

15 YEARS OF RESEARCH ON U.S. EMPLOYMENT AND THE MINIMUM WAGE

December 2015

ABSTRACT

After a quiet period in the 1980s, statistical analysis of the minimum wage and particularly its effect on employment has become very active. Since the exchange between David Neumark and William Wascher, and David Card and Alan Krueger in the *American Economic Review* (December 2000), more than 60 analyses of U.S. data have appeared, of which 37 report results in a way that makes them easily comparable and suitable for meta-analysis. We find a moderate degree of publication bias in this literature, but no support for the proposition that the minimum wage in the United States has had a noticeable effect on employment.

Paul Wolfson
Tuck School of Business

Dartmouth College
paul.j.wolfson@tuck.dartmouth.edu

Dale Belman
School of Human Resources
& Labor Relations
Michigan State University
dalebelman@outlook.com

15 YEARS OF RESEARCH ON U.S. EMPLOYMENT AND THE MINIMUM WAGE

Reinvigorated by the 1991 Cornell conference on the minimum wage and a series of increases in Federal, State and Local minimum wages, research on the effects of the minimum wage has blossomed over the last quarter century. More than 800 scholarly articles in English have been published on some aspect of the minimum wage in the years since the conference, more than 600 since the exchange between David Neumark and William Wascher (2000) and David Card and Alan Krueger (2000) at the end of 2000. A substantial plurality considers the effect of the minimum wage on some aspect of employment.

The very richness of estimates poses a challenge both for researchers and policy makers. Because the results, particularly the employment estimates, range widely, summarizing the findings is not straightforward. There is no immediate answer to the often asked question, “What is the effect of the minimum wage on employment?” One approach to answering this question has been to choose the estimates which the writer finds particularly convincing because of methodology, data, authorship and, perhaps, the results. Although picking and choosing the “best” studies has a long history in reviews of many literatures, it is particularly unsatisfactory here because of the large number of high quality studies, the richness of methods, data, outcome measures and time periods considered and because of ongoing controversies over the measured effect. Too often the choice of “best” appears to coincide with the author’s views on the minimum wage, undermining the legitimacy of the choice. What is then required is a more disciplined, transparent and reproducible method of obtaining a common estimate of the effect of the minimum wage across studies.

Meta-analysis, and in particular meta-regression, is a means of estimating a mean effect

across studies. Used widely in the medical and physical sciences, the techniques are well established. Meta-regression allows researchers to control for heteroscedasticity, for publication bias, and for heterogeneity in studies. This latter feature is particularly important in application to studies of the minimum wage as there has been a blossoming of methods, data sets and periods under study, each of which might account for differences in findings.

Card and Krueger (1995) may well contain the earliest formal meta-analysis of a minimum wage literature. Their primary concern was whether publication bias could explain the clear response of teenage employment in what was still the dominant literature, studies of the minimum wage based on aggregate U.S. time-series data. Doucouliagos and Stanley (2009: DS hereafter) demonstrated more rigorously that publication bias was indeed an issue in a superset of the studies that Card and Krueger (1995) considered. Of at least as much interest, they produced an unbiased meta-estimate of that employment effect. The meta-sample that Belman and Wolfson (2014, ch. 4: BW hereafter) examine overlaps that of DS but both begins and ends later than theirs, and includes results from several other OECD countries.

While limiting consideration only to the U.S., the present analysis otherwise extends that of BW with a dataset that is more than twice as large, from almost twice as many studies of U.S. data. The expansion has two sources. One is the inclusion of studies that have appeared in the meantime. The other is the inclusion of several earlier studies that did not explicitly report all the data needed for a meta-analysis, but reported enough information to calculate what was missing.

This study builds on techniques presented in DS and Stanley and Doucouliagos (2012: SD hereafter). Foremost among these is the formal statistical technique of meta-regression with

which it is possible to address several conditions that would otherwise make it difficult to derive reliable meta-estimates of the employment effect. We extend the methods used to address heterogeneity by relying on the LASSO to develop a specification for the final meta-regression equation. This technique better accounts for issues of multi-collinearity in a model with many explanatory variables than prior approaches used for meta-regression.

After accounting for heteroscedasticity, publication bias, and heterogeneity and non-independence of the estimates, the results for this literature are similar to those that DS and BW reported: some evidence of publication bias in favor of negative employment effects, and employment elasticities that are both negligible and far from statistically non-significant.

Some Econometric Issues of Meta-analysis

The analysis below relies on two complementary techniques of meta-analysis. For each, the unit of observation is an estimated effect, the meta-variable to be analyzed and explained: in this analysis, the elasticity of employment with respect to the minimum wage. One technique, the funnel plot, is graphical. It is a scatterplot of elasticity estimates against the inverse of the standard error. The other, meta-regression, has the appearance of a formal quantitative technique but, in the end, is also descriptive. It involves the specification and estimation of a regression model that describes the distribution of estimates. As both address similar issues, and these are more conveniently introduced in the context of meta-regression, we will start there.

Meta-regression is not only descriptive but also generates the result desired of meta-analysis, aggregating the results of several studies while correcting several issues that frequently characterize bodies of empirical work.¹ The simplest way to aggregate a body of results is to

¹What follows draws very heavily on SD, especially chapters four through six.

average them, in which case each estimate can be written as in equation (1):

$$Effect_i = \overline{Effect} + u_i = b_0 + u_i \quad (1.1)$$

This immediately suggests a regression equation, whence “meta-regression”. SD identify three issues in estimating (1.1):

1. Heteroscedasticity
2. Sample Selection bias due to publication bias
3. Heterogeneity

Heteroscedasticity: In the pair $\{Elasticity_i, SE_i\}$, $Elasticity_i$ is understood as the estimated mean of a population from which sample I is drawn, and SE_i is the standard deviation of its sampling distribution. That implies that in equation 1.1, $u_i \sim N(0, SE_i^2)$, a conventional case of heteroscedasticity. Neglecting this in calculating the average effect results in a larger standard error than otherwise. The obvious solution is to weight each observation by the reciprocal of SE_i^2 :

$$\frac{Elasticity_i}{SE_i} = \frac{\overline{Elasticity}}{SE_i} + \frac{u_i}{SE_i} \quad (1.2)^2$$

(where $Elasticity$ replaces the more general $Effect$ of equation 1.1). With this reweighting, the dependent variable is transformed into the t-statistic from the regression that generated the estimate. Defining a new variable, *precision*, as the reciprocal of SE , $precision_i = 1/SE_i$, this becomes:

²This reweighting, dividing each observation by SE_i , is identical to weighted least squares, multiplying the main diagonal of the variance-covariance matrix by SE_i^{-2} .

$$t_i = b_0 \bullet precision_i + v_i, v_i \sim N(0,1) \quad (1.3)$$

Notice that the estimate of the coefficient b_0 is the estimated mean effect size or elasticity.

Sample Selection Bias: In many research areas, suspicion arises that for one reason or another, not all analyses see the light of day, and that this outcome is systematic. One widely suspected reason is that journal editors are not interested in results that do not reject the null hypothesis, or that authors believe this to be true and do not submit such results for publication.³ Card and Krueger (1995: chapter 6) discuss a second possibility, the difficulty of publishing estimates that contradict widely held theoretical priors. In an ideal literature that consists only of studies that are well designed and executed, and where the only differences between studies are the specific sample used, one would expect the complete set of estimates for a specific effect to be symmetrically distributed about the true value of that effect. With publication bias for either of the reasons mentioned, one would expect the distribution about the mean to be asymmetric.⁴

In a conventional selection problem, some data is available on those who are not selected and more is available on those who are. For example, the researcher interested in the factors affecting women's wages may observe all women, but only observe a wage for women who are working. In this situation, where both a selection and a measurement equation can be estimated,

³This suspicion is so widely mentioned that specific citations to it are beside the point. Instead, google "publication bias null hypothesis" or "publication bias statistical significance", or read the Wikipedia article "Publication Bias".

⁴The sample we examine includes estimates from analyses that were published in one of two ways: either formally in a scholarly journal or book, or as a working paper perhaps prior to formal publication elsewhere. The latter group further includes several analyses labeled "forthcoming," so accepted for formal publication somewhere. Consequently, "publication bias" takes in both decisions of editors and referees, as well as those of researchers about what to write up in working papers.

the error term from equation (1) would consist of two parts, one of which is perfectly correlated with the selection error term (Davidson and MacKinnon 2004, 11.7):

$$Elasticity_i = b_0 + u_i = b_0 + \rho \cdot \sigma \cdot e_{1i} + e_{2i} \text{ where } e_{2i} \sim N(0, \sigma_2^2) \quad (2.1)$$

with homoscedasticity, where e_{1i} is the error term from the selection equation and σ is the constant standard deviation of u_i . Combining this with a heteroscedastic u_i , this becomes

$$Elasticity_i = b_0 + u_i = b_0 + \rho \cdot \sigma_i \cdot e_{1i} + e_{2i} = b_0 + \rho \cdot SE_i \cdot e_{1i} + e_{2i} \quad (2.2)$$

where σ_i is the variable standard deviation of e_i ; i.e., SE_i .

The inverse Mills ratio (IMR), e_{1i} in equations 2.1 and 2.2, is the expected value of the selection error term conditional on both being selected and the values of the regressors in the selection equation. The coefficients on the IMR proxy the unobserved factors relating to selection that also affect the measurement equation. Inclusion of the IMR removes coefficient bias associated with selection.

Publication bias poses a different and more challenging situation, as it is difficult -- likely impossible -- to assemble information on unpublished studies. Selection correction would be an obvious means of correcting coefficients for publication bias, but absent information on the unpublished work, the selection equation regressors, much less their values, cannot be known and the IMR can not be calculated. If the IMR were (roughly) constant, it could be folded into the estimation of equation 2.1 or 2.2, but that is clearly a convenient rather than an accurate assumption. SD (section 6.3) develop two approximations, and, based on simulations, recommend one for situations in which the true effect size (b_0) is zero and the other when the true effect size is non-zero. The former, for no effect, is (and hereafter omitting the subscripts that index the observation, i.e., using vectors for the variables in the equation)

$$Elasticity = b_0 + b_1 \bullet SE + u \quad (2.3)$$

With the correction for heteroscedasticity (as in equation 3), this becomes

$$t = b_1 + b_0 \bullet precision + v \quad (2.4)$$

The second approximation, for a situation of a non-zero effect, is

$$Elasticity = b_0 + b_1 \bullet SE^2 + u \quad (2.5)$$

and its heteroscedasticity-corrected counterpart is

$$t = b_1 \bullet SE + b_0 \bullet precision + v \quad (2.6)$$

In each of these equations, 2.3-2.6, the estimated value of coefficient b_1 is a measure of the publication bias and b_0 is the meta-estimate of the minimum wage elasticity of employment in this literature.

Of course, these equations look much like 2.1 and 2.2, derived under the assumption of a roughly constant IMR. We will revisit this issue following the discussion of heterogeneity.

Heterogeneity: If the differences in the estimated effects were due to nothing but their being calculated from different samples drawn from the same population, heterogeneity would not be an issue. That is rarely the case, especially in the literature of the *New Minimum Wage Research*, which draws not only on a much greater variety of data sources than the older literature, but uses a variety of statistical frameworks and techniques, and econometric specifications. All of these contribute to heterogeneity, and ideally, should be modeled to understand the source of different estimates of the effect of interest. Tests for heterogeneity exist, the most widely recognized being Cochran's Q statistic, which is the sum of squared residuals from estimating equation (1.3). Under the null hypothesis that the sample of estimates does not exhibit heterogeneity, the Q-statistic is distributed $\chi^2(N-1)$, where N indicates the

sample size, the number of estimates in the sample.

If heterogeneity is found, an obvious solution is to include regressors that control for the heterogeneity such as dummy variables for different characteristics suspected of contributing to it. What might these characteristics be in the case of the employment elasticity of the minimum wage? Possibilities include differences in groups studied; whether the estimate is for a specific industry such as the restaurant industry or a specific demographic group (teenagers, females), the data source or survey, and so on. For analyses that provide more than one estimate, latent common factors may be present that affect all of them, including the average quality of the study. In this case, it is appropriate to include a dummy variable for each study, much like panel fixed effects.⁵ Considering heterogeneity alone, this would modify equation 1.1 to 3.1:

$$Elasticity = b_0 + \mathbf{Z}\underline{B}_0 + u \quad (3.1)$$

where \mathbf{Z} is a matrix of variables used to control for the heterogeneity and \underline{B}_0 is a coefficient vector. Next, combining equations 3.1 and 1.3 to correct for both heterogeneity and heteroscedasticity gives us equation 3.2:

$$t = b_0 \bullet precision + \mathbf{diag}(precision) \bullet \mathbf{Z}\underline{B}_0 + v \quad (3.2)$$

where $\mathbf{diag}(precision)$ indicates the diagonal matrix of appropriate rank with elements of precision on the diagonal. If we were next to fold in the first pass at resolving publication bias, i.e., under the admittedly questionable assumption of a constant IMR, we would combine (3.2) with (2.4)

$$t = b_1 + b_0 \bullet precision + \mathbf{diag}(precision) \bullet \mathbf{Z}\underline{B}_0 + v \quad (3.3)$$

or (2.6)

⁵See Stanley and Doucouliagos 2012, section 6.2, especially 6.2.2.

$$t = b_1 \bullet SE + b_0 \bullet precision + \mathbf{diag}(precision) \bullet \mathbf{Z}\underline{B}_0 + v \quad (3.4)$$

However, a variable IMR, that is heterogeneity in the IMR can be addressed in the same way as heterogeneity in the studies themselves. With this recognition, equation 3.3 becomes

$$t = b_1 + \mathbf{K}\underline{B}_1 + b_0 \bullet precision + \mathbf{diag}(precision) \bullet \mathbf{Z}\underline{B}_0 + v \quad (3.5)$$

and 3.4

$$t = b_1 \bullet SE + \mathbf{diag}(SE) \bullet \mathbf{K}\underline{B}_1 + b_0 \bullet precision + \mathbf{diag}(precision) \bullet \mathbf{Z}\underline{B}_0 + v \quad (3.6)$$

where \mathbf{K} is a matrix of another set of variables, perhaps overlapping those in \mathbf{Z} , and \underline{B}_0 is the corresponding coefficient vector. This is the approach that SD recommend, equation 3.5 when there is no detectible effect (i.e. $b_0=0$ and $\underline{B}_0=\underline{0}$), and 3.6 when an effect can be detected. Of course, one cannot know which is appropriate in advance, so it makes sense to estimate both and compare the results.

Funnel Plots: Funnel plots are a scatterplot of two variables central to meta-analysis: the effect central to a specific meta-analysis, in this case the minimum wage elasticity of employment, and the estimated standard-error or its reciprocal, in this case precision.⁶ The name comes from the shape of the plot in a well-behaved sample, like an upright or inverted funnel, depending, respectively, on whether the y-axis is the standard error or precision. In a well-behaved sample, one typically sees that precise estimates (with small standard errors) are near the mean or median values of the effect size and that the range of estimates around the central tendency rises as precision falls (and standard error rises).

As with other statistical graphs, funnel plots are useful for becoming familiar with the data and for identifying errors, especially important since the data for meta-analyses are often

⁶Occasionally, the y-axis displays the degrees of freedom, the sample size, or its square root.

collected and recorded as part of the study. Specifically for meta-analysis, funnel-plots are used to look for suggestions of publication bias. Sampling error implies that the results from a random sample of analyses will be symmetrically distributed about the mean. A large deviation from symmetry in a funnel plot is a strong hint that publication bias has affected the sample.

Data

Gathering estimates of the minimum wage elasticity of employment begins with a search for empirical analyses of the minimum wage. We entered the phrase “minimum wage” into Google Scholar and several electronic databases of published articles or working papers: ISI-Web of Science, Econlit and the NBER. We limited ourselves to analyses of U.S. data that were published after 2000, either as articles in journals or as working papers. The beginning of the *New Minimum Wage Research* can be dated to a 1991 conference at Cornell University, and the exchange between Neumark and Wascher (2000) and Card and Krueger (2000) in the December 2000 issue of the *American Economic Review* marks the end of its first period. Our study includes analyses after this first period. We identified sixty analyses, working papers and published articles, that satisfied these criteria and included at least one estimate of the effect of the minimum wage on employment. Thirty six of these analyses either reported one or more elasticities and their standard errors (either explicitly or as estimates in a log-log equation) or we were able to calculate them from the estimates of a semi-log equation in which the dependent variable was a binary indicator of employment status, in combination with other data included in the analysis. The number of estimates per analysis varied greatly.⁷

⁷Appendix 1 lists these studies and the number of estimates garnered from each. The mean number of estimates in a study is 26 (s.d.=31) and the five-number summary is {1, 6, 15, 27.5, 125}. The total number of estimates in our study is 939.

Table 1 contains the information that we recorded to control for heterogeneity. The last two variables in Table 1 require some explanation. First, especially as the debate turns toward issues around identifying the effect of the minimum wage, it is becoming common to demonstrate that one could, if one chose, derive estimates similar to those found in earlier work, thus demonstrating that the results to come are not due to peculiarities of the particular data studied. Typically, this is done to make clear the source of differences with prior work, and not because the analyst has any confidence in these estimates. While they list several possible criteria for selecting estimates from each study, SD recommend including all estimates in the meta-analysis, and recording information about each to differentiate them.⁸ Thus we have included these estimates but wish to identify them so that we can examine their effect on our results.

Second, Gittings and Schmutte (2015) propose a hypothesis that the employment effect of the minimum wage varies systematically according to the tightness of the labor market. They examine it by arranging different industries into deciles (or more aggregated quantile) according to measures of the tightness of their labor markets, and estimate separate sets of employment elasticity for each quantile. It is not evident to us that these estimates here are comparable to the rest in our meta-study, and with 116 of their 125 estimated in this way, we feel a need to identify and examine the effects of these estimates as well (and will refer to this subsample as the GS subsample).

Variable Selection

⁸SD, sections 2.4.4 and 2.6. See also DS, and de Linde Leonard, Stanley, and Doucouliagos 2014.

SD and DS recommend Hendry's General-to-Specific (G2S) approach for variable selection; this can be roughly described as several backward stepwise regressions, each with a different variable eliminated at the first stage, and with a variety of specification and residual tests performed after each step to ensure model validity. In what may be the penultimate stage, there are several competing models which are compared to see if any encompasses the others. If not, one can restart the process with the union of the regressors in the competing model to see if a more spare specification can be achieved.⁹

There are a number of problems in applying this technique to meta-regression. To begin, the G2S approach was developed to capture the salient characteristics of an underlying data-generation process (DGP) in a reasonably spare specification. It is not obvious that thinking in terms of DGPs is useful in meta-analysis, a formal method of describing or characterizing a body of research. Does a body of research have a well structured DGP? At a more practical, less abstract level, the variables coded to capture the heterogeneity in this literature are multi-collinear; it is not possible to include all of them simultaneously in a regression, and the choice of which to exclude initially is necessarily arbitrary and may bias the outcome. Trying a sufficiently large (random) sample of these variables to feel confident that this has not occurred would be a mammoth undertaking.

We have instead used a version of the LASSO technique (Least Absolute Shrinkage and Selection Operator: Tibshirani 1996) that takes account of the clustering of our data by the analysis from which it was derived (Belloni, Chernozhukov, Hansen, and Kozbur 2015).¹⁰ The

⁹Hoover and Perez (1999) provide a careful discussion of this technique.

¹⁰We thank Christian Hansen for making available his stata program for this procedure.

LASSO is similar to ridge regression in that both minimize the sum of squared residuals subject to a summing up constraint concerning the regression coefficients. For ridge regression, the sum of squared coefficients must not exceed some value while for the LASSO, it is the sum of the absolute value. After standardizing the variables in the set to be considered, \mathbf{X} , the LASSO solves the following constrained optimization problem:

$$\min \|y - X\beta\| \text{ s.t. } \sum_{j=1}^N |\beta_j| \leq t \quad (4)$$

for some value of t . For sufficiently small t , the optimal values of $\{\beta_j\}$ will include some values equal to 0, that is some variables in the X matrix will not be selected. This typically results in a sparser equation specification than ridge regression. Informally, the LASSO minimizes the sum of the squared residuals by selecting first the independent variable that is most highly correlated with the dependent variable, then the one that is most correlated with a residualization of the dependent variable on the first variable selected, and so forth. It continues until inclusion of the next variable would violate the constraint.

The LASSO can be applied in different ways. Belloni, Chernozhukov and Hansen (2014) describe a *naive* approach, running the LASSO on the equation in question after specifying any variables of interest that must remain in, or be locked into, the equation. When the analysis concerns the effect of a treatment variable, for example the minimum wage,

[t]he problem with this approach can be seen by noting that LASSO and any other high-dimensional modeling device targets prediction, not learning about specific model parameters. From the standpoint of prediction, any variable that is highly correlated to the treatment variable will tend to be dropped since including such a variable will tend not to add much predictive power for the outcome given that the treatment is already in the model. Of course, the exclusion of a variable that is highly correlated to the treatment will lead to substantial omitted-variables bias if the [corresponding coefficient] is nonzero. Such omissions will happen routinely in any procedure that looks just at the equation above. (p. 36).

More simply put, the *naive* LASSO may well omit confounding variables, leading to a biased estimate of the effect of the treatment variable. Their preferred alternative, the *double* LASSO, is to apply the LASSO to both the response and treatment variables, and define the list of regressors in the equation ultimately to be examined as the regression of the response variable on the treatment variable and all the variables that appear in either set of LASSO results. We present results both for the *naive* LASSO with locked-in variables and a *multiple* LASSO, run on not one but several indicator variables for the employment responses of specific groups: teenagers, Eating and Drinking Establishments, those without a high school degree, males, and females.

Research into the employment effect of the minimum wage has at times been contentious and provoked unseemly, even defamatory, comments about participants from some who have not themselves been otherwise involved.¹¹ The previously mentioned exchange between Neumark and Wascher (2000), and Card and Krueger (2000) is decorous in comparison, with the former merely insinuating about poor technical choices (“We have not focused on CK's methods of analysis, which have drawn negative reactions from others (e.g., ...), although our impression of the data presented in CK's Reply ... is that they raise some questions about the validity of the assumptions needed to interpret the difference-in-differences estimates as a natural experiment.”) and the latter merely alluding to questionable statistical practice (“Instead, we suspect the

¹¹James M. Buchanan’s remark that “we have not yet become a bevy of camp-following whores.” in an Wall Street Journal opinion column (April 25, 1996) is the most famous. Writing about this work for Bloomberg a year earlier (April 23, 1995), Paul Craig Roberts concluded that “Political correctness seems to have crept into the inner sanctum of the AEA, discrediting its scholarly journal and debasing its top prize. Unless the association cleans up its act, it can kiss its credibility good-bye.” Leonard (2000) cites less vituperous but nevertheless peculiar comments by two other prominent economists.

common denominator is that representative samples show statistically insignificant and small differences in employment growth between New Jersey and eastern Pennsylvania, while the nonrepresentative sample informally collected for Berman produces anomalous results.”)

The passions that this topic has aroused suggests the importance of controlling for the source of estimates. This is reinforced by the large variation in the number of estimates from each analysis. With three analyses contributing only one estimate each, and three others each contributing over 100 estimates, the latter would swamp the former in our calculations.

Although the LASSO technique took account of the clustering, in no case did the *naive* LASSO select any of the analysis fixed effects when left to itself; consequently we present results where they were locked in prior to running the *naive* LASSO.¹² The *multiple* LASSO, on the other hand, consistently selected between five and ten of these variables, so we did not overrule it by locking in any variables other than precision and the term correcting for publication bias (i.e., the constant or se). The variables that the LASSO may include in the **K** and **Z** matrices are the same. Once a regressor list is specified, we redefine the dummies in deviation form so that the point estimate of b_0 can be interpreted as the equally weighted average effect from all the analyses (rather than from equally weighting all the estimates, neglecting for the moment that the estimates are already differently weighted to correct for heteroscedasticity).

¹²Fixed effects in the equations with *Elasticity* as the dependent variable correspond to fixed effects interacted with precision in the equations with *tstat* as the dependent variable, i.e., heteroscedasticity-corrected equations. For brevity, we will refer in the text to both, henceforth, as fixed effects, and be quite explicit when we are referring to fixed effects in the *tstat* equations that are not interacted with precision.

Funnel Plots and Other Graphs

Figure 1 displays the data in a funnel plot and two marginal box plots. Elasticity is displayed on the horizontal axis, with the tick marks indicating the minimum and maximum values in the data, -2.235 and 1.510; the unweighted mean, -0.050; one standard deviation either side of the mean, -0.342 and 0.243, and elasticities of ± 1 . The vertical dashed lines indicate the mean and the mean \pm one standard deviation. Precision, SE^{-1} , is shown on the vertical axis, which is drawn to a logarithmic scale. The tick marks, labeled not with the precision values which have no obvious meaning but with the SE value from which they are calculated, are of the same unitless scale as elasticity. The most precisely estimated elasticity in the data, at the top, has a standard error of 0.005 and the least precise has a standard error of 1.414. The mean value of precision is $0.056^{-1} = 17.8$ and one standard deviation above the mean is $0.025^{-1} = 40$ (one standard deviation below the mean is negative, so is not shown). The horizontal dashed line indicates the 90th percentile of precision.

The horizontal box-plot at the bottom right of figure 1 displays the median and extreme values of elasticity; the other two members of the five-number summary, not easily shown on this graph, are listed to the left, along with the sample size. The median elasticity, -0.032, is slightly larger than the mean, indicating perhaps a very slight skew to the left. The vertical box-plot on the upper left display the distribution of precision, with tick marks labeled by its reciprocal, SE. The five-number summary is (1.414⁻¹, 0.141⁻¹, 0.083⁻¹, 0.051⁻¹, 0.005⁻¹), and the 90th and 99th percentiles are 0.027⁻¹ and 0.009⁻¹. The funnel plot is not symmetric about its mean, with substantially fewer estimates on the lower right than the lower left, suggesting some

publication bias. However, the funnel plot is not truncated (certainly not severely so), so any estimate of publication bias is unlikely to be severe.

Figure 2 highlights the estimates for the teenage, minimum wage, employment elasticity. The black Xs are the GS labor-market tightness decile estimates. The dark gray triangles are the rest of the teenage elasticity estimates, and the small faint gray circles are the rest of the estimates, those not for teenagers. Two things stand out. First, the GS decile estimates are as a group quite different from the rest of the teenage estimates. Second, publication bias appears more clearly for the group comprising the non-GS teenage estimates, with asymmetry much in evidence in the lower parts of this plot.

Table 2 displays descriptive statistics for several subsamples (the first line shows the whole sample, for ease of comparison). Row 2 shows the distribution of the sample estimates for the employment elasticity of teenagers, roughly half the sample. Rows 3 and 4 further disaggregate this subsample into the GS subsample, roughly a quarter of the teenage estimates, and the remainder of the teenage sample. These three rows make very clear that the GS subsample is shifted to the right of the rest of the sample for teenagers. The other important point from this table comes from the information in the last three columns: the correlation coefficient between point estimates and standard errors of the elasticities, the standard errors of these correlation coefficients, and the p-value (for the null hypothesis that the correlation is zero in the population). In a sample without publication bias there should be no (statistically significant) correlation between these two statistics.¹³ Publication bias leads to truncation, most likely of the

¹³This statement does not apply to the GS sample since their hypothesis is that the employment elasticity varies systematically across labor-market tightness deciles.

estimates farthest from the mean of the untruncated sample, and this will result in the correlation between the point estimate and the estimated standard error. Only two of the correlation coefficients in the table are not statistically significant (for the Male employment elasticities and those for Eating and Drinking Establishments). What these different correlations suggest is that publication bias is likely a problem, but the differences in statistical significance across subsamples, and among those that are statistically significant, the variation in sign and size indicate that its effects are not homogenous within the sample.

Figure 3 shows box-plots summarizing the distributions of elasticity estimates within each of the analyses, for the whole sample in the top of the figure and for teenagers in the bottom. The box-plots are sorted by the median value for each analysis, from low median at the left to high median at the right. Several analyses at each end are labeled (see appendix 2 for a mapping of labels to analysis) as are those with the most extreme individual estimates. Although there is much overlap of distributions within the middle of the overall distribution, there is a very strong hint that analyses at both ends of these figures are different from those elsewhere.¹⁴

Meta-Estimates

Table 3 contains some simple meta-estimates of the elasticity for the full sample of 933 observations. First, for comparison, is the conventional mean, a meta-estimate uncorrected for any of the three issues discussed in the first part of the paper. The precision weighted – heteroscedasticity corrected – mean, -0.024, is half as large, and its standard error is less than a quarter as large. The unweighted median, -0.032, is about two thirds the corresponding mean,

¹⁴A similar plot sorted by year does not indicate any systematic variation over time in the distribution of estimates.

indicating some left skew and the IQ range occupies the interval[-0.127, 0.045]. The precision weighted (heteroscedasticity corrected) median is half the size of the precision weighted mean and the weighted IQ range is less than a quarter as large as the unweighted one. The skewness indicated by the medians relative to the means is consistent with the asymmetry of the funnel plots.

Table 4 displays estimates corrected for both heteroscedasticity and publication bias. Both the linear correction for publication bias (eq. 2.4) and the quadratic correction (eq. 2.6) indicate elasticities that are both statistically significant and very small, roughly 0.02, about the same size as the estimate that corrects only for heteroscedasticity. This last comparison anticipates the next result, publication bias that is not statistically significant by either correction. This last result is, however not definitive in light of Cochran's Q statistic which strongly indicates heterogeneity. Under the null hypothesis of homogeneity, it is distributed as $\chi^2(939-1)$, and with a value of 3569 it strongly rejects the null.

Table 5 displays results for three specifications each of which corrects for heteroscedasticity and heterogeneity, and two of which correct for publication bias. In the third, the coefficient b_1 and vector \underline{B} of equations 3.5 and 3.6 are constrained to be zero, preventing any correction for publication bias. Comparison of this set of estimates with the other two makes it possible to gauge the extent and effect of publication bias. For each of these three specifications, two sets of results are reported, the *naive* LASSO and the *multiple* LASSO.

The most striking characteristic of these results, before turning to the specifics of the estimates, is the stability of the standard error of each type of elasticity estimate within each row, that is across all the specifications. For instance, the standard error for precision lies entirely in

the interval [0.070, 0.099], while that for precision + teen lies in [0.084, 0.100]. The next broad pattern is that for both types of publication-bias correction, each elasticity point estimate from the *naive* LASSO lies to the right of the corresponding estimates of the *multiple* LASSO. With one exception (NoHS, equation 3.5), this means that the *naive* LASSO estimate is smaller in magnitude than the corresponding estimate from the *multiple* LASSO. Finally, in no case is any pair of estimates from publication-bias corrected specifications so far apart that we would conclude that they are not equal. For the moment, consider only the results in the first four columns of the table, those taken from specifications that include corrections for publication bias.

Because all terms interacted with precision are deviation coded, precision is a balanced-panel average of the estimated employment elasticities. It ranges between 0.008 and -0.090. If, as presumed, the *multiple* LASSO is more reliable than the *naive* LASSO, then the meta-estimates nearest zero are less reliable, suggesting an overall average effect that is most likely in the interval [-0.090, -0.076], none of which is statistically significant. The remainder of this discussion will refer solely to the meta-estimates in which the *multiple* LASSO was used.

What happens when we consider the response of a specific well studied group or industry? The point meta-estimates of the employment elasticity for teenagers lie in a slightly broader range, [-0.099, -0.070], closer to zero than any part of the old consensus range defined in Brown, Gilroy and Kohen (1983). Furthermore, neither is statistically significant.¹⁵ The interval for the elasticity of Eating and Drinking Establishments is similar to that for teenagers, [-0.096, -0.060],

¹⁵It is worth observing here that a variable identifying the estimates from the GS specification that allows the elasticity to vary systematically according to labor market tightness was available for the LASSO to select, but that in no case was it selected.

slightly closer to zero, but again, neither of the point estimates is statistically significant. The meta-estimates for those without a high school degree is contained in those for teenagers and Eating and Drinking Establishments, $[-0.097, -0.086]$, with much larger standard errors than seen so far. For males, the range, $[-0.067, -0.010]$, barely overlaps that for Eating and Drinking Establishments (and does not overlap that for teenagers) and the standard errors are comparable in size to all the prior ones except No HS degree. For females, the range of the meta-estimates is both greater and further to the left than for the other groups: $[-0.189, -0.108]$, and the meta-estimate further from zero is statistically significant at a 0.1 level

In the rightmost two columns of Table 5 are estimates when a correction for publication bias is not included in the set of variables for the LASSO to consider. The point estimates lacking the correction are very similar to those that include the correction.¹⁶ This pattern is consistent with that of the results in the bottom rows of the table, which indicate that the correction is not statistically significant.

One concern is the large standard errors in all the estimates. This begs the question, is it a mistake to combine these estimates, some for teenagers, some for Eating and Drinking Establishments, and some for others? Is this sample too heterogeneous even with the controls? To examine this question, Table 6 presents meta-estimates calculated by applying the *multiple* LASSO separately to subsamples consisting only of estimates for teenagers (and not including the GS Decile estimates), and for Eating and Drinking Establishments.

¹⁶At least for the *multiple* LASSO. For the *naive* LASSO, this is not at all the case, another reason for suspecting that the *multiple* LASSO generates more reliable results, or at least more robust results..

Treating the subsamples separately leads to point meta-estimates that are noticeably smaller than the corresponding figures in Table 5. For teenagers, the point estimates are roughly one-third as large as previously and for Eating and Drinking establishments about one-tenth as large. In addition, the standard errors of the meta-estimates are also considerably smaller.¹⁷

What we are left with is that whether one analyzes all the estimates together, controlling for heterogeneity within the sample or distinguishes them based on the group whose employment response is measured, the meta-estimates are very small. The larger meta-estimates, derived from considering all the estimates together, are not precisely estimated. Some of the smaller ones, with elasticities between -0.008 and -0.037 appear to be, that is to be statistically significant, but the point values are so small as to be economically not significant.

Conclusion

We have identified 60 analyses of U.S. data that have been completed since the exchange between Card and Krueger, and Neumark and Wascher in the December 2000 issue of the *AER*. From the 36 of these that either report elasticities of employment with respect to the minimum

¹⁷In considering these standard errors, it is important to recognize that their accuracy is sensitive to the number of clusters; it is widely recognized that clustering effectively captures sampling variation only when the number of clusters is sufficiently large. Opinions concerning the minimum size for “sufficiently large” vary (much like opinions about the minimum number of degrees of freedom for which it is reasonable to use the normal rather than the *t* distribution, where some are comfortable with *df*=30 and others prefer *df*=120: $t(30, 0.975) = 2.042$ and $t(120, 0.975) = 1.98$). Bertrand, Duflo and Mullainathan (2004) provide simulation evidence that 10 clusters are too few but 20 may be sufficient in the typical state-year or state-quarter panel. Hansen (2007) shows, also through simulations, that the minimum sufficient number of clusters varies with the degree of deviation from being *iid*, and the greater the deviation, the more clusters are necessary. Angrist and Pischke (2009) suggest, not quite tongue in cheek, that with (a minimum of) 42 clusters one can be confident that the standard errors are reasonably accurate. The upshot is that 20 clusters that go into calculating the standard errors for the teenage employment elasticities may well be sufficient, but that the 13 used for the E&D standard errors may be too few and the standard errors biased toward zero.

wage and their standard errors, or provide sufficient information to calculate them, we have gathered 933 estimates. While publication bias has a noticeable effect on the magnitude of the point meta-estimates, inducing a left shift, the estimated standard errors of the meta-estimates are hardly, if at all, affected. Furthermore, the magnitude of the shift of the point-estimates is small once these standard errors are taken into consideration.

After correcting for heteroscedasticity, publication bias and heterogeneity, the meta-estimate of the employment elasticity for is about one-third of what used to be the smallest elasticity in the consensus interval: -0.034 and -0.1, respectively. That for Eating and Drinking Establishments is about half the size of that for teenagers: -0.017. The meta-estimates for males and for those without a high school degree are both positive: 0.021 and 0.066. And the meta-estimate for females is almost zero: -0.003. None of these is anywhere near statistical significance, with cluster standard errors typically between 0.08 and 0.10 (and for those without a high school degree, about 0.15). As a whole, this literature provides no support for the position that minimum wage policy in the United States has had any detectible effect on employment, either negative or positive.

These estimates, which incorporate very recent research, parallel those of prior meta-analysis of the employment effects of the minimum wage. The current research adds to prior work not only in incorporating newer research but, more importantly, in introducing use of the LASSO to improve the specification of the model correcting for heterogeneity. This is a further step toward developing a disciplined, transparent and reproducible methodology summarizing results where there is a large body of research

References

- Angrist, J., and Pischke, J.-S. 2009. *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton, NJ: Princeton University Press.
- Belloni, A., Chernozhukov, V. and Hansen, C. 2014. "High-Dimensional Methods and Inference on Structural and Treatment Effects," *Journal of Economic Perspectives* 28(2): 29-50.
- Belloni, A., Chernozhukov, V., Hansen, C. and Kozbur, D. 2015. "Inference in High Dimensional Panel Models with an Application to Gun Control," WP (Oct. 6)
- Belman, D. and Wolfson, P. 2014. *What Does the Minimum Wage Do?* Kalamazoo, Michigan: W.E. Upjohn Institute for Employment Research.
- Bertrand, M., Duflo, E. and Mullainathan, S. 2004. "How Much Should We Trust Differences-in-Differences Estimates?" *Quarterly Journal of Economics* 119(1): 249–275.
- Brown, C., Gilroy, C. and Kohen, A. 1982. "The Effect of the Minimum Wage on Employment and Unemployment." *Journal of Economic Literature* 20(2): 487–528.
- Card, D., and Krueger, A. 1995. *Myth and Measurement: The New Economics of the Minimum Wage*. Princeton, NJ: Princeton University Press.
- Card, D., and Krueger, A. 2000. "Minimum Wages and Employment: A Case Study of the Fast-Food Industry in New Jersey and Pennsylvania: Reply." *American Economic Review* 90(5): 1397–1420.
- Doucouliafos, H., and Stanley, T. 2009. "Publication Selection Bias in Minimum-Wage Research? A Meta-Regression Analysis." *British Journal of Industrial Relations* 47(2): 406–428.
- Hansen, C. 2007. "Asymptotic Properties of a Robust Variance Matrix Estimator for Panel Data When T Is Large." *Journal of Econometrics* 141: 670–694.

Neumark, D., and Wascher, W. 2000. “Minimum Wages and Employment: A Case Study of the Fast Food Industry in New Jersey and Pennsylvania: Comment.” *American Economic Review* 90(5): 1362–1396.

Stanley, T. and Doucouliagos, H. 2012. *Meta-Regression Analysis in Economics and Business*. New York: Routledge.

Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal. Statistical Society B* 58(1): 267-288.

FIGURE 1

Employment and Hours Elasticities vs. Precision (SE^{-1})




FIGURE 2

Teen Employment and Hours Elasticities vs. Precision (SE^{-1})




FIGURE 3





TABLE 1

| <u>TYPE OF HETEROGENEITY</u> | | N |
|-------------------------------------|--|----------|
| 1. | Source of the estimate: analysis and authors | |
| 2. | The estimate refers to the employment | |
| | of members of a demographic group (e.g., teenagers, females) | 406 |
| | in an industry (e.g., eating and drinking establishments) | 354 |
| | of members of a demographic in an industry (e.g., teenagers in retail) | 138 |
| | (other) | 35 |
| 3. | The source of the data for this estimate | |
| | CPS | 387 |
| | QCEW | 187 |
| | QWI | 153 |
| | other | 206 |
| 4. | Whether the elasticity was for | |
| | hours of employment | 163 |
| | employment (# jobs, # people employed, employment-population ratio, binary indicator that an individual was employed) | 770 |
| 5. | The frequency of the data used for this estimate | |
| | Bi-Annual | 20 |
| | Annual | 250 |
| | Quarterly | 506 |
| | Monthly | 143 |
| | Other | 16 |
| 6. | The geographic reach of the study that provided this estimate | |
| | National | 827 |
| | Multi-state (less than national, but more than 1 state in the treatment group) | 13 |
| | State (a single state in the treatment group) | 57 |
| | City (a single city in the treatment group) | 35 |
| | Other | 1 |
| 7. | The analysis containing this estimate has been published | 487 |
| 8. | Explicit thought given to defining a control group in calculating this estimate | 355 |
| 9. | Either the data structure for this estimate is not one that would raise suspicions about the standard errors, or they have been resolved (e.g., by clustering) | 867 |

| | | |
|-----|---|-----|
| 10. | This estimate refers specifically to | |
| | the employment of teenagers | 476 |
| | employment in eating and drinking establishments | 205 |
| | the employment of those whose schooling ended before receiving a HS degree | 29 |
| | the employment of females | 36 |
| | the employment of males | 27 |
| 11. | This estimate is reported to replicate prior work of others | 77 |
| 12. | This estimate is from Gittings and Schmutte's labor-market deciles analysis | 116 |

TABLE 2
DESCRIPTIVE STATISTICS OF THE SAMPLE AND OF SUBSAMPLES

| sample | nobs | mean | sd | min | q1 | median | q3 | max | $\rho_{\text{elasticity, SE}}$ | se_{ρ} | p_{ρ} |
|-------------------------|------|--------|-------|--------|--------|--------|--------|-------|--------------------------------|-------------|------------|
| 1) Whole Sample | 939 | -0.050 | 0.292 | -2.235 | -0.129 | -0.032 | 0.049 | 1.510 | -0.213 | 0.032 | 0.000 |
| 2) Teens | 482 | -0.060 | 0.256 | -2.235 | -0.153 | -0.037 | 0.047 | 0.779 | -0.498 | 0.040 | 0.000 |
| 3) Gittings & Schmutte | 116 | 0.014 | 0.373 | -1.369 | -0.164 | 0.119 | 0.249 | 0.779 | -0.741 | 0.063 | 0.000 |
| 4) ~Gittings & Schmutte | 366 | -0.084 | 0.201 | -2.235 | -0.151 | -0.055 | 0.003 | 0.421 | -0.561 | 0.043 | 0.000 |
| 5) E&D | 205 | -0.019 | 0.242 | -0.750 | -0.079 | -0.016 | 0.022 | 1.510 | 0.176 | 0.069 | 0.118 |
| 6) Males | 27 | 0.006 | 0.114 | -0.214 | -0.060 | 0.008 | 0.079 | 0.202 | 0.188 | 0.196 | 0.349 |
| 7) Females | 36 | 0.072 | 0.314 | -0.313 | -0.095 | -0.005 | 0.097 | 1.180 | 0.672 | 0.127 | 0.000 |
| 8) NoHS | 29 | -0.211 | 0.588 | -1.010 | -0.701 | -0.282 | 0.106 | 1.180 | -0.607 | 0.153 | 0.000 |
| 9) Replications | 81 | -0.104 | 0.091 | -0.435 | -0.163 | -0.100 | -0.053 | 0.147 | -0.239 | 0.109 | 0.032 |

TABLE 3
Meta-Estimates of the Employment Elasticity with Respect to the Minimum Wage
Uncorrected and Corrected for Heteroscedasticity

| Full Sample | | |
|------------------------------|-----------|-------------------|
| N | 939 | |
| | Mean (se) | Median (IQ Range) |
| Uncorrected | -0.050 | -0.032 |
| standard error | (0.024) | [-0.129, 0.049] |
| Heteroscedasticity Corrected | -0.024 | -0.012 |
| standard error | (0.008) | [-0.034, -0.004] |

Standard Errors are calculated by clustering by title

TABLE 4
Meta-Estimates of the Employment Elasticity with Respect to the Minimum Wage
Corrected for Heteroscedasticity and Publication Bias (Not Heterogeneity)

| | Linear Correction for Publication Bias | Quadratic Correction for Publication Bias |
|---------------------|---|--|
| Elasticity | -0.019 | -0.023 |
| standard error | (0.009) | (0.008) |
| pub-bias correction | -0.227 | -0.671 |
| standard error | (0.273) | (0.606) |
| Cochran's Q | 3569 | |
| p | 0.000 | |

Standard Errors are calculated by clustering by title

TABLE 5
Meta-Estimates of the Employment Elasticity with Respect to the Minimum Wage
Corrected for Heteroscedasticity and Heterogeneity (and Sometimes for Publication Bias)

| | Eq. 3.5: Linear Correction for Publication Bias | | Eq. 3.6: Quadratic Correction for Publication Bias | | Eq. 3.x: No Correction for Publication Bias | |
|--------------------|--|-----------------------|---|-----------------------|--|-----------------------|
| | <i>Naive</i> Lasso | <i>Multiple</i> Lasso | <i>Naive</i> Lasso | <i>Multiple</i> Lasso | <i>Naive</i> Lasso | <i>Multiple</i> Lasso |
| precision | 0.008 | -0.076 | -0.034 | -0.090 | -0.070 | -0.080 |
| (clustered) se | 0.083 | 0.099 | 0.077 | 0.077 | 0.070 | 0.076 |
| precision + teen | -0.034 | -0.070 | -0.076 | -0.099 | -0.113 | -0.085 |
| (clustered) se | 0.097 | 0.100 | 0.091 | 0.086 | 0.084 | 0.084 |
| precision + E&D | -0.017 | -0.060 | -0.050 | -0.096 | -0.085 | -0.076 |
| (clustered) se | 0.096 | 0.099 | 0.093 | 0.086 | 0.086 | 0.085 |
| precision + NoHS | 0.067 | -0.097 | 0.000 | -0.086 | -0.053 | -0.089 |
| (clustered) se | 0.151 | 0.187 | 0.145 | 0.143 | 0.135 | 0.144 |
| precision + Male | 0.022 | -0.010 | -0.020 | -0.067 | -0.055 | -0.050 |
| (clustered) se | 0.084 | 0.108 | 0.078 | 0.080 | 0.070 | 0.079 |
| precision + Female | -0.002 | -0.189 | -0.047 | -0.108 | -0.084 | -0.100 |
| (clustered) se | 0.084 | 0.098 | 0.078 | 0.077 | 0.070 | 0.077 |
| Mean Pub Bias | -0.549 | 0.746 | -1.064 | -0.604 | - | - |
| (clustered) se | 0.264 | 0.584 | 0.778 | 0.659 | - | - |

VARIABLES INCLUDED IN EACH SPECIFICATION by LASSO

All *Naive* Lasso: (none but those locked in)

Multiple Lasso (stars on variables have conventional meanings about statistical significance at 0.1, 0.05 and 0.01 levels)

Eq. 3.5: Dummies for E&D, Male & Female***, and *precision* interacted with dummies for analyses 8, 19, 20***, 26 & 37***

Eq. 3.6: *precision* interacted with indicators for CPS, QCEW, reliable SEs***, Employment (not Hours), and analyses 7, 11*, 12, 15, 29**, 31***, 32, 33 & 36**

Eq. 3.x: *precision* interacted with indicators for CPS, QCEW, Nation-wide Data, reliable SEs, Employment (not Hours), and analyses 2, 7, 12, 15, 20**, 22***, 29**, 31, 32**, 33

TABLE 6
Meta-Estimates of the Employment Elasticity with Respect to the Minimum Wage for 2 Subsamples
Corrected for Heteroscedasticity and Heterogeneity (and Sometimes for Publication Bias)
(All Estimates based on *Multiple* LASSO variable selection)

| | | Eq. 3.5: Linear Correction for Publication Bias | Eq. 3.6: Quadratic Correction for Publication Bias | Eq. 3.x: No Correction for Publication Bias |
|----------------|----------------|--|---|--|
| Teen Subsample | precision | -0.022 | -0.037 | -0.040 |
| | (clustered) se | 0.017 | 0.017 | 0.017 |
| | Mean Pub Bias | -0.618 | -1.832 | - |
| | (clustered) se | 0.401 | 1.222 | - |
| E&D Subsample | precision | -0.008 | -0.013 | -0.014 |
| | (clustered) se | 0.004 | 0.003 | 0.004 |
| | Mean Pub Bias | -0.389 | 0.002 | - |
| | (clustered) se | 0.289 | 0.312 | - |

VARIABLES INCLUDED IN EACH SPECIFICATION by LASSO

Teen Subsample:

- Eq. 3.5: indicator for estimates in which explicit thought given to a control group*** (positively signed, 0.5)
- Eq. 3.6: none
- Eq. 3.x: none

E&D Subsample

- Eq 3.5: none
- Eq 3.6: *se* interacted with the indicator for estimates that are replications of prior work***, (negatively signed, 4 orders of magnitude larger than correction)
- Eq3.x: none

***Statistically significant at a 0.001 level

APPENDIX 1 STUDIES USED

| Title Dummy | # Obs | Analysis |
|----------------|-------|---|
| 1 | 4 | Bazen, S.; and Marimoutou, V. 2002. "Looking for a needle in a haystack? A re-examination of the time series relationship between teenage employment and minimum wages in the United States." <i>Oxford Bulletin of Economics and Statistics</i> . 64. 699-725. |
| 2 | 6 | Dodson, M. E. 2002. "The impact of the minimum wage in West Virginia: A test of the low-wage-area theory." <i>Journal of Labor Research</i> . 23(1). 25-40. |
| 3 | 24 | Orazem, P. F.; and Mattila, J. P. 2002. "Minimum wage effects on hours, employment, and number of firms: The Iowa case." <i>Journal of Labor Research</i> . 23(1). 3-23. |
| 4 | 1 | Neumark, D.; Schweitzer, M.; and Wascher, W. 2004. "Minimum wage effects throughout the wage distribution." <i>Journal of Human Resources</i> . 39(2). 425-450. |
| 5 | 5 | Potter, N. 2006. "Measuring the Employment Impacts of the Living Wage Ordinance in Santa Fe, New Mexico." Bureau of Bus. & Econ. Rsrch/ UNM. |

- 6 30 Dube, A.; Naidu, S.; and Reich, M. 2007. "The economic effects of a citywide minimum wage." *Industrial & Labor Relations Review*. 60(4). 522-543.
- 13 6 Neumark, D.; and Nizalova, E. 2007. "Minimum wage effects in the longer run." *Journal of Human Resources*. 42(2). 435-452.
- 7 8 Singell, L. D.; and Terborg, J. R. 2007. "Employment effects of two northwest minimum wage initiatives." *Economic Inquiry*. 45(1). 40-55.
- 8 36 Orrenius, P.; and Zavodny, M. 2008. "The Effect of Minimum Wages on Immigrants' Employment and Earnings." *Industrial & Labor Relations Review*. 61(4). 544-563.
- 9 9 Sabia, J.J. 2008. "Minimum Wages And The Economic Well-Being Of Single Mothers." *Journal of Policy Analysis and Management*. 27(4). 848-866.
- 10 23 Addison, J.; Blackburn, M.; and Cotti, C. 2009. "Do MWs raise employment? Evidence from the U.S. retail-trade sector." *Labour Economics*. 16(4). 397-408.
- 11 14 Allegretto, S.; Dube, A; and Reich, M. 2009. "Spatial Heterogeneity and Minimum Wages: Employment Estimates for Teens Using Cross-State Commuting Zones." IRLE Working Paper No. 181-09. Berkeley, CA: University of California, Berkeley, Institute for Research on Labor and Employment.

- 12 13 Keil, M.; Robertson, D.; and Symons, J. 2009. "Univariate Regressions of Employment on Minimum Wages
in the Panel of U.S. States." Robert Day School Working Paper No. 2009-03.
- 14 6 Sabia, J.J. 2009. "Identifying Minimum Wage Effects: New Evidence from Monthly CPS Data." *Industrial
Relations*. 48(2). 311-328.
- 15 16 Sabia, J.J. 2009. "The Effects of Minimum Wage Increases on Retail Employment and Hours: New Evidence
from Monthly CPS Data." *Journal of Labor Research*. 30(1). 75-97.
- 16 68 Belman, D.; and Wolfson, P. 2010. "The Effect of Legislated Minimum Wage Increases on Employment and
Hours: A Dynamic Analysis." *LABOUR*. 24(1). 1-25.
- 17 27 Dube, A., Lester, T. W.; and Reich, M. 2010. "Minimum Wage Effects Across State Borders: Estimates Using
Contiguous Counties." *The Review of Economics and Statistics*. 92(4). 945-964.
- 18 20 Persky, J.; and Baiman, R. 2010. "Do State Minimum Wage Laws Reduce Employment? Mixed Messages
from Fast Food Outlets in Illinois and Indiana." *Journal of Regional Analysis and Policy*. 40(2). 132-142.
- 19 104 Allegretto, S.; Dube, A; and Reich, M. 2011. "Do Minimum Wages Really Reduce Teenage Employment -
Accounting for Heterogeneity and Selectivity in State Panel Data." *Industrial Relations*. 50(2). 205-240.

- 20 12 Neumark, D.; and Wascher, W. 2011. “Does a Higher Minimum Wage Enhance the Effectiveness of the
 Earned Income Tax Credit.” *Industrial & Labor Relations Review*. 64(4). 712-746.
- 21 1 Addison, J.T.; and Ozturk, O.D. 2012. “Minimum Wages, Labor Market Institutions, and Female
 Employment: A Cross-Country Analysis.” *Industrial & Labor Relations Review*. 65(4). 779-809.
- 22 6 Addison, J. T. , Blackburn, M.L.; and Cotti, C.D. . 2012. “The Effect of MinimumWages on Labour Market
 Outcomes: County-Level Estimates from the Restaurant-and-Bar Sector.” *British Journal of Industrial
 Relations*. 50(3). 412-435.
- 24 15 Sabia, J.J.; Burkhauser, R.V.; and Hansen, B. 2012. “Are the Effects of Minimum Wage Increases Always
 Small? New Evidence from a Case Study Of New York State.” *Industrial & Labor Relations Review*. 65(2).
 350-376.
- 25 24 Addison, J.; Blackburn, M.; and Cotti, C. 2013. “Minimum Wage Increases in a Recessionary Environment. .
 Grupo de Estudos Monetários e Financeiros WP. 8.
- 29 28 Clemens, J.; and Wither, M. 2014. “The Minimum Wage and the Great Recession: Evidence of Effects on the
 Employment and Income Trajectories Of Low-skilled Workers. NBER WP 20724 (Dec).

- 23 8 Dube, A., Lester, T. W.; and Reich, M. 2014. “MW Shocks, Employment Flows & Labor Market Frictions.
IRLE Working Paper No. 149-13. Berkeley, CA: University of California, Berkeley, Institute for Research on
Labor and Employment.
- 30 30 Even, W.E; and Macpherson, D.A. 2014. “The Effect of Tip Credits on Earnings and Employment in the U.S.
Restaurant Industry.” *Southern Economic Journal*. 80(3). 633-655.
- 31 21 Hoffman, S.D. 2014. “Employment Effects of the 2009 Minimum Wage Increase: New Evidence from State-
Based Comparisons of Workers by Skill Level.” *B.E. Journal of Economic Analysis & Policy*. 14(3). 695-721.
- 34 8 Neumark, D.; Salas, J.M.I; and Wascher, W. 2014. “More on Recent Evidence on the Effects of Minimum
Wages in the United States.” NBER WP 20619 (Oct).
- 32 27 Neumark, D.; Salas, J.M.I; and Wascher, W. 2014. “Revisiting the Minimum Wage-Employment Debate:
Throwing Out the Baby with the Bathwater.” *Industrial & Labor Relations Review*. 67(Supplement). 608-648.
- 33 78 Totty, E. 2014. “Effect of Minimum Wages on Employment - a Factor model approach JULY.” KSM WP
1278 (July).
- 26 107 Allegretto, S; Dube, A.; Reich, M.; and Zipperer, B.. 2015. “Credible Research Designs for Minimum Wage
Studies.” IREL WP 116-29 (9/29).

- 27 1 Dube, A.; and Zipperer, B. 2015. “Pooled Synthetic Control Estimates for Recurring Treatments. IZA-DP
8944 (March).
- 35 125 Gittings, R.K.; and Schmutte, I.M. 2015. “Getting Handcuffs on an Octopus: Minimum Wages, Employment,
and Turnover.” SSRN WP (6/2/15).
- 36 12 Liu, S.; Hyclak, T.J.; and Regmi, K. 2015. “Impact of the Minimum Wage on Youth Labor Markets.”
LABOUR. Early View (10/28/15).
- 28 23 Meer, J.; and West, J. 2015. “Effects of the Minimum Wage on Employment Dynamics.” Upjohn Institute
WP 15-233.
- 37 6 Powell, D. 2015. “Synthetic Control Estimation Beyond Case Studies - Does the Minimum Wage Reduce
Employment.” BEPress WP (January).

APPENDIX 2

Figure A1, complementing Figure 2, the teen elasticities other than the Decile-based estimates of GS in dark gray triangles on a background of small, faint light gray circles for the rest of the sample. Comparing the two figures, it is obvious that the distribution of this part of the sample is more negative than the sample including the GS estimates. The mean (-0.084) and the upper four parts of the five-number summary are all noticeably shifted to left compared to the full teen sample, and the standard deviation is about 80% as large. Taken together, the most interesting point of these two graphs is that the GS sample looks very different from the rest of the teen elasticities, reinforcing the notion that it should be treated separately from the rest of the teen sample.

Figure A2 displays the estimates for Eating and Drinking Establishments. They are tightly distributed about their mean (-0.019), which is very close to the median value (-0.016). Half of the estimates lie between -0.079 and 0.022, and more than 90% are within one standard deviation of the mean (-0.261, 0.223). Ten lie to the left of this range, and four far to the right, with positive elasticities greater than 1.¹⁸

Figure A3 displays the handful of observations for each of three other demographic groups of interest: males, females and those without a high-school degree. The tick marks on both axes are calculated for the whole sample and are the same as in Figure 1. The estimates for males are tightly distributed within one standard deviation of the overall mean, and that is true of most of the estimates for females. A handful of those for females, however, are more positive than the upper end of this region. More than a handful of the estimates for those who did not complete

¹⁸These four all come from Persky and Baiman (2010), a two-state difference-in-differences analysis similar to the well known Card and Krueger (1994, 2000) studies. More than half of the 20 estimates from this study are negative.

high school are outside this band, most to the left but a few to the right. What is most striking about these estimates, however, is that they are generally less precisely estimated than those for either males or females.

The last of the funnel plots, Figure A4, highlights the estimates that are attempts to replicate others' prior work. They tend to be centrally located, both as point estimates of the elasticity, and in how precisely they are estimated. To the extent that including them affects the meta-analysis, it will likely be to reduce the standard error of the meta-estimates.

FIGURE A1

Employment and Hours Elasticities vs. Precision (SE^{-1}) Teens (Not Decile Estimates)




FIGURE A2

E&D Employment & Hours Elasticities vs. Precision (SE^{-1})




FIGURE A3




FIGURE A4

