American Economic Journal: Applied Economics 2019, 11(4): 268–304
https://doi.org/10.1257/app.20170085

Minimum Wages and the Distribution of Family Incomes[†]

By Arindrajit Dube*

There is robust evidence that higher minimum wages increase family incomes at the bottom of the distribution. The long-run (3 or more years) minimum wage elasticity of the non-elderly poverty rate with respect to the minimum wage ranges between −0.220 and −0.459 across alternative specifications. The long-run minimum wage elasticities for the tenth and fifteenth unconditional quantiles of family income range between 0.152 and 0.430 depending on specification. A reduction in public assistance partly offsets these income gains, which are on average 66 percent as large when using an expanded income definition including tax credits and noncash transfers. (*JEL* D31, I32, I38, J31, J38)

At least since Gramlich (1976), economists have recognized that the ability of minimum wage policy to aid lower income families depends on the joint distribution of wage gains, potential job losses, and other sources of family income. The poverty-reducing effects of the minimum wage are expected to be small if job losses from a minimum wage increase are sizable, or if most minimum wage workers are higher up in the family income distribution—say, because they are teens from higher income families or because their spouses are paid well. Therefore, the extent to which minimum wages raise family incomes at the bottom of the distribution is not clear a priori, and has to be studied empirically.

However, while there is a large and active literature on the employment effects of minimum wages—see Belman and Wolfson (2014) for a recent review—there are relatively fewer studies that empirically estimate the impact of the policy on family incomes. Compounding the problem, the existing papers often suffer from a number of key shortcomings including insufficient attention to the validity of the research design, and use of small samples sometimes with limited minimum wage variation—all of which tend to produce somewhat erratic and imprecise estimates.

In this paper, I use individual-level data from the March Current Population Survey (CPS) between 1984 and 2013 to provide a more complete assessment of how US

* Crotty Hall, University of Massachusetts, Amherst, MA 01002 and NBER (email: adube@econs.umass.edu). Alexandre Mas was coeditor for this article. I thank Charles Brown, David Card, Richard Dickens, Dayanand Manoli, Suresh Naidu, Michael Reich, Timothy Smeeding, Chris Taber, and participants at the 2015 ASSA/LERA session, 2015 UCL/LSE-CEP Low Pay Workshop, 2014 Building Human Capital Conference at University of Wisconsin-Madison and preconference at the University of Texas-Austin, UMass Boston, and the US Department of Labor conference on labor standards. I thank Ihsaan Bassier, Doruk Cengiz, Thomas Peake, Simon Sturn, Owen Thompson, and Ben Zipperer for excellent research assistance. I received funding from the University of Wisconsin-Madison Institute for Research on Poverty (grant number 454K252), US Department of Labor, and Washington Center for Equitable Growth (grant number 116-0417) for the completion of this research.

[†] Go to https://doi.org/10.1257/app.20170085 to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

minimum wage policies shift the distribution of family incomes for the non-elderly population than has been to date.[1] This paper makes three important advances over the existing literature. First, I fully characterize how minimum wage increases shift the cumulative distribution of family incomes, and subsequently use this to estimate unconditional quantile partial effects (UQPE) of the policy. This provides a fuller picture of the distributional effects of the policy than available to date. Second, I show how the income effects by quantile vary when we consider broader income definitions to include tax credits and noncash transfers. I am able to quantify the extent to which offsets through reduced public assistance affects the distributional impact of minimum wages. Third, I provide estimates that are more credibly causal than in the literature by utilizing a wide range of specifications with alternative sets of controls for time-varying heterogeneity, and by assessing the assumption of parallel trends using an array of falsification tests across the income distribution. Specifically, I use distributed lag specifications that both include leads and up to three years of lags in the minimum wage, allowing for a delayed impact of the policy; and I allow for time varying heterogeneity by region, business cycles, and state-specific trends.

Overall, I find robust evidence that higher minimum wages lead to increases in incomes at the bottom of the family income distribution. For the poverty rate, long-run minimum wage elasticities between -0.220 and -0.459 obtain from 8 specifications ranging from the classic two-way fixed effects model to a model with a rich set of controls for trends, regional shocks, and business cycle heterogeneity. All of these estimates are statistically distinguishable from zero at conventional levels.[2] There is a reduction in the shares below income cutoffs between 50 and 125 percent of the federal poverty threshold (FPT), with the largest proportionate reductions occurring around 75 percent of the poverty threshold (elasticities ranging between -0.186 and -0.566). When I look across demographic groups using the preferred (most saturated) specification, the poverty rate elasticities are similar when restricted to younger (under 30) and lower credentialed individuals without a high school degree. However, the elasticity is relatively larger for black and Latino individuals (-0.872), while relatively smaller for single mothers (-0.327). The poverty reducing effects are not substantially different after 1996, in the period following welfare reform and EITC expansion.

By estimating how a minimum wage increase affects the shares of population with incomes below various multiples of the federal poverty threshold we identify the effect of the policy on the cumulative distribution function (CDF) of equivalized family incomes. Subsequently, we can invert the impact of the policy on the CDF to estimate the impact on an income quantile, i.e., the UQPE. I use the recentered influence function (RIF) regression approach of Firpo, Fortin, and Lemieux (2009), which performs this inversion using a local linear approximation to the

[1] In this paper, when I refer to the 1984–2013 period, I am referring to the survey years for the March CPS. Note, however, that respondents in the March 2013 CPS survey are asked about their income during the year 2012.

[2] Most results in this paper are for the non-elderly population; so, when I refer to "the poverty rate," I am referring to the poverty rate among those under 65 years of age. Also, as a matter of terminology, in this paper, virtually all elasticities are elasticities with respect to the minimum wage. For brevity, I will sometimes refer to "the elasticity of the poverty rate with respect to the minimum wage" as either "the minimum wage elasticity for the poverty rate" or simply "the poverty rate elasticity." The same is true for elasticities of other outcomes with respect to the minimum wage, such as family income quantiles, the proportion under one-half poverty line, etc.

counterfactual CDF. This local inversion entails rescaling the marginal effect of minimum wages on the share above an income cutoff by the probability density of income at that cutoff.[3]

I find clear, positive effects of minimum wages on family incomes below the twentieth quantile. The largest impact occurs at the tenth and fifteenth quantiles, where estimates from most specifications are statistically significant, and the long-run minimum wage elasticities for these family income quantiles range between 0.152 and 0.430 depending on control sets. In the preferred specification, the family income elasticities with respect to the minimum wage are around 0.359 and 0.332 for the tenth and fifteenth quantiles, respectively, and diminish close to 0 by the thirtieth quantile. Since the conventional income definition used for official poverty calculations does not include tax credits such as the Earned Income Tax Credit (EITC), or noncash transfers such as the Supplemental Nutritional Assistance Program (SNAP), I also estimate the impact using an expanded income definition. After accounting for tax credits and noncash transfers, the minimum wage effect on the level of family incomes (i.e., the semi-elasticities) is about 66 percent as large for the bottom 30 percent of the distribution. Overall, the evidence clearly points to at least moderate income gains for low-income families resulting from minimum wage increases. At the same time, there is evidence for some substitution of government transfers with earnings as evidenced by the somewhat smaller income increase accounting for tax credit and noncash transfers such as SNAP and EITC. The estimates are similar from the more recent period (e.g., post 1996) when the policy environment is more homogeneous following welfare reform and EITC expansion.

While this paper substantially improves upon existing research on the topic of minimum wages, family income distribution and poverty, the existing research is consistent with the proposition that minimum wages likely reduce poverty. In online Appendix A, I quantitatively assess estimates from the 13 key papers in the literature by collecting or constructing nearly every minimum wage elasticity for the poverty rate from those studies for a variety of demographic groups. Seventy-two out of the 78 elasticities have a negative sign, and a simple average of the 78 elasticities produces a poverty rate elasticity of -0.18, while the average of the median elasticities from each study is -0.17. For 8 of the 13 studies that actually report an estimate for overall poverty (as opposed to narrower subgroups), the average of poverty rate elasticities is -0.13. These averages are broadly consistent with the range of findings in this paper in pointing toward a poverty-reduction effect of minimum wages, though they tend to be smaller in magnitude than what I find.

At the same time, the existing evidence is clouded by serious shortcomings in these studies: insufficient controls for state-level heterogeneity; short time periods, sometimes with little minimum wage variation; over-statement of precision due to improper methods of statistical inference; and the use of idiosyncratic sets of

[3] It is useful to contrast the UQPE with estimates from the more familiar (conditional) quantile regression. The quantile regression provides us with an estimate of the impact of minimum wages on, say, the tenth conditional quantile of family incomes. This tells us how the policy affects those with unusually low income within their demographic group, e.g., a college graduate with an income that is low *relative* to others in her educational category. However, here we are interested in the effect of the policy on those with low incomes in an *absolute* (or unconditional) sense, while controlling for covariates such as education. This is exactly what UQPE measures.

outcomes and target groups. In contrast, I use 30 years of data from a period with a large amount of cross-state minimum wage variation. I also pay close attention to the nonrandom nature of minimum wage policies (Allegretto et al. 2017). Since there is disagreement in the literature on the appropriateness of particular specifications, I show results using eight different specifications with alternative controls for state-level heterogeneity that subsumes much of the approaches used in the literature. Starting with the classic two-way fixed effects model, I progressively add regional controls, state specific trends, and state-specific business cycle effects—all of which have been shown to be important in the extant minimum wage literature (e.g., Allegretto et al. 2017, Zipperer 2016). The resulting set of estimates produces a credible range regardless of one's priors on, say, the desirability of using regional controls. At the same time, I assess the internal validity of various specifications using a host of falsification tests including estimating effects higher up in the income distribution, as well as analyzing leading effects (preexisting trends) across specifications. I show that the inclusion of controls for such state-level heterogeneity tends both to improve performance on falsification tests and to increase the magnitude of the estimated elasticity of the poverty rate with respect to minimum wages.

This paper adds to a growing literature that estimates distributional effects of policies using a quasi-experimental design. For example, Havnes and Mogstad (2015) also uses RIF regressions in a difference-in-difference setting to study the distributional impact of universal child care and find that a small average effect masks the more sizable increases in adult earnings at the bottom quantiles. The distributional analysis in this paper is also similar to Hoynes and Patel (2016), who show how EITC expansion shifted the lower end of the family income distribution. Finally, Autor, Manning, and Smith (2016) estimates the effect of minimum wages on the hourly wage distribution, refining the earlier analysis by Lee (1999).

The rest of the paper is structured as follows. In Section I, I describe the data and research design, including the RIF estimation of unconditional quantile partial effects. Section II presents my empirical findings on the effect of minimum wages on the proportions below various low-income cutoffs as well as on income quantiles. Section III concludes with a discussion of the policy implications.

I. Data and Research Design

A *Data and Sample Construction*

I use individual-level data from the UNICON extract of the March Current Population Survey (CPS) between 1984 and 2013. I augment the CPS data with information on state EITC supplements,[4] state per capita GDP, and state unemployment rates from the University of Kentucky Center for Poverty Research, and state and federal minimum wages from the US Department of Labor. I take the average of the effective minimum wage (maximum of the state or federal minimums) during the year for which respondents report incomes. For example, I match the effective

[4] Many states specify a percentage of the federal EITC as a supplement to be paid to state taxpayers. I use this state EITC supplement rate in my analysis as a control variable.

monthly minimum wage averaged over January through December of 2011 in a given state to respondents from that state in the 2012 March CPS.

There is extensive variation in minimum wages over the 30-year period studied in this paper. Online Appendix Figure B.1 plots the nominal federal minimum wage, as well as the fiftieth, seventy-fifth, and ninetieth percentiles of the effective nominal minimum wages (weighted by population). As the figure shows, the effective minimum wage varied substantially across different states over this period. It is also the case that there has been much more variation in minimum wages since 2000. Therefore, the inclusion of more recent data is particularly helpful as it allows us to estimate the effects of the policy more precisely.

The primary goal of this paper is to characterize how minimum wage changes affect the entire distribution of family incomes. For this reason, most of the analysis is performed for the non-elderly population as a whole.[5] The exclusion of the elderly is motivated by the fact that they have much lower rates of poverty than the rest of the population, in large part due to Social Security. For example, CPS data from March 2013 shows that 9.1 percent (2.7 percent) of the elderly had incomes under the poverty line (one-half the poverty line), whereas the corresponding proportions for the non-elderly population were 15.9 and 7.2 percent, respectively. For this reason, we are unlikely to learn very much about the impact of minimum wages on the bottom quantiles of the family income distribution from studying the elderly. Finally, a focus on the non-elderly is also common in the literature (e.g., Burkhauser and Sabia 2007, Sabia and Nielsen 2015). However, I also show that the overall estimates are quite similar if we include the elderly population.

As I discuss in online Appendix A, a number of researchers have studied the impact of minimum wages on children and single mothers (e.g., Morgan and Kickham 2001, DeFina 2008, Gundersen and Ziliak 2004). Several studies have also considered younger adults, or adults with limited education; these include Neumark (2016), Addison and Blackburn (1999), and Sabia and Nielsen (2015). Besides estimating the effect of minimum wages on the incomes of the non-elderly population overall, I also show key results by demographic groups similar to those that have been studied in the literature. These include: (1) all individuals without a high school degree, (2) individuals younger than 30 without a high school degree, (3) individuals with high school or less schooling, (4) black or Latino individuals, (5) children under 18 years of age, (6) all adults, and (7) single (unmarried) mothers with children.

B. *Outcomes and Research Design*

This paper focuses on equivalized[6] real family income, defined as multiples of the federal poverty threshold: $y_{it} = Y_{it}/FPT(N_i, Children_i, t)$. This is also sometimes referred to as the income-to-needs (ITN) ratio. As is standard, y_{it} is the ratio between family income, Y_{it}, and the federal poverty threshold $FPT(N_i, Children_i, t)$—which

[5] Official poverty measures do not include unrelated individuals under 15 years of age; for this reason, I exclude them from the sample as well.

[6] Equivalized family income is adjusted for family size and composition.

depends on family size (N_i) and the number of children, and varies by year (t). I start with the conventional definition of family income as is used for official poverty measurement: pretax cash income, which includes earnings and cash transfers, but does not include noncash benefits or tax credits.[7] I use the term family to refer to primary families, subfamilies, as well as single (non-family) householders. Family income and poverty thresholds are computed for each family type separately.[8]

The official poverty measure excludes noncash transfers such as SNAP and housing assistance, as well as tax credits such as EITC. As Fox et al. (2014) shows, exclusion of these categories of income tends to understate the economic gains made by families in the lower end of the income distribution during the 1990s and 2000s. Moreover, there is evidence that minimum wages may reduce some types of transfers such as SNAP (Reich and West 2015). Therefore, I also show income quantile elasticities using an expanded income definition that includes monetized value of noncash transfers (SNAP, housing assistance, school lunch) and refundable tax credits (EITC, child tax credit, and additional child tax credit). The values for the noncash transfers are self-reported, while tax credits are calculated by the Census Bureau based on the respondents' reported income.

Poverty Rate and Shares under Multiples of Federal Poverty Threshold.—To estimate the impact of minimum wages on the share with income under c times the federal poverty threshold, I use a linear probability model where the dependent variable is an indicator for whether individual i is in a family whose equivalized income y_{it} falls below c: $I_{cit} = \mathbf{1}\left(y_{it} < c\right)$. As an example, the share with $I_{1it} = 1$ corresponds to the official poverty rate.

A distributed lag version of the classic two-way (state and time) fixed effects regression specification is as follows:

$$(1) \qquad I_{cit} = \sum_{k=-1}^{3} \alpha_{ck} \ln\left(MW_{s(i)t-k}\right) + X_{it}\Lambda_c + W_{s(i)t}\Psi_c + \mu_{cs(i)} + \theta_{ct} + \epsilon_{cit}.$$

Here, $\mu_{cs(i)}$ is the state fixed effect, θ_{ct} is the time fixed effect, and ϵ_{cit} is the regression error term. The regression coefficients and the error components are all indexed by c to clarify that they are from separate regressions for each cutoff c as a multiple of the federal poverty threshold. The coefficient $\alpha_c^{LR} = \alpha_{c0} + \alpha_{c1} + \alpha_{c2} + \alpha_{c3}$ is the long-run semi-elasticity of the proportion under the income-to-needs cutoff, c, with respect to the minimum wage, $MW_{s(i)t}$, indexed by the state of residence $s\left(i\right)$ of individual i and time t. The coefficient $\alpha_c^{MR} = \alpha_{c0} + \alpha_{c1} + \alpha_{c2}$ is the "medium-run" semi-elasticity inclusive of the effect up to two years after the policy change. These are converted to elasticities by dividing by the sample mean of the population share below the cutoff, $F(c)$. For example, $\beta_c^{LR} = \left(1/F(c)\right) \times \left(\alpha_{c0} + \alpha_{c1} + \alpha_{c2} + \alpha_{c3}\right)$

[7] Eligible income includes earnings (excluding capital loss or gains), unemployment compensation, workers' compensation, social security, Supplemental Security Income, cash-based public assistance such as Temporary Assistance to Needy Families (TANF), veterans' payments, survivor benefits, pension or retirement income, interest, dividends, rents, royalties, income from estates, trusts, educational assistance, alimony, child support, assistance from outside the household, and other miscellaneous sources of cash income.

[8] I use the FAMINC and POVCUT variables from the UNICON extract.

is the long-run elasticity. The explicit inclusion of the lagged treatment variable may be of particular relevance when the specification includes a state-specific linear trend. With state trends, but without lagged treatment included as a regressor, a delayed impact can lead to a misestimation of the state trends, attenuating the measured effect of the treatment (e.g., Wolfers 2006, Meer and West 2016). Explicit inclusion of up to three years of lags in minimum wages mitigates this problem.

The inclusion of a leading minimum wage term allows us to net out any level differences in the outcome between treated and control states just prior to treatment. As a consequence, the medium- and long-run elasticities do not reflect such differences, unlike in a static model. Additionally, the "leading value" $(\beta_{c,-1})$ provides us with a falsification test to discern the reliability of a research design. A statistically significant or sizable leading value, $\beta_{c,-1}$, indicates that the specification may not be able to account for a violation of the parallel trends assumption prior to treatment, and, hence, may provide misleading estimates.[9] To clarify, the $\beta_{\tau,-1}$ term is not included in β_c^{MR} or β_c^{LR}. At the same time, a $\beta_{\tau,-1} \neq 0$ indicates that the treated and control units had some deviation in trends in the past, which raises questions about whether there may be additional deviations in trends in the future. For this reason, I subject all the specifications to the leading value falsification test, and use this information as a criteria for model selection. In some cases, I also show the cumulative response with three years of leads and three years of lags in minimum wages to better evaluate the assumption of parallel trends right before the minimum wage increases.

All regressions control for individual-level covariates X_{it} (quartic in age, and dummies for gender, race and ethnicity, education, family size, number of own children, and marital status), as well as state-level covariates $W_{s(i)t}$ (unemployment rate, state EITC supplement, and per capita GDP). We can calculate the minimum wage elasticity for the proportion under c, γ_c, by dividing α_c by the sample proportion under c. Therefore, γ_1 corresponds to the elasticity of the poverty rate with respect to the minimum wage. The state-level unemployment rate and per capita GDP are time-varying controls to account for aggregate economic shocks in the state that are unlikely to be affected by the policy, given the small share of minimum wage workers in the workforce. All regressions and summary statistics in this paper are weighted by the March CPS sample weights. Finally, the standard errors are clustered by state, which is the unit of treatment.

A problem with the two-way fixed effects model is that there are many potential time varying confounders when it comes to the distribution of family incomes. As shown in Allegretto et al. (2013), high- versus low-minimum wage states over this period are highly spatially clustered, and tend to differ in terms of growth in income inequality and job polarization, and the severity of business cycles. To account for such confounders, I also estimate specifications that allow for arbitrary regional shocks by the nine census divisions, by incorporating division-specific year effects $\theta_{cd(i)t}$. This is motivated by the finding in Allegretto, Dube, and Reich (2011) and Dube, Lester, and Reich (2010) of the importance of spatial heterogeneity in estimating minimum wage effects on employment, and these papers utilize

[9] For example, Dube, Lester, and Reich (2010) and Allegretto et al. (2017) show that the two-way fixed effects model often fails this falsification test when it comes to minimum wage impact on teen and restaurant employment.

division-specific time effects as well. Additionally, I will consider specifications with state-specific linear trends, $\sigma_{s(i)}t$, to account for long-run trend differences between states.

Given the importance of the business cycle as a determinant of family incomes and movements in the poverty rate, I pay special attention to the issue in this paper. The inclusion of the state unemployment rate and year dummies are the usual means of accounting for cyclical factors. However, there are strong prior reasons to worry about business cycle heterogeneity across states when it comes to poverty and minimum wages. Allegretto et al. (2013) shows that minimum wage increases are not uniformly distributed throughout the business cycle—they tend to occur more frequently during the second half of economic expansions. That paper also shows that states with higher minimum wages over the 1990–2012 period experienced sharper business cycle fluctuations. Moreover, states with higher minimum wages may systematically differ with respect to other attributes (such as unemployment insurance generosity) that may affect how a given change in the state unemployment rate translates into changes in family incomes or the incidence of poverty. As another illustration, Zipperer (2016) shows that low-wage employment was more reliant on the construction sectors in states bound by the federal minimum wage increase during 2007–2009; as a result, these states suffered a much steeper relative decline in low-wage employment during this period quite independent of minimum wage policy. For these reasons, I also consider specifications that include state-specific, Great Recession-year indicators, $\rho_{cr(t)s(i)}$, whereby a dummy for each Great Recession year is interacted with a dummy for the state; that is, state fixed effects interacted with separate dummies for each recessionary year: 2007, 2008, 2009.[10] This specification allows state-level outcomes to respond arbitrarily to each recession, but as a consequence of the inclusion of the state-specific recession-year dummies, the identifying variation in such specifications is largely limited to non-recessionary periods.[11]

The most saturated specification is as follows:

$$(2)\qquad I_{cit} = \sum_{k=-1}^{3} \alpha_{ck} \ln\left(MW_{s(i)t-k}\right) + X_{it}\Lambda_c + W_{s(i)t}\Psi_c + \mu_{cs(i)} + \theta_{cd(i)t} + \rho_{cr(t)s(i)} + \sigma_{s(i)}t + \epsilon_{cit}.$$

Besides equations (1) and (2), I also show results from all six of the intermediate specifications with combinations of the three sets of controls (division-specific year effects, state-specific recession-year effects, and state-specific linear trends), and discuss the full range of estimates.[12] Additionally, I assess the relative contribution

[10] These correspond to CPS survey years 1991, 1992, 2002, and 2008–2010.

[11] An added concern raised by Neumark, Salas, and Wascher (2014) is that recessionary periods can influence the estimation of state-specific trends. As Allegretto et al. (2017) argues, this too can be handled by the inclusion of state-specific, recession-year dummies. In studying minimum wage effects on welfare caseloads, Page, Spetz, and Millar (2005) also uses state-specific business cycle controls, although they interact the unemployment rate with state dummies.

[12] Using quadratic instead of linear trends in the saturated model produces similar results.

of each of the three sets of controls in explaining the difference between estimates from equations (1) and (2).

I estimate a series of regressions for alternative income cutoffs. In the main tables, I report the impact of minimum wages on the shares below the following cutoffs: 0.5, 0.75, 1, 1.25, 1.5, and 1.75 times the federal poverty threshold. In the figures, I show the effects for each threshold between 0.5 and 3.5 times the poverty threshold in increments of 0.1. Note that 3.5 times the threshold exceeds the median income-to-needs ratio in the sample (3.04). Consequently, these estimates characterize the impact of the policy on the bottom half of the equivalized family income distribution. The estimates for cutoffs near the middle of the distribution are also useful as falsification tests, since we do not expect the minimum wage to substantially affect incomes in that range.

Unconditional Quantile Partial Effects.—When we estimate the impact of a policy on the proportion of individuals below various income cutoffs, and do so for a large number of such cutoffs, the semi-elasticities show the effect of a log point increase in the minimum wage on the cumulative distribution function (CDF) of family incomes. This is an example of *distribution regressions* as discussed in Chernozhukov, Fernández-Val, and Melly (2013). Moreover, if we have estimates for the impact of the policy on the CDF for all values of an outcome y, we can then invert the impact of the policy on the CDF to estimate the effect of the policy on a particular quantile Q_τ of y. Figure 1 illustrates the concept. $F_A(y)$ is the actual CDF of the outcome y, say equivalized family income. The function $F_B(y)$ represents the counterfactual CDF, showing the distribution that would occur were there to be a small increase in the minimum wage. Under the assumption of conditional independence of the treatment, $F_B(y)$ is estimable using distribution regressions such as equations (1) or (2) of the outcome $I_c = \mathbf{1}(y)$ on the treatment, along with a set of covariates, for every value of c. The resulting estimates would fully characterize the impact of the treatment on the CDF of y, i.e., $F_B(y) - F_A(y)$, and, hence, form an estimate of the counterfactual distribution $F_B(y)$.

Say we are interested in the effect of the policy on the τth quantile of the outcome y. The unconditional quantile partial effect (UQPE) estimand is defined as $Q_{B,\tau} - Q_{A,\tau} = F_B^{-1}(\tau) - F_A^{-1}(\tau)$. It is a partial effect of minimum wages, since the distribution regressions used to estimate the counterfactual, $F_B(y)$, hold other covariates constant. It is an unconditional quantile effect because it measures the impact of the policy on quantiles of the unconditional (or marginal) distribution of y, which in the minimum wage context is of greater policy relevance than the conditional quantile partial effect (CQPE) that is the estimand associated with the quantile regression (Koenker and Bassett 1978). The latter represents the impact of the treatment on the τth quantile of the distribution of y conditional on covariates. For example, the CQPE informs us of the impact of minimum wages on those with low family incomes within their educational group—be they college graduates or without a high school degree. However, when thinking about distributional effects of minimum wages, we are not as interested in the impact of minimum wages on college graduates with unusually low family incomes—i.e., who are poor relative to other college graduates. We are more interested in the impact on those with low incomes

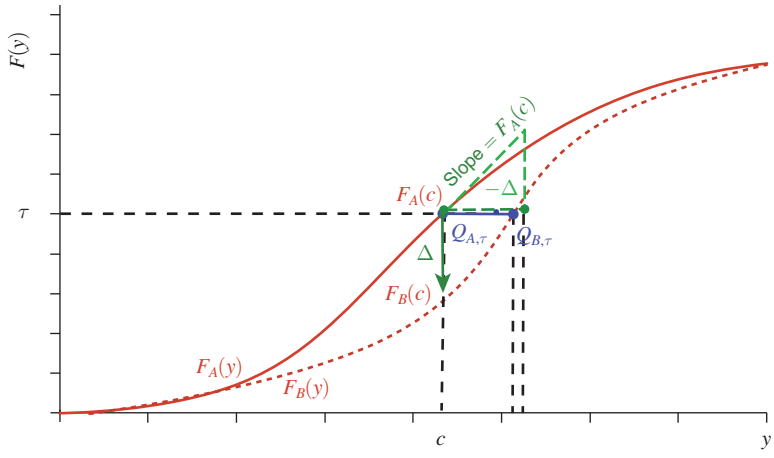


FIGURE 1. UNCONDITIONAL QUANTILE PARTIAL EFFECTS: LOCALLY INVERTING THE COUNTERFACTUAL DISTRIBUTION

Notes: The figure shows how the unconditional quantile partial effect (UQPE) is approximately estimated for a treatment such as a small increase in the minimum wage. $F_A(y)$ represents the actual distribution of outcome y, while $F_B(y)$ is the counterfactual distribution with a higher level of treatment. Under the assumption of conditional independence, the counterfactual distribution can be estimated using distribution regressions of the impact of the policy on the share below cutoffs c for all cutoffs. The UQPE for the τth quantile is $Q_{B,\tau} - Q_{A,\tau}$, represented as the solid (blue) segment. The recentered influence function (RIF) regression approximates the UQPE by inverting the counterfactual CDF $F_B(y)$ using a local linear approximation. After defining a cutoff c such that $F_A(c) = \tau$ using the actual distribution $F_A(y)$, it uses the impact on the proportion below c, i.e., $F_B(c) - F_A(c)$, and the slope of the CDF, $f_A(c)$, to estimate UQPE $\approx -\big(F_B(c) - F_A(c)\big)/f_A(c)$. The dashed (green) triangle shows the geometry of the RIF approximation to the UQPE, which is represented by the length of the triangle's base.

in an absolute (or unconditional) sense.[13] We do wish to *control for* factors like education, but do not wish to *condition* the distributional statistic on (e.g., define "low income" based on) those factors. The UQPE, $Q_{B,\tau} - Q_{A,\tau}$, controls for covariates, but does not define the quantiles based on them; hence, it captures the effect of the policy on the bottom quantiles of the marginal distribution.

It is possible to estimate the UQPE for the τth quantile by (i) estimating the effect of the policy on the proportions under a large set of cutoffs, c, and forming an estimate for the counterfactual distribution $F_B(\tau)$; and then (ii) globally inverting that distribution function and obtain an estimate for $F_B^{-1}(\tau)$ and, hence, an estimate for $F_B^{-1}(\tau) - F_A^{-1}(\tau)$. This procedure is feasible, and outlined in Chernozhukov, Fernández-Val, and Melly (2013). However, it is computationally demanding as it requires estimating a very large number of distribution regressions to globally invert $F_B(y)$ and estimate the quantile effects. As described in Firpo, Fortin, and Lemieux (2009) and Fortin, Lemieux, and Firpo (2011), we can instead invert the counterfactual distribution function using a local linear approximation. Figure 1 provides

[13] To be clear, both the UQPE and CQPE measure the effect of the treatment on low-income quantiles, and not specifically on people who would have earned low incomes (in either a conditional or an unconditional sense) absent the policy. The two concepts coincide only under the additional assumption of rank invariance, i.e., that the treatment does not alter the ranking of individuals.

the intuition behind this approach. We begin by defining a cutoff c associated with quantile τ such that $F_A(c) = \tau$ using the actual distribution. Next, we estimate the effect of the policy on the proportion below c using a single distribution regression. The effect on the proportion is graphically represented as $\Delta = \big(F_B(c) - F_A(c)\big)$ in Figure 1. Now, the quantity $Q_{B,\tau} - Q_{A,\tau}$ can be locally approximated by the product of the vertical distance $-\Delta = -\big(F_B(c) - F_A(c)\big)$ divided by the slope of the distribution function at $F_A(c) = \tau$, which is just the PDF of y at the τth quantile: $f_A\big(F_A^{-1}(\tau)\big)$. The green dashed triangle shows the geometry of this local linear approximation, which can be written as $UQPE \approx -\big(F_B(c) - F_A(c)\big)/f_A(c)$. While the global inversion would require us to estimate a large number of regressions for different values of c in order to obtain the estimate for a single quantile Q_τ, only one regression is needed for each quantile when inverting locally.

The key simplification of taking a linear approximation to the counterfactual CDF works well for a relatively continuous treatment with a substantial variation in treatment intensity, and less well for lumpy or discrete treatments. Given the fairly continuous variation in minimum wages, the approximation error is unlikely to be a major concern here. Later in this section, I discuss a few additional features of the data that further reduce the scope of the approximation error.

To operationalize the estimation, following Firpo, Fortin, and Lemieux (2009), I use as the dependent variable the recentered influence function of income. Since y is a multiple for the federal poverty threshold, for interpretational ease I rescale it by the average real value of the poverty threshold in my sample for the bottom 50 percent of the sample: $y'_{it} = \overline{FPT} \times y_{it} = \$22{,}476 \times y_{it}$. The RIF for the τth quantile, Q_τ, is as follows:

$$(3) \qquad RIF(y'_{it}, Q_\tau) = \left[Q_\tau + \frac{\tau}{f(Q_\tau)}\right] - \frac{\mathbf{1}(y'_{it} < Q_\tau)}{f(Q_\tau)} = k_\tau - \frac{\mathbf{1}(y'_{it} < Q_\tau)}{f(Q_\tau)}.$$

Since the first term in the bracket is a constant, the regression estimate for the UQPE at the τth quantile is simply a rescaled effect of the impact on the proportion under $c(\tau) = Q_\tau$, where the scaling factor is $-1/f_A(Q_\tau)$. This corresponds to the graphical demonstration of the technique in Figure 1.

I estimate a series of regressions for alternative quantiles, Q_τ. Again, I use a range of controls for time-varying heterogeneity across eight different specifications. The most saturated specification is as follows:

$$(4) \qquad RIF(y'_{it}, Q_\tau) = \sum_{k=-1}^{3} \gamma_{\tau,k} \ln\big(MW_{s(i),t-k}\big) + X_{it}\Lambda_\tau + W_{s(i)t}\Phi_\tau + \pi_{\tau s(i)} + \theta_{\tau d(i)t} + \sigma_{\tau s(i)}\, t + \rho_{\tau r(t)s(i)} + \epsilon_{\tau it}.$$

Here, $\gamma_\tau^{LR} = \sum_{k=0}^{3} \gamma_{\tau k}$ is the minimum wage long-run semi-elasticity for the UQPE at the τth quantile of equivalized family income. Note that $\gamma_\tau^{LR} = \big(\overline{FPT}/f\big(c(\tau)\big)\big) \times \alpha_{c(\tau)}^{LR}$, so there is a one-to-one correspondence between the estimates from equations (2) and (4). To obtain the minimum wage elasticity

for the τth income quantile, we divide γ_τ^{LR} by $Q_\tau = c(\tau) \times \overline{FPT}$, so $\eta_\tau^{LR} = \gamma_\tau^{LR} / \left(c(\tau) \times \overline{FPT}\right)$.[14]

A number of features of the data make it attractive for the application of the RIF-UQPE approach. Online Appendix Table B.1 and Figure B.2 show the CDF for the income-to-needs ratio. I note that the CDF is nearly linear in the bottom half of the distribution, especially between income-to-needs ratios of 0.75 and 2.50, which roughly correspond to the tenth and fortieth percentiles: in this range the PDF is essentially flat.[15] The linearity of the actual CDF (in combination with a continuous treatment) reduces the scope of the error from a linear approximation when inverting the counterfactual CDF using the RIF approach.

Additionally, online Appendix Figure B.3 shows that the income quantiles at the bottom of the distribution have been fairly stationary over the past three decades, although they do exhibit cyclical tendencies. Online Appendix Figure B.3 also shows that the probability densities at the associated income-to-needs cutoffs $f_A\left(c(\tau)\right)$ have also been fairly stable over time, with the possible exception of the fifth quantile. The relative stability of the income-to-needs quantiles and densities is relevant for interpreting the UQPE estimates. The estimation of the UQPE for a particular quantile, τ, is based on changes in the proportion below the income-to-needs cutoff $c(\tau)$ associated with that quantile, along with the probability density of the income-to-needs ratio at that cutoff, $f_A\left(c(\tau)\right)$. Both $c(\tau)$ and $f_A(y)$ are calculated by averaging over the entire sample. The relative stability of the mapping between c and τ over this period suggests that the estimated impact on income around a given cutoff c is referring to roughly the same quantile over this full period.

Finally, the use of the full-sample distribution to estimate the cutoff $c(\tau)$ and the density $f_A(c)$ may be an issue if the treatment and control units had very different income distributions. However, all states receive treatment at some point during the sample, and the variation in minimum wages is fairly continuous and widespread; therefore, the sample-averaged cutoffs and densities are broadly representative of where the minimum wage variation is coming from. Overall, the nature of both the treatment as well as the outcome facilitate the application of the RIF approach to a repeated cross-sectional setting.

C. *Descriptive Statistics*

Table 1 shows shares of the non-elderly population, as well as 8 key demographic groups under alternative income cutoffs, which range between 0.5 and 1.75 times the federal poverty threshold. To clarify, row 3 shows the poverty rates, while row 5 shows the share with income below 1.5 times the poverty threshold.

[14] I find that the estimation error in $f(Q_\tau)$ contributes trivially to the overall variance of the η_τ^{LR}. Using the delta method, accounting for the estimation of the density around the cutoff increases the standard error of η_τ^{LR} typically by less than 0.1 percent. For this reason, the results in this paper do not explicitly adjust the standard errors for the estimation of $f(Q_\tau)$.

[15] The kernel density estimation uses an Epanechnikov kernel and the STATA default bandwidth based on Silverman's rule of thumb.

Table 1—Population Shares below Multiples of Federal Poverty Threshold—By Demographic Group

Family income (multiples of FPT)	Age < 65	All, including elderly (65+)	<HS, under 30	<HS	≤HS	Black and Latino	Children	Adults	Single mothers
0.50	0.058	0.054	0.091	0.091	0.076	0.114	0.085	0.047	0.155
0.75	0.095	0.090	0.148	0.152	0.126	0.188	0.139	0.077	0.246
1.00	0.136	0.132	0.206	0.214	0.179	0.263	0.195	0.112	0.327
1.25	0.178	0.178	0.264	0.276	0.233	0.335	0.250	0.148	0.401
1.50	0.222	0.225	0.322	0.337	0.289	0.405	0.305	0.188	0.468
1.75	0.267	0.273	0.377	0.395	0.344	0.469	0.360	0.229	0.531
Observations	4,662,781	5,238,862	1,646,917	1,947,133	3,019,445	1,211,366	1,500,559	3,162,220	240,609

Notes: Each cell contains the proportion of the sample (or the mean of a poverty measure) under the family income cutoff. The sample consists of all non-elderly individuals in March CPS surveys for years between 1984 and 2013. All columns use the implicit equivalence scale used to calculate the official poverty rate. Columns 1 through 9 report the proportions for the following subsamples: non-elderly; all individuals, including elderly; individuals younger than 30 without a high school degree; all individuals without a high school degree; individuals with high school or lesser education; black and Latino individuals; children under 18; adults; and single mothers. All calculations use the March CPS person weights.

For non-elderly adults as a whole, the poverty rate averaged at 0.136 during the 1984–2013 period. Including elderly in the sample decreases it slightly to 0.132. Individuals without a high school degree (0.214) as well as those with high school or lesser schooling (0.179) had greater rates of poverty. The poverty rate for single mothers (0.327), black/Latino individuals (0.263), and children (0.195) were all higher than the average. The poverty rate among adults (0.112), on the other hand, is smaller than other groups. These patterns are as expected, and are qualitatively similar when we consider 50 percent or 150 percent of the poverty threshold.

II. Empirical Findings

A. *Main Results for Shares below Multiples of the Poverty Threshold*

Table 2 provides the estimates for the impact of minimum wages on the shares below various multiples of the federal poverty threshold. For ease of interpretation, I report the estimates as elasticities β_c^{MR} and β_c^{LR} by dividing the sum of regression coefficients by the sample proportion under each cutoff; this is true both for the point estimate and the standard errors. I show estimates using the specification that includes (i) division-specific year effects, (ii) state-specific Great Recession-year dummies, and (iii) state-specific linear trends. In Section IIID, I show that the general findings of poverty reduction are robust to virtually all forms of controls for time-varying heterogeneity used in this literature, including the simple two-way fixed effects model. However, as I also show in Section IIID, the saturated model performs the best in terms of falsification tests, which is why I use it as my preferred specification.

Overall, there is strong evidence that minimum wage increases reduce the share of individuals with low family incomes. The medium-run poverty elasticity is −0.399 (standard error = 0.126), while the long-run estimate is −0.446 (standard error

TABLE 2—MINIMUM WAGE ELASTICITIES FOR SHARE OF INDIVIDUALS WITH FAMILY INCOME BELOW MULTIPLES OF FEDERAL POVERTY THRESHOLD

Family income cutoff	Medium-run (2 year lagged)	Long-run (3+ year lagged)
0.50	−0.351	−0.455
	(0.236)	(0.247)
0.75	−0.537	−0.461
	(0.180)	(0.186)
1.00	−0.399	−0.446
	(0.126)	(0.137)
1.25	−0.048	−0.294
	(0.131)	(0.103)
1.50	0.060	−0.156
	(0.127)	(0.093)
1.75	0.027	−0.167
	(0.095)	(0.084)
Observations	4,662,781	4,662,781

Notes: The reported estimates are minimum wage elasticities for share with family income under various multiples of the federal poverty threshold. These estimates are from linear probability models that regress an indicator for having family income below multiples (between 0.50 and 1.75) of the federal poverty threshold on distributed lags of log minimum wage and covariates. The medium-run and long-run elasticities are calculated by summing the contemporaneous up to two- and three-year lagged log minimum wage coefficients, respectively, and then dividing by the sample proportion under the family income cutoff. Controls include state fixed effects, division-by-year fixed effects, state-specific linear trends, state-specific indicator for each Great Recession year, state-level covariates (GDP per capita, EITC supplement, and unemployment rate), and individual demographic controls (quartic in age, as well as dummies for race, marital status, family size, number of children, education level, Hispanic status, and gender), and are weighted by March CPS person weights. Additional division-by-year fixed effects, state-specific linear trends, and state-specific indicators for each Great Recession year are indicated in the table. State-cluster-robust standard errors are in parentheses.

$= 0.137$), both being statistically significant at the 1 percent level. More generally, there is a reduction in the shares of individuals with incomes below family income cutoffs between 50 and 125 percent of the official poverty threshold: the long-run elasticities range between -0.294 and -0.461, and these are all statistically significant at the 5 percent level (except for the 50 percent cutoff, which is significant at the 10 percent level). Figure 2 provides visual evidence on how minimum wages affect the bottom half of the family income distribution. The figure reports elasticities for each multiple of the poverty threshold between 0.5 and 3.5 in increments of 0.1—a total of 31 regressions. The top panel shows that by the second year after the policy change, there are large, statistically significant reductions in the share below cutoffs between 0.5 and 1.2 times the poverty threshold with minimum wage elasticities exceeding -0.5 in magnitude. The elasticities for cutoffs exceeding 1.2 are mostly close to zero, and never statistically significant. The bottom panel of the figure shows that that long run (3+ year) minimum wage elasticities are sizable and statistically significant for cutoffs between roughly 0.6 and 1.4 times the poverty threshold. However, there is some additional reduction in the shares below cutoffs between 1.4 and 2 times the poverty threshold. The elasticities subsequently decline in magnitude, and are close to zero for cutoffs of 2.75 times the poverty threshold or

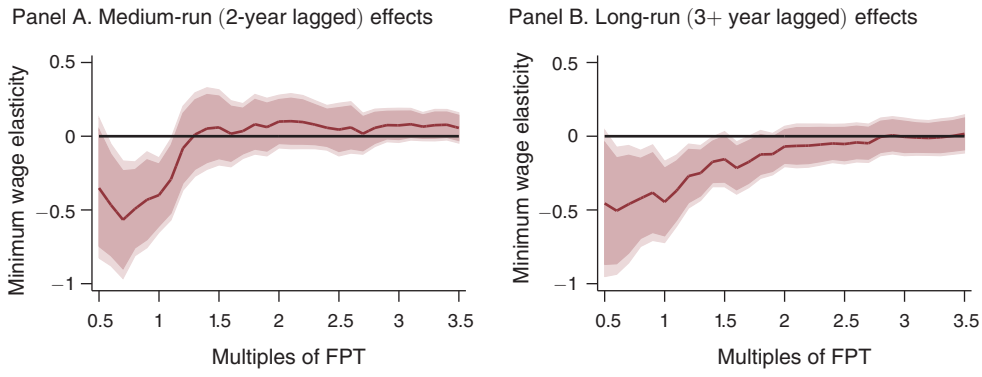


FIGURE 2. MINIMUM WAGE ELASTICITIES FOR SHARE UNDER ALTERNATIVE FAMILY INCOMES IN MULTIPLES OF FEDERAL POVERTY THRESHOLD

Notes: The reported estimates are long-run minimum wage elasticities for share with family income under various multiples of the federal poverty threshold. These estimates are from a series of linear probability models that are estimated by regressing an indicator for having family income below cutoffs between 0.5 and 3.5 times the federal poverty threshold (in increments of 0.05) on distributed lags of log minimum wage and covariates. The medium-run (2 year) and long-run (3+ year) elasticities are calculated by dividing the respective coefficients (or sums of coefficients) by the sample proportion under the cutoff. All specifications include state fixed effects, division-by-year fixed effects, state-specific linear trends, state-specific indicator for each Great Recession year, state-level covariates (GDP per capita, EITC supplement, and unemployment rate), and individual demographic controls (quartic in age as well as dummies for race, marital status, family size, number of children, education level, Hispanic status, and gender), and are weighted by March CPS person weights. The dark shaded area represents 90 percent and the light shaded area 95 percent state-cluster-robust confidence intervals.

greater. Overall, there is an indication that the medium run (2 year) effects are more concentrated at lower incomes than longer run effects. The lagged effects higher up in the family income distribution could be consistent with a lagged spillover effect of minimum wages on the wage distribution.

While the baseline specifications use one year of lead and three years of lags, to further assess the issue of parallel trends, I also estimate a longer distributed lag model with three leads and three lags. Figure 3 reports the cumulative response elasticities by successively adding the leads and lags, normalizing the elasticity to be 0 at event time equal to -1. The top panel shows the changes for the share below the federal poverty threshold: the share is stable in the three years prior to the policy change, but starts falling starting the year of the minimum wage increase, and reaches its maximum impact two years after the change. In contrast, the share below three times the federal poverty threshold shows no systematic movement: the cumulative response is close to zero for the three years prior to the policy change, and straddles zero subsequently. Overall, this evidence on timing provides additional validation for the research design.

To assess possible heterogeneity in the effect of the policy, Table 3 shows minimum wage elasticities for the shares below alternative family-income cutoffs for various demographic groups defined using education, age, gender, race, presence of kids, and marital status. For all groups, I find sizable reductions in the proportions under 50, 75, 100, and 125 percent of the poverty threshold. The long-run estimates of these 32 elasticities (8 groups with 4 cutoffs each) range between -0.225 and

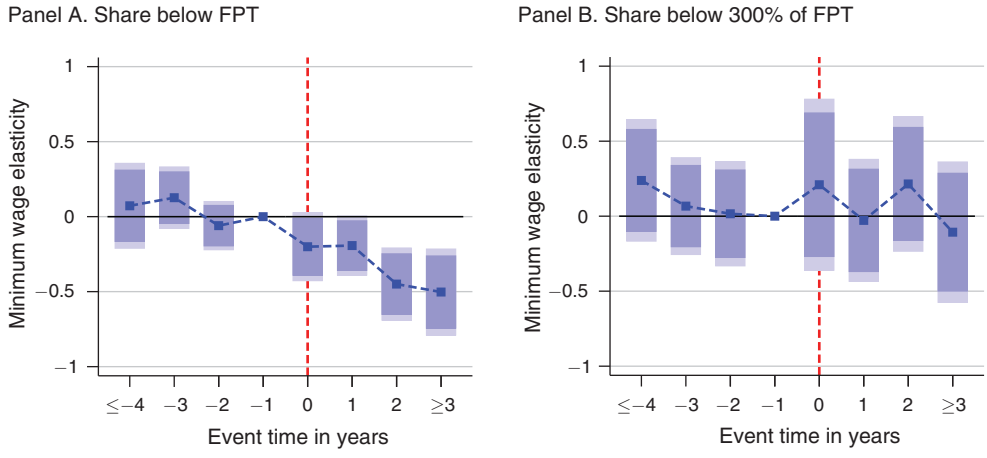


FIGURE 3. CUMULATIVE RESPONSE OF SHARE BELOW POVERTY THRESHOLD, AND THREE TIMES POVERTY THRESHOLD TO A LOG POINT INCREASE IN MINIMUM WAGES

Notes: The reported estimates are minimum wage elasticities for share with family income below one or three times the federal poverty threshold by event date. The elasticities are from linear probability models estimated by regressing an indicator for having family income below cutoffs of one and three times the federal poverty threshold on covariates and a distributed lags window including up to three leads and lags of log minimum wage. The elasticity at event date −1 is normalized to 0. The remaining elasticities are calculated by summing up the joint effect and dividing it by the sample proportion under the cutoff. Specifications include state fixed effects, division-by-year effects, state-specific linear trends, state-specific indicator for each Great Recession year, state-level covariates (GDP per capita, EITC supplement, and unemployment rate), and individual demographic controls (quartic in age and as well as dummies for race, marital status, family size, number of children, education level, Hispanic status, and gender). Regressions are weighted by March CPS person weights. The dark shaded area represents 90 percent, the light shaded area 95 percent state-cluster-robust confidence intervals.

−0.872, and 28 are statistically significant at least at the 10 percent level. As compared to our previous estimates for non-elderly individuals (−0.446), the poverty rate elasticities are quite similar for children (−0.499), those without a high school degree (−0.481), those with high school or lesser schooling (−0.474), all individuals including elderly (−0.436), all adults (−0.375), and individuals under 30 without a high school degree (−0.432). In contrast, they are smaller in magnitude for single mothers (−0.172), and substantially larger in magnitude for black and Latino individuals (−0.872). The reductions in low-income shares extend somewhat further up the distribution for black and Latino individuals, children under 18, and those under 30 without a high school degree; for these groups, there are substantial and statistically significant reductions for up to 175 percent of the poverty threshold. The key conclusion from these findings is that when we focus on disadvantaged groups such as black or Latino individuals, or those with lesser education, the antipoverty impact of minimum wages appears to be similar or somewhat greater; however, for another disadvantaged group (single mothers), the impact is somewhat smaller.

Next, I compare my findings with what the existing research suggests about heterogeneous impact by age, single mother status, education, and race, as summarized in online Appendix Table A.1. First, if we take the poverty rate elasticities for groups under 20 years of age in the literature, Morgan and Kickham (2001)

Table 3—Minimum Wage Elasticities for Share of Individuals with Family Income below Multiples of Federal Poverty Threshold by Demographic Subgroup

Family income cutoff	All, including elderly (65+)	<HS, under 30	<HS	≤HS	Black and Latino	Children	Adults	Single mothers
Panel A. Medium-run (2-year lagged) estimates								
0.50	−0.305	−0.523	−0.622	−0.487	−0.541	−0.430	−0.249	−0.501
	(0.242)	(0.291)	(0.279)	(0.244)	(0.421)	(0.264)	(0.226)	(0.391)
0.75	−0.503	−0.671	−0.768	−0.671	−0.715	−0.651	−0.411	−0.471
	(0.178)	(0.259)	(0.228)	(0.179)	(0.287)	(0.253)	(0.151)	(0.233)
1.00	−0.345	−0.442	−0.535	−0.424	−0.596	−0.387	−0.368	−0.350
	(0.115)	(0.191)	(0.159)	(0.116)	(0.168)	(0.188)	(0.122)	(0.144)
1.25	−0.029	−0.069	−0.124	−0.037	−0.319	0.004	−0.044	−0.078
	(0.110)	(0.185)	(0.158)	(0.137)	(0.143)	(0.188)	(0.113)	(0.112)
1.50	0.097	−0.024	−0.004	0.041	−0.137	0.015	0.120	0.182
	(0.098)	(0.197)	(0.170)	(0.140)	(0.119)	(0.201)	(0.104)	(0.145)
1.75	0.056	−0.024	0.031	0.026	−0.130	0.015	0.060	−0.049
	(0.075)	(0.146)	(0.129)	(0.104)	(0.097)	(0.151)	(0.082)	(0.099)
Panel B. Long-run (3+ year lagged) estimates								
0.50	−0.465	−0.522	−0.527	−0.543	−0.752	−0.551	−0.337	−0.584
	(0.242)	(0.294)	(0.287)	(0.271)	(0.375)	(0.314)	(0.210)	(0.286)
0.75	−0.469	−0.445	−0.488	−0.525	−0.831	−0.500	−0.390	−0.277
	(0.176)	(0.223)	(0.202)	(0.195)	(0.279)	(0.241)	(0.167)	(0.224)
1.00	−0.436	−0.432	−0.481	−0.474	−0.872	−0.499	−0.375	−0.172
	(0.126)	(0.168)	(0.140)	(0.139)	(0.178)	(0.191)	(0.123)	(0.210)
1.25	−0.285	−0.316	−0.309	−0.263	−0.745	−0.356	−0.225	−0.184
	(0.096)	(0.120)	(0.108)	(0.107)	(0.121)	(0.139)	(0.111)	(0.161)
1.50	−0.124	−0.228	−0.173	−0.136	−0.567	−0.286	−0.050	−0.081
	(0.090)	(0.124)	(0.118)	(0.111)	(0.100)	(0.136)	(0.098)	(0.170)
1.75	−0.115	−0.204	−0.129	−0.144	−0.528	−0.232	−0.110	−0.162
	(0.086)	(0.101)	(0.101)	(0.097)	(0.104)	(0.107)	(0.096)	(0.134)
Observations	5,238,862	1,646,917	1,947,133	3,019,445	1,211,366	1,500,559	3,162,220	240,607

Notes: The reported estimates are minimum wage elasticities for share with family income under various multiples of the federal poverty threshold. These estimates are from linear probability models that regress an indicator for having family income below multiples (between 0.50 and 1.75) of the federal poverty threshold on distributed lags of log minimum wage and covariates. The medium-run and long-run elasticities are calculated by summing the contemporaneous up to two- and three-year lagged log minimum wage coefficients, respectively, and then dividing by the sample proportion under the family income cutoff. The regression specification includes state fixed effects, division-specific year effects, state-specific indicator for each Great Recession year, state-specific linear trends, state-level covariates (GDP per capita, EITC supplement, and unemployment rate), and individual demographic controls (quartic in age as well as dummies for race, marital status, family size, number of children, education level, Hispanic status, and gender), and are weighted by March CPS person weights. <HS refers to individuals without a high school degree, while <HS under 30 further restricts to individuals under 30 years of age. ≤HS refers to individuals with high school or lesser schooling. State-cluster-robust standard errors are in parentheses.

and Addison and Blackburn (1999) also find a substantial effect of, respectively, −0.39 and −0.39 (averaged across specifcations for teens). While my estimate for all children are somewhat larger in magnitude (−0.499), both existing work and results in this paper point toward a substantial poverty reducing impact of minimum wages among children.

Second, for single mothers, I find elasticities for the proportion under the poverty line of −0.172, and under one-half poverty line of −0.584. As noted above, the poverty elasticity for single mothers is smaller than for the population overall, and is not

statistically significant at conventional levels. The implied elasticities in Neumark (2016) for 21–44-year-old single women with kids range between −0.14 and −0.30 depending on specification. These estimates are broadly similar to what I find, though I note the imprecision in both sets of estimates. Sabia (2008) finds a range of elasticities between −0.28 and −0.17 for single mothers, depending on the mother's education level. Burkhauser and Sabia (2007) finds poverty rate elasticities for single mothers between −0.21 and −0.07, depending on specification. DeFina (2008) finds poverty rate elasticities in woman headed households with kids of −0.42 (−0.35 when restricting to mothers without a college education). Finally, Gundersen and Ziliak (2004) finds very small effects for woman-headed households (−0.02). If we take an average of the poverty rate elasticities for single mothers (or woman heads of households) across these five studies, we get an average elasticity of −0.18, which is close to my estimate of −0.17.

The third comparison concerns heterogeneity in the effect by levels of education. I find poverty elasticities for those without a high school degree (−0.481) or with high school or lesser schooling (−0.474) that are comparable to the overall estimate (−0.446). Neumark (2016) provides estimates for those 21 years of age or older by education: the unrestricted estimates range between −0.11 and −0.15, while for those with high school or lesser education, they range between −0.01 and −0.19; none of these are statistically significant. Sabia (2008) finds somewhat larger reductions in the poverty rate for single mothers without a high school degree (−0.28) than those with (−0.17), although neither estimate is statistically significant. In contrast, restricting to those with less education tends to slightly diminish the effects in DeFina (2008), though they continue to be sizable (changing from −0.42 to −0.35). The estimates in Sabia and Nielsen (2015) are highly imprecise, and the impact of conditioning on education levels is contradictory across specifications. Finally, while Addison and Blackburn (1999) do not provide comparable estimates by levels of education, averages across their specifications do suggest a somewhat large elasticity (−0.45) for individuals who did not complete junior high school. While the estimates in the literature do not paint a clear picture, on balance they do not suggest that the poverty reducing effect of minimum wages is smaller among those with less education. That is consistent with what I find here.

The fourth, and final, comparison concerns heterogeneity by race. Here, I find clear evidence of substantially stronger reduction in poverty, and near poverty, among black or Latino individuals as compared to the population as a whole. Gundersen and Ziliak (2004) also finds a slightly larger effect in the black population—though the magnitude is still very small (−0.06). Finally, the estimates in Sabia and Nielsen (2015) are, again, imprecise and qualitatively differ by specification. This paper provides sharper evidence than available in existing work that minimum wages tend to raise bottom incomes more strongly for African Americans and Latinos.

B. *Effect on Family Income Quantiles*

We can use the impact of minimum wages on the cumulative distribution of family income to estimate the impact on family income quantiles. Figure 4 shows the sample average CDF of family income, along with the counterfactual distribution.

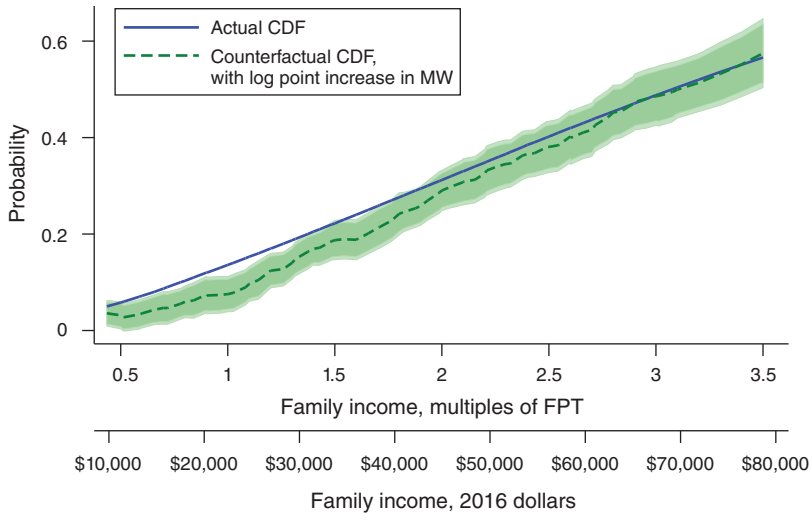


FIGURE 4. ACTUAL AND COUNTERFACTUAL CUMULATIVE DISTRIBUTION FUNCTIONS OF FAMILY INCOME IN MULTIPLES OF FEDERAL POVERTY THRESHOLD

Notes: The solid (blue) line shows the actual cumulative distribution function (CDF) of family income, estimated as the sample average over the full period. The dashed (green) line shows the counterfactual CDF of family income with a log point increase in the minimum wage. The counterfactual is calculated by adding to the actual CDF the regression-based, long-run effects of minimum wage on the share below cutoffs. These long-run estimates come from a series of linear probability models that are estimated by regressing an indicator for having a family income below cutoffs (between 0.40 and 3.50) on distributed lags of log minimum wage and covariates. The long-run estimates are calculated as the sum of the contemporaneous up to three-year lagged log minimum wage coefficients. All specifications include state fixed effects, division-by-year fixed effects, state-specific linear trends, state-specific indicator for each Great Recession year, state-level covariates (GDP per capita, EITC supplement, and unemployment rate), and individual demographic controls (quartic in age as well as dummies for race, martial status, family size, number of children, education level, Hispanic status, and gender), and are weighted by March CPS person weights. The dark shaded area represents 90 percent, the light shaded area 95 percent state-cluster-robust confidence intervals. The additional x-axis shows the average family income values in 2016 dollars corresponding to the multiples of federal poverty threshold.

In the top scale of the figure, income is shown as multiples of the federal poverty threshold, y. The bottom scale shows income in dollar amounts obtained by multiplying y by the sample average of the real federal poverty threshold ($22,476). Here, the counterfactual distribution at income c is calculated as $F(c) + \alpha_c^{LR}$, where $F(c)$ is the CDF averaged over the entire sample, and α_c^{LR} is the 3+ year minimum wage semi-elasticity estimated from equation (2). In other words, it is the counterfactual CDF with a log point increase in the minimum wage. The figure is based on 31 separate regressions estimated for cutoffs c between 0.5 and 3.5 in increments of 0.1, using the preferred specification. This is just a different way of reporting the estimates shown in Figure 2, which instead showed the impact as elasticities. In both figures, we can see that there is a statistically significant and quantitatively substantial shift in the CDF for incomes between 0.5 and 1.75 times the poverty threshold. Overall, an increase in the minimum wage seems to improve the family income distribution in the sense of first-order stochastic dominance. By juxtaposing the actual and counterfactual CDF, we can also visually assess the impact by quantiles (the vertical axis). By considering the horizontal distance between the actual

and counterfactual CDFs, we can infer that there is clear, statistically significant, increases in incomes between the tenth and twentieth percentiles. And that above the fortieth percentile, the changes (horizontal distance) are mostly close to zero.

This exercise illustrates the key idea behind estimating the unconditional quantile partial effects. In principle, we can estimate the quantiles by globally inverting the CDF as we did visually when inspecting Figure 4. This is similar to the approach in Chernozhukov, Fernández-Val, and Melly (2013). While feasible, this requires estimating a large number of regressions, even when estimating the effect for a single quantile. Therefore, the global inversion can be computationally demanding, especially as we wish to estimate it for many specifications with a large set of fixed effects and for different income definitions, and using individual level data. For this reason, I use the RIF approach proposed by Firpo, Fortin, and Lemieux (2009). We approximate the horizontal distance between the actual and counterfactual CDFs at quantile q by dividing the vertical distance between the two by the PDF at $q = F^{-1}(c)$. This requires us to estimate only a single regression per quantile. The unconditional quantile partial effects (β_τ) are estimated using equation (3), or analogous regressions for the less saturated specifications. To convert the UQPEs into elasticities (η_τ), they are subsequently divided by the income-to-needs cutoffs corresponding to a given quantile.

Figure 5 shows the long-run minimum wage elasticities for each family income quantile between 5 and 50 using the preferred specification. We find substantial and statistically significant effects for quantiles between the seventh and twentieth, declining sharply by the thirty-fifth quantile. The elasticities around the median of the family income distribution straddle zero. The figure provides compelling visual evidence that minimum wage policies tend to raise income at the bottom of the distribution, consistent with evidence on the CDF in Figure 4. Column 2 of Table 4 reports the magnitudes of the medium and long-run elasticities for the baseline (cash) income definition. The long-run elasticities are 0.273, 0.359, 0.332, 0.152, and 0.164 for the fifth, tenth, fifteenth, twentieth and twenty-fifth quantiles of equivalized family incomes, respectively. The tenth, fifteenth, and twenty-fifth quantiles are statistically significant at least at the 5 percent level, while the twentieth quantile is significant at the 10 percent level.[16,17]

So far, our income defintion has corresponded to that used in the official poverty definition, which limits to cash-based incomes. Next, we expand the income definition, and estimate minimum wage elasticities for quantiles of family income

[16] As noted above, globally inverting the CDF is computationally intensive, especially for conducting inference. However, the estimates are broadly similar. For example, the elasticities for tenth, fifteenth, twentieth, thirtieth, fortieth, and fiftieth quantiles when inverting globally are 0.44, 0.26, 0.18, 0.09, 0.05, and 0.08, while the locally inverted estimates are 0.36, 0.33, 0.15, 0.09, 0.05, and 0.02, respectively. (Results not reported in tables.)

[17] These are based on point-wise confidence intervals. Additionally, I implement the bootstrap-based approach of Lehrer, Pohl, and Song (2016)—which identifies the ranges of the outcome distribution that exhibit positive treatment effects accounting for possible multiple testing concerns. The more conservative inference of Lehrer, Pohl, and Song (2016) controls for the family-wise error rate, which here is the probability that we mistakenly declare a positive UQPE for at least one quantile in the considered range. When I consider the quantiles 10, 15, 20, 25, and 30, I find that the tenth and fifteenth quantile effect are statistically significant at the 5 percent level. This contrasts with the point-wise CI, which are significant for the tenth, fifteenth, and twenty-fifth quantiles. Overall, the conclusion that there are positive effects at the tenth and fifteenth quantiles do not appear to be driven by multiple-testing concerns.

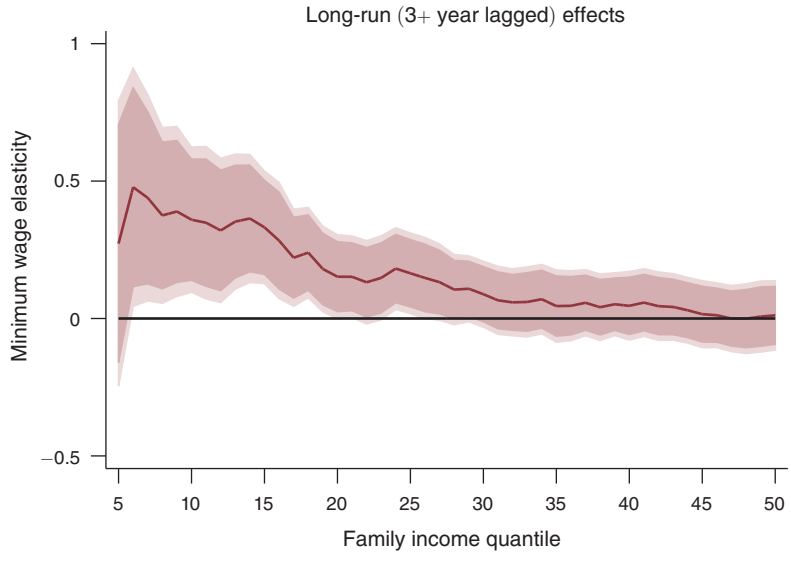


FIGURE 5. MINIMUM WAGE ELASTICITIES FOR UNCONDITIONAL FAMILY INCOME QUANTILES

Notes: The reported estimates are long-run minimum wage elasticities by unconditional quantiles of family income. The estimates are from a series of linear probability models that are estimated by regressing an indicator for having a family income below cutoffs associated with alternative quantiles (between 5 and 50) on distributed lags of log minimum wage and covariates. The long-run coefficients are calculated as the sum of the contemporaneous up to three-year lagged log minimum wage coefficients. Unconditional quantile partial effects (UQPE) are calculated by dividing the coefficient on log minimum wage by the negative of the family income density at the appropriate quantile. The UQPE estimates are subsequently divided by the family income cutoff for the quantile to transform the estimates into elasticities. All specifications include state and year fixed effects, division-by-period fixed effects, state-specific linear trends, state-specific indicator for each Great Recession year, state-level covariates (GDP per capita, EITC supplement, and unemployment rate), and individual demographic controls (quartic in age as well as dummies for race, marital status, family size, number of children, education level, Hispanic status, and gender), and are weighted by March CPS person weights. The dark shaded area represents 90 percent, the light shaded area 95 percent state-cluster-robust confidence intervals.

inclusive of tax credits and noncash transfers. These elasticities are reported in column 4 of Table 4, and are generally smaller than those using pretax cash income (column 2). Specifically, averaged across the tenth and the fifteenth quantiles—where there is a sizable effect—the elasticities are 69 percent and 60 percent as large from using the expanded income definition for the medium and long run, respectively. Figure 6 shows the elasticities visually for each quantile between 5 and 50 using the preferred specification, where we continue to find sizable and statistically significant effects between seventieth and twentieth quantiles. However, the elasticities using expanded income are about 60 percent as large as those in Figure 5 for those quantiles.

While the elasticities are useful in understanding how responsive income is to minimum wages for various quantiles, they are less useful for assessing how much of the cash income gain is being offset by a loss in tax credits or noncash transfers. This is because the proportionate change in cash versus expanded incomes is also driven by the fact that the expanded incomes are larger: so a $1 increase in income due to minimum wages is a smaller share of expanded income, just because the $1 is

TABLE 4—MINIMUM WAGE EFFECTS ON UNCONDITIONAL QUANTILES OF FAMILY INCOMES, FOR ALTERNATIVE INCOME DEFINITIONS

Income	Cash income		+ Noncash transfers		+ Noncash transfers tax credits	
Family income quantile	Semi-elasticity	Elasticity	Semi-elasticity	Elasticity	Semi-elasticity	Elasticity
Panel A. Medium-run (2-year lagged) estimates						
5	$1,923	0.196	$884	0.070	−$31	−0.002
	(2,364)	(0.242)	(2,657)	(0.210)	(2,750)	(0.203)
10	$7,794	0.445	$7,955	0.409	$7,757	0.370
	(2,358)	(0.135)	(2,599)	(0.134)	(2,457)	(0.117)
15	$6,912	0.284	$3,612	0.141	$4,255	0.155
	(2,781)	(0.114)	(2,918)	(0.114)	(2,412)	(0.088)
20	−$1,222	−0.040	−$511	−0.016	−$580	−0.017
	(3,584)	(0.116)	(3,108)	(0.098)	(2,966)	(0.089)
25	−$846	−0.023	−$1,758	−0.047	−$1,362	−0.035
	(2,816)	(0.076)	(2,813)	(0.075)	(2,628)	(0.067)
30	−$2,611	−0.060	−$3,125	−0.071	−$3,731	−0.083
	(3,441)	(0.079)	(3,220)	(0.074)	(3,040)	(0.068)
Panel B. Long-run (3+ year lagged) estimates						
5	$2,669	0.273	$1,746	0.138	$1,079	0.080
	(2,528)	(0.258)	(2,364)	(0.187)	(2,656)	(0.197)
10	$6,294	0.359	$6,258	0.322	$5,370	0.256
	(2,302)	(0.131)	(2,599)	(0.134)	(2,439)	(0.116)
15	$8,077	0.332	$6,200	0.243	$4,377	0.160
	(2,484)	(0.102)	(2,339)	(0.092)	(2,548)	(0.093)
20	$4,696	0.152	$5,002	0.158	$3,794	0.114
	(2,345)	(0.076)	(2,152)	(0.068)	(2,208)	(0.066)
25	$6,121	0.164	$5,051	0.134	$4,155	0.107
	(2,719)	(0.073)	(2,496)	(0.066)	(2,363)	(0.061)
30	$3,838	0.088	$4,108	0.094	$2,254	0.050
	(2,599)	(0.060)	(2,490)	(0.057)	(2,371)	(0.053)
Observations	4,662,781	4,662,781	4,662,781	4,662,781	4,662,781	4,662,781

Notes: The reported estimates are minimum wage elasticities for unconditional quantiles of equivalized family income using alternative income definitions. The estimates are from linear probability models that regress an indicator for having a family income below cutoff associated with a quantile (between 5 and 30) on distributed lags of log minimum wage and covariates. The first two columns only consider pretax cash income, columns 3 and 4 augment it by adding noncash transfers (SNAP, NLSP, and housing subsidies), and the last two columns further add tax credits (EITC and child tax credits) to the family income definition. Medium- and long-run unconditional quantile partial effects (UQPE or semi-elasticities) for family incomes are calculated by summing of the contemporaneous up to two- and three-year lagged log minimum wage coefficients, respectively, and dividing by the negative of the family income density at the appropriate quantile. To calculate elasticities, the UQPE estimates are subsequently divided by the family income cutoff for the quantile. The regression specification includes state fixed effects, division-specific year effects, state-specific indicator for each Great Recession year, state-specific linear trends, state-level covariates (GDP per capita, EITC supplement, and unemployment rate), and individual demographic controls (quartic in age as well as dummies for race, marital status, family size, number of children, education level, Hispanic status, and gender), and are weighted by March CPS person weights. State-cluster-robust standard errors are in parentheses.

divided by a larger denominator. For example, the fifteenth quantile in cash income is $24,123. When we expand the income definition to include noncash transfers and tax credits, it is $27,423. Therefore, even if the income gains using the two definitions are identical, the elasticity using the expanded income definition would be around 12 percent smaller.

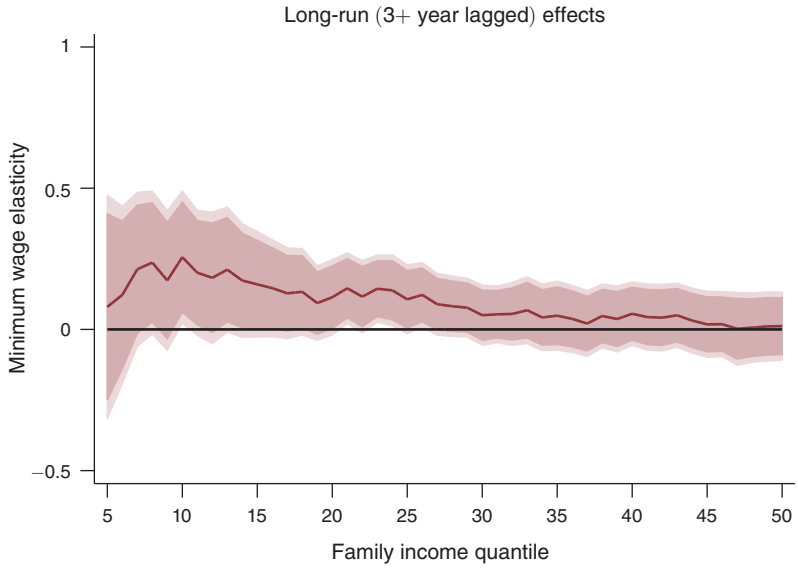


FIGURE 6. MINIMUM WAGE ELASTICITIES FOR UNCONDITIONAL FAMILY INCOME QUANTILES; EXPANDED INCOME WITH TAX CREDITS AND NONCASH TRANSFERS

Notes: The reported estimates are long-run minimum wage elasticities by unconditional quantiles of an expanded definition of family income, which includes tax credits (EITC, child tax credit) and noncash transfers (SNAP, NSLP, and housing subsidy). The estimates are from a series of linear probability models that are estimated by regressing an indicator for having an expanded family income below cutoffs associated with alternative quantiles (between 5 and 50) on distributed lags of log minimum wage and covariates. The long-run coefficients are calculated as the sum of the contemporaneous up to three-year lagged log minimum wage coefficients. Unconditional quantile partial effects (UQPE) are calculated by dividing the coefficient on log minimum wage by the negative of the family income density at the appropriate quantile. The UQPE estimates are subsequently divided by the expanded income cutoff for the quantile to transform the estimates into elasticities. All specifications include state and year fixed effects, division-by-period fixed effects, state-specific linear trends, state-specific indicator for each Great Recession year, state-level covariates (GDP per capita, EITC supplement, and unemployment rate), and individual demographic controls (quartic in age as well as dummies for race, marital status, family size, number of children, education level, Hispanic status, and gender), and are weighted by March CPS person weights. The dark shaded area represents 90 percent, the light shaded area 95 percent state-cluster-robust confidence intervals.

To better capture the extent of offset, Table 4 reports the minimum wage effect on the level of real incomes (i.e., the semi-elasticities, γ_τ) for progressively more expansive income definitions. When we consider all 5 quantiles (fifth, tenth, fifteenth, twentieth, twenty-fifth, and thirtieth), the increase in expanded income including tax credits and noncash transfers is around 66 percent as large. About a third of the loss is due to reduced noncash transfers (column 3), while the remaining two thirds is due to a reduction in tax credits (column 5). However, there is an interesting difference in where in the income distribution these two losses occur. To see this more easily, I plot the semi-elasticities from the three income definitions in Figure 7. The loss due to reduced noncash transfers is highly concentrated between the fourteenth and sixteenth quantiles, around or just above the federal poverty threshold. This is consistent with the nature of eligibility requirements of noncash assistance programs; for instance, typically only individuals with income

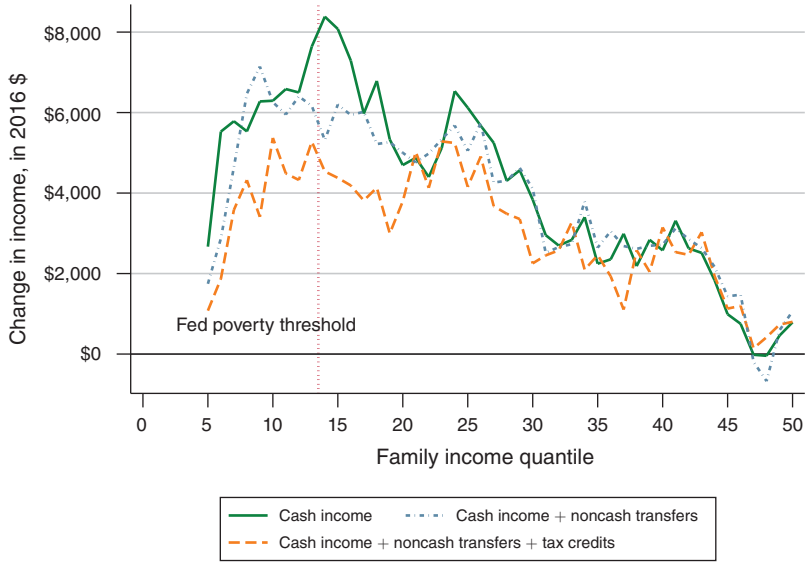


FIGURE 7. MINIMUM WAGE SEMI-ELASTICITIES FOR UNCONDITIONAL QUANTILES OF FAMILY INCOMES, FOR ALTERNATIVE INCOME DEFINITIONS

Notes: The reported estimates are semi-elasticities, showing the increase in income (in 2016 dollars) from a log point increase in the minimum wage. The solid line is for the change in cash income, the dash-dot line augments it by adding noncash transfers (SNAP, NSLP, and housing subsidies) to the income definition, and the dashed line is for the expanded income that includes both noncash transfers and tax credits (EITC and child tax credits). The estimates are from a series of linear probability models that are estimated by regressing an indicator for having family income below cutoffs associated with alternative quantiles (between 5 and 50) on distributed lags of log minimum wage and covariates. The long-run coefficients are calculated from the sum of the contemporaneous up to three-year lagged log minimum wage coefficients. Unconditional quantile partial effects (UQPE) are calculated by dividing the coefficient on log minimum wage by the negative of the family income density at the appropriate quantile. The UQPE estimates are subsequently multiplied by the average cash income in 2016 dollars corresponding to the federal poverty threshold to transform the estimates into changes in income. All specifications include state and year fixed effects, division-by-period fixed effects, state-specific linear trends, state-specific indicator for each Great Recession year, state-level covariates (GDP per capita, EITC supplement, and unemployment rate), and individual demographic controls (quartic in age as well as dummies for race, marital status, family size, number of children, education level, Hispanic status, and gender), and are weighted by March CPS person weights. The vertical dotted line shows the quantile corresponding to the sample averaged federal poverty threshold.

at or lower than 130 percent of the federal poverty threshold are eligible for SNAP. As a result, some individuals with incomes just below the eligibility threshold can lose the benefits due to higher earnings. Figure 7 provides compelling evidence on the importance of these offsets. At the fifteenth quantile, there is a 40 percent offset to the income gain due to loss in public assistance, mostly due to reduced noncash transfers. In contrast, at the tenth or twentieth quantiles, the offset in income gains due to loss in public assistance was only around 15 percent. Moreover, while the loss in noncash transfers is relatively more concentrated around the fifteenth quantile, the loss due to tax credits is more uniform across the income distribution. This likely reflects the fact that the EITC phase-out range is more spread out.[18]

[18] That the income gains at the fifteenth quantile is 40 percent smaller when using an expanded income definition does not necessarily imply that an individual at the fifteenth quantile of the cash incomes experiences a 40 percent

Finally, Table 4 also shows that the offsets in public assistance appear to be larger in the long run. This could potentially reflect a behavioral response in utilization; for example, higher incomes could reduce SNAP take-up rate and not the eligibility rate, and changes in take-up behavior can take longer to adjust than the mechanical effects of changes in eligibility.

While use of the 30 years of data provides us with greater precision, one may wonder whether the patterns are similar in more recent times. In particular, when it comes to tax credits and transfers, EITC expansion of 1992 and welfare reform in 1996 are important, and may interact with the minimum wage policy. Moreover, there is evidence that target efficiency of minimum wages has improved over time (Lundstrom 2017). For these reasons, in Table 5, I show the income elasticities by quantiles for cash income and expanded income including tax-credits and transfers for subsamples beginning in 1990 as well as 1996. The estimates are quite similar to those from the full sample. For the sample beginning in 1996, the long-run elasticities for the tenth and fifteenth quantiles when considering cash incomes are 0.447 and 0.503, respectively, and both are statistically significant at the 5 percent level. For comparison, these elasticities for the full sample are 0.359 and 0.332, respectively (online Appendix Table B.4). When using the expanded income definition, the elasticities for these two quantiles are 0.352 and 0.275 in the 1996–2013 sample, as compared to 0.256 and 0.160 in the full sample. Estimates from the sample beginning in 1990 are similar. As expected, the estimates for the smaller samples are generally less precise, but the magnitudes are similar.

C. *Robustness to Alternative Assumptions about Unobserved Heterogeneity*

To assess how controls for unobserved heterogeneity influence the findings, Figure 8 plots the long-term minimum wage elasticities for shares below various multiples of the federal poverty threshold. (The full set of medium- and long-run coefficients and standard errors are reported in online Appendix Table B.3.) I use eight different regression specifications that range from the classic two-way fixed effects model (panel A, column 1 in Figure 8) to the most saturated specification (panel B, column 4), which includes (a) division-specific year effects, (b) state-specific Great Recession-year dummies, and (c) state-specific linear trends. The six other specifications exhaust all intermediate combinations of controls and provide us with evidence on how the inclusion of various types of time-varying controls affects the estimates.

Overall, there is robust evidence that minimum wage increases reduce the share of individuals with low family incomes. For family income cutoffs between 0.50 and 1.25 (i.e., between 50 and 125 percent of the official poverty threshold), and across the eight specifications, all 32 of the long-run estimates (sum of the

offset due to lost public assistance. Individuals at these quantiles need not be the same (a) before versus after the policy change and (b) when considering different types of incomes. We would need to make a rank-invariance assumption to assure that our findings about quantiles can be interpreted as reflecting impacts on individuals. As an empirical matter, ranks in the conventional cash income and the expanded incomes are highly correlated. For those in the bottom half of the cash income distribution, the correlation coefficient between the two ranks is 0.99. However, this does not tell us whether the policy affects the rank of individuals in the two distributions.

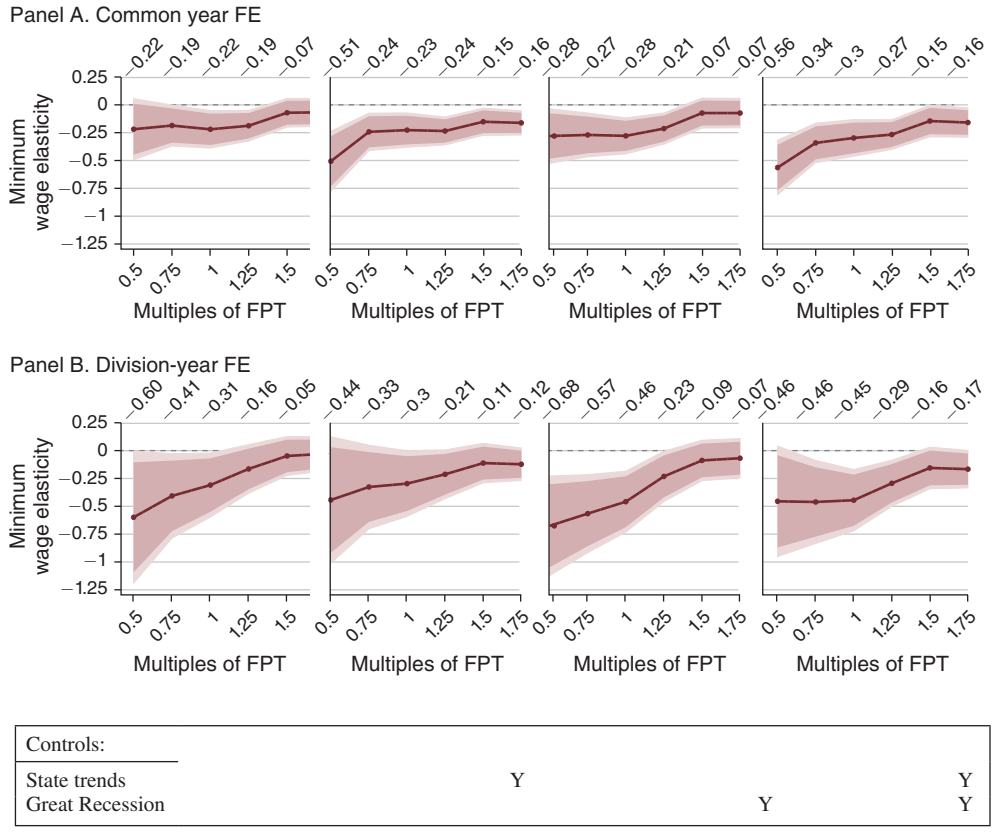
Table 5—Minimum Wage Elasticities for Unconditional Quantiles of Family Incomes for Samples Starting in 1990 and 1996

	1990				1996			
	Pre-TC&T		After-TC&T		Pre-TC&T		After-TC&T	
Family income quantile	Medium-run	Long-run	Medium-run	Long-run	Medium-run	Long-run	Medium-run	Long-run
5	0.299	0.261	0.161	0.233	0.321	−0.094	0.218	0.144
	(0.241)	(0.259)	(0.201)	(0.182)	(0.268)	(0.316)	(0.222)	(0.235)
10	0.536	0.405	0.458	0.308	0.662	0.447	0.570	0.352
	(0.129)	(0.135)	(0.111)	(0.118)	(0.152)	(0.189)	(0.141)	(0.159)
15	0.364	0.407	0.206	0.216	0.461	0.503	0.227	0.275
	(0.121)	(0.117)	(0.092)	(0.101)	(0.139)	(0.169)	(0.110)	(0.132)
20	0.034	0.250	0.038	0.171	−0.028	0.236	−0.006	0.199
	(0.111)	(0.088)	(0.084)	(0.075)	(0.102)	(0.115)	(0.083)	(0.090)
25	0.035	0.200	−0.001	0.142	−0.018	0.198	−0.004	0.165
	(0.081)	(0.083)	(0.068)	(0.067)	(0.079)	(0.110)	(0.067)	(0.092)
30	−0.023	0.131	−0.046	0.087	−0.083	0.113	−0.097	0.112
	(0.086)	(0.066)	(0.070)	(0.061)	(0.065)	(0.096)	(0.058)	(0.081)
Observations	3,836,209	3,836,209	3,836,209	3,836,209	3,021,983	3,021,983	3,021,983	3,021,983

Notes: The reported estimates are minimum wage elasticities for unconditional quantiles of equivalized family income using alternative income definitions. The estimates are from linear probability models that regress an indicator for having a family income below cutoff associated with a quantile (between 5 and 30) on distributed lags of log minimum wage and covariates. The first four columns start the estimation samples from 1990 and the last from 1996. Columns 1, 2, 5, and 6 only consider cash income; and columns 3, 4, 7, and 8 include noncash transfers (SNAP, NLSP, and housing subsidies) and tax credits (EITC and child tax credits). Medium- and long-run unconditional quantile partial effects (UQPE) for family incomes are calculated by summing of the contemporaneous up to two- and three-year lagged log minimum wage coefficients, respectively, and dividing by the negative of the family income density at the appropriate quantile. The UQPE estimates are subsequently divided by the family income cutoff for the quantile to transform the estimates into elasticities. The regression specification includes state fixed effects, division-specific year effects, state-specific indicator for each Great Recession year, state-specific linear trends, and state-level covariates (GDP per capita, EITC supplement, and unemployment rate), and individual demographic controls (quartic in age as well as dummies for race, marital status, family size, number of children, education level, Hispanic status, and gender), and are weighted by March CPS person weights. State-cluster-robust standard errors are in parentheses.

contemporaneous plus one, two, and three year lags in minimum wage) are negative in sign, and 29 are statistically significant at least at the 10 percent level. This includes all of the poverty rate elasticities, which are statistically significant at the 5 percent level. While all of the estimates are sizable, the more saturated specifications tend to produce estimates with somewhat larger magnitudes. For example, the long-run poverty elasticity estimate of −0.220 from the two-way fixed effects estimate (panel A, column 1 of Figure 8) is the smallest in magnitude among all eight specifications. In contrast, the estimate with state-trends and division-period FE (panel B, column 3) is the largest in magnitude (−0.459); the preferred specification used in the paper (panel B, column 4) produces an elasticity of −0.446.

Figure 9 presents analogous evidence on equivalized family income elasticities for quantiles, and shows robust evidence across all 8 specifications that minimum wages lead to at least moderate increases in incomes for the bottom 20 percent of the equivalized family income distribution. (online Appendix Table B.4 reports the full set of estimates and standard errors.) Of the 32 estimates, all are positive in sign, and 24 are statistically significant at least at the 10 percent level. The 16 estimates for the tenth and the fifteenth quantiles range between 0.152 and 0.430, and all are statistically significant at least at the 10 percent level. As before, the two-way fixed



Controls:			
State trends	Y		Y
Great Recession		Y	Y

FIGURE 8. MINIMUM WAGE ELASTICITIES FOR SHARE UNDER ALTERNATIVE FAMILY INCOMES IN MULTIPLES OF FEDERAL POVERTY THRESHOLD: ALTERNATIVE SPECIFICATIONS

Notes: The reported estimates are long-run minimum wage elasticities for share with family income under various multiples of the federal poverty threshold. These estimates are from a series of linear probability models that are estimated by regressing an indicator for having family income below cutoffs between 0.5 and 1.75 times the federal poverty threshold (in increments of 0.25) on distributed lags of log minimum wage and covariates. The long-run (3+ year) elasticities are calculated by dividing the respective coefficients (or sums of coefficients) by the sample proportion under the cutoff. All specifications include state and year fixed effects, state-level covariates (GDP per capita, EITC supplement, and unemployment rate), and individual demographic controls (quartic in age as well as dummies for race, marital status, family size, number of children, education level, Hispanic status, and gender), and are weighted by March CPS person weights. Additional controls used for each specification are as indicated above. The dark shaded area represents 90 percent, the light shaded area 95 percent state-cluster-robust confidence intervals.

effects specification (panel A, column 1) provides the smallest estimated magnitudes, and the inclusion of division-specific year effects and state-specific Great-Recession controls tend to increase the size of the estimates. These patterns are as expected, since the elasticities for the family income quantiles are simply rescaled semi-elasticities for the proportions below alternative income-to-needs cutoffs, and the rescaling factors are common across specifications.[19]

[19] Online Appendix Table B.5 additionally presents family income elasticities using the broader income definition including tax credits and noncash transfers. Averaging across all specifications, the long-run estimates for these quantiles using the expanded definition are about 60 percent as large as when using pretax cash income, which is similar to the 66 percent estimate I find in the preferred specification above.

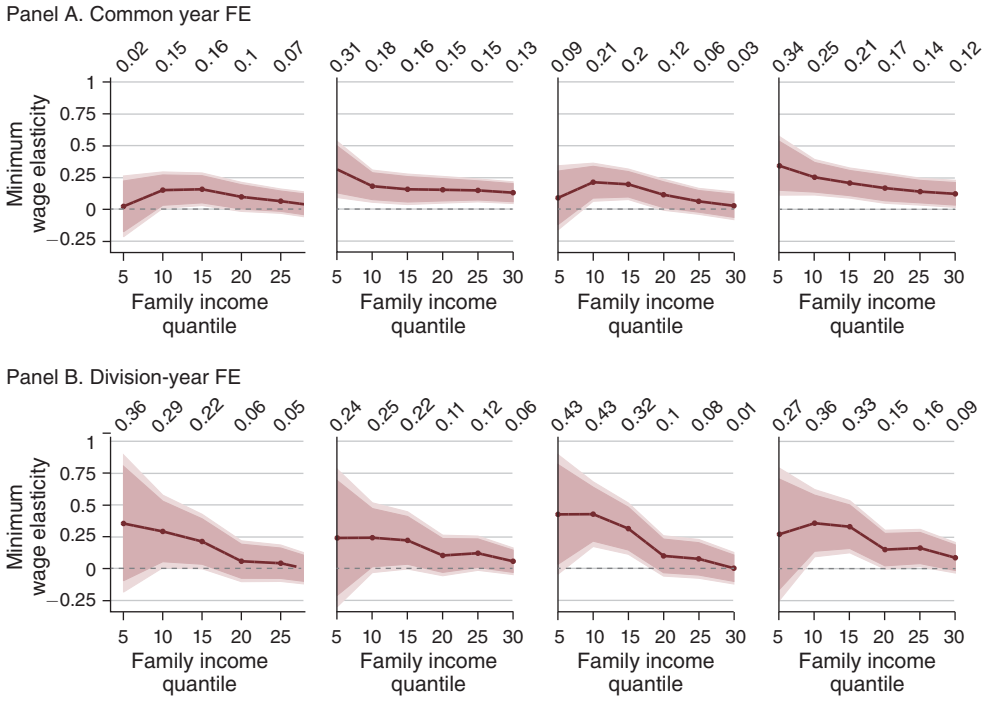


FIGURE 9. MINIMUM WAGE ELASTICITIES FOR UNCONDITIONAL FAMILY INCOME QUANTILES: ALTERNATIVE SPECIFICATIONS

Notes: The reported estimates are long-run minimum wage elasticities by unconditional quantiles of family income. The estimates are from a series of linear probability models that are estimated by regressing an indicator for having a family income below cutoffs associated with alternative quantiles (between 5 and 50) on distributed lags of log minimum wage and covariates. The long-run coefficients are calculated as the sum of the contemporaneous up to three-year lagged log minimum wage coefficients. Unconditional quantile partial effects (UQPE) are calculated by dividing the coefficient on log minimum wage by the negative of the family income density at the appropriate quantile. The UQPE estimates are subsequently divided by the family income cutoff for the quantile to transform the estimates into elasticities. All specifications include state and year fixed effects, state-level covariates (GDP per capita, EITC supplement, and unemployment rate), and individual demographic controls (quartic in age as well as dummies for race, marital status, family size, number of children, education level, Hispanic status, and gender), and are weighted by March CPS person weights. Additional controls used for each specification are as indicated above. The dark shaded area represents 90 percent, the light shaded area 95 percent state-cluster-robust confidence intervals.

The range of estimates raises the issue of model selection. There is an a priori case for using richer controls that better account for time-varying heterogeneity across states. Allowing for regional shocks and state-specific trends receives strong support in existing work. For example, Allegretto et al. (2017) shows that the inclusion of these controls mitigates contamination from preexisting trends when it comes to estimating the effect of minimum wages on teen employment. They also provide evidence that synthetic control methods tend to put substantially more weight on nearby states in constructing a control group, providing additional validity to the

intuition that nearby states are better controls. One argument against using more saturated models is the loss of precision: that they may lack the statistical power to detect an effect.[20] However, while the standard errors for the most saturated specification 8 are larger, the 95 percent confidence intervals for that specification does not contain the point estimates using specification 1 for 75 and 100 percent of the federal poverty threshold. In other words, there is strong indication that the differences in estimates across the two specifications are not driven primarily by the imprecision of the more saturated model.

Beyond this, I consider two types of falsification exercises for model selection. First, I use the leading values as a falsification test, analogous to tests used in Dube, Lester, and Reich (2010); Allegretto, Dube, and Reich (2011); and Allegretto et al. (2017). To be clear, the medium- and long-run estimates do not include the leading coefficient, so they net out any level differences in the outcome just prior to the minimum wage increase. However, since a nonzero leading coefficient provides an indication of the violation of parallel trends, such a test can help us assess the validity of particular specifications. The results are shown in Table 6. They indicate that specifications 1, 2, 5, and 6 all produce positive leads for the proportion below one-half the poverty line, and these are statistically significant at the 5 percent level. Specification 1 also produces positive leads for 1.75 times poverty threshold. In contrast, specifications 3, 4, and 7 tend to produce negative leads for 0.75 times the poverty threshold. Considering the full range of cutoffs, the most saturated specification 8 usually performs the best when it comes to the leading values falsification test: it is the only specification where none of the leading coefficients are statistically significant; moreover, this is mostly driven by the small size of the coefficients and not from a greater imprecision. For example, the absolute value of the mean of all leads in specification 8 is 0.081, while the other specifications range between 0.083 and 0.194. Overall, this suggests the most saturated specification may be able to guard against preexisting trends throughout the income distribution better than less saturated ones.

In addition, we should not expect minimum wages to affect the proportion earning under 3, 3.5, or 4 times the poverty threshold, which roughly corresponds to the fiftieth, fifty-seventh, and sixty-fourth percentiles of the family income distribution in the national sample. Therefore, reliable specifications should produce estimates for these cutoffs that are small or close to zero. Panel B of Table 6 reports the long-run elasticities for shares below these middle and upper income cutoffs. In general, none of the specifications produce statistically significant long-run effects, which is reassuring. However, the most saturated specification performs particularly well: averaged over all the cutoffs, the mean absolute value of the estimates is 0.14 for specification 8, as compared to a range between 0.014 and 0.070 for the other specifications. Overall, the most saturated specification 8 performs very well on the falsification exercises, while the results from the intermediate specifications vary.[21]

[20] A second rationale for excluding covariates is that some of them are "bad controls" in the sense of blocking a causal pathway between the treatment and the outcome. As I discussed above, the state-specific linear trends may constitute a problem if there are delayed effects of the policy, but I address this issue by including three annual lags in minimum wage.

[21] Additional results (not reported in the tables) also suggest that the saturated specification 8 is more robust to exclusion of the state-level controls. Excluding these controls produces a long-run poverty elasticity (standard

TABLE 6—LEADING ESTIMATES AND UPPER INCOME FALSIFICATION TESTS: SHARE OF INDIVIDUALS WITH FAMILY INCOME BELOW MULTIPLES OF FEDERAL POVERTY THRESHOLD

Family income cutoff	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A. 1-year leading estimates								
0.50	0.496	0.367	−0.093	0.003	0.579	0.423	0.011	0.102
	(0.151)	(0.115)	(0.137)	(0.147)	(0.149)	(0.111)	(0.144)	(0.151)
0.75	0.041	−0.011	−0.298	−0.262	0.125	0.054	−0.199	−0.176
	(0.098)	(0.116)	(0.090)	(0.107)	(0.098)	(0.117)	(0.100)	(0.115)
1.00	0.072	0.033	−0.134	−0.142	0.128	0.068	−0.055	−0.080
	(0.084)	(0.097)	(0.072)	(0.085)	(0.089)	(0.102)	(0.080)	(0.091)
1.25	0.070	0.009	−0.095	−0.128	0.084	−0.020	−0.060	−0.121
	(0.085)	(0.090)	(0.080)	(0.087)	(0.093)	(0.094)	(0.093)	(0.095)
1.50	0.097	0.022	−0.082	−0.122	0.092	−0.022	−0.068	−0.125
	(0.081)	(0.081)	(0.069)	(0.070)	(0.088)	(0.087)	(0.080)	(0.078)
1.75	0.161	0.077	−0.022	−0.081	0.154	0.034	−0.022	−0.091
	(0.074)	(0.073)	(0.061)	(0.059)	(0.076)	(0.074)	(0.070)	(0.065)
Panel B. Long-run (3+ year lagged) estimates of upper income								
3.00	0.055	−0.020	0.071	0.038	0.043	−0.035	0.017	−0.004
	(0.038)	(0.038)	(0.051)	(0.053)	(0.039)	(0.039)	(0.063)	(0.064)
3.50	0.055	0.007	0.069	0.052	0.032	−0.020	0.024	0.016
	(0.038)	(0.037)	(0.052)	(0.050)	(0.041)	(0.041)	(0.066)	(0.064)
4.00	0.022	−0.016	0.069	0.048	0.006	−0.031	0.025	0.024
	(0.031)	(0.030)	(0.043)	(0.039)	(0.032)	(0.035)	(0.056)	(0.052)
Observations	4,662,781	4,662,781	4,662,781	4,662,781	4,662,781	4,662,781	4,662,781	4,662,781
Div-period FE			Y	Y			Y	Y
State trends		Y		Y		Y		Y
GR-state dummies					Y	Y	Y	Y

Notes: The reported estimates are either leading minimum wage elasticities (panel A) or long-run minimum wage elasticities (panel B) for share with family income under various multiples of the federal poverty threshold. These estimates are from linear probability models that regress an indicator for having family income below multiples (between 0.50 and 1.75) of the federal poverty threshold on distributed lags of log minimum wage and covariates. The leading elasticities are calculated by dividing the one-year leading minimum wage coefficients by the sample proportion under the family income cutoff. The long-run elasticities are calculated by dividing the sum of the contemporaneous up to three-year lagged log minimum wage coefficients by the sample proportion under the family income cutoff. All specifications include state and year fixed effects, state-level covariates (GDP per capita, EITC supplement, and unemployment rate), and individual demographic controls (quartic in age as well as dummies for race, marital status, family size, number of children, education level, Hispanic status, and gender), and are weighted by March CPS person weights. Additional division-by-year fixed effects, state-specific linear trends, and state-specific indicators for each Great Recession year are indicated in the table. State-cluster-robust standard errors are in parentheses.

Therefore, based both on a priori grounds as including the richest set of controls for time-varying heterogeneity, as well as its performance on the falsification tests, I use the saturated model to be the preferred specification. At the same time, I recognize that reasonable observers may disagree on which specification is the most reliable, and may place somewhat different weights on the evidence associated with

error) of −0.397 (0.157) for the saturated model, close to the original elasticity of −0.446 (0.137). In contrast, the two-way FE estimate is substantially smaller in magnitude when these controls are left out: −0.062 (0.087) instead of −0.220 (0.084). To the extent unobservables are correlated with observable controls, these results suggests that more saturated specification is better able to account for them.

each specification. For this reason, in this paper, I report the range of key estimates across all eight specifications, which allows readers to put more weight on various specifications based on their own priors.

While the specifications in Table 2 include leads and lags in minimum wages, in much of the literature, estimates come from static models with just a contemporaneous minimum wage. To assess the difference this might make, in Table 7 I report estimates from a static model for the poverty rate. In every set of controls, the minimum wage elasticity from the static model is smaller in magnitude than either the medium- or long-run elasticities from the dynamic model. On average, the contemporaneous elasticity from the static model is less than half the size of the long-run estimate. For example, the static two-way fixed effects specification (column 1) that is most commonly used in the literature produces an estimate of -0.074, not statistically distinguishable from zero. It also has a positive leading coefficient of 0.072, although not statistically significant. This subjects the static elasticity to a positive bias. Indeed, if we estimate a regression with a leading minimum wage in addition to the contemporaneous minimum wage, the magnitude of the minimum wage elasticity doubles to -0.144 with a standard error of 0.064, rendering it statistically significant (results not shown in the table). The addition of lags raise the magnitude further, with the long run estimate being -0.220. To summarize, the static two-way FE estimate appears to be biased towards zero both due to preexisting trends, and due to missing some of the lagged effects of the policy.

What are the relative contributions of: (1) using a dynamic model, (2) including division-year effects, (3) including Great Recession-state effects, and (4) including state-specific trends in explaining the gap between the the static, two-way fixed effects estimate (-0.074) and the long-run estimate from the most saturated model (-0.446)? Any decomposition of these four factors will depend on the order in which they are implemented. There are exactly $4! = 24$ different orderings for incrementally changing each of these four factors going from the static version of specification 1 to the dynamic version of specification 8, and each of these orderings provides a different decomposition. Averaging across these 24 orderings, the average incremental contribution of these four factors (in the same order as above) are: (1) 44 percent, (2) 41 percent, (3) 18 percent, and (4) -3 percent. Use of the dynamic model (both netting out the lead, and considering long-run effects) is the biggest source of difference, while controls for regional and business cycle heterogeneity are important as well. Use of state-specific trends appears to be the least important of these, at least when considering the poverty rate elasticity.

III. Discussion

There is robust evidence that minimum wage increases over the past 30 years have boosted pretax-and-transfer incomes at the bottom of the income distribution. Across specifications with alternative assumptions about the appropriate counterfactual, I find that minimum wages reduce the shares with family incomes below 50, 75, 100, and 125 percent of the federal poverty threshold. The long-run poverty rate elasticities range between -0.220 and -0.459, and all of these are statistically significant. I also find that a higher minimum wage improves the family income

Table 7—Comparing Minimum Wage Elasticities for the Poverty Rate: Static and Dynamic Specifications

Family income cutoff	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Contemporaneous effect	−0.074	−0.069	−0.250	−0.196	−0.075	−0.082	−0.293	−0.232
	(0.062)	(0.058)	(0.095)	(0.094)	(0.062)	(0.066)	(0.098)	(0.107)
Dynamic specification:								
Lead	0.072	0.033	−0.134	−0.142	0.128	0.068	−0.055	−0.080
	(0.084)	(0.097)	(0.072)	(0.085)	(0.089)	(0.102)	(0.080)	(0.091)
Medium-run	−0.161	−0.145	−0.302	−0.274	−0.234	−0.178	−0.475	−0.399
	(0.101)	(0.099)	(0.129)	(0.124)	(0.092)	(0.089)	(0.147)	(0.126)
Long-run	−0.220	−0.227	−0.310	−0.296	−0.279	−0.298	−0.459	−0.446
	(0.084)	(0.077)	(0.143)	(0.147)	(0.081)	(0.082)	(0.136)	(0.137)
Observations	4,662,781	4,662,781	4,662,781	4,662,781	4,662,781	4,662,781	4,662,781	4,662,781
Div-period FE			Y	Y			Y	Y
State trends		Y		Y		Y		Y
GR-state dummies					Y	Y	Y	Y

Notes: The reported estimates are minimum wage elasticities for share with family income under the federal poverty threshold. The estimates are from linear probability models where the outcome is an indicator for having a family income below the federal poverty threshold. The static specification includes only the contemporaneous log minimum wage. The elasticity is calculated by dividing the log minimum wage coefficient by the sample proportion under the family income cutoff. The dynamic specification includes a distributed lags window of one lead and up to three lags of log minimum wage. The leading, medium-run, and long-run elasticities are calculated from the leading log minimum wage coefficient, the sum of the contemporaneous up to two-year lagged log minimum wage coefficients, and the sum of the contemporaneous up to three-year lagged log minimum wage coefficients, respectively, and then divided by the sample proportion under the cutoff. All specifications include state and year fixed effects, state-level covariates (GDP per capita, EITC supplement, and unemployment rate), and individual demographic controls (quartic in age as well as dummies for race, marital status, family size, number of children, education level, Hispanic status, and gender), and are weighted by March CPS person weights. Additional division-by-year fixed effects, state-specific linear trends and state-specific indicators for each Great Recession year are indicated in the table. State-cluster-robust standard errors are in parentheses.

distribution in the sense of first-order stochastic dominance. The clearest increases are for the tenth and fifteenth quantiles, with elasticities ranging between 0.152 and 0.430 across alternative models. While these results hold across a range of specifications, the preferred specification with a rich set of controls suggests an elasticity of around 0.359 for the tenth quantile of family incomes. Overall, around a third of the income increases resulting from minimum wage changes accrue to those in the bottom 15 percent of the family income distribution (roughly corresponding to those with incomes under the federal poverty level), while around 80 percent accrue to those in the bottom third of the family income distribution. At the same time, I also find evidence that there are sizable offsets in public assistance in response to minimum wage gains. For those in the bottom fifth of the distribution, around 30 percent of the income gains are offset through reductions in noncash transfers as well as tax credits. These offsets seem to be the largest near the federal poverty threshold—consistent with losses in eligibility in programs such as SNAP.

What do these estimates imply about the likely impact on poverty from an increase in the federal minimum wage from the current $7.25/hour to, say, $12/hour? Taking into account the state minimum wages as of January 2017 suggests that this policy would raise the effective minimum wage—i.e., the maximum of federal or state standard—by 41 percent (or 0.34 log points). Using the long-run minimum wage elasticity for the poverty rate of −0.446 from the preferred

specification and a 12.7 percent poverty rate among the non-elderly population in 2016 suggests a 1.9 percentage point reduction in the poverty rate from this minimum wage increase. Given the roughly 270 million non-elderly Americans in 2016, this translates into 6.16 million fewer individuals living in poverty. We can also expect the same minimum wage increase to raise family incomes by 12.2 percent at the tenth quantile of the equivalized family income distribution in the long run. For the tenth quantile, this translates into an annual income increase of $2,140; after accounting for the offset due to reduced tax credits and transfers, this amounts to an increase of $1,826.[22] Therefore, the increase in the federal minimum wage can play an important role in reducing poverty and raising family incomes at the bottom. To put this in context, Hoynes, Page, and Stevens (2006) estimates that EITC reduces non-elderly poverty rate by around 1.7 percentage points. They also find that (means tested and non-means tested) cash transfers reduce it by as much as 3.8 percentage points, while noncash transfers (other than Medicaid) reduce it by around 0.9 percentage points.[23] In other words, a substantial increase in the minimum wage would likely have an impact on the non-elderly poverty rate comparable to means tested public assistance programs.

The income estimates in this paper do not factor in changes in prices due to minimum wage increase. However, the expected increase in the overall price level from this policy is quite small when compared to the income gains for those at the bottom of the income distribution. For example, in a simulation study using input-output tables, MaCurdy (2015) calculated that the bottom quintile faced a 0.5 percent increase in prices from the 21 percent increase in minimum wage in 1996—an elasticity of around 0.02. The preferred long-run elasticity for the tenth quantile after public offsets is around 0.256, or more than an order of magnitude larger. Therefore, netting out the likely price increase would not substantially affect the estimates of the real income gains for the bottom quintile of the income distribution. However, it does mean that some (possibly much) of these income gains at the bottom are likely borne by middle and upper income consumers through small increases in prices.

The estimates in this paper tend to suggest bigger increases in bottom incomes in response to minimum wage increases than some simulation based studies, such as Sabia and Burkhauser (2010), CBO (2014), or MaCurdy (2015). There are a number of problems in using the cross-sectional relationship between reported wages and family incomes to simulate how the gains from a minimum wage increase will be distributed. Most obviously, we would need to make assumptions about how behavior changes: this includes employer hiring and firing behavior, as well as workers' job search behavior and labor force participation, which could vary by family income and other characteristics. In addition, simulations such as these face a number of challenges that tend to suggest a weaker link between low wages and

[22] If we take the range of estimates from all specifications, the proposed minimum wage changes can be expected to reduce the poverty rate among the non-elderly population by 0.950 and 1.98 percentage points, hence, reducing the number of non-elderly individuals living in poverty by somewhere between 3.03 and 6.33 million. For the tenth quantile of family incomes, this translates to an annual income increase ranging between 5.2 and 14.5 percent, or between $905 and $2,937. After accounting for offsets due to lost public assistance, the income increases would range between $657 and $2,790.

[23] The estimates for means tested and non-means tested are reported separately in Hoynes, Page, and Stevens (2006); therefore, the sum of the effects (3.8 percentage points) may overstate the net impact.

low family income than is truly the case. A key concern is measurement error in both wages and other sources of incomes (which includes wage and salary incomes of other family members). It is a straightforward point that measurement error in reported wages leads to an attenuation in the measured relationship between workers' wages and family incomes.[24] As a result, simulating wage changes for those earning around the minimum wage will typically suggest smaller effects on poverty and smaller income increases at the bottom quantiles than would occur in reality. This is because (1) some of the individuals with high reported wages in low-income families are actually low-wage earners, and (2) some of the low-wage earners reporting high levels of other sources of income (including spousal wage and salary income) in reality are in poorer families. A related practical issue that arises from this is the treatment of sub-minimum wage workers. For example, in their simulations of raising the minimum wage from \$5.70 to \$7.25, Sabia and Burkhauser (2010) assumes that all those with reported hourly earnings below \$5.55 will receive no wage increases because they are in the "uncovered sector." Moreover, they assume that no one above \$7.25 will get a raise. These particular assumptions seem implausible due to both measurement error issues, as well as the well-known "lighthouse effect" phenomenon whereby even uncovered sector workers' wages are affected by minimum wages (Card and Krueger 1995; Boeri, Garibaldi, and Ribeiro 2011). Moreover, as Autor, Manning, and Smith (2016) shows, effects of the minimum wage extend up to the twentieth percentile of the wage distribution, which would be unlikely absent some spillovers.[25] Therefore, results from simulation studies may not provide reliable guidance in assessing the impact of minimum wages on bottom incomes, making it critical for us to consider actual evidence from past minimum wage changes when analyzing policy proposals.

While minimum wages tend to raise family incomes at the bottom, they also tend to substitute earnings for public assistance. Were we to assess public policies strictly based on their efficacy in reducing post-tax and transfer poverty, the offsets through reduced tax credits and noncash transfers further suggests a lower effectiveness of minimum wages in raising post-tax-and-transfer incomes. Policies like cash

[24] Consider the relationship between own-wage income, W, and family income $F = W + I$, where I represents other incomes (possibly others' wages). The linear approximation to the true relationship is represented by the population regression $F = \beta W + u$. Note that $\beta = \text{cov}(W + I, W)/V(W) = 1 + \left(\sigma_{WI}/\sigma_W^2\right)$. So if wages are at all positively correlated with other sources of family incomes, I, as is likely, then $\beta > 1$.

Now consider the case where W is measured with error, so that $\tilde{W} = W + e$, and $\tilde{F} = W + I + e$ are the observed wage and family income. This is slightly different from the textbook classical measurement error case because the measurement error, e, affects both the independent and dependent variables. Substituting the reported values into the true regression equation produces $\tilde{F} - e = \beta\left(\tilde{W} - e\right)$. Rearranging, we have $\tilde{F} = \beta\tilde{W} + (1 - \beta)e = \beta\tilde{W} + \tilde{u}$.

Note that $\tilde{\beta} = \text{cov}\left(\tilde{F}, \tilde{W}\right)/V\left(\tilde{W}\right)$ is the estimate from a population regression of $\tilde{F}$ on $\tilde{W}$. Substituting $\tilde{F} = \beta\tilde{W} + (1 - \beta)e$ into the expression for $\tilde{\beta}$ we have $\tilde{\beta} = \beta + (1 - \beta)\left(\sigma_e^2/\left(\sigma_w^2 + \sigma_e^2\right)\right)$, which will be attenuated toward zero if $\beta > 1$, which is true if wages are at all positively correlated with other sources of family incomes. The attenuation will also be proportionate to the share of wage variance that is due to error.

[25] Autor, Manning, and Smith (2016) also highlights how measurement error in wages and wage spillovers have similar implications about the effects of the minimum wage on the observed wage distribution. This is an interesting point that affects the interpretation of the effects on higher wage quantiles. But for our purposes here, regardless of the interpretation of these effects as true spillovers or measurement error spillovers, ignoring them will tend to downward bias the predicted effects of minimum wages on poverty in simulation studies.

transfers, SNAP, and tax credits are better targeted to raise incomes for those at the very bottom of the income distribution.

At the same time, there are positive aspects of a policy that reduces public assistance via increased earnings. First, a reduction in public benefits like SNAP can be efficiency enhancing, since in principle these programs are funded using taxation that can have deadweight losses. Hendren (2014) finds that lower pretax incomes at the bottom have an efficiency cost due to such transfers: $1 of surplus accruing to the bottom quintile of the income distribution can be transformed into a $1.15 surplus for everyone. A corollary to that argument is that by reducing such transfers, minimum wages can have an efficiency-enhancing attribute.[26] Alternatively, if there are deadweight losses from the minimum wage policy, the reduced public assistance will tend to mitigate such distortions. Additionally, there is increasing recognition that individuals may not see labor earnings and tax credits/transfers as perfect substitutes, and may prefer to receive a higher compensation for their work than what is often perceived as a government transfer. Relatedly, voters may prefer "pre-distributive" policies like minimum wages over "redistributive" ones using tax and transfers, even when they are concerned with inequality—in part because of low levels of trust in government (Kuziemko et al. 2015). All of these factors add to the cost of pretax inequality and raises the benefits of a policy like minimum wage that affects the pretax-and-transfer income distribution.

Overall, motivations of voters and policymakers in raising minimum wage policies tend to go beyond reducing poverty. For example, using data from the 2016 YouGov Common Content Pre-Election Survey of voters ($N = 1{,}000$), I find that voters' preferred level of federal minimum wage barely changes (from $10.45 to $10.18) if there were also to be sufficient tax credits to eliminate poverty among working families. (See online Appendix Figure B.4.) This is consistent with the observation that the popular support for minimum wages is in part fueled by a desire to raise earnings of low- and moderate-income families more broadly, and by concerns of fairness that seek to limit the extent of wage inequality (Green and Harrison 2009), or employers' exercise of market power (Fehr and Fischbacher 2004; Kahneman, Knetsch, and Thaler 1986). The findings from this paper suggest that attaining these goals is also consistent with at least a moderate increase in family incomes, and a reduced reliance on public assistance for those at the bottom of the income distribution.

REFERENCES

Addison, John T., and McKinley L. Blackburn. 1999. "Minimum Wages and Poverty." *Industrial and Labor Relations Review* 52 (3): 393–409.

Allegretto, Sylvia A., Arindrajit Dube, and Michael Reich. 2011. "Do Minimum Wages Really Reduce Teen Employment? Accounting for Heterogeneity and Selectivity in State Panel Data." *Industrial Relations* 50 (2): 205–40.

Allegretto, Sylvia, Arindrajit Dube, Michael Reich, and Ben Zipperer. 2013. "Credible Research Designs for Minimum Wage Studies." IZA Discussion Paper 7638.

[26] Moreover, the public savings afforded by a higher minimum wage could be used to expand program generosity, further enhancing income at the bottom.

Allegretto, Sylvia, Arindrajit Dube, Michael Reich, and Ben Zipperer. 2017. "Credible Research Designs for Minimum Wage Studies: A Response to Neumark, Salas, and Wascher." *ILR Review* 70 (3): 559–92.

Autor, David H., Alan Manning, and Christopher L. Smith. 2016. "The Contribution of the Minimum Wage to US Wage Inequality over Three Decades: A Reassessment." *American Economic Journal: Applied Economics* 8 (1): 58–99.

Belman, Dale, and Paul J. Wolfson. 2014. *What Does the Minimum Wage Do?* W.E. Upjohn Institute.

Boeri, Tito, Pietro Garibaldi, and Marta Ribeiro. 2011. "The Lighthouse Effect and Beyond." *Review of Income and Wealth* 57 (S1): S54–78.

Burkhauser, Richard V., and Joseph J. Sabia. 2007. "The Effectiveness of Minimum-Wage Increases in Reducing Poverty: Past, Present, and Future." *Contemporary Economic Policy* 25 (2): 262–81.

Card, David, and Alan B. Krueger. 1995. *Myth and Measurement: The New Economics of the Minimum Wage.* New Jersey: Princeton University Press.

Chernozhukov, Victor, Iván Fernández-Val, and Blaise Melly. 2013. "Inference on Counterfactual Distributions." *Econometrica* 81 (6): 2205–68.

Congressional Budget Office (CBO). 2014. *The Effects of a Minimum-Wage Increase on Employment and Family Income.* Congressional Budget Office, February.

DeFina, Robert H. 2008. "The Impact of State Minimum Wages on Child Poverty in Female-Headed Families." *Journal of Poverty* 12 (2): 155–74.

Dube, Arindrajit. 2019. "Minimum Wages and the Distribution of Family Incomes: Dataset." *American Economic Journal: Applied Economics*. https://doi.org/10.1257/app.20170085.

Dube, Arindrajit, T. William Lester, and Michael Reich. 2010. "Minimum Wage Effects across State Borders: Estimates Using Contiguous Counties." *Review of Economics and Statistics* 92 (4): 945–64.

Fehr, Ernst, and Urs Fischbacher. 2004. "Third-Party Punishment and Social Norms." Evolution and Human Behavior 25 (2): 63–87.

Firpo, Sergio, Nicole M. Fortin, and Thomas Lemieux. 2009. "Unconditional Quantile Regressions." *Econometrica* 77 (3): 953–73.

Fortin, Nicole, Thomas Lemieux, and Sergio Firpo. 2011. "Decomposition Methods in Economics." In *Handbook of Labor Economics*, Vol. 4A, edited by Orley Ashenfelter and David Card, 1–102. Amsterdam: North Holland.

Fox, Liana, Irwin Garfinkel, Neeraj Kaushal, Jane Waldfogel, and Christopher Wimer. 2014. "Waging War on Poverty: Historical Trends in Poverty Using the Supplemental Poverty Measure." NBER Working Paper 19789.

Gramlich, Edward M. 1976. "Impact of Minimum Wages on Other Wages, Employment, and Family Incomes." *Brookings Papers on Economic Activity* 6 (2): 409–61.

Green, David A., and Kathryn Harrison. 2009. "Minimum Wage Setting and Standards of Fairness." Institute for Fiscal Studies Working Paper W10/09.

Gundersen, Craig, and James P. Ziliak. 2004. "Poverty and Macroeconomic Performance across Space, Race, and Family Structure." *Demography* 41 (1): 61–86.

Havnes, Tarjei, and Magne Mogstad. 2015. "Is Universal Child Care Leveling the Playing Field?" *Journal of Public Economics* 127: 100–114.

Hendren, Nathaniel. 2014. "The Inequality Deflator: Interpersonal Comparisons without a Social Welfare Function." NBER Working Paper 20351.

Hoynes, Hilary W., and Ankur J. Patel. 2016. "Effective Policy for Reducing Poverty and Inequality? The Earned Income Tax Credit and the Distribution of Income." https://gspp.berkeley.edu/assets/uploads/research/pdf/Hoynes-Patel-EITC-Income-11-30-16.pdf.

Hoynes, Hilary W., Marianne E. Page, and Ann Huff Stevens. 2006. "Poverty in America: Trends and Explanations." *Journal of Economic Perspectives* 20 (1): 47–68.

Kahneman, Daniel, Jack L. Knetsch, and Richard Thaler. 1986. "Fairness as a Constraint on Profit Seeking: Entitlements in the Market." *American Economic Review* 76 (4): 728–41.

Koenker, Roger, and Gilbert Bassett, Jr. 1978. "Regression Quantiles." *Econometrica* 46 (1): 33–50.

Kuziemko, Ilyana, Michael I. Norton, Emmanuel Saez, and Stefanie Stantcheva. 2015. "How Elastic Are Preferences for Redistribution? Evidence from Randomized Survey Experiments." *American Economic Review* 105 (4): 1478–1508.

Lee, David S. 1999. "Wage Inequality in the United States during the 1980s: Rising Dispersion or Falling Minimum Wage?" *Quarterly Journal of Economics* 114 (3): 977–1023.

Lehrer, Steven F., R. Vincent Pohl, and Kyungchul Song. 2016. "Targeting Policies: Multiple Testing and Distributional Treatment Effects." NBER Working Paper 22950.

Lundstrom, Samuel M. 2017. "When Is a Good Time to Raise the Minimum Wage?" *Contemporary Economic Policy* 35 (1): 29–52.

MaCurdy, Thomas. 2015. "How Effective Is the Minimum Wage at Supporting the Poor?" *Journal of Political Economy* 123 (2): 497–545.

Meer, Jonathan, and Jeremy West. 2016. "Effects of the Minimum Wage on Employment Dynamics." *Journal of Human Resources* 51 (2): 500-522.

Morgan, David R., and Kenneth Kickham. 2001. "Children in Poverty: Do State Policies Matter?" *Social Science Quarterly* 82 (3): 478–93.

Neumark, David. 2016. "Policy Levers to Increase Jobs and Increase Income from Work after the Great Recession." *IZA Journal of Labor Policy* 5 (6): 1–38.

Neumark, David, J.M. Ian Salas, and William Wascher. 2014. "Revisiting the Minimum Wage-Employment Debate: Throwing Out the Baby with the Bathwater?" *ILR Review* 67 (3S): 608–48.

Page, Marianne E., Joanne Spetz, and Jane Millar. 2005. "Does the Minimum Wage Affect Welfare Caseloads?" *Journal of Policy Analysis and Management* 24 (2): 273–95.

Reich, Michael, and Rachel West. 2015. "The Effects of Minimum Wages on Food Stamp Enrollment and Expenditures." *Industrial Relations* 54 (4): 668–94.

Sabia, Joseph J. 2008. "Minimum Wages and the Economic Well-Being of Single Mothers." *Journal of Policy Analysis and Management* 27 (4): 848–66.

Sabia, Joseph J., and Richard V. Burkhauser. 2010. "Minimum Wages and Poverty: Will a $9.50 Federal Minimum Wage Really Help the Working Poor?" *Southern Economic Journal* 76 (3): 592–623.

Sabia, Joseph J., and Robert B. Nielsen. 2015. "Minimum Wages, Poverty, and Material Hardship: New Evidence from the SIPP." *Review of Economics of the Household* 13 (1): 95–134.

Wolfers, Justin. 2006. "Did Unilateral Divorce Laws Raise Divorce Rates? A Reconciliation and New Results." *American Economic Review* 96 (5): 1802–20.

Zipperer, Ben. 2016. "Did the Minimum Wage or the Great Recession Reduce Low-Wage Employment? Comments on Clemens and Wither (2016)." Washington Center for Equitable Growth Working Paper 120616.

This article has been cited by:

1. David Neumark, Maysen Yen. 2021. Effects of recent minimum wage policies in California and nationwide: Results from a pre-specified analysis plan. *Industrial Relations: A Journal of Economy and Society* **50**. . [Crossref]
2. Jefferson T. Barrett, Lois K. Lee, Michael C. Monuteaux, Caitlin A. Farrell, Jennifer A. Hoffmann, Eric W. Fleegler. 2021. Association of County-Level Poverty and Inequities With Firearm-Related Mortality in US Youth. *JAMA Pediatrics* **320**. . [Crossref]
3. Caitlin E. Caspi, Molly De Marco, Emily Welle, Claire Sadeghzadeh, Leah Chapman, Lisa J. Harnack, Rebekah Pratt. 2021. A Qualitative Analysis of SNAP and Minimum Wage Policies as Experienced by Workers with Lower Incomes. *Journal of Hunger & Environmental Nutrition* **129**, 1-19. [Crossref]
4. Todd Morris. 2021. The unequal burden of retirement reform: Evidence from Australia. *Economic Inquiry* **88**. . [Crossref]
5. Amy Shanafelt, Claire Sadeghzadeh, Leah Chapman, Molly De Marco, Lisa Harnack, Susan Gust, Melvin Jackson, Caitlin Caspi. 2021. Recruitment of Low-wage Workers for a Time-sensitive Natural Experiment to Evaluate a Minimum Wage Policy: Challenges and Lessons Learned. *Field Methods* **33**:3, 268-286. [Crossref]
6. Michael Lokshin, Martin Ravallion. Would Mexican Migrants be Willing to Guarantee Americans a Basic Income? . [Crossref]
7. Jose Garcia-Louzao, Linas Tarasonis. 2021. A First Glance at the Minimum Wage Incidence in Lithuania Using Social Security Data*. *Ekonomika* **100**:1, 31-53. [Crossref]
8. Young Cheol Jung, Adian McFarlane, Anupam Das. 2021. The effect of minimum wages on consumption in Canada. *The Economic and Labour Relations Review* **32**:1, 65-89. [Crossref]
9. Jhacova A. Williams, Trevon D. Logan, Bradley L. Hardy. 2021. The Persistence of Historical Racial Violence and Political Suppression: Implications for Contemporary Regional Inequality. *The ANNALS of the American Academy of Political and Social Science* **694**:1, 92-107. [Crossref]
10. Tatsushi Oka, Ken Yamada. 2019. Heterogeneous Impact of the Minimum Wage: Implications for Changes in Between- and Within-Group Inequality. *SSRN Electronic Journal* . [Crossref]