

NBER WORKING PAPER SERIES

IS THE TIME-SERIES EVIDENCE ON
MINIMUM WAGE EFFECTS
CONTAMINATED BY PUBLICATION BIAS?

David Neumark
William Wascher

Working Paper 5631

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 1996

The views expressed in this paper do not necessarily represent those of the Federal Reserve Board or its staff. We thank Jeff Biddle, Harry Holzer, Alan Krueger, Robert Lalonde, Kevin Lang, Peter Schmidt, John Strauss, Jeff Wooldridge, and seminar participants at Michigan State, the Milken Institute, UCSD, and the University of Washington for helpful comments, David Card and Alan Krueger for making their data available, and Daniel Hansen for research assistance. This paper is part of NBER's research program in Labor Studies. Any opinions expressed are those of the authors and not those of the National Bureau of Economic Research.

© 1996 by David Neumark and William Wascher. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

IS THE TIME-SERIES EVIDENCE ON
MINIMUM WAGE EFFECTS
CONTAMINATED BY PUBLICATION BIAS?

ABSTRACT

Publication bias in economics may lead to selective specification searches that result in overreporting in the published literature of results consistent with economists' priors. In reassessing the published time-series studies on the employment effects of minimum wages, some recent research has reported evidence consistent with publication bias, and concluded that the most plausible explanation of this evidence is "editors' and authors' tendencies to look for negative and statistically significant estimates of the employment effect of the minimum wage," (Card and Krueger, 1995a, p. 242).

We present results indicating that the evidence is more consistent with a change in the estimated minimum wage effect over time than with publication bias. More generally, we demonstrate that existing approaches to testing for publication bias may generate spurious evidence of such bias when there are structural changes in some parameters. We then suggest an alternative strategy for testing for publication bias that is more immune to structural change. Although changing parameters may be uncommon in clinical trials on which most of the existing literature on publication bias is based, they are much more plausible in economics.

David Neumark
Department of Economics
Michigan State University
East Lansing, MI 48824
and NBER

William Wascher
Board of Governors of the
Federal Reserve System
20th and Constitution Avenue, NW
Washington, DC 20052

I. Introduction

In recent years, some economists have expressed concern that the journal review process has had the unintended effect of biasing published empirical results towards the rejection of economic null hypotheses. In other disciplines, meta-analyses of published papers--generally focusing on clinical trials--provide some evidence consistent with a bias towards publication of significant results (see, for example, Berlin, et al., 1989), under the (reasonable) presumption that the parameter being estimated across studies, such as the effect of a drug, is constant. In economics, however, it is considered more likely that parameters may vary across samples--especially when the samples cover different time periods--either because of changed behavior (e.g., Lucas, 1976) or model misspecification. The possibility that parameters change makes it difficult to apply standard meta-analysis techniques to the economics literature, because these techniques draw inferences regarding publication bias from changes in estimates across samples.¹

This problem is illustrated by revisiting recent work by David Card and Alan Krueger, published in the American Economic Review Papers and Proceedings, and in their book Myth and Measurement (1995a, 1995b). They apply standard meta-analysis techniques to the time-series literature on the employment effects of minimum wage laws, and show that estimated t-statistics have fallen rather than risen with sample sizes, and that t-statistics tend to be clustered around two. Based on this evidence, Card and Krueger conclude that “the studies in the literature have been affected by specification searching and publication biases, induced by editors’ and authors’

¹In the clinical trial literature--where there is little room for manipulating specifications--the typical assertion is that it is difficult to publish insignificant results (the “file-drawer” problem). If a similar problem exists in empirical economics, one response by researchers could be specification search to look for significant or “expected” results (e.g., Leamer, 1978). For ease of exposition, in this paper we use the term publication bias to refer to the file-drawer problem as well as the selective specification search that publication bias may induce, although the latter is the issue with which we are most concerned.

tendencies to look for negative and statistically significant estimates of the employment effect of the minimum wage” (1995a, p. 242), implying that earlier time-series studies of the employment effects of minimum wages should be discounted in assessing the overall evidence.²

In this paper, we show that this conclusion suffers from an incomplete application of meta-analysis techniques to economic studies. In particular, the linear reduced-form employment equations that were typically estimated in the time-series minimum wage literature apparently did not capture changes over time in the econometric relationship between minimum wages and employment. Although the patterns of coefficient estimates and t-statistics in published minimum wage studies are consistent with what would be expected in the event of publication bias, similar patterns are also evident in regression estimates generated from performing rolling regressions on a consensus time-series specification that emerged subsequent to most of the published studies, and on specifications that Card and Krueger advocate.

In addition to presenting this qualitative evidence, we suggest a method of testing the null hypothesis of no publication bias when parameters are changing. If there is a specification that is plausibly free of selective specification search, then that specification can be used as a benchmark with which to compare changes in estimates in the published studies. We suggest a statistical test for whether the changes in estimates in the published studies are significantly different from those for the benchmark specification estimated over successively longer sample periods. Both the qualitative evidence and our formal statistical test indicate that the time-series evidence on the employment effects of minimum wages is not contaminated by publication bias.

²As Card and Krueger note, De Long and Lang (1992) have extended the meta-analysis approach to study publications in economics more generally. However, De Long and Lang find evidence of a “reverse” form of publication bias, namely that “journals tend to publish papers that fail to reject their null hypotheses only when the null hypotheses are likely to be false” (p. 1257).

II. Evidence of Publication Bias

The approach used by Card and Krueger to test for publication bias in the minimum wage literature is relatively straightforward. First, in the absence of publication bias and if parameters are stable, then estimated t-statistics for employment effects of minimum wages should rise as the sample period is extended; more precisely, the log of the absolute t-ratios should rise one-for-one with the log of the square root of the degrees of freedom. This is closely related to the approach of Berlin, et al. (1989), who test for publication bias in research from clinical trials by asking whether the estimated magnitude of the effects in published studies is negatively related to the sample size. The intuition is that if only significant results get published, the estimated effect will be bigger when it is based on a smaller sample. Second, a tendency for published t-statistics for employment effects to be clustered around two can be interpreted as an indication of publication bias.

The evidence that Card and Krueger report from their meta-analysis of the minimum wage literature is consistent with publication bias. Focusing on the existing time-series studies that use quarterly data, they first report that a regression of the log of the absolute t-statistics from these studies on the log of the square root of the degrees of freedom yields a negative coefficient estimate, rather than the positive coefficient (of one) implied by the null hypothesis of no publication bias. Second, they provide a graph suggesting that the published absolute t-statistics are clustered around two; more precisely, the graph shows that a scatterplot of estimated elasticities and standard errors of these estimates fits a line with slope two quite well.

Alternatively, such evidence could be interpreted as indicating structural change over time, perhaps associated with model misspecification. If, for example, the absolute size of the minimum wage coefficient in the time-series employment equations has been declining over time, the absolute t-statistic may have fallen as more years were added to the time-series regressions, even though the

standard errors were also falling; a declining coefficient would also explain why, as shown below, the Berlin, et al. (1989) test generates evidence consistent with publication bias. Nevertheless, for reasons that are not explicitly stated, Card and Krueger conclude that publication bias is the more "likely" or "plausible" explanation (1995a, p. 242; 1995b, p. 192).³

III. Assessing the Evidence

Is publication bias the most likely explanation of this evidence? To examine this question, we attempt to mimic the following historical "experiment." Suppose researchers at different times in the past had approached the time-series data with a regression specification that was not influenced by selective specification search aimed at producing significant negative effects of minimum wages. Would the coefficients from this "benchmark" specification, estimated over varying numbers of years, have been different from those reported in the literature? A positive conclusion--in particular, if the findings of these hypothetical researchers would have generated stable estimates of minimum wage effects (and rising absolute t-statistics)--would support the publication bias explanation. A negative answer, in contrast, would point to changing parameters rather than publication bias as the explanation for the observed pattern of t-statistics and coefficient estimates in the literature.

This approach is related to the so-called "gold standard" approach in the medical literature: using a high-quality, large study with which to compare other published studies (Chalmers, et al.,

³Their conclusion may be based in part on their interpretation of the exchange between Welch (1974, 1976, and 1977) and Siskind (1977)--discussed in Chapter 6 of Myth and Measurement--as providing an example of publication bias. We leave it to the reader to examine the original articles and assess this interpretation. However, even if one accepts Card and Krueger's interpretation of this exchange, this anecdotal evidence provides no basis for interpreting other studies similarly. Regardless of whether or not one interprets their conclusions as coming down firmly on the side of publication bias, the contribution of our paper is to distinguish between publication bias and structural change.

1987). Of course, the existence of such a study--or a "benchmark" specification in our case--is not always a given, and the selection of such a study or specification from a set of alternatives is a subjective matter. Consequently, the approach we pursue in this particular context may not always be applicable and may be influenced by the subjective decisions of the researcher. In the present case, however, we believe that a good set of plausible benchmark specifications is available.

The first benchmark specification we consider is one suggested by Solon (1985), which Card and Krueger use in their book to update the existing time-series evidence. In this specification, the log of the teen employment rate is regressed on the log of the Kaitz index (an industry-coverage-weighted average of the minimum wage relative to the average industry wage), the log of the unemployment rate of adult males, the fraction of 16-19 year-olds who are 16-17, the fraction of 16-19 year-olds in the armed forces, the log of the fraction of 16-19 year-olds in the population, quarterly seasonal dummies (since the data used are not seasonally adjusted), a time trend and its square, and interactions of the time trend and its square with each of the quarter dummies. We use the same data as Card and Krueger, which they have made available via the Internet.

This "consensus" specification reflects some features that are relatively common across the entire array of time-series minimum wage studies, such as the Kaitz index specification, the inclusion of an aggregate demand indicator, and the inclusion of some supply indicators. However, it excludes the school enrollment rate, which was included in some earlier studies but has been criticized as endogenous. It also excludes lagged values of the Kaitz index, which were also included in some early studies, and which appear to yield slightly larger disemployment effects (Hamermesh, 1981). In addition, this specification includes some features that were only adopted in work that followed most of the earlier time-series studies. For example, although Brown, et al. (1983) note that correction for serial correlation was rare in the work they review, later research

corrected for first-order serial correlation with the Cochrane-Orcutt method. Similarly, Solon (1985) notes that seasonal effects are not constant over time, necessitating the inclusion of the time-quarter interactions (in part, to get to a model consistent with an AR(1) specification). Because this specification reflects current “wisdom” regarding the correct time-series specification, while excluding variables (such as the inclusion of enrollment rates or lagged minimum wage effects) that are contentious, it seems a potentially useful one with which to carry out our historical experiment.

The first panel of Figure 1 reports the estimated elasticities of teen employment with respect to the minimum wage that we obtain when we repeatedly reestimate this specification after adding an additional quarter of data (“rolling regressions”). We begin with a sample period ending in 1968, the earliest ending year of the sample periods in the time-series studies that Card and Krueger consider. We continue through 1988--the last year of data in the published studies that they cover. We also show the upper end of the 95-percent confidence interval, so that the reader can see which estimates are significantly different from zero. The graph shows that, in fact, the estimated elasticity has generally fallen in absolute value over time.

Next, we regressed the log of the absolute t-statistics generated by this procedure against the log of the square root of the degrees of freedom; OLS estimates of this meta-analysis regression are reported in column (1) in panel A of Table 1. The estimated coefficient is negative (-.41) and significant.⁴ Also reported in Table 1 is the regression of the absolute value of the estimated coefficient on the square root of the degrees of freedom (the test proposed by Berlin, et al., 1989). This estimated slope is also negative (-.004) and significant (see column (1), panel B). Finally, the second panel of Figure 1 shows that the absolute t-statistics from the rolling regressions are

⁴As Card and Krueger (1995b) note, these standard errors cannot be taken literally, since the samples on which the estimates are based overlap substantially.

clustered around two.

What are the implications of these results? Card and Krueger interpreted similar results for published studies as evidence of publication bias. However, the results in Figure 1 and in column (1) of Table 1 suggest that studies done by unbiased researchers at different points in time, approaching the data with a pre-specified or benchmark regression equation (in this case, an equation that evolved after much of the published work), would also have generated evidence consistent with publication bias. That is, the findings from the published studies may reflect nothing but objective researchers estimating a relationship that was, in fact, changing over time.⁵

To differentiate between publication bias and structural change, we need to test whether the estimated coefficients from the meta-analysis regressions based on the published studies are significantly different from those based on the estimates from the rolling regressions. One approach to assessing the significance of the differences in the meta-analysis results is to use the Hausman test framework. In particular, under the null of no publication bias, there are two consistent estimates of changes in the minimum wage effect over time: the estimate in Table 1, based on the benchmark specification (b_B); and the estimate based on the published studies (b_P).⁶ Under the

⁵There is no obvious reason why the t-statistics from the rolling regressions should be clustered around two. However, if the t-statistics for the shorter samples tended to be near two and the estimated coefficients fell over time, then such clustering might occur.

⁶That is, we are thinking of b_B and b_P as alternative estimates of the parameter β in either of the two meta-analysis regressions that we consider:

$$\ln(t) = \alpha + \beta \ln(N^{1/2}) + \epsilon$$

or

$$g = \alpha + \beta N^{1/2} + \epsilon ,$$

alternative hypothesis of publication bias, only b_B is consistent. Thus, a test of the statistical significance of the difference between b_P and b_B provides a test for publication bias. More specifically, as long as the estimate b_B from the benchmark specification is efficient, then the results in Hausman (1978) imply that under the null hypothesis of no publication bias, the statistic

$$\frac{(b_P - b_B)^2}{Var(b_P) - Var(b_B)}$$

is distributed as $\chi^2(1)$, and the null hypothesis is rejected at the α -percent level if this statistic exceeds the α -percent critical value; the efficiency of b_B eliminates the covariance between the two estimates from the test statistic.⁷

Card and Krueger report the estimated coefficient from a regression of the log absolute values of the t-statistics from the published studies on the log square root of the degrees of freedom. This estimate, which corresponds to the regression in panel A of Table 1, is -.81, with a standard error of .70. As noted above, implementation of the above test statistic also requires an asymptotically efficient estimate (under the null) of the corresponding coefficient from the time-

where t and g are the time series of t-statistics or estimated minimum wage effects from the published studies (in which case the estimates of β are denoted b_P) or the rolling regressions (in which case they are denoted by b_B).

⁷We are grateful to Kevin Lang for suggesting this framework for evaluating the significance of the difference between the estimates from the rolling regressions and those from the published studies. Note that this case differs slightly from the application of Hausman tests to model specification. In specification tests, the estimate that is efficient under the null is inconsistent under the alternative (which is the reason the specification test is of interest). However, the distributional results in Hausman (1978) for the statistic in the text do not require this restriction, and as we know b_P is inefficient under the null hypothesis (both because it is based on a limited number of possible sample periods and because it is subject to random variation in specifications across published studies), the application of the test to the question of publication bias instead requires the estimate from the benchmark specification (b_B) to be efficient under the null.

series of estimated minimum wage effects from the benchmark specification (b_B). LM tests indicated significant evidence of first-order serial correlation (but no higher-order serial correlation), as well as heteroscedasticity corresponding to increased precision of the estimates with more degrees of freedom. We therefore reestimated the meta-analysis regression correcting for first-order serial correlation and for heteroscedasticity of this form.⁸

The corresponding estimate is reported in column (1) of panel A of Table 2. The efficient estimate of the coefficient from the t-statistic regression is $-.39$, with a standard error of $.44$. As the second-to-last row of panel A reports, the p-value from the Hausman test of the null hypothesis of no publication bias is $.44$, indicating no evidence against this null hypothesis.⁹ This implies that the deviations of the elasticities in the published studies from the regression line produced by the rolling regression, although in the direction of publication bias, could easily have been produced by chance.¹⁰ A similar analysis in Panel B, based on our estimates of the regression of the absolute

⁸Because of the overlapping samples, one could in principle solve for the error structure in the meta-analysis regressions using the rolling regression estimates, with some assumptions. However, the first-order serial correlation structure fits the residuals very well, and therefore seems a good approximation to the true process.

⁹The standard error of the Card and Krueger estimate is undoubtedly biased downward because of overlapping samples, but a larger standard error would only strengthen our failure to reject the null, since the denominator of the test statistic would be larger.

¹⁰The estimates of ρ near unity in panel A suggest that the residual may be nearly $I(1)$, which might argue for using a first-difference estimator instead of the AR correction used in Table 1. However, in a closely-related estimation problem (with unknown ρ), Canjels and Watson (1994) show that the Prais-Winsten estimator is the preferred estimator when ρ is likely to be near one, but may not be exactly one. (In particular, this estimator is more robust to variation in ρ than is the first-difference estimator.) Using our data, the Prais-Winsten estimates were very similar to the maximum likelihood estimates in Table 2. For example, the estimated coefficient (standard error) of the log square root of degrees of freedom corresponding to column (1) of Table 2 was $-.47$ ($.32$). The Prais-Winsten estimates were also similar to the maximum likelihood estimates in the other columns of this panel, discussed below. Straight first-difference estimates (correcting the differenced regression for heteroscedasticity) also were negative, with larger absolute values than the estimates from the published studies reported by Card and Krueger (which is inconsistent with publication bias); however, the first-difference estimates

coefficient estimates on the square root of the degrees of freedom, leads to the same conclusion. Indeed, in this case the p-value is .99, as the estimated change in the minimum wage effect from the rolling regressions is almost identical to that from the published studies.

Our analysis to this point hinges on accepting Solon's specification as free of publication bias, and hence as a valid benchmark specification. If the benchmark specification itself reflects publication bias, then a comparison of estimates from the benchmark specification with those from published studies does not provide a valid test for publication bias. To see this, suppose that researchers (Solon included) tended to select specifications by looking for negative effects of minimum wages. They would tend to report those specifications with relatively large negative effects, for the sample period they were using. However, because the estimated effects are partly random, we would expect the estimated minimum wage effect to be attenuated as the sample period is extended, reflecting regression to the mean in the estimated minimum wage effect.¹¹ In other words, rolling regression estimates of a specification that itself reflects publication bias might generate the same qualitative patterns in the estimated coefficients or t-statistics that publication bias generates in a sequence of published studies, hence biasing our test towards finding no evidence of publication bias.

We address this concern in two ways. First, we repeat the previous analysis for two alternative benchmarks based on specifications advocated by Card and Krueger (1995b), who we assume did not engage in selective specification search to find negative minimum wage effects.¹²

were quite imprecise.

¹¹This is a potential issue because Solon's paper used data only through 1979.

¹²Given the large body of research by these authors indicating either no employment effects or positive employment effects of minimum wages, their preferred specifications should provide valid benchmarks.

The first adds the log of the adult male employment-to-population ratio as an independent variable. The second, for which they offer a priori theoretical arguments, also adds the log of the average real manufacturing wage. If these specifications were not considered by earlier researchers, or were discarded as part of their selective specification search, then the estimated minimum wage effects from these specifications should not show the same attenuation that earlier published specifications would show if they had been chosen with the goal of generating negative minimum wage effects for the sample periods under study.

The results are reported in columns (2) and (3) of Tables 1 and 2, and in Figure 2. For both specifications, the estimated minimum wage effects decline over time and the absolute values of the t-statistics are clustered around two. This is evidence that--were it to come from published studies--would point to publication bias. Moreover, this spurious evidence of publication bias is even stronger for these specifications than for the Solon specification, as the absolute values of the estimated slope coefficients in both tables are larger for these specifications. In fact, in panel B the evidence is not even in the direction of publication bias, as the estimated coefficients from the rolling regression estimates are more negative than that from the published studies (reported in the note to Table 2).

We also re-computed the tests for publication bias using data through 1993, although it seems to us that these additional observations are less relevant to the question of publication bias. Including the more recent data may provide a more efficient estimate of the change in the minimum wage effect over time, but it seems more appropriate to base our conclusions on a comparison between the published studies and what these studies would have been obtained had they used the benchmark specification. Regardless, the qualitative conclusions, as reported in the last row in

each panel of Table 2, are unchanged; there is still no significant evidence of publication bias.¹³

The second way we address the concern that the consensus specification used as a benchmark (in Figure 1) reflects publication bias is to calculate these various tests using only the estimates for samples that included observations for the years after 1979, the last year used in the literature employing the consensus specification. In particular, if that specification itself is contaminated by publication bias, we might expect the attenuation of the estimated minimum wage effect to emerge only in subsequent years. As shown in Table 3, the estimated minimum wage effects using this more limited sample of estimates do decline faster than in the longer sample. However, the same is true for the two specifications advocated by Card and Krueger, for which the estimated relationships are at least as strongly negative, suggesting that the sharper decline in estimated minimum wage effects in this latter period using the Solon specification reflects parameter changes, rather than publication bias in that specification.

Finally, while the formal test for publication bias hinges on a number of statistical assumptions, we regard this test as a useful heuristic device for assessing the statistical strength of the evidence suggesting that Card and Krueger's results from the published studies do not, in fact, reflect publication bias. In our view, however, the key result in this paper is that we replicate the qualitative results in Card and Krueger's work--which they interpret as consistent with publication bias--using benchmark specifications that are presumably free of publication bias.

IV. Why Did Minimum Wage Effects Decline?

A natural question that arises from our analysis is why “objective” estimates of minimum wage effects apparently declined as later years were added to the samples. While it is not the focus

¹³The differences in p-values indicate that the evidence is a bit more in the direction of publication bias in the t-statistic regressions when data through 1993 are included, and a bit less so in the regressions for the absolute value of the estimated coefficient.

of this paper, our case against publication bias in this particular context would be bolstered if these parameter changes could be explained. We can think of three possible sources of parameter change, all of which are in part related to a misspecification of the basic model.¹⁴

First, the Kaitz index fell over much of the 1980s, implying that the percentage of teens bound by the minimum wage may have fallen. Because the time-series estimate of the minimum wage effect in the simple linear employment equation can be viewed roughly as an average of the minimum wage effect prevailing in each year, a nonlinearity in the true relationship between employment and the Kaitz index could produce changes over time in the estimated coefficient in the misspecified linear model.¹⁵ In particular, a positive relationship between the absolute size of the minimum wage effect and the level of the Kaitz index could explain the declining minimum wage effect found when data from the 1980s were added to the sample. As it turns out, however, the data are inconsistent with this explanation. Because the value of the Kaitz index in the 1980s was still well above that experienced in the 1950s or 1960s, the sample average of the Kaitz index continues to rise as observations from the 1980s are added. This suggests that a nonlinearity of the type that could have produced an increase in the estimated minimum wage coefficient in the 1970s also would have produced a further increase as the 1980s data were included, rather than the decline we observe.

Second, early increases in the Kaitz index partly reflected changes in coverage by minimum wage laws, whereas later changes were generated mainly from changes in minimum wage levels or

¹⁴Williams and Mills (1995) offer an alternative explanation associated with the time-series properties of teenage employment rates and the Kaitz index. In particular, they claim that the decline in the minimum wage coefficient in standard regressions occurs because these two variables are integrated of different orders.

¹⁵Such a nonlinearity could arise for reasons discussed in Neumark and Wascher (1994).

average wages. If the effects of coverage changes are larger than are the effects of changes in the relative minimum wage, the effective absence of further increases in coverage after the early 1970s could generate a declining employment effect.

Finally, if the workers for whom minimum wage workers are most substitutable are not those earning the average wage, but those earning a wage lower in the distribution (such as the 25th centile), then the well-documented rise in wage dispersion over the 1980s suggests that a Kaitz index measured using the average wage would overstate movements in a Kaitz index measured using the 25th centile wage. The use of this "incorrect" Kaitz index then produces an underestimate of the absolute minimum wage elasticity. A complete assessment of this hypothesis entails a major undertaking, because it requires micro-data on wage distributions over a long period. Nonetheless, given evidence that minimum wage effects declined more sharply in the 1980s, we think this is the most plausible of the three explanations that we have offered.¹⁶

Given our findings that declines in minimum wage effects in published studies estimated over longer sample periods do not stem from publication bias, but from one of these sources of parameter change (or perhaps some other explanation), we would advise researchers not to view the earlier time-series evidence as biased. Rather, the accumulated evidence should be interpreted as showing that time-series estimates of minimum wage effects are negative and significant for some sample periods, but also that these effects fall with the addition of data from the 1980s (and early 1990s). Studying why these effects vary over alternative sample periods is likely to enhance our understanding of the employment effects of minimum wages. In addition, our perspective on the accumulated evidence leaves open the possibility that time-series regressions some years hence will

¹⁶This explanation was suggested by Lawrence Katz at an AEI conference on minimum wages, and is supported in recent unpublished empirical evidence reported by Donald Deere, Kevin Murphy, and Finis Welch.

again reveal a stronger negative effect of minimum wages on teen employment.

V. Conclusion

This paper presents an approach to testing for publication bias in the economics literature, where it may be difficult to distinguish the effects of publication bias from the effects of changes in parameters. The basic approach is to identify a benchmark specification that is unlikely to have been influenced by selective specification search, and to compare changes in the estimate of interest across published studies using different sample periods to changes in estimates from the benchmark specification estimated over similar periods. This approach contrasts with the clinical trial literature, in which the null hypothesis of no publication bias is based on a comparison of estimates across published studies using different sample sizes with a parameter assumed to be constant.¹⁷

We apply this approach to evaluate the claims by Card and Krueger (1995a, 1995b) that the time-series evidence on minimum wage effects reflects selective specification search motivated by authors' and editors' predisposition toward finding and publishing results that conform to conventional economic theory. It is true that the absolute magnitude of estimated minimum wage effects in published studies fell as the time series were extended, consistent with publication bias. However, it is also true that the absolute magnitude of the estimated effects falls over time for specifications that emerged in the literature subsequent to all or most of the published work that Card and Krueger criticize, including specifications that they advocate on a priori theoretical grounds. In addition, a statistical test indicates that the reduction over time in the estimated

¹⁷In this paper we discuss the approach as applied to time-series studies. In principle, a similar approach could be applied to cross-sectional studies, where true parameters may vary by location, demographic group, etc., perhaps leading to corresponding variation in parameter estimates across published studies. However, because we would not necessarily expect variation in sample size to be correlated with variation in the representation of different locations or demographic groups in the data, such parameter variation is less likely to point to publication bias when meta-analysis methods from the clinical trial literature are used.

minimum wage effect in the published studies is not significantly different from the reduction in the estimates from the benchmark specifications using successively longer sample periods.

Thus, we conclude that the reduction in estimated minimum wage effects in published studies using more recent data is due to structural change or misspecification of the employment equation, rather than to publication bias. While we do not make any claim that these findings generalize to other areas of empirical research in economics, provisionally, at least, our evidence should make researchers less skeptical of published evidence that is consistent with widely-held priors.

References

- Berlin, Jesse A., Colin B. Begg, and Thomas A. Louis." 1989. "An Assessment of Publication Bias Using a Sample of Published Clinical Trials." Journal of the American Statistical Association, 84, 381-92.
- Brown, Charles, Curtis Gilroy, and Andrew Kohen. 1983. "Time-Series Evidence on the Effect of the Minimum Wage on Youth Employment and Unemployment." Journal of Human Resources, 18(1), 3-31.
- Canjels, Eugene, and Mark W. Watson. 1994. "Estimating Deterministic Trends in the Presence of Serially Correlated Errors." Mimeograph, Northwestern University.
- Card, David, and Alan B. Krueger. 1995a. "Time-Series Minimum-Wage Studies: A Meta-Analysis." American Economic Review, 85(2), 238-43.
- _____. 1995b. Myth and Measurement: The New Economics of the Minimum Wage (Princeton, NJ: Princeton University Press).
- Chalmers, T.C., H. Levin, H.S. Sacks, D. Reitman, J. Berrier, and R. Nagalingam. 1987. "Meta-analysis of Clinical Trials as a Scientific Discipline, I: Control of Bias and Comparison With Large Cooperative Trials." Statistics in Medicine, 6, 315-325.
- De Long, J. Bradford, and Kevin Lang. 1992. "Are All Economic Hypotheses False?" Journal of Political Economy, 100(6), 1257-72.
- Hamermesh, Daniel S. 1981. "Minimum Wages and the Demand for Labor." NBER Working Paper No. 656.
- Hausman, Jerry. 1975. "Specification Tests in Econometrics." Econometrica, 46, 1251-71.
- Leamer, Edward P. 1978. Specification Searches: Ad Hoc Inference with Nonexperimental Data (New York: John Wiley and Sons, Inc.).
- Lucas, Robert E. Jr. 1976. "Econometric Policy Evaluation: A Critique." In Brunner and Meltzer, eds., The Phillips Curve and Labor Markets (Amsterdam: North Holland).
- Neumark, David, and William Wascher. 1994. "Minimum Wage Effects and Low-Wage Labor Markets: A Disequilibrium Approach." NBER Working Paper No. 4617.
- Siskind, Frederic. 1977. "Minimum Wage Legislation in the United States: Comment." Economic Inquiry, 135-38.
- Solon, Gary. 1985. "The Minimum Wage and Teenage Employment: A Reanalysis with Attention to Serial Correlation and Seasonality." Journal of Human Resources, 20(2), 292-7.

Welch, Finis. 1974. "Minimum Wage Legislation in the United States." Economic Inquiry, 12, 285-318.

_____. 1976. "Minimum Wage Legislation in the United States." In Orley Ashenfelter and James Blum, eds., Evaluating the Labor Market Effects of Social Programs (Princeton, NJ: Princeton University Press).

_____. 1977. "Minimum Wage Legislation in the United States: Reply." Economic Inquiry, 15, 139-42.

Williams, Nicolas, and Jeffrey A. Mills. 1995. "The Minimum Wage and Teenage Employment: Is There a Relationship?" Manuscript, University of Cincinnati.

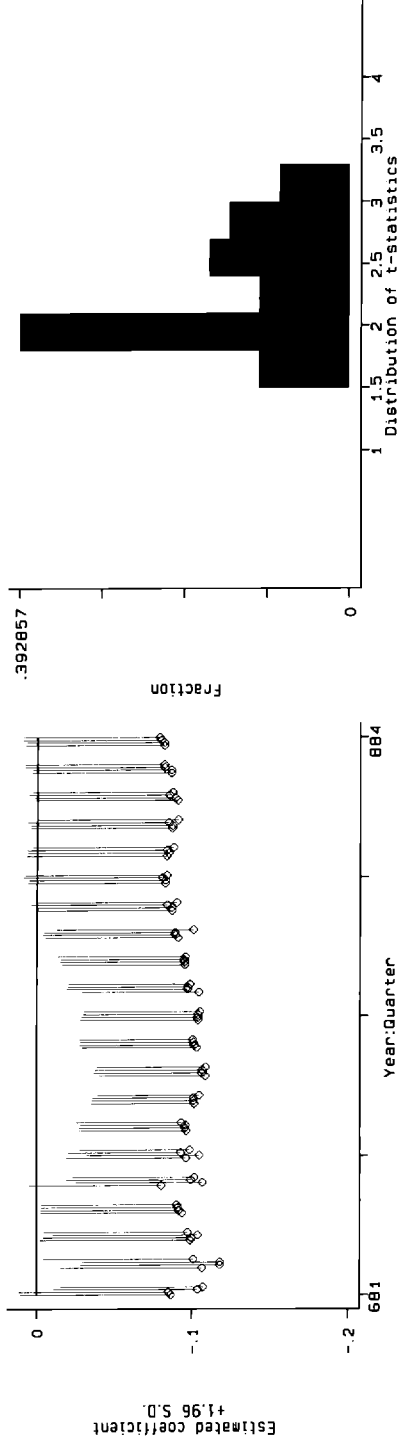
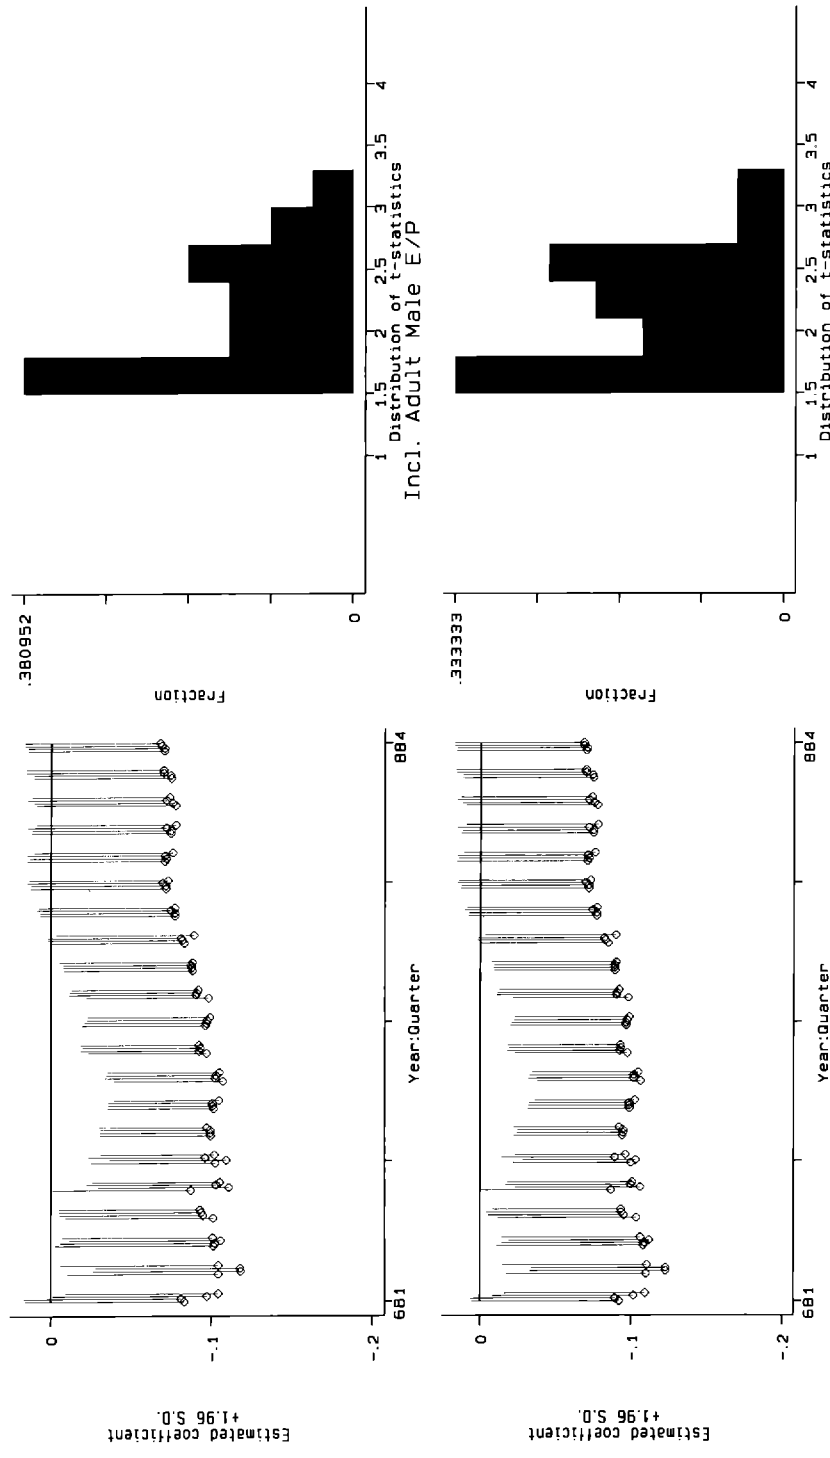


Figure 1--Consensus (Solon) Specification, 1968--1988



Incl. Log Real Manu. Wage and Adult Male E/P

Figure 2--Card/Krueger Specifications, 1968-1988

Table 1: Meta-Analysis OLS Regressions for Time-Series Estimates of Employment Effects of Minimum Wages, Quarterly Data, 1968-1988

A. Regressions of Log Absolute Values of T-Statistics on Log Square Root of Degrees of Freedom

	<u>Consensus</u> <u>Specification</u> (1)	<u>Incl. Adult</u> <u>Male Empl./Pop.</u> (2)	<u>Also Incl. Log Real</u> <u>Manu. Wage</u> (3)
Constant	1.69 (.25)	2.18 (.27)	2.31 (.23)
Log square root of degrees of freedom	-.41 (.12)	-.66 (.12)	-.73 (.11)
R ²	.13	.26	.36

B. Regressions of Absolute Coefficient Estimates on Square Root of Degrees of Freedom

Constant	.13 (.005)	.16 (.005)	.16 (.005)
Square root of degrees of freedom	-.004 (.0006)	-.008 (.0006)	-.008 (.0005)
R ²	.38	.67	.75

Standard errors are reported in parentheses. There are 84 observations used in each column. Minimum wage employment equations from which the coefficients and t-statistics are obtained are described in the text.

Table 2: Meta-Analysis Regressions for Time-Series Estimates of Employment Effects of Minimum Wages, Quarterly Data, 1968-1988, AR(1), Heteroscedasticity-Corrected Estimates

A. Regressions of Log Absolute Values of T-Statistics on Log Square Root of Degrees of Freedom

	Consensus <u>Specification</u> (1)	Incl. Adult <u>Male Empl./Pop.</u> (2)	Also Incl. Log Real <u>Manu. Wage</u> (3)
Constant	1.57 (1.00)	2.02 (1.12)	2.06 (.94)
Log square root of degrees of freedom	-.39 (.44)	-.64 (.49)	-.65 (.41)
ρ	.94 (.04)	.95 (.03)	.94 (.03)
P-value, H_0 : no publication bias	.44	.74	.77
Data through 1993	.24	.56	.61

B. Regressions of Absolute Coefficient Estimates on Square Root of Degrees of Freedom

	Consensus <u>Specification</u> (1)	Incl. Adult <u>Male Empl./Pop.</u> (2)	Also Incl. Log Real <u>Manu. Wage</u> (3)
Constant	.15 (.01)	.17 (.01)	.17 (.01)
Square root of degrees of freedom	-.006 (.001)	-.009 (.001)	-.009 (.001)
ρ	.74 (.08)	.65 (.09)	.69 (.08)
P-value, H_0 : no publication bias	.99	.80	.78
Data through 1993	.79	.98	.96

Standard errors are reported in parentheses. There are 84 observations used in each column. Minimum wage employment equations from which the coefficients and t-statistics are obtained are described in the text. The heteroscedasticity correction is based on a regression of the squared residuals on the log of the degrees of freedom in Panel A, and the degrees of freedom in Panel B. After transforming the data to obtain homoscedastic errors, maximum likelihood was used to estimate the AR(1) model. The test for publication bias is based on a Hausman test of the significance of the difference between the estimated coefficient from these regressions and the estimate from the published studies. For Panel A, the estimate from the published studies, reported by Card and Krueger, is -.81 (with a standard error of .70). For Panel B, the estimate (computed from the estimates in Table 6.1 of *Myth and Measurement*, and degrees of freedom supplied by Alan Krueger) is -.006 (with a standard error of .012).

Table 3: Meta-Analysis Regressions for Time-Series Estimates of Employment Effects of Minimum Wages, Quarterly Data, 1980-1988, AR(1), Heteroscedasticity-Corrected Estimates

A. Regressions of Log Absolute Values of T-Statistics on Log Square Root of Degrees of Freedom

	<u>Consensus</u> <u>Specification</u>	<u>Incl. Adult</u> <u>Male Empl./Pop.</u>	<u>Also Incl. Log Real</u> <u>Manu. Wage</u>
	(1)	(2)	(3)
Constant	3.63 (.67)	3.89 (.75)	3.96 (.77)
Log square root of degrees of freedom	-1.28 (.28)	-1.44 (.32)	-1.47 (.33)
ρ	.69 (.12)	.70 (.12)	.72 (.12)

B. Regressions of Absolute Coefficient Estimates on Square Root of Degrees of Freedom

Constant	.16 (.02)	.17 (.02)	.17 (.02)
Square root of degrees of freedom	-.007 (.002)	-.009 (.002)	-.009 (.002)
ρ	.47 (.15)	.59 (.14)	.61 (.13)

Standard errors are reported in parentheses. There are 36 observations used in each column. See notes to Table 2 for details.