Estimating the Impact of Public Policy with Unobserved Variation

Aaron Sullivan*

November 15, 2024

[Click here for latest version]

Abstract

In recent years, many sub-state jurisdictions in the United States have introduced local economic policies previous only seen at the state or federal level. Because most publicly available data does not identify individuals at these local levels, estimating the impact of these policies is difficult as local variation will bias state-level estimates. In this paper I propose a solution that combines an intention-to-treat approach with two-sample IV where aggregates are sufficient for the first-stage estimation and identifies the local average treatment effect (LATE) of the policy. Using the recent prevalence of local minimum wage policy changes as an example, I show that estimating the impact of these policies on workers with this method provides statistically distinct results, though differences are economically small and qualitatively similar to previous studies.

Keywords: Non-Compliance, Intention-to-Treat, Two-Sample IV, Minimum Wage
JEL Codes: J08, J21, J42

*UC Irvine, a.sullivan@uci.edu

1 Introduction

In recent years, city- and municipal-level governments in the United States have introduced broad economic policies previous only seen at the state or federal level. For example, as of March 2023 61 sub-state jurisdictions have binding minimum wages that surpass the federal- or state-level minimum wage. Figure 1 highlights the relative scale of these policies in terms of the number of individuals facing state-level minimum wage changes and sub-state-level minimum wage changes. However, examining the effects of these and similar policies can be difficult as privacy restrictions often prevent common public social surveys from identifying individuals or groups in these sub-state jurisdictions. Indeed, if researchers wish to estimate the effects of these policies, but fail to control for this sub-jurisdiction variation, then common methods of estimation will fail to identify the policy effect of interest.

In this study I propose a standard intention-to-treat instrumental variables approach that uses observed state-level policy as an instrument for the binding policy in each jurisdiction. The typical method for this estimation requires the researcher to have information on the actual treatment status of individuals along with their assigned treatment. Since public data may lack this information due to privacy restrictions about where survey respondents live, I propose using aggregated population and treatment information in a two-sample instrumental variables (TSIV) approach to recover the local average treatment effect (LATE) of the policy.[1] I show that an aggregated first-stage can recover the TSIV estimate, provided that the policy is constant within the unit of observation and the researcher has information on the size of the population within each unit.

To highlight these facts, I first outline the nature of the bias associated with the unobserved variation in the context of the minimum wage (though this exercise is general to many forms of unobserved sub-jurisdiction policy variation). I then characterize the intention-to-

[1] Data that does identify individuals in Census "Place" jurisdictions or at the PUMA level, such as the American Community Survey (ACS), may be used in some scenarios to estimate the effects of these policies, but limit the scope of such analysis to repeated cross sections and longer time-frames, since these data typically do not identify the month the sample was taken in and does not follow individuals over time.

Figure 1: Treated Population and Population Distribution by Jurisdiction

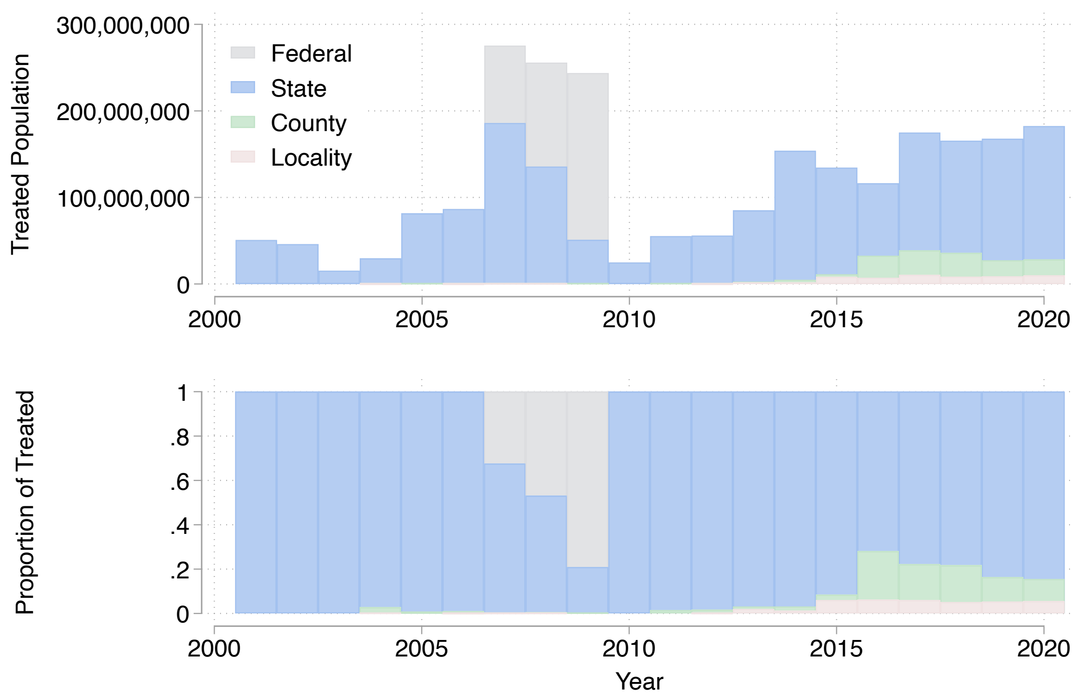

Source: Data on jurisdiction-level minimum wage changes are from Vaghul and Zipperer (2021) and population levels are from the Census Bureau's Intercensal and Postcensal population estimates.
Note: In the top panel the figure depicts the total population for jurisdictions that face a minimum wage change in the given year. Bars are stacked and so the total height of the bar equals the total population of all jurisdictions that faced a minimum wage change in the year. The bottom panel shows the same data in proportion of the total population treated within a given year.

treat approach nested within the TSIV procedure and highlight a special case where data in the first stage is aggregated while the second stage is disaggregated. Following this, I implement this TSIV procedure in the context of the minimum wage over three periods of varying unobserved variation, 2000-2020, 2000-2014, and 2015-2020, showing that this unobserved variation does indeed produce statistically distinct effects, but are economically and qualitatively small.

In these results I highlight differences between state-level, county-level, and city-level variation and show that for low levels of unobserved variation, state-level estimates are relatively unchanged by the correction, while those with higher levels become statistically distinct. These differences are, however, economically small, reaching approximately 0.015 percentage points in terms of the elasticity of the minimum wage with respect to employment or earnings. Qualitatively, these results are within the distribution of estimates found in the body of minimum wage research,[2] and so are unlikely to meaningfully alter the interpretation of any previous research that do not control for this unobserved variation.

In general, this work relates directly to research on the empirical estimation of the effects of the minimum wage (Card and Krueger, 1994; Dube et al., 2010; Allegretto et al., 2011, 2017; Neumark and Shirley, 2021; Jha et al., 2022, and others), as well as to those that examine city-level minimum wages directly (Jardim et al., 2022; Dube and Lindner, 2021; Neumark and Yen, 2021), and expands on them by providing a novel framework for estimating the effects of these state-level policy changes despite unobserved sub-state variation. This work also relates to the growing body of literature on place based policies (Busso et al., 2013; Kline and Moretti, 2013; Neumark and Simpson, 2015), highlighting the effect of these policies on the measurement of the effects of state- and national-level policies when these policies overlap. The main contribution of this work is to provide a framework to estimate the effects of state- and national-level policies with the public data currently available. Using this framework, I show that estimates for the effect of state-level minimum wage policies on earnings and employment are statistically distinct from those derived in traditional methods, though the differences are qualitatively and economically small.

2 The Effects of Unobserved Treatment Variation

As is well known (e.g., Durbin (1954); Imbens and Angrist (1994); Angrist et al. (1996); Frangakis and Rubin (1999)), misspecification—whether through measurement error or non-compliance—biases estimates of the parameter of interest in OLS. In this section I motivate this idea with an example where I model the unobserved policy variation as treatment

[2] For a meta analysis of minimum wage studies see Neumark and Shirley (2021).

non-compliance in a structural framework. I then outline the standard intention-to-treat approach using instrumental variables (IV) that can obtain a consistent estimate of the local average treatment effect (LATE) under standard TSIV assumptions, as in Imbens and Angrist (1994) and Angrist et al. (1996). In this approach treatment assignment is used as an instrument for actual treatment received in the instrumental variables estimation. In the context of minimum wages, the state-level minimum wage policy is the assigned treatment and the relevant binding minimum wage—city, county, state, or federal—is the treatment received. If the setting satisfies the standard assumptions for TSIV, then this procedure obtains a consistent estimate of the causal effect of interest.

This intention-to-treat approach requires that the researcher has information on who actually receives treatment, not just those who are assigned to treatment. This presents a problem for our setting, as data that identifies individuals in these sub-jurisdictions is often not available. In the context of the minimum wage, the availability of data that identifies individuals in these jurisdiction is sparse and will vary based on the data frequency and source. For example the American Community Survey (ACS) can identify individuals in "Census Place" designations and at the PUMA level, but is restricted to monthly cross-sections that are aggregated to the annual level, making it unclear when in fact the survey took place relative to the policy change within the year, and impossible to follow individuals over time. The solution I propose here is that, under certain conditions, population aggregates are sufficient for estimation of the first-stage in the two-stage estimation procedure. Essentially, this implies a TSIV approach, where the first-stage estimation is performed using the aggregated data sample and the second-stage with the original sample of preference. Since data on sub-jurisdiction aggregates are generally publicly available, for example both minimum wages and various population strata are freely available for US cities and counties from the US Census, this approach allows for the use of individual level data that otherwise would not be adequate.

To motivate this setup, consider a structural model of the form,

$$Y = X\beta + \varepsilon, \tag{1}$$

where Y is an $n \times 1$ vector of outcomes, X is an $n \times 1$ vector of values from the regressor, and ε is an $n \times 1$ vector of errors. Further assume that X can be decomposed into endogenous and exogenous components such that,

$$X = Z^s + Z^c. \tag{2}$$

Where Z^s is an $n \times 1$ vector of observable treatment components and Z^c is an $n \times 1$ vector of unobservable treatment components. In the context of the minimum wage we might think of Z^s as the state-level minimum wage policy and Z^c as any unobserved sub-state-level increase in the minimum wage. In this context Z^c would be zero when the state-level policy is binding. Suppose we attempt to estimate the effects of this policy, but only with the observed component of X, Z^s. In expectation the estimator for this endogenous regression obtains,

$$\mathbb{E}[\hat{\beta}] = \beta + \left(Z^{s\prime} Z^s\right)^{-1} Z^{s\prime} Z^c \beta, \tag{3}$$

where the second term is the bias associated with omitting the Z^c portion of X.

One common approach to deal with this sort of endogeneity is to observe that (3) may be rewritten more simply as

$$\mathbb{E}[\hat{\beta}] = \left(Z^{s\prime} Z^s\right)^{-1} Z^{s\prime} X \beta. \tag{4}$$

This formulation is useful because it makes clear that the bias is simply the OLS estimator for a regression of X on Z^s. So if, in the context of the minimum wage, state-level policies are positively correlated with local policies then this bias term will be positive, pushing estimates of the effect of interest away from zero. Given that state-level minimum wage policies preempt local policies when they surpass them, and in general these policies only increase in nominal terms, then this it is likely the case that these are positively correlated.

However, since the original premise is that we can only observe an individual's treatment assignment, Z^s, but not their actual treatment, X, we have no means to estimate this parameter directly. Fortunately, even if we cannot observe X in our sample of that includes outcomes, it is still possible to estimate a regression of X on Z^s, provided we have data with information on population aggregates for both X and Z^s. More specifically, we observe an individual's outcome, Y, and their assignment, Z^s, in the original sample, but have information on aggregate *levels* of treatment assignment, Z^s, and actual treatment, X, in a separate sample. The use of this aggregated sample implies a two-sample IV approach.

2.1 Two-Sample Instrumental Variables

The premise for two-sample IV stems from the fact that for a given instrumental variables estimation the construction of the first and second stage estimates could in principle be obtained from distinct samples (Angrist and Krueger, 1992). If we consider the structural setup in (1), we can denote the sample of original with subscript s, then in addition to standard assumptions for IV estimation, TSIV only requires that (a) $\text{plim}_{n_s \to \infty}(Z_s' X_s / n_s) = \Sigma_{ZX}$ and (b) statistical moments from the two samples are independent, which implies that $\lim_{n_1, n_2 \to \infty}(n_1/n_2) = k$, where k is a positive constant. The first assumption implies that the statistical moments for Z and X converge in expectation and so have the same limiting distribution. The second assumption, while not strictly required, allows for some simplifications with the covariance matrix.

Given the bivariate setup above, and assuming homoskedastic standard errors, we can write the TSIV estimator and covariance matrix as

$$\hat{\beta}_{IV} = \left(Z_2' X_2 / n_2\right)^{-1} Z_1' Y_1 / n_1, \tag{5}$$

and

$$\hat{\Omega} = \hat{\sigma}^2 \left(Z_2' X_2/n_2\right)^{-1} Z_1' Z_1/n_1 \left(X_2' Z_2/n_2\right)^{-1}, \tag{6}$$

where the inverted left term in (5) and the "bread" in (6) are moment estimators from the sample of Z and X and the right term in (5) and the "meat" in (6) are moment estimators from the sample of Y and Z. This construction is missing one final piece. Since we do not have a dataset that observes Z^c for individuals, we need some alternative way to estimate the moments in (5) and (6). This, however, is a comparatively simple task. Since X is constant for all individuals within the some jurisdiction, the moments above can be estimated directly using population aggregates.

To illustrate this, observe that within each jurisdiction the policy is constant for a given time period. Define an $m \times 1$ matrix, P, that is composed of the $j = 1, 2, \ldots, m$ unique entries in X, such that $x_i = x_{i'} = p_j$. If we further group the unique values of $x_i = x_{i'}$ in X, we can define n_j as the column dimension of the jth block representing the number of entries for p_j corresponding to X. Using this observation we can define the moment estimator for (5) and (6) as equivalently,

$$\sum_i^n z_i^s x_i = \sum_j^m n_j z_j^s x_j,$$

where the j subscript on the right-hand-side is used to identify the unique p_j values arising from this equivalence in Z^s and X, respectively. This leads to a feasible estimate of β_{IV}, and corresponding covariance matrix, that takes the form

$$\hat{\beta}_{IV} = \left(Z_2^{s\prime} \mathbf{n}_j X_2 / \sum_J n_j\right)^{-1} Z_1^{s\prime} Y_1/n_1, \tag{7}$$

and

$$\hat{\Omega} = \hat{\sigma}^2 \left(Z_2^{s\prime} \mathbf{n}_j X_2 / \sum_J n_j\right)^{-1} Z_1^{s\prime} Z_1^s/n_1 \left(X_2' \mathbf{n}_j Z_2^s / \sum_J n_j\right)^{-1}, \tag{8}$$

where now Z_2^s and X_2 are used to denote the aggregated P matrix as described above, $\mathbf{n}_j$ is a diagonal weighting matrix where the entries represent the m corresponding columndimensions of Z_2^s and X_2, n_j. Here $\hat{\sigma}^2$ can be estimated using a two-sample 2SLS approach. Importantly, this process also works if we substitute in averages from our original dataset, or a third dataset, for these jurisdictions in the estimation procedure.

Fortunately, this type of estimator is a direct representation of a "frequency" weighted estimator in common statistical software. For example, in Stata this regression would simply be a regression of minimum wages for all jurisdictions on state minimum wages weighted by the jurisdiction population using "fweights".

3 Data

Data for this study comes from four primary sources. First-stage estimates use postcensal and intercensal census population estimates from the Census Bureau's Population Estimates Program (PEP), and minimum wage data from the Historical State and Sub-state Minimum Wages series from Vaghul and Zipperer (2021). Second-stage estimates employ data from the Current Population Survey (CPS) for monthly state-level panel estimates and the Quarterly Census of Employment and Wages (QCEW) for quarterly county-level estimates. Summary statistics for all of the data employed is available in Table 1.

3.1 Minimum Wages

Data on minimum wages are obtained from the Historical State and Sub-State Minimum Wages series from Vaghul and Zipperer (2021). Minimum wages are available in all sub-jurisdictions for the entire period studied (2000-2020) at up to the daily frequency. For estimates in the aggregated first-stage estimates, which has a yearly frequency, the annual series is used. In the second-stage estimates, the minimum wage frequency that matches the data frequency is used. For the CPS this is the monthly frequency series and for the QCEW this is at the quarterly level. For all samples with a quarterly or yearly frequency, the average minimum wage over the period is used.

Over the period studied the minimum wage ranges from \$5.15 per hour to almost \$17 (averaged over the period in question), with the average increasing from \$5.20 per hour to \$8.84, a 70 percent increase. In general, the minimum wage is weakly increasing year-over-year in nominal terms.[3] This, combined with the fact that sub-jurisdiction wage floors are preempted by super-jurisdiction wage floors when they might fall below them, creates a positive relationship between state and sub-state minimum wages.

3.2 Postcensal and Intercensal City Population Estimates

Estimates of city-level population levels comes from the annual Postcensal and Intercensal Population Estimates from the US Census Bureau. Postcensal population estimates are annual population levels for incorporated and residual populated regions following a decennial census that are adjusted with location birth, death, and migration data. Intercensal population estimates are similar to postcensal estimates, in that the baseline annual estimates are generated with these same population records, but are then smoothed by extrapolating between decennial census periods. The estimates are adjusted such that the ratio of the intercensal and postcensal estimates follows a geometric progression over the decade.[4] In some cases this method can not be used, for instance if the population for the unit of observation declines to zero. For those cases a simple linear interpolation is used instead. This process makes these data particularly flexible. Originally the data is organized at the annual level, but by using this same process it is possible to adjust the population estimates to be at the quarterly or monthly frequency. This adjustment may limit the extent that aggregation plays a roll in estimation bias when observations vary within the unit of observation.

[3]Over the 21 year period, only five jurisdictions, across three states, have seen their nominal minimum wage decline year-over-year. In 2010 Colorado reduced it's minimum wage due to a deflationary period that activated a state-wide indexing provision to reduce the minimum wage in kind. In 2017 the Kentucky Supreme Court found that sub-state jurisdictions lacked the authority to institute their own local minimum wages, reverting increases enacted in Louisville and Lexington. Lastly, in 2017 the Iowa legislature and Governor enacted legislation that restricted sub-state jurisdictions from enacting their own local minimum wages, reverting previous minimum wage increases in Linn, Johnson, and Wapello County to the state level.

[4]This method is often referred to as the "Das Gupta" method and takes the form: $P_t = Q_t(P_{3652}/Q_{3652})^{t/3652}$, where P_t is the intercensal estimate for day t, Q_t is the postcensal estimate for day t, Q_{3652} is the end day postcensal estimate, and P_{3652} is the new decennial census estimate.

Table 1: Summary Statistics

	Mean	Std. Dev.	Min	Max	N
Census					
State Population	6,017,605	6,783,270	336,465	39,437,610	1,071
County Population	97,845	315,221	0	10,094,865	65,868
Locality Population	7,541	48,072	0	3,982,885	854,627
Minimum Wages					
State MW	7.1	1.6	5.1	15	1,071
County MW	10	2.4	5.2	16	237
Locality MW	12	1.6	6.8	17	3,182
QCEW					
Weekly Wage: All	681	180	380	1,963	1,071
Weekly Wage: Restaurants	302	120	117	847	1,071
Employment: All	4,474	5,778	320	44,845	1,071
Employment: Restaurants	58,352	77,772	4,487	539,477	1,071
CPS					
Employment Rate	.35	.48	0	1	1,764,164
Hourly Wage	10	17	.0012	1,800	143,966
Age	17	1.1	16	19	1,766,750
Share Male	.51	.5	0	1	1,766,750
Share Hispanic	.17	.37	0	1	1,766,750
Unemployment Rate	.17	.079	0	.72	1,766,750
Share Teen	.056	.0064	.019	.084	1,766,750

Source: Data on jurisdiction-level minimum wage changes are from Vaghul and Zipperer (2021) and population levels are from the Census Bureau's Intercensal and Postcensal population estimates.
Note: The table reports summary statistics for population levels and minimum wages by jurisdiction level from the year 2000 to 2020. For the CPS panel, the data is representative of those 16 to 19 years old, where the share teen category is the proportion 16 to 19 in the weighted sample population.

The unit of observation in the sample is the census designated "primitive geography level". These units represent cities, towns, residual unincorporated areas, or other indivisible sub-jurisdiction within a county. In total there are 40,722 of these primitive geographic locations over 3,143 counties. Table 2 reports population levels by jurisdiction for individuals within a location subject to a minimum wage change. When a location is either both a city and a county (San Francisco) or when the city may be completely identified by a combination of several counties (New York City), the unit is categorized as a county. This classification is intended to capture the extent to which various jurisdiction policy variation is observable in the data.

Table 2: Treated Individuals by Jurisdiction Level and Year

	Total	Federal	State	County	Locality
2001	50,804,834	0	50,804,834	0	0
2002	46,032,984	0	46,032,984	0	0
2003	15,397,143	0	15,397,143	0	0
2004	29,563,642	0	28,742,266	750,133	71,243
2005	81,578,189	0	80,829,343	748,846	0
2006	86,604,935	0	85,780,293	751,431	73,211
2007	275,365,441	89,321,603	184,766,796	758,348	518,694
2008	255,738,597	120,053,912	134,393,052	767,067	524,566
2009	243,754,397	192,779,519	50,200,531	774,347	0
2010	24,883,722	0	24,883,722	0	0
2011	55,337,218	0	54,521,524	815,694	0
2012	55,943,360	0	55,032,221	828,963	82,176
2013	84,982,751	0	82,292,994	958,615	1,731,142
2014	153,865,774	0	149,003,922	3,042,779	1,819,073
2015	134,519,523	0	123,183,572	3,215,024	8,120,927
2016	116,360,512	0	83,705,875	25,493,029	7,161,608
2017	175,016,171	0	136,223,422	28,343,463	10,449,286
2018	165,575,790	0	129,699,122	27,345,781	8,530,887
2019	167,820,836	0	140,382,076	18,467,837	8,970,923
2020	182,497,468	0	154,164,645	18,354,586	9,978,237

Source: Data on jurisdiction-level minimum wage changes are from Vaghul and Zipperer (2021) and population levels are from the Census Bureau's Intercensal and Postcensal population estimates.
Note: The table reports the number of individuals within a jurisdiction-level that are subject to binding jurisdiction-level minimum wage policy changes. In all years the policy represents nominal increases in the minimum wage except in five cases: Colorado in 2010, Lexington and Louisville localities (Kentucky) in 2017, and Linn, Johnson, and Wapello County (Iowa) in 2018. For San Francisco, which is both a city and a county, the variation is treated as county-level.

3.3 Quarterly Census of Employment and Wages

The QCEW reports employment and wages for over 95 percent of all workers and is derived from filings firms provide for workers covered by unemployment insurance. Previously called titled the ES-202, this dataset allows for identification of individuals within specific employment sectors. Using these data the sample is restricted to all individuals employed in the restaurant industry. Specifically the sample is restricted to workers within industry NAICs code 7221 and 7222 for “full-service” and “limited-service” restaurants from 2000-2010 and 7225 for “restaurants and other eating places” after 2010, due to NAICs coding changes.

This restriction is similar those used by other minimum wage studies that seek to restrict the sample to workers most likely to be affected by the minimum wage policy (e.g. Dube et al. (2010)) so that comparisons can be made.

3.4 Current Population Survey

The CPS is a monthly household survey administered by the Bureau of Labor Statistics and the US Census. The data used for this study begins with the January 2000 survey and extends to Decemeber 2020. The survey is administered in a staggered-cohort fashion, where-by individuals are surveyed during four consecutive month, are then left out for eight months, and then surveyed again the following year for the same consecutive four months. This provides that each individual surveyed is only present for the same four months one year apart, but may start those four months in any month, January to December. For example, if a respondent’s survey cohort begins in October 2010, they will be surveyed in October, November, December, in that year, and then January in 2011. Eight months later, they will again be surveyed in October, November, and December in 2011, and then again in January 2012. This structure allows for a limited panel of individuals over time.

Data for earnings is derived from the CPS’s Outgoing Rotation Group (ORG), otherwise known as the “earner study”. Questions that generate these data are asked only during the fourth and eighth month of a respondents survey sequence, their “outgoing” months. Estimates for earnings are further restricted to those with positive earnings in a period and so reduce the size of the sample considerably.

The CPS data are obtained from the Intregrated Public Use Microdata Series (IPUMS) (Flood et al., 2021). For the purposes of this study, information on a respondents state of residents, employment and labor force status, earnings, hours worked, as well as demographics details such as age, marriage status, sex, and education are obtained. Demographic weights are also obtained for both the basic monthly sample and the ORG.

4 Empirical Exercise

From Section 2 it's clear that estimates that fail to control for sub-jurisdiction variation will in general be biased.[5] In the context of the minimum wage, estimates face an upward bias. This is because standard estimation methods use super-jurisdiction policies (state-level), which preempt local policies when they are larger. Thus, all unobserved variation is positively correlated with the observed treatment.

In this section I outline the estimation procedure in the context of two prominent minimum wage studies. The first follows the traditional state-level procedure using aggregate data from the QCEW from Dube et al. (2010). While the purpose of this paper was to highlight the contiguous-border design the authors employ, state-level estimates are used as a benchmark and so seem appropriate for one here as well. The second, using the individual-level data from CPS and ORG, follows Allegretto et al. (2017). Using these methods as benchmarks, I provide empirical evidence of the impact of unobserved variation on state-level estimates of the effects of the minimum wage.

For first-stage regressions, estimates follow

$$\ln(\mathrm{mw}_{lt}) = \xi + \pi \ln(\mathrm{mw}_{st}) + \mathbf{X}_{st}\gamma + \tau_t + \mu_s + \varepsilon_{st}, \tag{9}$$

where $\ln(\mathrm{mw}_{lt})$ is the binding nominal minimum wage in location l in year t, $\ln(\mathrm{mw}_{st})$ is the binding nominal minimum wage for l's super-jurisdiction state s in year t, $\mathbf{X}_{st}$ is a set of covariates that are structured to match the second-stage equation as closely as possible, τ_t and δ_s are state and year fixed effects, and ε_{st} is an exogenous error term. The covariates included in each first-stage are the closest aggregate version of the second-stage equivalents.

[5] For a simulation exercise highlighting these results see the Appendix.

In general, these are state averages or total sums.

Estimates using the QCEW follow the reduced-form equation (1) from Dube et al. (2010), restricted to the state-level:

$$\ln(y_{sq}) = \alpha + \eta \ln(\text{mw}_{sq}) + \delta \ln(y_{sq}^{TOT}) + \gamma \ln(pop_{sq}) + \phi_q + \mu_s + \kappa_{sq}, \tag{10}$$

where the outcome of interest, $\ln(y_{sq})$ is log of employment or earnings in state s, in quarter q, $\ln(y_{sq}^{TOT})$ is the log of the total private sector employment or earnings in state s, in quarter q, and $\ln(pop_{sq})$ is the log of the population for state s in quarter q. For these estimates data is restricted to those in the restaurant industry (industry NAICs code 7221 and 7222 for "full-service" and "limited-service" restaurants from 2000-2010 and 7225 for "restaurants and other eating places" after 2010). For estimates using the CPS, the reduced-form equation (1) from Allegretto et al. (2017) is used. This equation follows:

$$y_{ism} = \beta \ln(\text{mw}_{sm}) + \mathbf{H}_{ism}\Gamma + \lambda \cdot \text{unemp}_{sm} + \psi_m + \mu_s + \nu_{ism}, \tag{11}$$

where y_{ism} is the outcome of interest and takes the form of either individual i's log hourly earnings or a dichotomous variable that takes the value one if the individual is employed and zero if they are unemployed or not in the labor force in month m. $\mathbf{H}$ is a matrix of individual characteristics including age, marital status, sex, education, race, and Hispanic status, and unemp_{sm} is the state-level unemployment rate in state s for month m. As in Allegretto et al. (2017), estimates are restricted to teens age 16-19 and also include in the equation the share of the population within this age group.

For each dataset and outcome three estimates are performed. One for the whole period, 2000-2020, one for a "placebo" period with relatively low levels of unobserved variation, 2000-2014, and a third for a high-unobserved-variation period, 2015-2020, during which an average of 20 percent of the minimum wage policy variation (in terms of population) occurred in sub-state regions, 6 percent of which were within cities.

Figure 2: First-Stage IV estimates by Year

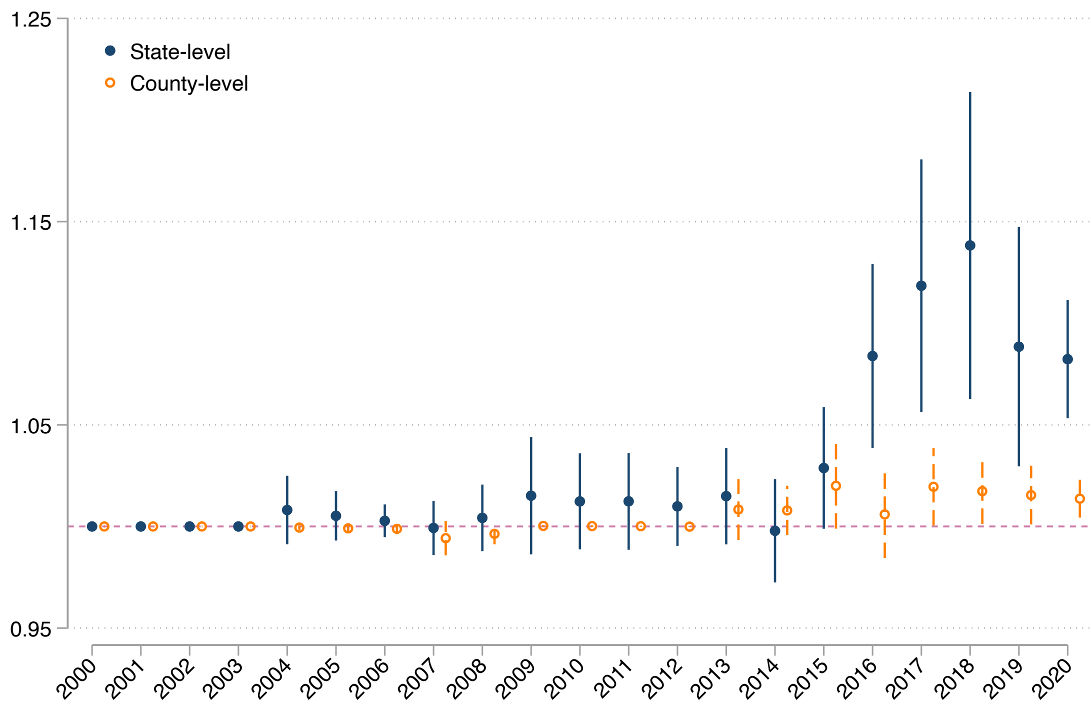

Note: Estimates follow the first-stage regression from (9) and are restricted to each year over the period. In all years the covariates for second-stage matching are omitted.

5 Results

To better motivate the empirical results, estimation of the first-stage equation for each year is given in Figure 2. From this figure we can see that during the 2015-2020 period the relative size of π, the coefficient from (9), increases dramatically. If we think of this coefficient in the classical IV case, that is as the term premultiplying β in (4), then it's a natural analog to the level of relative bias introduced by the unobserved variation. This implies that as we estimate models that focus on this region of time we should see larger biases in those estimates.

Table 3 displays the consolidated results from the empirical exercise. Columns 2-7 represent estimates of the elasticity of the minimum wage with respect to earnings for the period in question. The first two columns represent those elasticities for the whole period of this study, 2000-2020, with the second section representing the placebo period, 2000-2014, and the third two columns the high-unobserved-variation period, 2015-2020. Columns 8-13 are similarly grouped and represent estimates of the elasticity of the minimum was with respect to the employment to population ratio. The first row of results employ equation (10) and uses data from the QCEW for the second stage. The second row follows equation (11) and uses the CPS for it's second stage estimation. CPS estimates further divide the coefficient and standard errors by the mean of the outcome variable in order to generate the correct elasticities. First-stage results are in the bottom panel along with the subsequent F-test statistic.

The first thing to notice is that for all samples and periods the F-statistic is large (over 500). Additionally, the first-stage coefficient is large and positive for both the baseline and high-unobserved-variation periods, but is close to one for the placebo period. This isn't surprising, but still supports the claims made earlier. Outside of the placebo columns, all IV estimates are closer to zero when compared to their OLS counterparts (consistent with the Monte Carlo simulation in the Appendix). In all but the CPS employment estimates for 2015-2020, the IV estimates are statistically distinct from their OLS counterparts. Additionally, placebo estimates are not statistically distinguishable from the OLS estimates. These results support the idea that failing to control for this unobserved variation makes estimating these models infeasible.

Despite the clear evidence of bias in reduced form estimates from this period, the magnitude of this bias is relatively small. These results imply a difference of approximately 0.015 percentage points between IV and OLS estimates for both the employment and earnings elasticities, which is both economically and qualitatively small. That is, these differences are not likely to produce different interpretations of these effects if the research had estimated the original models. Furthermore, quantitatively, all of the biased results are within the historical distribution of minimum wage estimates (Neumark and Shirley, 2021).

6 Conclusion

In this paper I provide a novel framework for estimating the effects of state- and national-level policy when some portion of the populations' treatment is unknown. I examine the relative scale of unobservable sub-state policy variation and how that variation may inhibit identifying the policy effect of interest, using sub-state minimum wages as a motivating case. Between 2000-2020, minimum wage policy changes in locations not observable in common public datasets has increased considerably, with as much as 20 percent of the policy variation, in terms of population, residing in locations with unobservable minimum wage changes. In response to this, I propose a general framework that can obtain the local average treatment affect of interest. That is, by pairing intention-to-treat IV with two-sample IV, consistent estimates of the parameter of interest can be obtained, despite not observing individuals with these local jurisdictions.

Over the period of this study, standard difference-in-differences estimates of the effect of minimum wage increases on labor outcomes do not identify the causal effect of interest. Differences in standard state-level estimates and the framework proposed here, while statistically significant, are economically small. Differences between OLS and IV estimates for both the elasticity for earnings with respect to the minimum wage and for the elasticity of the employment to population ratio with respect to the minimum wage are about 0.015 percentage points. This difference is relatively small in the universe of minimum wage studies, with, for example, estimates for employment for young adults ranging from over -0.5 to 0.5 (Neumark and Shirley, 2021).

Table 3: Comparison of TS2SLS and OLS Minimum Wage Estimates

	ln(Earnings)						ln(Emp/Pop)					
	2000-2020		2000-2014		2015-2020		2000-2020		2000-2014		2015-2020	
	IV	OLS	IV	OLS	IV	OLS	IV	OLS	IV	OLS	IV	OLS
QCEW												
$\ln(MW)$	0.176***	0.191***	0.310***	0.308***	0.252***	0.270***	-0.202**	-0.219**	-0.011	-0.010	-0.174***	-0.186***
	(0.064)	(0.069)	(0.066)	(0.065)	(0.027)	(0.027)	(0.077)	(0.083)	0.050	0.049	(0.053)	(0.056)
CPS												
$\ln(MW)$	0.186***	0.200***	0.126***	0.125***	0.198***	0.212***	-0.095**	-0.102**	-0.116***	-0.115***	0.035	0.037
	(0.027)	(0.028)	(0.024)	(0.024)	(0.052)	(0.055)	(0.042)	(0.045)	(0.042)	(0.041)	(0.063)	(0.068)
First-Stage												
	QCEW Controls						CPS Controls					
	2000-2020		2000-2014		2015-2020		2000-2020		2000-2014		2015-2020	
$\ln(MW_s)$	1.082***		0.994***		1.070***		1.076***		0.994***		1.072***	
	(0.002)		(0.003)		(0.004)		(0.035)		(0.003)		(0.038)	
F-Stat	774.18		94,200.60		562.22		935.59		116,106.82		781.33	
N_{qcew}	259,512	259,512	184,826	184,826	74,686	74,686	259,512	259,512	184,826	184,826	74,686	74,686
N_{cps}	142,922	142,922	111,199	111,199	31,723	31,723	1,764,164	1,764,164	1,326,264	1,326,264	438,656	438,656
N_{fs}	841,010		600,459		240,551		841,010		600,459		240,551	

Standard errors in parentheses

* $p < .1$, ** $p < .05$, *** $p < .01$

Note: The table reports the effects of minimum wage changes on employment and earnings. Columns 2-5 report earnings elasticities with respect to the minimum wage, while columns 6-9 report employment elasticities. Results in the first row use data from the Quarterly Census of Employment and Wages (QCEW) for the restaurant sector, results in the second row are from the Current Population Survey (CPS) for individuals age 16-19. The bottom panel displays the first-stage regression for the IV columns and subsequent F-statistic. All estimates control for time period and state fixed effects. Estimates using the QCEW control for log employment in the private sector and log population. Estimates using the CPS control for individual demographics as well as the teenage population share and state unemployment rate.

References

Allegretto, Sylvia A., Arindrajit Dube, and Michael Reich, "Do Minimum Wages Really Reduce Teen Employment? Accounting for Heterogeneity and Selectivity in State Panel Data," *Industrial Relations: A Journal of Economy and Society*, April 2011, *50* (2), 205–240.

Allegretto, Sylvia, Arindrajit Dube, Michael Reich, and Ben Zipperer, "Credible Research Designs for Minimum Wage Studies," *ILR Review*, May 2017, *70* (3), 559–592.

Angrist, Joshua D. and Alan B. Krueger, "The Effect of Age at School Entry on Educational Attainment: An Application of Instrumental Variables with Moments from Two Samples," *Journal of the American Statistical Association*, June 1992, *87* (418), 328–336.

_ , Guido W. Imbens, and Donald B. Rubin, "Identification of Causal Effects Using Instrumental Variables," *Journal of the American Statistical Association*, June 1996, *91* (434), 444–455.

Busso, Matias, Jesse Gregory, and Patrick Kline, "Assessing the Incidence and Efficiency of a Prominent Place Based Policy," *American Economic Review*, April 2013, *103* (2), 897–947.

Card, David and Alan Krueger, "Minimum Wages and Employment: A Case Study of the Fast Food Industry in New Jersey and Pennsylvania," *American Economic Review*, 1994, *84* (4), 772–793.

Dube, Arindrajit and Attila Lindner, "City Limits: What Do Local-Area Minimum Wages Do?," *Journal of Economic Perspectives*, February 2021, *35* (1), 27–50.

_ , T. William Lester, and Michael Reich, "Minimum Wage Effects Across State Borders: Estimates Using Contiguous Counties," *Review of Economics and Statistics*, November 2010, *92* (4), 945–964.

Durbin, J., "Errors in Variables," *Revue de l'Institut International de Statistique / Review of the International Statistical Institute*, 1954, *22* (1/3), 23.

Flood, Sarah, Miriam King, Renae Rodgers, Steven Ruggles, J. Robert Warren, and Michael Westberry, "Integrated Public Use Microdata Series, Current Population Survey: Version 9.0 [dataset]," 2021. Place: Minneapolis.

Frangakis, Constantine E. and Donald E. Rubin, "Addressing complications of intention-to-treat analysis in the combined presence of all-or-none treatment-noncompliance and subsequent missing outcomes," *Biometrika*, June 1999, *86* (2), 365–379.

Imbens, Guido W. and Joshua D. Angrist, "Identification and Estimation of Local Average Treatment Effects," *Econometrica*, March 1994, *62* (2), 467.

Jardim, Ekaterina, Mark C Long, Robert Plotnick, Emma van Inwegen, Jacob Vigdor, and Hilary Wething, "Minimum-Wage Increases and Low-Wage Employment: Evidence from Seattle," *American Economic Journal: Economic Policy*, May 2022, *14* (2), 263–314.

Jha, Priyaranjan, David Neumark, and Antonio Rodriguez-Lopez, "What's Across the Border? Re-Evaluating the Cross-Border Evidence on Minimum Wage Effects," Technical Report 2022.

Kline, Patrick and Enrico Moretti, "Place Based Policies with Unemployment," *American Economic Review: Papers & Proceedings*, May 2013, *103* (3), 238–243.

Neumark, David and Helen Simpson, "Place-Based Policies," in "Handbook of Regional and Urban Economics," Vol. 5, Elsevier, 2015, pp. 1197–1287.

_ **and Maysen Yen**, "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*, December 2021, (October), 1–28.

_ **and Peter Shirley**, "Myth or Measurement: What Does the New Minimum Wage Research Say about Minimum Wages and Job Loss in the United States?," Technical Report, National Bureau of Economic Research, Cambridge, MA January 2021.

Vaghul, Kavya and Ben Zipperer, "Historical State and Sub-state Minimum Wages," 2021.

A Appendix

A.1 Monte Carlo Simulation

The following section outlines the Monte Carlo simulation employed. These estimates support the results outlined in Section 2.1. The simulation assumes a 20 percent rate of non-compliance and that the elasticity of the minimum wage with respect to wages is 0.20 percent. The simulation performs 10,000 repetitions of a simple model that relates the log of the worker wage to the log of the minimum wage.

Using the data from the QCEW, the sample is generated by attributing a minimum wage increase to a 20 percent sub-sample of counties and then estimating a two-sample two stage least squares model where the first stage contains information on local minimum wages and the second stage only contains information on aggregate state-level minimum wages. This increase is generated such that it follows a $N(2, 1)$ distribution and is then added to the minimum wage for counties within the selected sub-sample. Wage outcomes are generated such that the effect of the "true" minimum wage increase follows a $N(0.2, 0.025)$, which is similar in magnitude to previous studies. Figure 3 depict the results of this process. The red dashed line depicts the imposed "true" effect. The black line depicts the reduced-formed estimates that would obtain had the researcher ignored the sub-sample variation and estimated their model regardless. These estimates have a mean of 0.2023 compared to 0.2 from the true effect. The gray dashed curve depicts the TS2SLS estimates, where the first-stage uses the original disaggregated data and the second stage employs an aggregated version at the state-level. The average estimate from this method is 0.1997, considerably closer than the reduced-form estimates to the true effect.

As is clear from the figure, the reduced-form estimates overstate the effect of the minimum wage policy. Consider equation (3), where the estimated parameter faces some additive bias associated with the relationship between state and sub-state treatment. Since this relationship is weakly positive—state minimum wages preempt local minimum wages when

Figure 3: Monte Carlo Simulation: Standard and IV Methods

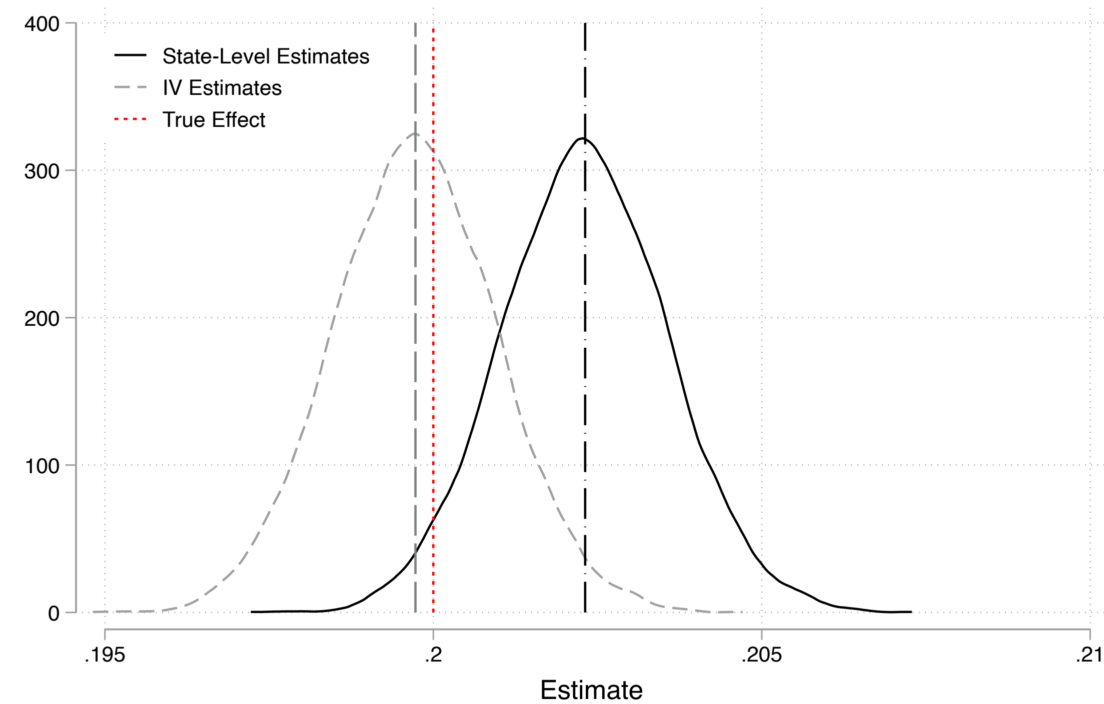

Notes: The figure displays a density plot of the resulting coefficient estimates from 10,000 Monte Carlo repetitions of the effect of minimum wage changes on the average wage of workers where 20% of the sample faces unobserved sub-jurisdiction minimum wage changes.

they surpass them—this bias is weakly positive. In fact, only 328 (3.28 percent) of the repetitions are at or below the defined true effect. This result is in contrast to typical errors-in-variables bias, where the mean zero noise generates a downward bias in the estimates of the parameter of interest. However, the TS2SLS method is not without its own drawbacks. From the figure we can see that the average estimate for this model is biased downward. Furthermore, this bias is in the opposite direction of that of the reduced-form estimates. This procedure thus provides a "lower bound" for the true parameter of interest.