the fraction of the vote for the union.

When they do, they find that the event study estimates of the simple fact of union recognition are actually insignificant (in fact, are slightly positive) when the fraction of workers voting for unionization is included as a regressor. This fraction of union votes, however, is negatively related to firm equity value. One piece of evidence that argues for something besides union recognition as the driver of their results is the fact that the negative correlation between union vote-share and firm equity value persists even for vote-shares less than 50%—that is, it holds even in those elections that do not result in union recognition. Other interpretations of the event study analysis, however, are also consistent with this evidence. The problem is that for this part of their analysis, unlike

the RD estimates, it is hard (impossible) to conceive of something resembling “experiment.” While it is possible to imagine small changes in the threshold for representation (the experiment implied by the RD analysis) there is not an obvious analog for changes in the proportion supporting a union.

Simply put, “describing heterogeneous treatment effects” (the fact that the effect of unionization may be different for different plants) is inherently more difficult than even the very difficult effort to estimate average causal effects. The validity of estimates of average causal effects can be assessed by checking that the *ceteris paribus* conditions necessary for valid inference are indeed met. That is, the fact that “treatment” and “control” groups are identical in every way except union status confirms the experimental nature of the data and provides powerful evidence that only the causal impact of unionization is being estimated.

Describing heterogeneity is inherently a non “quasi-experimental” task, since the essence of issue is the attempt to extrapolate internally valid estimates to situations unlike the conditions in which the internally valid effect was estimated. Although Lee and Mas study the correlation between vote share and their event study estimates, the estimates from this part of their study may reflect other differences that are coincidentally correlated with vote share.

John DiNardo is Professor of Economics and Public Policy at the University of Michigan and research associate at the National Bureau of Economic Research. His research focuses on applied econometrics, labor economics, health economics, political science and econometrics. Recent work has included characterizing the finite sample performance of so-called semi-parametric treatment effect estimators, evaluating “accountability” standards for public schools, the labor market effects of Hawaii’s health insurance mandate, and the accuracy and reliability of pre-election polls, among other things. Some of his current projects include a chapter for the forthcoming Handbook of Labor Economics on Policy Evaluation Methods, a fifth edition of an econometric textbook, *Econometric Methods*, (Johnston and DiNardo), constructive proposal for attrition and non-response, and the health effects of obesity.

Endnotes

1. This is the source of the “regression discontinuity” label for this research design.
2. Lee and Mas (2008) has been cited by those claiming that unions have an adverse impact on firms, and their paper does indeed find a negative correlation between union vote share and firm equity values when they turn to another research design besides the regression discontinuity. For more details on the Lee and Mas (2008) paper, see Appendix B.
3. See Schmitt and Zipperer (2007) on employer opposition to unionization drives.

References

- Ashenfelter, Orley. 1978. “Union Relative Wage Effects, New Evidence, and a Survey of Their Implications for Wage Inflation,” in R. Stone and W. Peterson, eds., *Economic Contributions to Public Policy*, MacMillan Press.
- Card, David. 1992. *The Effect of Unions on the Distribution of Wages: Redistribution or Relabelling?* NBER Working Paper 4195, National Bureau of Economic Research, Cambridge, Mass.: NBER. October.
- DiNardo, John. 2007. “The Causal Effect of Labor Unions In the U.S. Context.” Presentation. http://www-personal.umich.edu/~jdinardo/dinardo_for_printing2.pdf
- DiNardo, John, and David S. Lee. 2002. *The Impact of Unionization on Establishment Closure: A Regression Discontinuity Analysis of Representation Elections*. Working Paper 8993, National Bureau of Economic Research. June.
- DiNardo, John, and David S. Lee. 2004. Economic impacts of new unionization on private sector employers: 1984-2001. *Quarterly Journal of Economics*. Vol. 119, No. 4, pp. 1383-1441. November.
- DiNardo, John, Nicole Fortin, and Thomas Lemieux. 1996. Labor market institutions and the distribution of wages, 1973-1993: A semi-parametric approach. *Econometrica*. Vol. 64, No. 5, pp. 1001-45. September.
- Freeman, Richard. 1984. Longitudinal analysis of the effects of trade unions. *Journal of Labor Economics*. Vol. 2, pp. 1-26.
- Krashinsky, Harry A. 2004. Do marital status and computer usage really change the wage structure? *Journal of Human Resources*. Vol. 3, No. 3, pp. 774-91. Summer.
- Lee, David S., and Alexandre Mas. 2008. *Long-Run Impacts of Unions on Firms: New Evidence from Financial Markets, 1961-1999*. Unpublished Working Paper, Princeton University 2008. Presented at National Bureau of Economic Research Summer Institute, July.
- Lemieux, Thomas. 1998. Estimating the effects of unions on wage inequality in a panel data model with comparative advantage and non-random selection. *Journal of Labor Economics*. Vol. 16, pp. 261-91.
- Mankiw, N. Gregory. 1997. *Macroeconomics*, 3rd edition. New York: Worth Publishers.
- Manning, Alan. 2003. *Monopsony in Motion: Imperfect Competition in Labor Markets*, Princeton, N.J.: Princeton University Press.
- Mincer, Jacob. *The Economics of Wage Floors*. 1981. NBER Working Papers 0804. Cambridge, Mass.: National Bureau of Economic Research, Inc. November.
- Schmitt, John, and Ben Zipperer. 2007. *Dropping the Ax: Illegal Firings During Union Election Campaigns*. CEPR Research Report. Washington, D.C.: Center for Economic and Policy Research. January.