Does regional variation in wage levels identify the effects of a national minimum wage?

Daniel Haanwinckel *

May 2024

Abstract

I evaluate the performance of estimators that exploit regional variation in wage levels to identify the employment and wage effects of national minimum wage laws. For the "effective minimum wage" design, I show that the identification assumptions in Lee (1999) are difficult to satisfy in settings without regional minimum wages. For the "fraction affected" design, I show that economic factors such as skill-biased technical change or regional convergence may cause parallel trends violations and should be investigated using pre-treatment data. I also show that this design is subject to misspecification biases that are not easily solved with changes in specification.

*Department of Economics, UCLA and NBER. Contact information: haanwinckel@econ.ucla.edu
This paper benefited from conversations with Martha Bailey, Adriana Lleras-Muney, Maurizio Mazzocco, Yotam Shem-tov, and many others at UCLA, as well as with participants of the 2023 Columbia Junior Micro/Macro Labor Conference, the 2023 NBER Labor Studies Fall Meeting, and the 2024 SOLE Meeting. Filipe Fiedler provided excellent research assistance.

1 Introduction

One approach to measuring the effects of minimum wage regulations involves leveraging variation in wage levels across different regions within a country. For instance, the influence of the US federal minimum wage is more pronounced in states like Mississippi or Arkansas than in Texas or Georgia due to the higher median wages in the latter states. Thus, one can use wage distributions to construct treatment intensity measures for a "differences-in-differences" analysis. Classic applications using cross-state variation from the US are Card (1992), who introduced the "fraction affected" design, and Lee (1999), who introduced the "effective minimum wage" design. This strategy is particularly convenient in countries with little or no spatial variation in minimum wage laws. Examples of such applications published by leading economic journals include Mexico (Bosch and Manacorda, 2010), South Africa (Dinkelman and Ranchhod, 2012), Germany (Dustmann et al., 2021), the US in the 1960s and 1970s (Bailey, DiNardo and Stuart, 2021), and Brazil (Engbom and Moser, 2022).

This paper examines the identification assumptions underlying those econometric approaches. Specifically, it uses a combination of economic theory and simulation exercises to investigate whether fraction affected and effective minimum wage designs can accurately capture the causal effects of the minimum wage on employment and wages, focusing on applications where the minimum wage is set at the national level.

For the effective minimum wage design, I show that two identification assumptions emphasized by Lee (1999) are crucial for unbiased estimation of causal effects but very difficult to satisfy in applications where there is no regional variation in minimum wage laws. The first assumption is that observed median log wages are a good proxy for the centrality of the latent log wage distribution—that is, the distribution of log wages that would prevail with no minimum wage in place. I show that employment effects lead to violations of that assumption by introducing correlated measurement errors in the design. The ensuing biases can be significant regardless of whether the average employment effects are negative, zero, or positive, and even if the correlation between the unobserved centrality of the latent wage distribution and the observed median log wage is above 0.99.

The second assumption in Lee (1999) is that overall wage levels should be uncorrelated with the dispersion of latent log wages across regions. This assumption may be violated if, for example, regions differ in the share of skilled workers, and there is more latent wage dispersion for skilled workers (Lemieux, 2006). I show that correlations as small as 0.07—no larger than what one would infer from US data—can introduce large biases.

I discuss the effectiveness of potential solutions to those problems with the effective minimum wage design. They include testing for spillovers in the upper tail as a diagnostics tool, using higher quantiles of the wage distribution to construct the effective minimum wage, or changing the set of fixed effects or trends included in the regression. The general message is that the design is unlikely to recover true causal effects unless there exists an economic factor that (i) changes the "bite" of the minimum wage in some regions compared to others; (ii) is otherwise unrelated to employment outcomes and the shape of the log wage distribution, conditional on the fixed effects and controls included in the regression; and (iii) has sufficient residual variance. If the data includes a variable that is a good candidate for being that economic factor, then an instrumental variables approach that directly exploits it is more likely to be successful.

Next, I discuss fraction affected and gap designs implemented at the regional level (this paper does not address estimators that compare firms within the same region based on the firm-level share of affected workers, e.g. Card and Krueger, 1994; Harasztosi and Lindner, 2019). The key identification assumption is transparent in those designs: in a counterfactual scenario without an increase in the national minimum wage, trends in outcomes would be orthogonal to the treatment intensity variable used, conditional on the controls and fixed effects used in the design. My analysis points out three potential issues. First, I show that structural factors that may seem unproblematic, such as a common trend in the dispersion of latent log wages affecting all regions in the same way (due, e.g., to skill-biased technical change), may cause violations of the parallel trends assumption. Second, the design is subject to bias from regression to the mean in regional wage statistics, originating from sampling variation when constructing those statistics and from region-specific productivity shocks. Third, the design is sensitive to the functional form chosen for the treatment intensity variable, with misspecification biases being more significant when the minimum wage is more binding or when it causes an increase in the number of workers earning a bit more than the minimum wage.

I also discuss possible diagnostics and solutions to those issues. I show that tests for differential pre-trends may effectively detect the first two issues when the data includes a "pre-treatment" period without significant changes in minimum wage laws, provided that the econometrician is careful when implementing the test. Similarly, if sufficient pre-treatment data is available, regression to the mean can be controlled using the procedure discussed in

Dustmann et al. (2021). For the misspecification issue, I consider three potential solutions: binary treatment intensity measures, quadratic specifications, and using one treatment intensity variable as an instrument for the other. Neither of these approaches solves the problem, and the binary version typically displays larger biases than the continuous fraction affected or gap specifications.

The paper is structured as follows. Section 2 introduces the setup, clarifies the scope of my analysis, and shows how it relates to other papers on the econometrics of minimum wage. Sections 3 and 4 analyze the effective minimum wage and the fraction affected estimators, respectively. In Section 5, I show that simulation results shown in Sections 3 and 4 also hold when using alternative economic models as the data-generating processes. The final section concludes with a summary of recommendations for researchers interested in measuring the effects of national minimum wage changes in contexts without regional variation in minimum wage laws.

2 Setup and relationship to literature

I consider two-period ($t \in \{0,1\}$) data-generating processes (DGP) of the following form:

$$\begin{aligned} \boldsymbol{y}_{r,t} &= f\left(mw_t, \boldsymbol{\theta}_{r,t}\right) \\ [\boldsymbol{\theta}'_{r,0}, \boldsymbol{\theta}'_{r,1}]' &\sim G \end{aligned} \tag{1}$$

where $r \in \{1, \ldots, R\}$ indexes regions, $\boldsymbol{y}_{r,t}$ is a vector of equilibrium outcomes (such as employment to population ratio or quantiles of the log wage distribution), mw_t denotes the logarithm of the national minimum wage, $\boldsymbol{\theta}_{r,t}$ is a vector of region-time-specific determinants of the outcomes of interest that differ across regions, and f is a function that outputs the equilibrium outcomes of a particular economic model (which I will later specify in the simulation exercises). The variables that compose the $\boldsymbol{\theta}_{r,t}$ vectors may display correlations across periods within regions, as determined by the distribution G, but they are independent across regions.

I assume henceforth that $mw_1 > mw_0$.[1] Given the data-generating process described above, there are two natural ways to define the ceteris paribus causal effects of the rise in the national

[1] This assumption is without loss of generality given the data-generating process above, as outcomes depend only on current values of mw_t and $\boldsymbol{\theta}_{r,t}$. It would not be without loss of generality in a model that accounted for nominal rigidities or other dynamic concerns. See the discussion at the end of this section.

minimum wage:

$$\begin{aligned}
\boldsymbol{ATE}_0 &= \mathbb{E}\left[f\left(mw_1,\theta_{r,0}\right) - f\left(mw_0,\theta_{r,0}\right)\right] \\
&= \mathbb{E}\left[f\left(mw_1,\theta_{r,0}\right)\right] - E_G\left[\boldsymbol{y}_{r,0}\right] \\
\boldsymbol{ATE}_1 &= \mathbb{E}\left[f\left(mw_1,\theta_{r,1}\right) - f\left(mw_0,\theta_{r,1}\right)\right] \\
&= \mathbb{E}\left[\boldsymbol{y}_{r,1}\right] - E_G\left[f\left(mw_0,\theta_{r,1}\right)\right]
\end{aligned}$$

where the expectation is taken with respect to the distribution of the $\theta_{r,t}$ variables. The first formulation, $\boldsymbol{ATE}_0$, requires evaluating a counterfactual where the minimum wage rises from mw_0 to mw_1 but other characteristics remain at their $t=0$ levels. The second formulation compares the outcomes as of $t=1$ to a counterfactual scenario where the minimum wage remained at the $t=0$ level. The two definitions are identical if the G distribution is time-invariant, an assumption that holds for only some of the DGPs I study. For consistency throughout the paper, I use the average of these two definitions as the object of interest to be recovered by the econometric designs:

$$\boldsymbol{ATE} = \frac{\boldsymbol{ATE}_0 + \boldsymbol{ATE}_1}{2}$$

All exercises in the paper impose an additional restriction on the data-generating process: there are no trends in overall wage levels. Suppose the minimum wage change is simultaneous with an unobserved shock to total factor productivity (TFP) affecting all regions. In that case, it is only possible to separately identify the average effects of the minimum wage by imposing further assumptions. To abstract from this "missing intercept" issue, I rule out common TFP shocks, though I explore idiosyncratic, mean-zero TFP shocks. In practice, econometricians should interpret estimates coming from these regressions as the impact of the minimum wage net of common TFP shocks.

Comparison to existing literature: Equation (1) is fairly general, but it imposes important constraints that limit the scope of my analysis. Discussing these limitations helps pinpoint how my findings differ from and complement existing literature.

Many papers discuss econometric challenges arising from the fact that minimum wage effects may take some time to materialize. Using Canadian data, Baker, Benjamin and Stanger (1999) document that employment negatively responds to low-frequency minimum wage variation. In contrast, the response to the high-frequency component of the variation is positive (though statistically insignificant). Sorkin (2015) builds a putty-clay model where firms cannot easily adjust labor inputs in the short run. Using this model, he argues that short- and long-run effects of the minimum wage can be very different and that it may be impossible to identify long-run effects through econometric designs if the data does not contain permanent shocks to the real minimum wage (which is often the case, as minimum wage increases are eroded by inflation over a few years). Meer and West (2016) argue that, in many economic models (including the putty-clay model of Sorkin, 2015), minimum wages cause changes in employment growth rates, not levels. Then, standard difference-in-differences designs may fail to capture the actual effects of the minimum wage on employment, especially if the designs include region-specific time trends. Vogel (2023) documents that minimum wage effects are concentrated on the bottom of the wage distribution on impact but "trickle up" over the next few years. Because I study a two-period model without dynamics, all of the issues I document in this paper are separate from those discussed above.

Equation (1) does not include measurement error, meaning that the issues I discuss in this paper are also different from the mechanical bias issue discussed in Autor, Manning and Smith (2016). These authors propose an instrumental variables estimator as an improvement over the effective minimum wage design of Lee (1999). Later in the paper, I will show that the estimator of Autor, Manning and Smith (2016) can also solve some of the issues I document. However, that strategy is only feasible with alternative data-generating processes that include region-specific changes in the minimum wage.

My model also imposes complete independence between regions, ruling out spillover effects coming from, e.g., migration responses. See Huang (2019), Chapter 2 for an investigation of whether such spillovers create biases in the context of the US.

Finally, the issues I discuss complement recent papers studying the econometrics of difference-in-differences models. The restriction to a two-period model means that issues related to staggered treatment and treatment effect heterogeneity across groups do not apply; see de Chaisemartin and D'Haultfœuille (2020) and Roth et al. (2023). Callaway, Goodman-Bacon and Sant'Anna (2024) study difference-in-differences models with continuous treatment intensity variables, a group that includes the fraction affected-style regressions I study in Section 4. Part of my results can be interpreted as showing that, under a range of economic models, the fraction affected design and related estimators do not satisfy the "strong parallel trends" assumption defined by Callaway, Goodman-Bacon and Sant'Anna (2024), causing biases that may be large in some cases and whose sign is difficult to predict.

3 The effective minimum wage design

3.1 Definition

Let $w_{q,r,t}$ denote quantile q of the log wage distribution in region r at time t. Now suppose that the econometrician is interested in two types of endogenous outcomes $y_{i,r,t}$: *quantile gaps* of the form $w_{q,r,t} - w_{0.5,r,t}$ and an employment measure such as the employment-to-population ratio. The effective minimum wage design uses the following ordinary least squares regression to estimate the effects of the national minimum wage increase:

$$y_{i,r,t} = \alpha_{i,r} + \delta_{i,t} + \beta_i \left[mw_t - w_{0.5,r,t}\right] + \gamma_i \left[mw_t - w_{0.5,r,t}\right]^2 + \varepsilon_{i,r,t}, \tag{2}$$

where i indexes the particular outcome of interest, such that each outcome corresponds to a separate regression. The term $mw_t - w_{0.5,r,t}$ is the (log) effective minimum wage. It measures how binding the national minimum wage is in a particular region and time, using the observed median wage as the benchmark. This baseline specification includes region and time fixed effects, following the bulk of the literature. I discuss alternative specifications later.

To calculate the predicted treatment effects of the minimum wage increase in each region r, the econometrician multiplies the changes in the effective minimum wage (and its square) by the estimated $\hat{\beta}$ and $\hat{\gamma}$ parameters. Those products can then be added up and averaged across regions, yielding the estimated average treatment effects of the national minimum wage increase following the definition from Section 2.[2]

This regression model was first introduced by Lee (1999), who focused on quantile gaps as the outcomes of interest. The design was later used to estimate employment effects as well; Engbom and Moser (2022) is one example.[3]

[2] That is, the predicted average treatment effects of the minimum wage are:

$$\widehat{ATE}_q = \frac{1}{R}\sum_r \left\{ \hat{\beta}_q \left[(mw_1 - w_{0.5,r,1}) - (mw_0 - w_{0.5,r,0}) \right] + \hat{\gamma}_q \left[(mw_1 - w_{0.5,r,1})^2 - (mw_0 - w_{0.5,r,0})^2 \right] \right\}$$

The econometrician could be interested not in the average change in region-specific quantile gaps but in the change in quantile gaps based on the national log wage distribution. I focus on the former definition because it is more closely linked to the regression model.

[3] Such regressions follow in the tradition of earlier papers that used variation in wage levels to measure employment effects of minimum wages. A well-known example is Neumark and Wascher (1992), who use as the treatment variable the nominal minimum wage in a state-year multiplied by the state-specific minimum wage coverage and divided by the state-specific average wage. My analysis focuses on the quantile-based effective minimum wage because it is more common in recent work.

In order to discuss identifying assumptions, it is helpful to introduce notation from Lee (1999). Suppose each region has a latent distribution of log wage in each period—that is, the distribution of log wages that would prevail with no minimum wage regulation. Assume that the cumulative distribution function for those latent log wages has the form:

$$F_t\left(\frac{w-\mu_{r,t}}{\sigma_{r,t}}\right)$$

where $\mu_{r,t}$ and $\sigma_{r,t}$ are the *centrality* (or location) and *dispersion* parameters, respectively.

Using this notation, Lee (1999) emphasizes two identification assumptions. First, the deflator used to construct the effective minimum wage—that is, the median wage $w_{0.5,r,t}$ in Equation (2)—should provide a good approximation for the centrality parameter $\mu_{r,t}$. Second, the location and dispersion parameters should be uncorrelated across regions conditional on t. When employment is the outcome of interest, one must also assume that latent employment is uncorrelated with location parameters, conditional on the fixed effects included in the regression.

The central message of my analysis below is that those assumptions are essential but unlikely to hold in practice for economic reasons. Later, I will also show that the issues are less severe in contexts with regional variation in the minimum wage, and in those contexts, the instrumental variables estimator of Autor, Manning and Smith (2016) is preferable.

3.2 Issue #1: Imperfect measurement of latent centrality

The effective minimum wage design is predicated upon minimum wage effects being stronger where it bites more into the *latent* wage distribution. Following that logic, the econometrician would ideally use $mw_t - \mu_{r,t}$ as the key regressor. But because $\mu_{r,t}$ is not observed, $mw_t - w_{0.5,r,t}$ is used instead. In this subsection, I argue that even minor deviations between $w_{0.5,r,t}$ and $\mu_{r,t}$ are enough to introduce economically significant biases.

3.2.1 Good and bad variation

To understand why even minor deviations between $w_{0.5,r,t}$ and $\mu_{r,t}$ can be problematic, consider a simple model with only two regions, A and B. In this model, a more binding minimum wage has small but positive effects on the median wage. These effects at the median can arise from strong spillovers (e.g., workers moving from low- to high-wage firms) or because the minimum wage causes disemployment effects in the lower tail, such that all quantiles of the log wage distribution mechanically move to the right.

I start by discussing the ideal source of variation for the effective minimum wage designs: random shocks to the location parameters $\mu_{r,t}$. In period $t = 0$, both regions have identical distributions of latent log wages. In period $t = 1$, two things happen. First, the national minimum wage rises. Second, the location parameter in Region B increases— for simplicity, assume that it increases by the same amount as the minimum wage. Panel A in Figure 1 illustrates this scenario.

In this scenario, Region A is the "treatment group" while B is the "control." In B, the minimum wage binds as much in period $t = 1$ as it did in $t = 0$. Thus, we should not expect any changes in the effective minimum wage or outcomes of interest. Thus, comparing A and B provides a valid quasi-experiment from which we can recover the causal effects of the minimum wage, even though the change in the national minimum wage was the same everywhere. The broader point is that, in the effective minimum wage design with region and time fixed effects, the ideal identifying variation comes from idiosyncratic shocks to the location parameters $\mu_{r,t}$.[4]

Next, I show how heterogeneity in dispersion parameters $\sigma_{r,t}$ can introduce "bad variation:" a spurious empirical link between the effective minimum wage and the outcomes of interest. Again, we will explore a scenario where Region A is affected by the increase in the national minimum wage, while Region B is not. But now assume that neither location nor dispersion parameters change over time; the reason why B is the "control" is because the time-invariant dispersion σ_B is minimal, as illustrated in Panel B of Figure 1. The minimum wage has no bite in that region because latent wages tightly concentrate around the median.

What would the regression recover in this scenario? The effective minimum wage rises in both regions as the minimum wage gets closer to the median wage. However, the relative increase is higher in Region B due to spillover effects on the median wage in Region A. Thus, if those permanent differences in the dispersion of latent wages are the only source of variation, then the predicted treatment effects would have the opposite sign compared to the actual causal effects of the minimum wage.

[4]One may wonder whether the small spillover effects in Region A generate bias, since the change in the effective minimum wage is smaller than the change in the national minimum wage. To see why this is not a problem, recall from Footnote 2 that after estimating Equation (2), the predicted treatment effects are calculated by multiplying the coefficients and changes in the effective minimum wage, not changes in the national minimum wage.

Panel A: Good variation arising from a shock to location, $\Delta\mu_B$

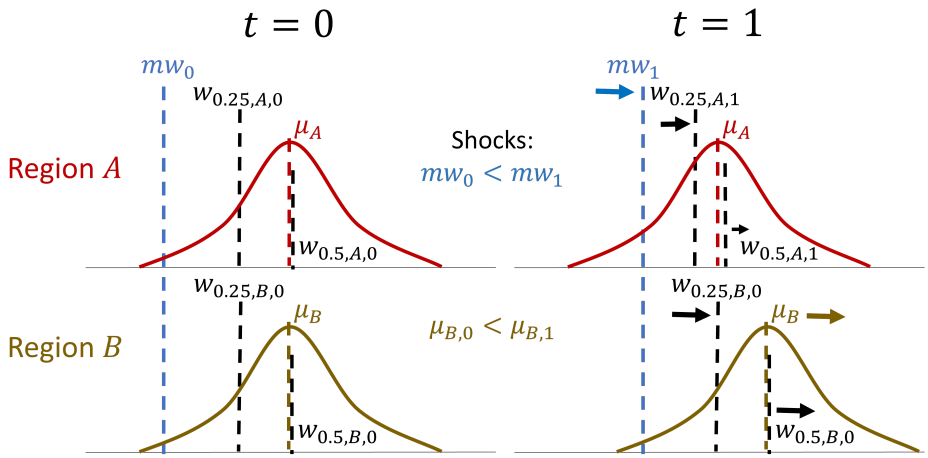

Panel B: Bad variation arising from differences in dispersion, $\sigma_A > \sigma_B$

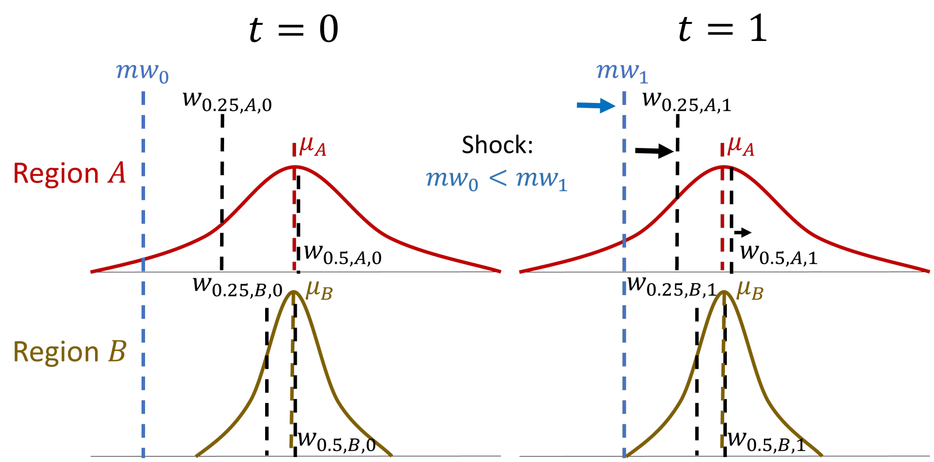

Panel C: Bad variation arising from a shock to dispersion, $\Delta\sigma_B$

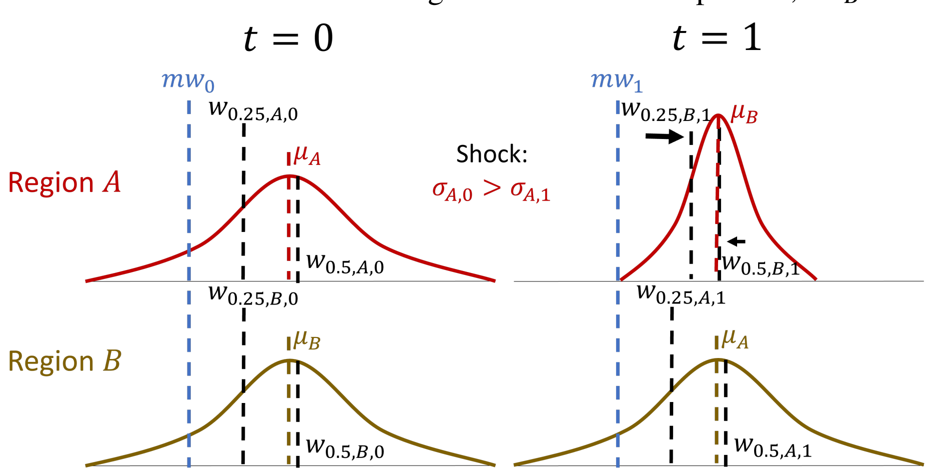

Figure 1: Good and Bad Variation in the Effective Minimum Wage Design

Shocks to location parameters $\sigma_{r,t}$ can also introduce bad variation. Panel C in Figure 1 shows a scenario where the only change over time is a fall in the dispersion parameter in Region A. That shock reduces inequality, moving percentile 25 of the log wage distribution closer to the median. Moreover, because the minimum wage becomes less binding, the median log wage falls, such that the effective minimum wage rises. Thus, comparing changes between Region A and Region B, the estimator will estimate a positive relationship between the effective minimum wage and the log wage gap $w_{0.25,r,t} - w_{0.5,r,t}$. However, the magnitude of this link is likely to be significantly overstated compared to the causal minimum wage effects. That is because the regression entirely attributes the inequality-reducing effects of the $\sigma_{r,t}$ change to that small change in the effective minimum wage induced by the spillovers at the median.

I finish this discussion with three remarks. First, since the scenarios in Panels B and C imply biases in opposite directions, it would only be possible to know the direction of the bias by making stronger assumptions about the data-generating process. Second, to see how this is fundamentally a measurement error issue, note that the issues I described above would not exist if the econometrician could observe $\mu_{r,t}$ and use it to construct the effective minimum wage. If that were the case, only the scenario shown in Panel A would generate identifying variation for the effective minimum wage design. Third, in all of those examples, there is no systematic relationship between location and dispersion parameters in all of those scenarios. I discuss problems arising from such correlations in the following subsection.

Simulations. To investigate the potential magnitude of the bias in empirical applications, I perform simulation exercises with parameters calibrated based on state-level data from the US Current Population Survey. I assume that latent log wages are Normally distributed in every region. There is a "markdown" parameter $m \in [0,1]$ such that the latent distribution is truncated at $mw_t + \log m$ and censored at mw_t. That is, workers who would earn less than the minimum wage times the markdown become disemployed, and those with latent wages above that cutoff but below the minimum earn exactly the minimum wage (that is, they create a minimum wage "spike" in the simulated log wage distribution). Unless otherwise noted, all simulations have a markdown parameter of $m = 0.7$.

Figure 2 illustrates the minimum wage effects in these simulations for two regions that differ in minimum wage bindingness. The truncation and censoring effects above correspond to the red and orange areas. That figure also includes the possibility of positive employment effects slightly above the minimum wage, illustrated in green. The first set of simulation

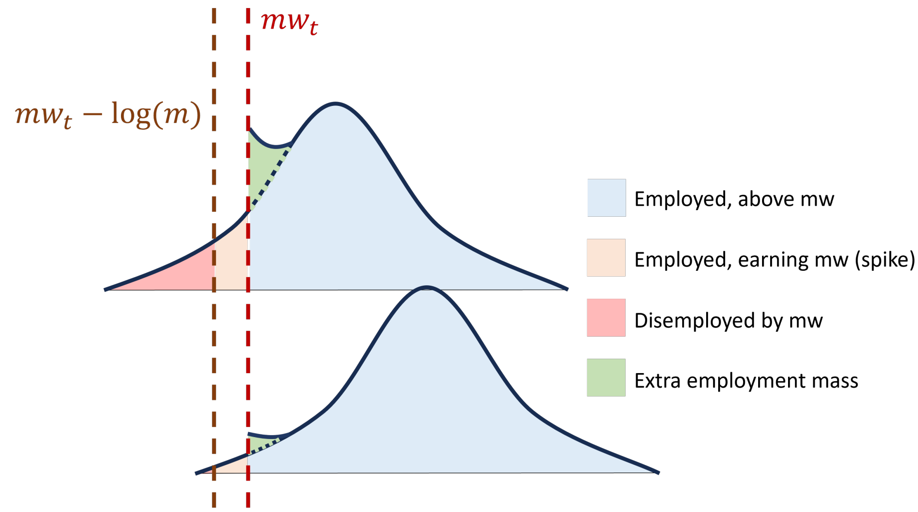

Figure 2: Minimum wage effects in the baseline simulation model

results reported below does not include those positive employment effects; I will return to this point at the end of this section.

Each region is described by a vector $[\mu_{r,0}, \sigma_{r,0}, \mu_{r,1}, \sigma_{r,1}]$, drawn from a multivariate Normal distribution. The parameters for that multivariate Normal are calibrated based on state-level data from the US Current Population survey, based on the years 1989 (corresponding to $t = 0$) and 2004 (corresponding to $t = 1$).[5] As explained below, Each particular simulation exercise makes different assumptions about that meta-distribution of parameters across regions. In all exercises shown, the data contains 200 regions. See Appendix A.1 for details.

Before showing the results, I make an important note. Since I ignore state-level minimum wage regulations and select particular years in the analysis, these exercises do not constitute an evaluation of the effective minimum wage design in the US context. Instead, I use the US data to argue that the econometric issues I describe could be significant in contexts similar to the US regarding latent log wage distributions and how heterogeneous they are across regions.

The first panel in Table 1 shows a case where all regions have the same dispersion parameters

[5] For each state, I calculate the mean log wage and the and standard deviation of log wages for each year. Next, I calculate the means, variances, and pairwise correlations for this four-element vector across states. I use those summary statistics to calibrate the simulations. I use 1989 and 2004 because the real federal minimum wage bottomed out in those years. In addition, unemployment rates are also similar in both years. Thus, the summary statistics based on these two years provide a reasonable approximation for how the latent distribution of log wages varies between states and over time.

Table 1: Effective minimum wage design: good vs. bad variation

	Outcome			
	Emp.	p10 - p50	p25 - p50	p90 - p50
Panel A: Regions differ only in location				
True average causal effect	-0.010	0.019	0.006	-0.004
Effective min. wage	-0.010	0.020	0.006	-0.004
	(0.000)	(0.001)	(0.000)	(0.000)
Panel B: Regions differ in location and dispersion				
True average causal effect	-0.010	0.020	0.006	-0.004
Effective min. wage	-0.007	0.034	0.015	-0.023
	(0.002)	(0.011)	(0.007)	(0.014)
Panel C: As above, but larger increase in min. wage				
True average causal effect	-0.032	0.078	0.017	-0.012
Effective min. wage	-0.014	0.117	0.046	-0.080
	(0.007)	(0.020)	(0.012)	(0.026)
Panel D: St. dev. of dispersion is 50% larger				
True average causal effect	-0.010	0.020	0.006	-0.004
Effective min. wage	-0.003	0.050	0.025	-0.047
	(0.003)	(0.017)	(0.010)	(0.021)

Notes: This table summarizes simulation results with 200 regions and two periods. The top row in each panel reports the average of 1,000 simulations of the true ATE_i for different outcomes i, corresponding to different columns. The second row shows estimated average treatment effects for each outcome based on the effective minimum wage regressions, averaged over the same 1,000 simulations. The third row shows the average over simulations of the corresponding standard errors, which are clustered at the region level in each simulation. The data-generating process includes truncation and censoring effects of the minimum wage. Each panel corresponds to different assumptions on the data-generating process. In **Panel A**, regions differ only in the location parameter $\mu_{r,t}$, with a correlation between the initial and final location of 0.89. **Panel B** includes differences in the dispersion parameter $\sigma_{r,t}$, with a correlation between initial and final dispersion of 0.46. **Panel C** is like Panel B but with an increase of the log minimum wage of 0.4 instead of 0.2. **Panel D** increases the between-region standard deviation of the $\sigma_{r,t}$ parameters by 50%. See Appendix A.1 for details on the calibration of the model.

$\sigma_{r,t} = \sigma_t$, but differ in location parameters—which are subject to changes over time. This model corresponds to the ideal scenario with only "good" variation. Correspondingly, the estimator performs very well. The true average causal effects from the model are nearly identical to the predicted effects from the regressions, averaged over 1,000 simulations (shown in the second row of each panel of the table). In addition, the confidence intervals are very tight, as implied by the standard errors reported in the third row (which are also averaged over the 1,000 simulations).

Panel B introduces differences in dispersion parameters. The data-generating process still satisfies the structural assumptions emphasized by Lee (1999): distributions only differ in location and dispersion parameters, not shape, and the location and dispersion parameters have zero correlation. In addition, the median wage is an excellent proxy for the latent $\mu_{r,t}$ location parameter: Table A1 in Appendix A shows that their correlation is 0.999. Still, the estimator displays economically meaningful biases. They arise because, though the differences between $w_{0.5,r,t}$ and $\mu_{r,t}$ constitute a tiny share of the variation, they are systematically correlated with the outcomes of interest.[6]

Panel C shows that biases are more significant when the simulated increase in the federal minimum wage is 40 log points instead of 20 log points. Even though the correlation between $\mu_{r,t}$ and $w_{r,t}$ remains above 0.99 (see Table A1), measured employment effects are only 40% as large in magnitude as the true causal effects. At the same time, the lower-tail spillovers are amplified by 50%. This exercise reinforces the idea that the "good" identifying variation comes not from the change in the minimum wage itself but from idiosyncratic shocks to $\mu_{r,t}$. Thus, a larger minimum wage shock does not help reduce bias. On the contrary, it amplifies biases because the effects on observed median wages become stronger.

Panel D highlights how differences in the dispersion of latent log wages between regions constitute the source of the bias. Making those differences 50% larger in magnitude while keeping the other parameters constant is enough to essentially double the average bias in the regressions.

Are there biases if average employment effects are zero or positive? To answer this question, I simulate an alternative model where the minimum wage can increase employment levels for individuals with latent wages just above it, corresponding to the green areas in Figure 2. Such effects could emerge from increased search effort by workers (see Adams, Meer and Sloan, 2022, for an empirical evaluation of this hypothesis in the context of the US). The details on that model are provided in Appendix A.1.3.

Results are reported in Appendix Tables A3 and A4. The simulations are analogous to those shown in Table 1, but the average employment effects of the minimum wage are zero (in Table A3) or positive at around one percentage point (in Table A4). These alternative models also display economically significant biases, though they are smaller than in Table 1. This exercise clarifies that the econometric problem arises from idiosyncratic, rather than average,

[6] The direction of the biases suggests that the bias coming from *shocks* to dispersion parameters $\mu_{r,t}$ are the main problem in the simulations, instead of the permanent differences in dispersion. Among the problems illustrated in Figure 1, the one in Panel C seems the most relevant. In unreported simulations, I confirm this intuition, noting that a model that maintains cross-sectional dispersion in $\sigma_{r,t}$ but makes it invariant over time is almost unbiased.

differences between $\mu_{r,t}$ and $w_{r,t}$.

3.3 Issue #2: Correlation between location and dispersion parameters

The second assumption emphasized by Lee (1999) is independence between the location and dispersion parameters, $\mu_{r,t}$ and $\sigma_{r,t}$, conditional on t. For an intuition of why this assumption is essential for measuring spillover effects, consider again the "good variation" example from the previous subsection. In that example, Region A was "treated" by the minimum wage because its location parameters $\mu_{A,t}$ are constant over time, while Region B is the "control" because $\mu_{B,1} - \mu_{B,0} = mw_1 - mw_0$. Now, suppose that, along with the increase in location, the dispersion parameter also increases for Region B. That would increase all quantile gaps $w_{q,B,t} - w_{0.5,B,t}$ in the control region. Thus, comparing changes in treatment versus control regions would no longer provide a valid estimate of the causal effects of the minimum wage.

A correlation between location and dispersion parameters is also problematic if the outcome is employment. The reason is that changes in dispersion parameters can make the minimum wage bind more or less in some regions, causing independent effects on the median wage in the presence of a minimum wage. For example, rising dispersion can add more probability mass in the lower tail of the latent log wage distribution, increasing the amount of truncation and, thus, the mechanical effects of the minimum wage on the median wage. That effect would magnify the correlated measurement error issues discussed in the previous subsection.

Below, I show through simulations that even a mild contemporaneous correlation between location and dispersion parameters can introduce significant biases in the effective minimum wage design. Next, I argue that there are plausible economic reasons why we should expect such correlations to occur.

Panel A in Table 2 shows a baseline scenario where regions differ in dispersion parameters, but dispersion and correlation parameters are uncorrelated. Panel B introduces a within-period correlation of 0.076, the value I find in US data for 1989 (the correlation for 2004 is 0.264). That mild correlation is enough to bring the estimated employment effects to almost zero and make estimated spillover effects much larger than the true ones. Note that this correlation does not significantly affect estimated standard errors; if anything, the estimates become more precise.

The US data also displays intertemporal correlations between location and dispersion.[7] Panel C

[7] Specifically, initial location has a significant correlation with final dispersion, and initial dispersion has a

Table 2: Correlation between location and dispersion parameters

	Outcome			
	Emp.	p10 - p50	p25 - p50	p90 - p50
Panel A: No correlation between location and dispersion				
True average causal effect	-0.010	0.020	0.006	-0.004
Effective min. wage	-0.007	0.033	0.014	-0.022
	(0.002)	(0.011)	(0.007)	(0.014)
Panel B: Contemporaneous correlation of 0.076				
True average causal effect	-0.010	0.020	0.006	-0.004
Effective min. wage	-0.002	0.077	0.040	-0.076
	(0.002)	(0.010)	(0.006)	(0.013)
Panel C: Full correlation matrix in US data				
True average causal effect	-0.010	0.019	0.006	-0.004
Effective min. wage	-0.014	-0.006	-0.009	0.027
	(0.002)	(0.011)	(0.006)	(0.014)

Notes: Each panel displays average results for 1,000 simulations, each with 200 regions and two periods, for different assumptions on the data-generating process (see the notes below Table 1 for an explanation of the table's structure). **Panel A** is identical to Panel B in Table 1: regions differ in location ($\mu_{r,t}$) and dispersion ($\sigma_{r,t}$) parameters, but they are orthogonal to each other. **Panel B** introduces a correlation of 0.076 between location and dispersion parameters within each period. **Panel C** uses the full set of correlations between $[\mu_{r,0}, \sigma_{r,0}, \mu_{r,1}, \sigma_{r,1}]$ observed in US data. See Appendix A.1 for details on the calibration of the model.

includes those correlations in the simulated model. The biases now have the opposite sign. That result shows that it may be difficult to predict the direction of the bias in empirical applications.

Correlations between the location and dispersion of latent log wages may emerge from economic reasons. One is the observation that if workers are split into education-age groups, higher-wage groups tend to display more within-group inequality. This fact is discussed in detail by Lemieux (2006), who argues that much of the increase in inequality observed in the US from 1973 to 2003 is a compositional effect deriving from increased educational achievement. The same result has been found in other contexts, such as Brazil (Ferreira, Firpo and Messina, 2017). Then, if regions differ in workforce composition, the correlation we discussed above may follow.

Education is not the only economic factor that can introduce problems for the effective minimum wage design. Regional differences in endowments, leading to heterogeneity in industrial composition, may also generate a correlation between location and dispersion parameters. That is because industries—or clusters of connected industries that tend to co-

mild correlation with final location.

locate—may differ in wage premiums and the breadth of occupations and skill levels used in production.

3.4 Fixed effects, trends, controls, and confounders

The baseline specification in Lee (1999) does not include region fixed effects. Concerning the inclusion of such fixed effects, he writes: "... *the reduced identifying variation resulting from eliminating the "permanent" state effects may magnify biases due to misspecification, in the same way biases stemming from measurement error in the independent variable are magnified when true variation in the independent variable is reduced.*" Using the language introduced in Subsection 3.2, the estimator without region fixed effects has another source of "good" variation: within-period differences in the location parameters of latent log wage distributions (instead of simply differential shocks to location). That may significantly reduce the influence of "bad" variation coming from correlated measurement error in the centrality measure, reducing the amount of bias.

Table 3 illustrates Lee's argument through simulations. The data-generating process for those simulations is the same as reported in Panel B of Table 1. I report predicted treatment effects using the default effective minimum wage design and two alternative specifications, the first being the estimator without region fixed effects. The comparison of the fifth row to the third shows that, by using more "good" variation coming from level differences in $\mu_{r,t}$, the estimator without region fixed effects can indeed perform better.

Still, it is easy to contemplate omitted variable biases that could cause problems for estimators without region fixed effects. One example would be that unregistered employment, not visible in the data, is more relevant in low-wage areas, generating a spurious negative correlation between measured employment-to-population and the effective minimum wage. Another is regional differences in the supply of skills, corresponding latent log wage distributions with different shapes (and thus different quantile gaps). In this paper, I focus on the version with region fixed effects because those concerns make it more popular in recent literature.

Indeed, the specifications in papers such as Bosch and Manacorda (2010), Autor, Manning and Smith (2016), and Engbom and Moser (2022) go beyond region fixed effects and include region-specific trends as well. These trends may absorb region-specific supply and demand shocks that affect the median wage and the outcomes of interest. One example is changes in educational composition, which, as discussed in the previous subsection, may affect both

Table 3: Effective minimum wage: alternative fixed effects specifications

	Outcome			
	Emp.	p10 - p50	p25 - p50	p90 - p50
True average causal effect	-0.010	0.020	0.006	-0.004
Effective min. wage	-0.007	0.033	0.014	-0.022
	(0.002)	(0.011)	(0.007)	(0.014)
Effective min. wage, no region FE	-0.010	0.022	0.007	-0.007
	(0.001)	(0.004)	(0.002)	(0.005)
Effective min. wage, no time FE	-0.007	0.052	0.024	-0.041
	(0.001)	(0.003)	(0.002)	(0.004)

Notes: Each panel displays average results for 1,000 simulations with 200 regions and two periods. See the notes below Table 1 for an explanation of the table's structure. The data-generating process corresponds to Panel B from Table 1; see Appendix A.1 for details. The table reports predicted average treatment effects for three estimators: the baseline effective minimum wage design and the same design without region effects or without time effects.

the location and the dispersion of latent log wages. Another example is demand-side shocks such as "the China syndrome" (Autor, Dorn and Hanson, 2013), whose wage effects are not uniform over the distribution and whose employment effects may generate mechanical shifts in the median wage.

However, it is not evident that region-specific trends and controls reduce biases; they may instead amplify them. The reason is analogous to Lee's discussion on including region fixed effects. By including the trends, the econometrician may throw out the "good variation" with the bathwater. The fixed effects, region-specific trends, and controls may absorb much of the $\Delta\mu_{r,t}$ shocks, such that the measurement errors become a larger share of the residual variation in the effective minimum wage. In addition, once those controls are included, it may be difficult to interpret where the variation in the effective minimum wage is coming from. This lack of intuition is problematic; ideally, the econometrician should be able to defend the assumption that there exists an economic factor, separate from all trends and controls, that shifts wage levels but does not affect employment levels or the shape of the log wage distribution in any way (other than making the minimum wage more or less binding).

Based on this discussion, one may wonder whether dropping the time effects from the design would add more "good" variation. Lee (1999) explains that this choice is only wise if the econometrician believes that latent log wages' shape and average dispersion do not change over time. It is not warranted in the presence of secular wage trends coming from technical change or international trade, for example.

Table 4: Effective minimum wage using percentile 90 as the deflator

	Outcome			
	Emp.	p10 - p90	p25 - p90	p50 - p90
True average causal effect	-0.010	0.024	0.009	0.000
Effective min. wage, p90	0.009	0.219	0.176	0.000
	(0.002)	(0.013)	(0.011)	(0.000)

Notes: This table has the same structure as Table 3, but reports regression results where the effective minimum wage is calculated based on percentile 90 of the observed log wage distribution.

Table 3 illustrates the sensitivity of that estimator to changes in the economic environment. The baseline data-generating process displays minor differences in the marginal distributions of $\mu_{r,t}$ and $\sigma_{r,t}$ between periods. The most salient differences are that average $\sigma_{r,t}$ falls from 0.54 to 0.51, and the standard deviation of $\sigma_{r,t}$ between regions increases from 0.026 to 0.049. Those small changes are enough to warrant the inclusion of time effects, as biases are much more prominent when the model does not include them.

3.5 Does using a higher quantile as the deflator help?

In some applications, the econometrician may have a strong prior that the minimum wage significantly impacts the median wage, making it a poor measure of centrality. In those cases, they may consider using a higher quantile of the wage distribution to construct the effective minimum wage. For example, Bosch and Manacorda (2010) use quantile 0.7 as the deflator in a study of Mexico, and Engbom and Moser (2022) use quantile 0.9 when studying Brazil.

Lee (1999) argues that the deflator should be a good approximation for centrality $\mu_{r,t}$ instead of merely an overall measure of wages. Otherwise, the regression may yield non-zero estimates even when the observed log wage distribution is identical to the latent wage distribution. The discussion regarding correlated measurement error introduces another reason to be wary of choosing higher quantiles of the wage distribution. While it is true that those higher quantiles may be less affected by the minimum wage, the effects will still not be zero if the minimum wage has employment effects, positive or negative. In addition, higher quantiles are likely to be more sensitive to cross-region differences in the dispersion of latent log wages. Due to those two issues, the biases may be more significant when a quantile other than the median is used as the deflator.

Table 4 evaluates the performance of an estimator based on quantile 90 of the log wage distribution using the baseline scenario with regional differences in location and dispersion parameters (the same from Table 3). The biases are significantly larger than those for other estimators previously discussed. In unreported simulations, I tested that estimator in a broader range of scenarios and found that it consistently underperforms relative to the estimator based on the median.

3.6 Is the standard diagnostic test effective?

Lee (1999) proposes estimating relative effects on high log wage quantiles $q > 0.5$ to validate the model. The justification for that approach is that, in many applications (such as in the US), the econometrician may have a strong prior that the minimum wage should have minimal effects on the upper tail of the wage distribution. Autor, Manning and Smith (2016) use the same specification test to validate their instrumental variables implementation of the effective minimum wage design.

However, the econometrician must know that such a test is subject to both false positives and false negatives. False positives—detecting a problem where none exists—may arise because many plausible mechanisms could lead to minimum wage spillovers that extend beyond the median wage. Engbom and Moser (2022) develop and estimate an on-the-job search model where minimum wages cause spillovers that extend far into the upper tail of the wage distribution, primarily due to worker reallocation from low- to high-wage firms. The model in Haanwinckel (2023) also includes endogenous changes in within-firm returns to skill in response to reallocation flows, firm entry responses, and price effects as mechanisms that can generate spillovers in the upper parts of the wage distribution. Those channels may be quantitatively important even when net disemployment effects are minor, as in Engbom and Moser (2022). Thus, an econometrician with a strict rejection rule based on effects in the upper tail may reject a valid model.

There are two concerns regarding false negatives. One is that the estimator may be biased in the lower tail but not in the upper tail. This may happen if, for example, the negative upper-tail bias illustrated in Table 1 is combined with positive bias arising from measurement error, as discussed by Autor, Manning and Smith (2016). Second, even if spurious upper tail spillovers are detected, the econometrician may still interpret those results as not indicating a problem if they have the prior that such spillovers are economically plausible in the specific application.

3.7 State-level minimum wages and instrumental variables

In a previous subsection, I argued that adding region fixed effects, trends, and controls can absorb much of the "good" variation that could be exploited and thus magnify biases coming from misspecification. One exception to that logic is where the data includes changes in region-specific minimum wage laws. In that scenario, the estimator can exploit variation from those regulatory changes while using a battery of controls and region-specific trends to net out the influence of other factors. Still, the effective minimum wage estimator may remain biased, as it uses both the good variation from state-specific minimum wages and the bad variation induced by measurement error and the residual correlation between location and dispersion parameters.

One may consider an instrumental variables (IV) estimator that isolates that source of good variation. One approach is to use the prevailing institutional minimum wage (and its square) as an instrument for the effective minimum wage (and its square). In their pursuit of an effective minimum wage estimator robust to measurement error, Autor, Manning and Smith (2016, henceforth AMS) propose an IV estimator along those lines but include a third instrument: the interaction of the log minimum wage with the average median wage in each region. Because it uses observed median wages in its construction, this third instrument may be subject to some of the abovementioned concerns.

Table 5 presents the outcomes of simulations that incorporate region-specific minimum wages and implement alternative instrumental variables estimators. As with the previous simulations, the parameters of the data-generating process are tailored to mirror the US context; for more details, refer to Appendix A.1.4. Panel A showcases the baseline model, where there is a slight correlation between location and dispersion parameters (as in Table 2). Panels B and C introduce region-specific minimum wages that surpass the national minimum wage. The distinction between the panels is the proportion of regions with local minimum wages exceeding the national minimum wage. In Panels B and C, I present results not only for the regular effective minimum wage design but also for instrumental variables specifications, with either two or three instruments.

From the table, three key findings emerge. First, the more variation derived from state-level minimum wages, the smaller the biases, even when using the ordinary least squares estimator. This is evident when comparing the "Effective min. wage" rows across panels, which gradually align with the corresponding "Mean causal effect" rows. However, some bias persists. Second, the use of instrumental variables approaches significantly mitigates

Table 5: State-level minimum wages and instrumental variables approaches

	Outcome			
	Emp.	p10 - p50	p25 - p50	p90 - p50
Panel A: No regional variation in minimum wage.				
True average causal effect	-0.010	0.020	0.006	-0.004
Effective min. wage	-0.002	0.076	0.040	-0.075
	(0.002)	(0.010)	(0.006)	(0.013)
Panel B: 20% of regions with local min. wage				
True average causal effect	-0.015	0.035	0.008	-0.005
Effective min. wage	-0.015	0.051	0.015	-0.018
	(0.001)	(0.005)	(0.003)	(0.006)
Two instruments	-0.015	0.036	0.009	-0.006
	(0.002)	(0.006)	(0.004)	(0.008)
Three instruments (AMS)	-0.017	0.042	0.009	-0.005
	(0.002)	(0.005)	(0.003)	(0.006)
Panel C: 40% of regions with local min. wage				
True average causal effect	-0.020	0.052	0.011	-0.007
Effective min. wage	-0.019	0.059	0.015	-0.016
	(0.001)	(0.004)	(0.002)	(0.005)
Two instruments	-0.020	0.051	0.011	-0.007
	(0.002)	(0.004)	(0.002)	(0.006)
Three instruments (AMS)	-0.020	0.052	0.011	-0.007
	(0.001)	(0.004)	(0.002)	(0.005)

Notes: Each panel displays average results for 1,000 simulations, each with 200 regions and two periods, for different assumptions on the data-generating process (see the notes below Table 1 for an explanation of the table's structure). Models in all panels are similar to those from Panel B in Table 2, where there is a small intra-temporal correlation between location ($\mu_{r,t}$) and dispersion ($\sigma_{r,t}$) parameters. Panels B and C introduce region-specific minimum wages. They differ in the share of regions with a local minimum wage higher than the national minimum wage. "Two instruments" corresponds to regressions that employ the nominal minimum wage and its square as instruments for the effective minimum wage and its square. "Three instruments (AMS)" adds a third instrument following Autor, Manning and Smith (2016). See Appendix A.1.4 for details.

this bias. Third, the biases are least pronounced when employing the estimator with two instruments, albeit at the expense of precision.

Therefore, the issues discussed in this section are an additional reason to adopt instrumental variables regressions in the style of AMS when the data includes regional-level variation in minimum wage laws. Such estimators circumvent previously discussed biases by eschewing the potentially endogenous variation from median wages. This section also provides a rationale for avoiding the "interaction" instrument in AMS if the minimum wage instruments alone offer sufficient identifying variation.

3.8 Taking stock

The core message of this discussion is that to evaluate whether the effective minimum wage strategy is likely to be successful, the econometrician should have a clear sense of what constitutes the identifying variation. Exogenous changes in state-level minimum wages are the clearest example of such variation. Because the effective minimum wage estimator might be biased even when that variation is available, instrumental variables approaches such as that in Autor, Manning and Smith (2016) are recommended.

If the data includes little or no variation in state-level minimum wage laws, does the variation in wage levels identify the effects of the national minimum wage? My analysis shows that this is only the case if, conditional on the fixed effects (or trends) used in the regression, differences in median wages come from an underlying structural factor that shifts the location of latent log wage distributions but has no independent effects on shape, dispersion, or employment. If the econometrician does not have an intuitive sense of what that structural factor could be, then the effective minimum wage design may not be warranted.

The existence of this structural factor is necessary but insufficient for unbiased estimation. That's because observed median wages are not a perfect proxy for the centrality of the latent log wage distribution, which gives rise to biases originating from correlated measurement error. Ideally, this underlying shifter of minimum wage bindingness should be observable, in which case the econometrician may use it as an instrument.

4 Fraction Affected and Gap estimators

4.1 Definition

Now, I study a difference-in-differences model with a time-invariant, continuous measure of treatment intensity based on the initial distribution of wages:

$$y_{i,r,t} = \alpha_{i,r} + \delta_{i,t} + \beta_i FA_r \cdot \mathbf{1}\{t=1\} + \varepsilon_{i,r,t} \tag{3}$$

where the subscript i indexes a specific equilibrium outcome, such that different o correspond to separate regressions. The treatment intensity variable FA_r is the "fraction affected," that is, the share of workers in the initial period earning less than mw_1.[8] The regressions include

[8] When the data includes non-compliance with minimum wage regulations, researchers typically define the fraction affected as the share earning between mw_0 and mw_1, but not always (see e.g. Bailey, DiNardo and

region and time fixed effects.

Given the linearity of this model, the estimated average treatment effect of the national minimum wage increase on outcome i is given by the product of the average of FA_r and the estimated $\hat{\beta}_i$ parameter.

Card (1992) first introduced the fraction affected design in an analysis of the 1990 increase in the federal minimum wage in the US. In that paper, he emphasizes that much of the identifying variation in his application comes from significant heterogeneity in the bindingness of state-level minimum wages in the preceding years. Since that original application, that estimator has been applied in other contexts with no regional variation in nominal minimum wages, such as the introduction of a federal minimum wage in Germany in 2015 (Ahlfeldt, Roth and Seidel, 2018; Fedorets and Shupe, 2021).

Identification comes from comparing the evolution of outcomes for "more treated" versus "less treated" units, where the treatment intensity variable only uses information from the initial period. This design is thus fundamentally different from the effective minimum wage one, which, as discussed in the previous section, relies on idiosyncratic shocks to the location parameter of latent log wage distributions (when the regression includes both region and time fixed effects). The core identification assumption is standard for differences-in-differences designs: absent the increase in the national minimum wage, outcomes in treatment and control regions would evolve similarly.

This design is ideal in scenarios where the minimum wage increases after at least a few years without adjustments. In those cases, the econometrician can use pre-treatment data to check for differential trends, which may provide support for the parallel trends identification assumption. In the next subsections, I will discuss the effectiveness of the parallel trends assumption in detecting each of the issues I highlight.

I also study other closely related designs. The main one is based on the "Gap measure:"

$$\begin{aligned} y_{i,r,t} =& \alpha_{i,r}^{Gap} + \delta_{i,t}^{Gap} + \beta_i^{Gap} Gap_r \cdot \mathbf{1}\{t=1\} + \varepsilon_{i,r,t}^{Gap} \\ Gap_r =& \frac{\sum_{j=1}^{J_r} \max\{\exp(mw_1) - \exp(w_{j,0}), 0\}}{\sum_{j=1}^{J_r} \exp(w_{j,0})} \end{aligned}$$

where $j \in \{1, \ldots, J_r\}$ indexes workers in the initial period and $w_{j,0}$ is their log wage in that

Stuart, 2021). In all simulations below, there is perfect compliance, so both approaches are equivalent.

period. Card and Krueger (1994) introduces the gap measure in a firm-level econometric design. It has later been extended to region-level designs like the ones studied in this paper. Dustmann et al. (2021) provides an example, again in the context of Germany. It would correspond to the resulting relative increase in the average wage in each region if all low-wage workers were to receive raises to comply with the new minimum. Other variations, such as transforming the treatment intensity variable into a binary indicator, are also explored as potential ways to solve issues discussed in the following subsection.

4.2 Issue #1: Sensitivity to functional form assumptions

The fraction affected design and its variations rely on specific functional forms that determine the sensitivity of individual regions to the national change in the minimum wage. If that sensitivity measure is misspecified, estimates of average treatment effects based on that estimator may be biased.

To investigate this possibility, I run simulations based on the same model used in the previous section: latent log wages are Normal in each region, and the minimum wage causes truncation, censoring, and potentially positive employment effects for workers a bit above the minimum wage (see Appendix A.1 for details). I consider four data-generating processes encompassing different values for the initial minimum wage and whether it causes positive employment effects. In the scenarios with the higher initial minimum wage, it bites more into the latent distribution after the increase of 20 log points, such that the disemployment and censoring effects become more significant.

To focus on the role of functional form assumptions, I designed those simulations to be ideal applications for the fraction affected and gap designs. First, the national minimum wage increase is the only time-varying factor in the model, preventing violations of the parallel trends assumption. Second, regions only differ in their time-invariant location parameter μ_r. Thus, to the extent that those simulations find misspecification issues, one may expect they would be even more severe if the data also includes heterogeneity in the dispersion and shape of latent log wage distributions.

Table 6 shows that even in this ideal scenario, biases arising from functional form misspecification may be statistically and economically significant. The direction of the bias may change depending on the model used, and in unreported simulations, I find that it also changes with the markdown parameter m. The biases are more significant when the minimum wage is more binding or causes positive employment effects. For example, in Panel C,

Table 6: Misspecification biases in the Fraction Affected and Gap designs

	Outcome				
	Emp.	p10	p25	p50	p90
Panel A: Small initial min. wage, truncation/censoring only					
True average causal effect	-0.006	0.016	0.008	0.004	0.002
Fraction affected	-0.008	0.020	0.010	0.006	0.003
	(0.000)	(0.001)	(0.001)	(0.001)	(0.001)
Gap measure	-0.006	0.015	0.007	0.004	0.002
	(0.000)	(0.001)	(0.001)	(0.001)	(0.001)
Panel B: Large initial min. wage, truncation/censoring only					
True average causal effect	-0.031	0.118	0.036	0.020	0.010
Fraction affected	-0.039	0.185	0.044	0.026	0.013
	(0.000)	(0.009)	(0.001)	(0.001)	(0.001)
Gap measure	-0.028	0.127	0.031	0.019	0.009
	(0.000)	(0.009)	(0.001)	(0.001)	(0.001)
Panel C: Small initial min. wage, positive emp. effects					
True average causal effect	0.010	-0.002	-0.012	-0.006	-0.003
Fraction affected	0.002	0.067	-0.002	-0.001	-0.000
	(0.000)	(0.003)	(0.001)	(0.001)	(0.001)
Gap measure	0.001	0.052	-0.001	-0.001	-0.000
	(0.000)	(0.002)	(0.001)	(0.001)	(0.001)
Panel D: Large initial min. wage, positive emp. effects					
True average causal effect	-0.003	0.149	0.039	0.002	0.001
Fraction affected	-0.038	0.132	0.143	0.026	0.012
	(0.001)	(0.007)	(0.004)	(0.002)	(0.001)
Gap measure	-0.028	0.090	0.102	0.019	0.008
	(0.001)	(0.006)	(0.003)	(0.001)	(0.001)

Notes: In all panels, the national minimum wage increases by 20 log points from the first period to the second. Regions differ only in the time-invariant location parameter $\mu_r \sim \mathcal{N}(0, 0.2^2)$. Each panel displays average results for 1,000 simulations, each with 200 regions. For each outcome, the numbers correspond to the mean true ATE across simulations, the mean estimates of causal effects based on the regressions listed on the left, and the average standard error associated with the estimates (in parentheses, clustered at the region level).

the minimum wage increases employment by one percentage point, but both the fraction affected and gap estimators yield precise estimates that are very close to zero.

Given this particular data-generating process, the measured employment and wage effects tend to be smaller in magnitude when estimated with the Gap measure. If one is exclusively interested in the ratio of employment effects to wage effects in the lower tail (proxying for the employment elasticity with respect to the worker's wage), then the estimators are remarkably similar to each other across panels. However, the estimated ratio is generally different from the true one.

In the remainder of this subsection, I consider three alternative econometric specifications and show that neither of them adequately solves the misspecification biases.

Binary design: Recent papers have explored a binary version of the fraction affected intensity measure (Derenoncourt et al., 2021; Parente, 2024). One may wonder whether such a strategy can help with misspecification issues. Callaway, Goodman-Bacon and Sant'Anna (2024) argue that, in difference-in-differences designs with continuous treatment variables, a binary design has a more straightforward interpretation when the control group in the binary design is composed of entirely untreated units. However, this is not the case for the data-generating processes studied here.

In Appendix Table A5, I show that a binary definition of treatment suffers from more substantial biases than the continuous versions. I split regions into treatment or control groups based on whether the initial median wages are below a given threshold. I choose thresholds such that either half or 90% of the regions are in the treatment group. Consistent with the logic of Callaway, Goodman-Bacon and Sant'Anna (2024), biases are smaller when 90% of the sample is in the treatment group since it makes the "zero" group closer to being entirely untreated. However, there is a precision loss, and a significant bias remains.

Instrumental variables approach: Another potential solution for the misspecification biases discussed above would be to use the regressor in the Fraction Affected design as an instrument for the regressor based on the Gap measure, or vice-versa. This strategy could be justified if each of those regressors were equal to an unobserved metric of propensity to be affected by the minimum wage plus some random noise and if those noise terms are uncorrelated with one another. Appendix Table A6 reports the results of using such strategies. They have no impact on the estimates compared to the basic OLS estimator.

Quadratic specification: One potential source of misspecification biases is failing to account for the fact that minimum wage effects can be very heterogeneous depending on how binding it is. Appendix Table A7 reports results for models that include a quadratic term in the regressions to account for such heterogeneity. Biases become smaller for some outcomes but larger for others. Thus, this approach is not a practical solution for the misspecification problems either.

4.3 Issue #2: Regression to the mean at the regional level

The fraction affected and gap measures are constructed based on extreme wage observations, in the sense that individual workers only contribute to those measures if their wages are below some threshold. Thus, these estimators may be subject to bias emerging from regression to the mean. This issue is well-known in minimum wage studies at the individual worker or firm levels. For example, in their worker-level analysis, Dustmann et al. (2021) use data from before the minimum wage was implemented to control for regression to the mean. However, that issue is typically not discussed in regional-level studies.

One potential source of regression to the mean at the regional level is sampling error. A region may have a high fraction affected because of an "unlucky" draw of workers in the survey in the year used to construct the treatment intensity variable. The ensuing bias is likely to be negligible with large sample sizes. However, one must be aware of this possibility in studies that define regions at a fine geographical level, especially if regions with the smallest samples also have low average wages.

Time-varying structural factors that determine regional wages may also introduce reversion to the mean, even when samples are large. Caliendo et al. (2017) document that regional-level productivity shocks are quantitatively significant in the United States. Gennaioli et al. (2014) collect time-series data on regional GDP for 83 countries and document within-country regional convergence. Their results mean that, in general, regions that have particularly low GDP per capita in a given period are likely to have stronger growth ex-post. Since these regional productivity shocks may affect both wages and employment, potential biases are not limited to regressions where wages are the dependent variable. That kind of regression to the mean will likely be more consequential for longer-run specifications.

The comparison between Panels A and B in Table A8 illustrates this issue in the context of the Normal-markdown model previously used in this paper. It reports results for the Gap design; Table A8 in Appendix B shows similar results for the Fraction Affected design. In Panel A, regions only differ in a time-invariant location parameter μ_r. Panel B introduces

Table 7: Sensitivity of the Gap design

	Outcome				
	Emp.	p10	p25	p50	p90
Panel A: Only permanent differences in location					
True average causal effect	-0.010	0.026	0.012	0.007	0.003
Gap measure	-0.009	0.027	0.011	0.006	0.003
	(0.000)	(0.002)	(0.001)	(0.001)	(0.001)
Panel B: Adding location shocks, stable distributions					
True average causal effect	-0.010	0.026	0.012	0.007	0.003
Gap measure	-0.007	0.043	0.030	0.026	0.023
	(0.001)	(0.005)	(0.006)	(0.006)	(0.006)
Panel C: Adding dispersion differences and shocks, stable distributions					
True average causal effect	-0.010	0.026	0.013	0.007	0.003
Gap measure	-0.007	0.052	0.034	0.024	0.009
	(0.001)	(0.006)	(0.006)	(0.006)	(0.009)
Panel D: Average dispersion falls over time					
True average causal effect	-0.010	0.027	0.013	0.007	0.003
Gap measure	-0.004	0.046	0.032	0.023	0.009
	(0.001)	(0.006)	(0.006)	(0.006)	(0.009)

Notes: All panels illustrate scenarios where the only time-varying factor is an increase in the national minimum wage of 20 log points. Each panel displays average results for 1,000 simulations with 200 regions and two periods. For each outcome, the numbers correspond to the mean true ATE across simulations, the mean estimates of causal effects based on the regressions listed on the left, and the average standard error associated with the estimates (in parentheses, clustered at the region level). See the notes for Tables 1 and 2 for a description of Panels A and B.

time-specific location parameters $\mu_{r,t}$, with each region's initial and final parameters being jointly Normal with a correlation of 0.894 (equal to the correlation between state-level mean log wages in the US for 1989 and 2004). As expected, the estimated wage effects become substantially more positive.

Panel C further to explore this issue by including time-varying heterogeneity in dispersion parameters $\sigma_{r,t}$ between regions. Those parameters are assumed to be independent of the location parameters but have an autocorrelation of 0.456 between periods. That magnifies the positive bias in the lower tail, though it reduces it in the upper tail.

Fortunately, biases arising from regression to the mean can be detected with tests for differential pre-trends if the context allows such tests. Appendix Table A9 illustrates this concept using a placebo exercise that parallels Table 7. The Gap measure is calculated as if the national minimum wage would increase by 0.2, but there is no increase from period 0 to 1. That placebo exercise shows positive employment effects that have about the same size as the biases discussed above.

For the pre-trends test to work well in detecting that issue, the econometrician needs to be careful in how to define the treatment intensity variable. As an example, consider Dustmann et al. (2021), who calculate the regional Gap measure for each year in the period before the new minimum wage and then use the average for all of those years as the main regressor. Such a definition may help with precision and reduce regression to the mean originating from sampling error. However, it does not solve regression to the mean originating from region-level shocks. More importantly, constructing the treatment variable in that way prevents regression to the mean from being detected in tests for differential pre-trends. Thus, using a single pre-period year is preferable for diagnostics purposes. Ideally, if sufficient pre-treatment data is available, the econometrician should directly control for regression to the mean; see the individual-level design in Dustmann et al. (2021) for details.

Another implication of regression to the mean is that using region-specific linear trends to control for deviations from the parallel trends assumption may not be a valid strategy. The reason is that, in an "event study" graph, regression to the mean implies a V-shaped pattern with the bottom located in the year used to construct the treatment variable. Thus, if one were to extrapolate the pre-trends into the post-period and use it as the counterfactual, one would increase the bias in the estimated treatment effects instead of attenuate it.[9]

4.4 Issue #3: Trends in the dispersion of latent wages

Suppose that, in all regions, there is a change in the dispersion of latent log wages occurring at the same time as the increase in the national minimum wage. Since that structural shock is the same in all regions, one may expect the time effects included in the regression to account for it. This subsection shows that this may not be the case when the minimum wage is not zero in the initial period.

Consider the model illustrated in Figure 3. Regions differ in a location parameter μ_r constant over time. The shape parameter σ_t is the same in all regions and decreases over time. All else equal, regions with lower μ_r have a higher fraction affected and a higher causal effect of the minimum wage. However, with the fall in σ_t, these regions may see a relative *increase* in employment. That is because the truncation effects of the minimum wage may decrease

[9] The recommendation to avoid region-specific trends was first given by Meer and West (2016), based on the fact that the minimum wage may cause changes in employment trends rather than a step change in employment levels.

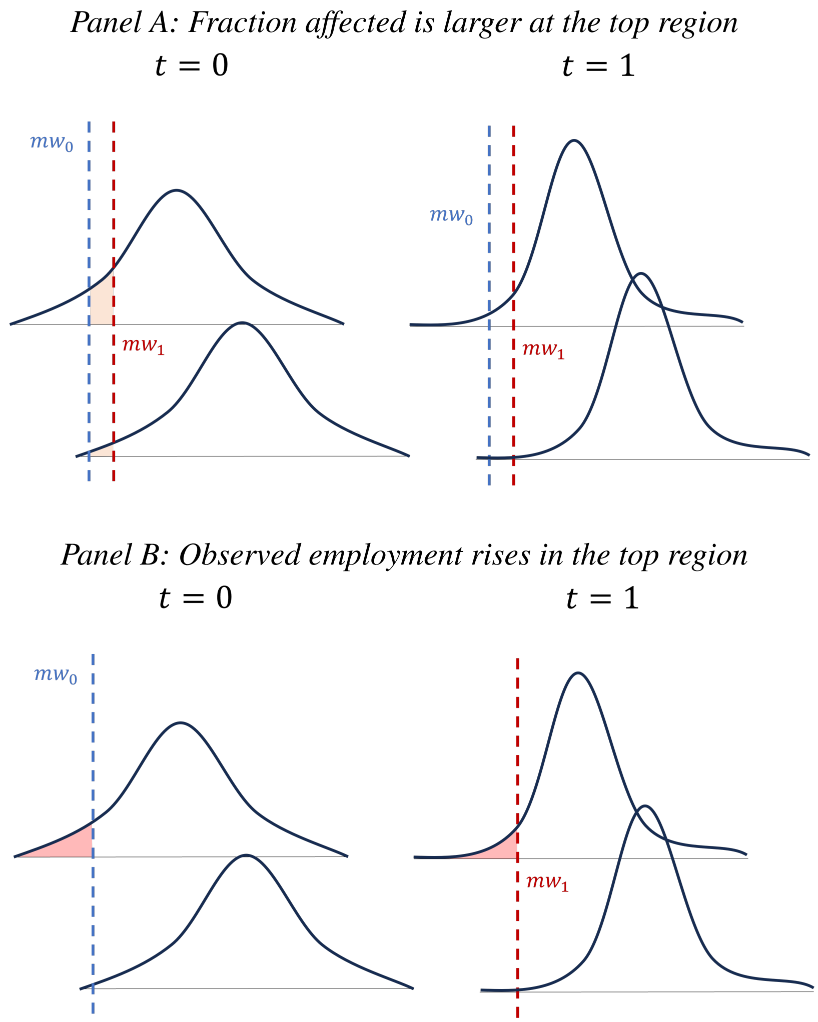

Figure 3: A truncation model with a secular decline in latent log wage dispersion

as latent wages become more concentrated around the median. As a result, the estimator may recover a positive relationship between the fraction affected and employment despite the causal effects being negative for all observations in all periods.

The comparison between Panels C and D in Table 7 illustrates this issue. Panel C was discussed in the previous subsection: it features idiosyncratic shocks to both location and dispersion parameters, but the distributions of those parameters are stable over time. Panel D differs from it by reducing the average dispersion parameter $\sigma_{r,t}$, from 0.54 in $t = 0$ to 0.51 in $t = 1$.[10] This small change is enough to reduce the measured employment effects by almost

[10] Those numbers are based on the previously discussed calibration based on state-level statistics from CPS data for 1989 and 2004. In the context of the US, the decline in dispersion could be partly explained by the higher prevalence of state-level minimum wages in 2004 compared to 1989, in which case it should not

half, making them statistically different from the true causal effects on average.

Tests for differential pre-trends can detect this issue if the trend in dispersion starts before the minimum wage change and has a roughly constant effect on employment over time. The placebo regressions reported in Appendix Table A9 illustrate that possibility. The comments made in the previous subsection apply: for those tests to be most effective, the treatment intensity variable should be calculated based on a single year instead of averaged over all pre-treatment years. In addition, one must be aware of the possibility that pre-trends look flat because they combine the effect of regression to the mean (discussed in the previous subsection) and other secular trends that have a stronger effect on treated regions.

4.5 Taking stock

I showed that the Fraction Affected and Gap designs explore a fundamentally different and more transparent source of variation compared to the effective minimum wage design. It relies on the well-understood parallel trends assumption. Part of my contribution is showing that some factors that do not seem problematic, such as national trends in the dispersion of latent log wages or region-level productivity shocks, can imply violations of the parallel trends assumption and cause quantitatively significant biases. That makes it all the more important to report tests for differential pre-trends. Those tests are most effective when the treatment variable is constructed using only one pre-treatment period.

I also showed that because those estimators fundamentally rely on the functional form used to construct the treatment variable, they are sensitive to model misspecification. Practitioners are encouraged to report predicted employment and wage effects from both the Fraction Affected and Gap designs to assess the possibility of misspecification biases. However, readers should be made aware that there is no theoretical guarantee that the true average treatment effect lies between those estimates. In addition, I showed that binarized versions of those estimators should be avoided as they display more significant misspecification biases.

5 Robustness to alternative data generating processes

One may wonder whether the conclusions from the previous sections were sensitive to the data-generating process used in the simulations. In this section, I evaluate the performance

be interpreted as a change in *latent* log wages. But as mentioned before, this exercise aims to illustrate an econometric issue using reasonable numbers, not evaluate the effects of the national minimum wage in the US.

of the estimators using two alternative frameworks: the canonical model of labor demand and the task-based monopsonistic model with firm heterogeneity from Haanwinckel (2023).

5.1 The canonical model of labor demand

Consider a competitive economy with two types of labor: low- or high-skilled. Each worker is characterized not only by their type but also by their amount of efficiency units of labor. The skill wage premium is pinned down by the ratio of marginal products of labor between the two labor types. Marginal products of labor are, in turn, determined by a representative production function with constant returns and constant elasticity of substitution. This setup mirrors the classic model of Katz and Murphy (1992).

Now, consider the implications of a binding minimum wage in that model. The representative firm will not employ workers whose productivity is below the minimum wage. Thus, one can think of this model as one where each worker has a latent log wage given by the sum of the log price of the efficiency unit of their type (low- or high-skill) and the log of their amount of efficiency units. The observed wage distribution is the truncated version of the latent one. See Appendix A.2 for details on the model.

The critical difference between this model and the one presented before is that the shape of the latent log wage distribution responds to the minimum wage. As the minimum wage causes more disemployment for low-skill workers, returns to skill are expected to fall. These price responses generate wage spillovers for other low-skill workers and attenuate the minimum wage's disemployment effects.

Appendix Table A10 reports simulations that assess the effectiveness of the Fraction Affected and Gap designs using the Canonical model as the data-generating process. In those simulations, regions only differ in the initial share of workers in the high-skill group, and the only time-varying factor is the minimum wage. In that sense, they parallel Table 6 above, illustrating ideal scenarios for the Fraction Affected and Gap designs. Also, similar to that previously discussed table, I report results for different model specifications, varying the bindingness of the initial minimum wage and the elasticity of substitution between low- and high-skill workers.

The takeaways from that exercise are the same as those from Table 6. The predicted average causal effects estimated using those approaches do not always equal the true ones, and those misspecification biases are more significant when the minimum wage is more binding.

Appendix Table A11 evaluates the effective minimum wage design using the same simulations. The baseline effective minimum wage specification displays severe biases when both region and time fixed effects are included. That is because the model does not include region-specific wage shocks, the source of "good" variation in that design. For that reason, I also report estimates for the model without region fixed effects (the baseline specification in Lee, 1999), which exploit differences in wage levels instead of changes. The model is almost ideal for that design in that there are no sources of "endogeneity" (like systematically lower latent employment levels in low-wage regions) and also in that the dispersion of efficiency units for both low- and high-skill workers are the same (which reduces the correlation between "location" and "dispersion" parameters of latent log wages). Even so, that model displays significant biases in some specifications.

5.2 A task-based, monopsonistic model with firm heterogeneity

The previous simulation exercises have two important limitations. First, the causal channels from the minimum wage to employment and wages are somewhat limited, excluding margins that have been shown to matter, such as the reallocation of workers from low- to high-wage jobs (Dustmann et al., 2021). Second, they relied on simplistic—and thus potentially inaccurate—functional form assumptions on the shape of the underlying worker productivity distribution and how it varies across regions.

To address those limitations, I use the model estimated in Haanwinckel (2023) as the data-generating process. In that model, workers belong to one of ten comparative advantage types and differ in the number of efficiency units of labor within groups. The quantity of workers of each type is associated with the educational composition of the region at each period. Firms differ in their demand for skills because they use task-based production technologies, and some firms may have a higher demand for high-complexity tasks that skilled workers perform better. Firms also differ in the wages they pay to each worker type; that is, the model features firm wage premiums. Minimum wages have a range of effects, including disemployment (truncation), mechanical wage increases (censoring), between-firm reallocation, changes in returns to skill within firms, in the distribution of firm types (that is, the share of high-wage, high-skill firms versus low-wage, low skill ones), in participation decisions of workers, and in prices for goods. Regions differ not only in educational composition but also in total factor productivity (TFP), three types of labor demand parameters (which may correlate with TFP and initial educational composition), and some parameters regulating employment

Table 8: Task-based, monopsonistic model with two-sided heterogeneity

Panel A: Fraction Affected and Gap Designs

	Outcome				
	Emp.	p10	p25	p50	p90
True average causal effect	-0.046	0.318	0.197	0.120	0.107
Fraction affected	-0.021	0.232	0.278	0.325	0.175
	(0.007)	(0.012)	(0.015)	(0.015)	(0.026)
Gap measure	-0.012	0.103	0.123	0.153	0.085
	(0.004)	(0.017)	(0.019)	(0.021)	(0.014)
Binary measure, 50% treated	-0.003	0.092	0.113	0.135	0.072
	(0.003)	(0.006)	(0.008)	(0.009)	(0.012)
Binary measure, 90% treated	0.002	0.227	0.260	0.271	0.102
	(0.008)	(0.021)	(0.025)	(0.029)	(0.032)

Panel B: Effective Minimum Wage Designs

	Outcome			
	Emp.	p10 - p50	p25 - p50	p90 - p50
True average causal effect	-0.046	0.198	0.077	-0.013
Effective min. wage	-0.015	0.218	0.122	0.070
	(0.012)	(0.011)	(0.013)	(0.037)
Effective min. wage, no region FE	-0.073	0.196	0.088	-0.016
	(0.006)	(0.005)	(0.006)	(0.014)
Effective min. wage, no time FE	0.113	0.212	0.121	-0.139
	(0.004)	(0.003)	(0.004)	(0.012)
AMS, no time FE	0.125	0.211	0.121	-0.159
	(0.005)	(0.003)	(0.003)	(0.011)

Panel C: Effective Minimum Wage based on percentile 90

	Outcome			
	Emp.	p10 - p90	p25 - p90	p50 - p90
True average causal effect	-0.046	0.211	0.090	0.000
Effective min. wage, p90	0.025	0.384	0.306	0.000
	(0.012)	(0.015)	(0.021)	(0.000)

Notes: This table compares the true average causal effects of the minimum wage from the model of Haanwinckel (2023) to the predicted average treatment effects from different estimators. The model has 151 regions (corresponding to microregions in Brazil) and two time periods (corresponding to 1998 and 2012). Estimated standard errors (clustered at the microregion level) are shown in parentheses.

rates (which capture unobserved heterogeneity related to the Brazilian informal sector).

Haanwinckel (2023) estimates that model using Brazilian matched employer-employee data, targeting a large set of moments at the region-time levels. There are 151 regions (corresponding to *microregions* as defined by the Brazilian Statistical Bureau, IBGE) and two time periods (corresponding to 1998 and 2012). Appendix Table A12 shows that the estimated model fits several dimensions of the data pretty well: inequality measures between and within educational groups, the contribution of firm wage premiums to inequality (measured using the methodology developed by Kline, Saggio and Sølvsten, 2018), the share of workers earning up to 30 log points of the minimum wage, and the minimum wage "spike" (measured as the share of workers earning up to 5 log points of the minimum wage).

Table 8 evaluates the performance of the estimators discussed in this paper when the fitted model of Haanwinckel (2023) is used as data. Each row labeled as an estimator in that table reports the predicted average treatment effects on a range of outcomes using the results of that single regression. The table also shows the corresponding standard errors in parenthesis. At the top row of each panel, I report the true average causal effect of the national minimum wage predicted by the fitted model, following the definition from Section 2.

The overall result is that all of the estimators display significant biases. With one exception, they all underestimate disemployment effects by at least 50%. Some versions of the effective minimum wage design predict positive, instead of negative, employment effects. Most estimators also fail to accurately capture the spillover effects on the wage distribution, though the effective minimum wage estimator without region fixed effects performs surprisingly well in that regard.

6 Conclusion

In this paper, I analyzed the performance of two classes of estimators of the employment and wage effects of minimum wages in contexts where that policy is set at the national level. I discussed the key source of identifying variation for each of them, showed that identification assumptions required for unbiased estimation are stronger than what existing literature documents, and discussed potential solutions via adjustments of the estimation procedures.

The main takeaway of my analysis is that if the data in a specific application includes several pre-treatment periods when the minimum wage was constant, then the "fraction affected" or "gap" designs should be the preferred choice. In that case, the specific recommendations coming from this paper are (i) construct the treatment intensity variable using only one pre-treatment year rather than averaging over years; (ii) check for differential pre-trends; (iii) control for regression to the mean using the procedure in Dustmann et al. (2021); (iv) check for trends in the dispersion of log wages in the pre-treatment period, which may cause bias if the initial minimum wage is not zero; and (v) do not use a binary version of the treatment intensity variable. Even if the econometrician takes such precautions, the estimator may be subject to significant misspecification biases arising from functional form assumptions, especially when the minimum wage increase is significant or generates positive employment effects. At a minimum, econometricians may consider reporting results of both the "fraction affected" and "gap" designs to partially assess the relevance of this problem in a particular application.

If the empirical context does not feature a stable pre-treatment period, then the validity of estimates from the "fraction affected" and "gap" designs cannot be evaluated. In such cases, one may consider the effective minimum wage design. However, my analysis shows that this estimator relies on assumptions unlikely to hold in any application without regional differences in minimum wage regulation. Among variations of the effective minimum wage design, the best-performing one includes both region and time fixed effects and constructs the effective minimum wage using the median wage (as opposed to a higher wage quantile). However, even in these cases, biases can be considerable. The effective minimum wage design should only be trusted if the econometrician believes there exists an economic factor that increases the bite of the minimum wage in some regions compared to others but that does not affect the shape of the wage distribution or employment levels in any other way (conditional on the controls included in the design). If such a variable is available, the econometrician should consider an instrumental variables design that directly exploits it as an instrument.

In contexts that are not ideal for the fraction affected/gap designs due to a lack of a pre-period with stable minimum wages and where there is no quasi-experimental shifter of minimum wage bindingness at the regional level, the econometrician may consider two alternative strategies. One is to use within-region difference-in-differences designs where the unit of analysis is either firms (Harasztosi and Lindner, 2019) or workers (Dustmann et al., 2021), and treatment and control groups are defined based on initial wages. In that case, the econometrician should be mindful that some concerns discussed above, such as regression to the mean, may still apply. In addition, careful interpretation of results is warranted; potential effects on firm or worker entry may be undetectable, and the control group may also be affected by the shock even if general equilibrium effects are believed to be small (for example, high-wage firms may be affected by the minimum wage if it induces worker reallocation, as in Dustmann et al., 2021).

The second alternative strategy is to estimate a parametric economic model; see, for example, Engbom and Moser (2022) and Haanwinckel (2023) in the Brazilian context. Such models can use information from the data to quantify and correct potential sources of bias. Those solutions come at the cost of higher complexity and the need to pre-specify the causal pathways through which the minimum wage affects the economy. A promising direction for further research is developing an econometric model that is simple to implement and agnostic about economic channels but adequately controls for data features that are problematic for the designs studied in this paper.

References

Adams, Camilla, Jonathan Meer, and CarlyWill Sloan. 2022. "The minimum wage and search effort." *Economics Letters*, 212: 110288.

Ahlfeldt, Gabriel M., Duncan Roth, and Tobias Seidel. 2018. "The regional effects of Germany's national minimum wage." *Economics Letters*, 172: 127–130.

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.

Autor, David H., David Dorn, and Gordon H. Hanson. 2013. "The China Syndrome: Local Labor Market Effects of Import Competition in the United States." *American Economic Review*, 103(6): 2121–68.

Bailey, Martha J., John DiNardo, and Bryan A. Stuart. 2021. "The Economic Impact of a High National Minimum Wage: Evidence from the 1966 Fair Labor Standards Act." *Journal of Labor Economics*, 39(S2): S329–S367.

Baker, Michael, Dwayne Benjamin, and Shuchita Stanger. 1999. "The Highs and Lows of the Minimum Wage Effect: A Time-Series Cross-Section Study of the Canadian Law." *Journal of Labor Economics*, 17(2): 318–350.

Bosch, Mariano, and Marco Manacorda. 2010. "Minimum Wages and Earnings Inequality in Urban Mexico." *American Economic Journal: Applied Economics*, 2(4): 128–49.

Caliendo, Lorenzo, Fernando Parro, Esteban Rossi-Hansberg, and Pierre-Daniel Sarte. 2017. "The Impact of Regional and Sectoral Productivity Changes on the U.S. Economy." *The Review of Economic Studies*, 85(4): 2042–2096.

Callaway, Brantly, Andrew Goodman-Bacon, and Pedro H. C. Sant'Anna. 2024. "Difference-in-Differences with a Continuous Treatment." arXiv:2107.02637.

Card, David. 1992. "Using Regional Variation in Wages to Measure the Effects of the Federal Minimum Wage." *Industrial and Labor Relations Review*, 46(1): 22–37.

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

de Chaisemartin, Clément, and Xavier D'Haultfœuille. 2020. "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects." *American Economic Review*, 110(9): 2964–96.

Derenoncourt, Ellora, François Gérard, Lorenzo Lagos, and Claire Montialoux. 2021. "Racial Inequality, Minimum Wage Spillovers, and the Informal Sector." Unpublished manuscript.

Dinkelman, Taryn, and Vimal Ranchhod. 2012. "Evidence on the impact of minimum wage laws in an informal sector: Domestic workers in South Africa." *Journal of Development Economics*, 99(1): 27–45.

Dustmann, Christian, Attila Lindner, Uta Schönberg, Matthias Umkehrer, and Philipp vom Berge. 2021. "Reallocation Effects of the Minimum Wage*." *The Quarterly Journal of Economics*, 137(1): 267–328.

Engbom, Niklas, and Christian Moser. 2022. "Earnings Inequality and the Minimum Wage: Evidence from Brazil." *American Economic Review*, 112(12): 3803–47.

Fedorets, Alexandra, and Cortnie Shupe. 2021. "Great expectations: Reservation wages and minimum wage reform." *Journal of Economic Behavior and Organization*, 183: 397–419.

Ferreira, Francisco H. G., Sergio P. Firpo, and Julián Messina. 2017. "Ageing Poorly?: Accounting for the Decline in Earnings Inequality in Brazil, 1995-2012." Inter-American Development Bank IDB Publications (Working Papers) 8220.

Gennaioli, Nicola, Rafael La Porta, Florencio Lopez De Silanes, and Andrei Shleifer. 2014. "Growth in regions." *Journal of Economic Growth*, 19(3): 259–309.

Haanwinckel, Daniel. 2023. "Supply, Demand, Institutions, and Firms: A Theory of Labor Market Sorting and the Wage Distribution." NBER Working Paper 31318.

Harasztosi, Peter, and Attila Lindner. 2019. "Who Pays for the Minimum Wage?" *American Economic Review*, 109(8): 2693–2727.

Huang, Chen. 2019. "Essays in Labor Economics." Doctoral dissertation, University of Arizona.

Katz, Lawrence F, and Kevin M Murphy. 1992. "Changes in Relative Wages, 1963-1987: Supply and Demand Factors." *The Quarterly Journal of Economics*, 107(1): 35–78.

Kline, P., R. Saggio, and M. Sølvsten. 2018. "Leave-out estimation of variance components." *ArXiv e-prints*.

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

Lemieux, Thomas. 2006. "Increasing Residual Wage Inequality: Composition Effects, Noisy Data, or Rising Demand for Skill?" *American Economic Review*, 96(3): 461–498.

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

Neumark, David, and William Wascher. 1992. "Employment Effects of Minimum and Subminimum Wages: Panel Data on State Minimum Wage Laws." *Industrial and Labor Relations Review*, 46(1): 55–81.

Parente, Rafael. 2024. "Minimum Wages, Inequality, and the Informal Sector." Unpublished manuscript.

Roth, Jonathan, Pedro H.C. Sant'Anna, Alyssa Bilinski, and John Poe. 2023. "What's trending in difference-in-differences? A synthesis of the recent econometrics literature." *Journal of Econometrics*, 235(2): 2218–2244.

Sorkin, Isaac. 2015. "Are there long-run effects of the minimum wage?" *Review of Economic Dynamics*, 18(2): 306–333.

Vogel, Jonathan. 2023. "The Race Between Education, Technology, and the Minimum Wage." NBER Working Paper 31028.

Appendix

A Details on data-generating processes for simulations

This appendix lists the parameters used in all simulations. Every simulation exercise is repeated 1,000 times. The increase in the log minimum wage is always 0.2, except when otherwise noted.

A.1 Normal-markdown model

A.1.1 Model description

Each region r and each time t has a Normal distribution of latent log wages $G^*_{r,t}(w^*) = \Phi\left(\frac{w^*-\mu_{r,t}}{\sigma_{r,t}}\right)$. The employment-to-population ratio and the distribution of observed wages depend on latent wages, the level of the national minimum wage, and a "markdown" parameter $m \in (0,1]$ as follows:

$$
\begin{aligned}
emp_{r,t} =& 1-\Phi\left(\frac{mw_t-\log m-\mu_{r,t}}{\sigma_{r,t}}\right) \\
G_{r,t}(w) =& \frac{\Phi\left(\frac{w-\mu_{r,t}}{\sigma_{r,t}}\right)-\Phi\left(\frac{mw_t-\log m-\mu_{r,t}}{\sigma_{r,t}}\right)}{1-\Phi\left(\frac{mw_t-\log m-\mu_{r,t}}{\sigma_{r,t}}\right)} \qquad \text{for } w \geq mw_t
\end{aligned}
$$

This model generates both truncation and censoring of the latent wage distribution. Workers whose latent log wages are below the minimum minus the log markdown become disemployed. For those with latent log wages above the log minimum wage, the observed wage is equal to the latent wage. Finally, those who remain employed but have latent log wages below the log minimum wage see a mechanical increase in their wage. The latter group corresponds to the minimum wage "spike" in the log wage distribution.

The model can be understood as reflecting an economy with an inelastic labor supply, exogenous worker productivities, and identical monopsonistic firms paying wages that are below the marginal products of labor unless mandated to pay higher wages via the minimum wage. When the markdown m is low, disemployment effects are smaller, and positive effects on wages are bigger. Unless otherwise noted, I use $m = 0.7$.

A.1.2 Calibration

The meta-parameters governing the distribution of region-specific parameters $[\mu_{r,0}, \sigma_{r,0}, \mu_{r,1}, \sigma_{r,1}]$ are based on data from the US Current Population Survey for 1989 (corresponding to period $t = 1$) and 2004 (corresponding to $t = 0$). I chose those years because the national minimum wage was small and approximately the same, in real terms, in both years and the unemployment rate was also approximately equal.

The data was processed using the same procedures as in Lemieux (2006). The sample is restricted to workers between 16 and 64 years of age, with positive potential experience, and whose wages and worked hours are reported by the respondent instead of inferred. Top-coded earnings are adjusted by a factor of 1.4.

Using this sample, I calculate the mean and standard deviation of log wages in each combination of state and year, weighting by the CPS sampling weights and worker hours. Then, I de-mean the $\mu_{r,t}$ elements using simple averages within the period so that the $\mu_{r,t}$ are mean zero in both periods. I treat those statistics as corresponding to the $[\mu_{r,0}, \sigma_{r,0}, \mu_{r,1}, \sigma_{r,1}]$ vector for each state. Thus, I calculate the corresponding covariance matrix of that vector and use it to calibrate the simulation models.

Finally, I calibrate the simulations using the estimated vector of means and covariance matrix. As stated in the main text, in each simulation, the vectors $[\mu_{r,0}, \sigma_{r,0}, \mu_{r,1}, \sigma_{r,1}]$ for each region r are drawn from a Multivariate Normal distribution. The parameters for that meta-distribution are created by either "shutting down" some of the correlations in the estimated covariance matrix, eliminating differences in dispersion parameters, increasing the correlation between some initial and final region parameters to one (to impose that those parameters are time-invariant), or averaging some meta-parameters between both periods so that the distributions are stable over time. Tables A1 and A2 report the meta-parameters used in every simulation exercise with the Normal-markdown model.

A.1.3 Positive employment effects

For some simulation exercises, I augment the Normal-markdown model to include the possibility of positive employment effects. I add two parameters to the model: P_{base} and P_{height}. The total employment mass added to the model is equal to $\frac{P_{base}P_{height}}{2}\phi\left(\frac{mw_t-\mu_{r,t}}{\sigma_{r,t}}\right)$, where the latter term corresponds to the density of the latent log wage distribution evaluated at the point where the minimum wage binds. The wage distribution for that extra mass is triangular, with

Table A1: Simulation meta-parameters: Normal-markdown model, effective minimum wage design

		Min. Wage		Means			Std. Dev.			Correlations						Corr.
Model	m	mw_0	mw_1	$\sigma_{r,0}$	$\sigma_{r,1}$	μ_0	σ_0	μ_1	σ_1	μ_0,σ_0	μ_0,μ_1	μ_0,σ_1	σ_0,μ_1	σ_0,σ_1	μ_1,σ_1	$\mu_{r,t}, w_{0.5,r,t}$
Tables 1, A3, and A4																
Panel A	0.7	-1.0	-0.8	0.542	0.510	0.123	0.000	0.112	0.000	0.000	0.894	0.000	0.000	0.000	0.000	1.000
Panel B	0.7	-1.0	-0.8	0.542	0.510	0.123	0.026	0.112	0.049	0.000	0.894	0.000	0.000	0.456	0.000	0.999
Panel C	0.7	-1.0	-0.6	0.542	0.510	0.123	0.026	0.112	0.049	0.000	0.894	0.000	0.000	0.456	0.000	0.994
Panel D	0.7	-1.0	-0.8	0.542	0.510	0.123	0.039	0.112	0.074	0.000	0.894	0.000	0.000	0.456	0.000	0.998
Table 2																
Panel A	0.7	-1.0	-0.8	0.542	0.510	0.123	0.026	0.112	0.049	0.000	0.894	0.000	0.000	0.456	0.000	0.999
Panel B	0.7	-1.0	-0.8	0.542	0.510	0.123	0.026	0.112	0.049	0.076	0.894	0.000	0.000	0.456	0.076	0.999
Panel C	0.7	-1.0	-0.8	0.542	0.510	0.123	0.026	0.112	0.049	0.076	0.894	0.366	0.063	0.456	0.264	0.999
Tables 3 and 4																
All panels	0.7	-1.0	-0.8	0.542	0.510	0.123	0.026	0.112	0.049	0.000	0.894	0.000	0.000	0.456	0.000	0.999
Table 5																
All panels	0.7	-1.0	-0.8	0.542	0.510	0.123	0.026	0.112	0.049	0.076	0.894	0.000	0.000	0.456	0.076	0.999

Notes: This table shows meta-parameters used to draw simulation samples for the Normal-markdown model. See Appendix A.1 for details.

Table A2: Simulation meta-parameters: Normal-markdown model, Fraction Affected/Gap designs

		Min. Wage		Means			Std. Dev.			Correlations						Corr.
Model	m	mw_0	mw_1	$\sigma_{r,0}$	$\sigma_{r,1}$	μ_0	σ_0	μ_1	σ_1	μ_0,σ_0	μ_0,μ_1	μ_0,σ_1	σ_0,μ_1	σ_0,σ_1	μ_1,σ_1	$\mu_{r,t},w_{0.5,r,t}$
Tables 6, A5, A6, and A7																
Panel A	0.7	-1.1	-0.9	0.526	0.526	0.118	0.000	0.118	0.000	0.000	0.999	0.000	0.000	0.000	0.000	1.000
Panel B	0.7	-0.7	-0.5	0.526	0.526	0.118	0.000	0.118	0.000	0.000	0.999	0.000	0.000	0.000	0.000	0.994
Panel C	0.7	-1.1	-0.9	0.526	0.526	0.118	0.000	0.118	0.000	0.000	0.999	0.000	0.000	0.000	0.000	1.000
Panel D	0.7	-0.7	-0.5	0.526	0.526	0.118	0.000	0.118	0.000	0.000	0.999	0.000	0.000	0.000	0.000	0.999
Tables 7 and A8																
Panel A	0.7	-1.0	-0.8	0.526	0.526	0.118	0.000	0.118	0.000	0.000	0.999	0.000	0.000	0.000	0.000	0.999
Panel B	0.7	-1.0	-0.8	0.526	0.526	0.118	0.000	0.118	0.000	0.000	0.894	0.000	0.000	0.000	0.000	0.999
Panel C	0.7	-1.0	-0.8	0.526	0.526	0.118	0.038	0.118	0.038	0.000	0.894	0.000	0.000	0.456	0.000	0.999
Panel D	0.7	-1.0	-0.8	0.542	0.510	0.118	0.038	0.118	0.038	0.000	0.894	0.000	0.000	0.456	0.000	0.999
Table A9																
Panel A	0.7	-1.0	-1.0	0.526	0.526	0.118	0.000	0.118	0.000	0.000	0.999	0.000	0.000	0.000	0.000	1.000
Panel B	0.7	-1.0	-1.0	0.526	0.526	0.118	0.000	0.118	0.000	0.000	0.894	0.000	0.000	0.000	0.000	1.000
Panel C	0.7	-1.0	-1.0	0.526	0.526	0.118	0.038	0.118	0.038	0.000	0.894	0.000	0.000	0.456	0.000	1.000
Panel D	0.7	-1.0	-1.0	0.542	0.510	0.118	0.038	0.118	0.038	0.000	0.894	0.000	0.000	0.456	0.000	1.000

Notes: This table shows meta-parameters used to draw simulation samples for the Normal-markdown model. See Appendix A.1 for details.

support $[mw_t, mw_t + P_{base}]$ and peak at the left extreme of the support.

Intuitively, that model corresponds to one where the minimum wage increases labor force participation of individuals with potential wages just above the minimum wage, in the interval $[mw_t, mw_t + P_{base}]$. This effect's overall intensity is assumed to be proportional to the density of latent log wages evaluated at the minimum wage level and to the P_{height} parameter. In this model, a small minimum wage is likely to have positive employment effects, which are initially increasing. However, at some point, the effects of disemployment start to become more significant. Eventually, the effects of minimum wage on employment will become negative.

A.1.4 The Normal-markdown model with state-level minimum wages

For the exercise shown in Table 5, I augment the Normal-markdown model to include the possibility of state-specific minimum wages that surpass the national minimum wage. I first choose the share of regions that, in each period, are selected to have a higher local minimum wage. Those shares are 0.2 or 0.4, depending on the Panel in Table 5. For reference, the share of states in the US that had local minimum wages at least 5 log points above the federal minimum wage was 0.23 in both 1989 and 2004.

When simulating the model for the initial period, I randomly draw the subset of regions with higher local minimum wages. Given the shares chosen above, those subsets have the same size in all simulations. Then, I draw a number from a Normal distribution with a mean of 0.25 and a standard deviation of 0.075. I assign a local log minimum wage that is equal to the federal minimum wage plus that number (or the federal minimum wage plus 0.05, whatever is higher). The numbers 0.25 and 0.075 above are chosen to match the mean and standard deviation of the log gap between local minimum wages and the federal minimum wage in 2004 for the subset of states for which the minimum wage is higher than the federal one (the corresponding numbers are 0.15 and 0.052 for 1989).

In the second period simulation, I follow the same procedure, except that I do not allow for reductions of local minimum wages between periods. So, the local minimum wage is either calculated from the procedure above or observed in the first period, whichever is higher.

The definition of the average treatment effect to be estimated is updated in the following way to account for the possibility of local minimum wages:

$$
\begin{aligned}
\boldsymbol{ATE}_0 &= \mathbb{E}\left[f\left(mw_{r,1}, \theta_{r,0}\right) - f\left(mw_{r,0}, \theta_{r,0}\right)\right] \\
&= \mathbb{E}\left[f\left(mw_{r,1}, \theta_{r,0}\right) - \boldsymbol{y}_{r,0}\right] \\
\boldsymbol{ATE}_1 &= \mathbb{E}\left[f\left(mw_{r,1}, \theta_{r,1}\right) - f\left(mw_{r,0}, \theta_{r,1}\right)\right] \\
&= \mathbb{E}\left[\boldsymbol{y}_{r,1} - f\left(mw_{r,0}, \theta_{r,1}\right)\right] \\
\boldsymbol{ATE} &= \frac{\boldsymbol{ATE}_0 + \boldsymbol{ATE}_1}{2}
\end{aligned}
$$

The only difference is that the counterfactuals being considered correspond to state-specific changes in the minimum wage induced by the minimum wage, caused either by the increase in the federal minimum wage or by a random draw of a higher local minimum wage in the latter period. Calculating the estimated average treatment effects is the same, using region-specific changes in the observed effective minimum wage. That is consistent with the updated definition, as changes in the effective local minimum wage reflect both national and local minimum wage changes.

A.2 The canonical model of labor demand

A.2.1 Model description

Consider a competitive economy where the only production factors are skilled ($i = 1$) and skilled ($i = 2$) labor, both of which have inelastic supply. A representative firm produces the numeraire good in the economy using a constant elasticity of substitution (CES) production function:

$$F(L_1, L_2) = \left[\alpha L_1^{\frac{E-1}{E}} + (1-\alpha) L_2^{\frac{E-1}{E}}\right]^{\frac{E}{E-1}}$$

The measure of workers is normalized to one, and the region-specific share of skilled workers is s_r. Each worker supplies $\exp(e)$ efficiency units of labor, where e has a Normal distribution with mean zero and standard deviation D. Workers whose log marginal product of labor falls below the log minimum wage mw_t are not employed by the representative firm. That is, the minimum wage truncates the worker productivity distribution. The equilibrium log prices per efficiency unit of labor, p_i, are then given by the solution to a system of two equations:

$$p_{i,r,t} = \log F_i\left(s_r E(p_{1,r,t}, mw_t), (1-s_r)E(p_{2,r,t}, mw_t)\right) \qquad i \in \{1,2\}$$

$$\text{where } E(p, mw) = \int_{mw_t - p}^{\infty} \exp(e)\phi\left(\frac{e}{D}\right) de$$

$$\text{and } F_i(L_1, L_2) = \frac{dF(L_1, L_2)}{dL_i}$$

In the expressions above, ϕ denotes the density of a standard Normal distribution. The function $E(\cdot)$ calculates the average amount of efficiency units supplied by workers of a given type, taking into account the disemployment effects of the minimum wage.

The resulting employment-to-population ratio in a given region and period is given by:

$$emp_{r,t} = s_r\left[1 - \Phi\left(\frac{mw_t - p_{1,r,t}}{D}\right)\right] + (1-s_r)\left[1 - \Phi\left(\frac{mw_t - p_{2,r,t}}{D}\right)\right],$$

and the corresponding cumulative distribution function for log wages is:

$$G_{r,t}(w) = s_r \frac{\Phi\left(\frac{w - p_{1,r,t}}{D}\right) - \Phi\left(\frac{mw_t - p_{1,r,t}}{D}\right)}{1 - \Phi\left(\frac{mw_t - p_{1,r,t}}{D}\right)} + (1-s_r)\frac{\Phi\left(\frac{w - p_{2,r,t}}{D}\right) - \Phi\left(\frac{mw_t - p_{2,r,t}}{D}\right)}{1 - \Phi\left(\frac{mw_t - p_{2,r,t}}{D}\right)} \qquad \text{for } w \geq mw_t$$

where Φ is the cumulative distribution function of a standard Normal.

A.2.2 Calibration

I also use US Current Population Survey data to calibrate the simulations. Using the same sample restrictions described in the previous subsection and data for 1989, I define a worker as belonging to the skilled group $i = 1$ if they have at least four years of college education. Then, I calculate the mean and standard deviation of log wages by skill group for each state and the share of workers in each group.

On the labor supply side, the (unweighted) average of the share of skilled workers across states is 0.224, and the standard deviation is 0.047. Then, in the simulations, I draw the share of skilled workers in each region from a Normal distribution with the corresponding mean and standard deviation, trimming the results so that the share of each worker type can never be below 0.01. The average standard deviation of log wages within states is close to 0.5 for both educational groups. Thus, I set $D = 0.5$.

On the demand side, the mean log wage gap between skilled and unskilled workers is also close to 0.5. Thus, I choose the α parameter such that the skill premium $p_{2,r,t}/p_{1,r,t}$ is 0.5 in an equilibrium of the model with the share of skilled workers equal to the cross-state average, and at the lowest initial value of the minimum wage used (see below). That corresponds to $\alpha = 0.563$ when the elasticity of substitution used in the simulation is $E = 3$, and $\alpha = 0.493$ for $E = 1.4$.

The simulations are run for six scenarios. They combine the two values for the elasticity of substitution in production and three initial values of the minimum wage: -2.2, -1.8, and -1.5. The corresponding initial employment-to-population ratios given the average share of skilled workers are around 0.995, 0.966, and 0.896, respectively.

B Additional tables

Table A3: Effective minimum wage design with zero employment effects

	Outcome			
	Emp.	p10 - p50	p25 - p50	p90 - p50
Panel A: Regions differ only in location				
True average causal effect	-0.000	0.020	0.000	-0.000
Effective min. wage	-0.000	0.020	0.000	-0.000
	(0.000)	(0.003)	(0.000)	(0.000)
Panel B: Regions differ in location and dispersion				
True average causal effect	-0.000	0.021	0.000	-0.000
Effective min. wage	0.001	0.025	0.004	-0.008
	(0.001)	(0.012)	(0.007)	(0.015)
Panel C: As above, but larger increase in min. wage				
True average causal effect	-0.008	0.106	0.009	-0.003
Effective min. wage	0.001	0.123	0.028	-0.043
	(0.006)	(0.019)	(0.013)	(0.028)
Panel D: St. dev. of dispersion is 50% larger				
True average causal effect	-0.001	0.022	0.001	-0.000
Effective min. wage	0.002	0.031	0.009	-0.019
	(0.002)	(0.017)	(0.011)	(0.022)

Notes: This table is analogous to Table 1, except that the markdown parameter is reduced from 0.7 to 0.65, and the data-generating process includes positive employment effects. The size of those effects is calibrated such that the average impact of a 20 log point increase in the minimum wage is zero in Panel A, corresponding to $P_{height} = 0.5$ and $P_{base} = 0.25$. See Appendix A.1.3 for details.

Table A4: Effective minimum wage design with zero employment effects

	Outcome			
	Emp.	p10 - p50	p25 - p50	p90 - p50
Panel A: Regions differ only in location				
True average causal effect	0.009	0.026	-0.005	0.003
Effective min. wage	0.009	0.027	-0.005	0.003
	(0.000)	(0.003)	(0.001)	(0.000)
Panel B: Regions differ in location and dispersion				
True average causal effect	0.009	0.028	-0.005	0.003
Effective min. wage	0.009	0.023	-0.007	0.007
	(0.001)	(0.012)	(0.008)	(0.015)
Panel C: As above, but larger increase in min. wage				
True average causal effect	0.015	0.137	0.004	0.005
Effective min. wage	0.020	0.140	0.019	-0.011
	(0.005)	(0.019)	(0.015)	(0.029)
Panel D: St. dev. of dispersion is 50% larger				
True average causal effect	0.008	0.030	-0.004	0.003
Effective min. wage	0.009	0.019	-0.008	0.009
	(0.001)	(0.018)	(0.012)	(0.022)

Notes: This table is analogous to Table 1, except that the markdown parameter is reduced from 0.7 to 0.6, and the data-generating process includes positive employment effects. The size of those effects is calibrated such that the average impact of a 20 log point increase in the minimum wage is close to one percentage point in Panel A, corresponding to $P_{height} = 1.0$ and $P_{base} = 0.25$. See Appendix A.1.3 for details.

Table A5: Difference-in-differences with binary treatment

	Outcome				
	Emp.	p10	p25	p50	p90
Panel A: Small initial min. wage, truncation/censoring only					
True average causal effect	-0.006	0.016	0.008	0.004	0.002
Binary measure, 50% treated	-0.003	0.007	0.003	0.002	0.001
	(0.000)	(0.001)	(0.000)	(0.000)	(0.000)
Binary measure, 90% treated	-0.004	0.011	0.006	0.003	0.002
	(0.000)	(0.001)	(0.001)	(0.001)	(0.001)
Panel B: Large initial min. wage, truncation/censoring only					
True average causal effect	-0.031	0.118	0.036	0.020	0.010
Binary measure, 50% treated	-0.009	0.053	0.010	0.006	0.003
	(0.001)	(0.002)	(0.001)	(0.000)	(0.000)
Binary measure, 90% treated	-0.017	0.084	0.020	0.012	0.006
	(0.001)	(0.004)	(0.001)	(0.001)	(0.001)
Panel C: Small initial min. wage, positive emp. effects					
True average causal effect	0.010	-0.002	-0.012	-0.006	-0.003
Binary measure, 50% treated	0.001	0.021	-0.001	-0.001	-0.000
	(0.000)	(0.002)	(0.000)	(0.000)	(0.000)
Binary measure, 90% treated	0.003	0.017	-0.004	-0.002	-0.001
	(0.000)	(0.003)	(0.001)	(0.001)	(0.001)
Panel D: Large initial min. wage, positive emp. effects					
True average causal effect	-0.003	0.149	0.039	0.002	0.001
Binary measure, 50% treated	-0.008	0.033	0.035	0.006	0.003
	(0.001)	(0.002)	(0.002)	(0.001)	(0.000)
Binary measure, 90% treated	-0.013	0.078	0.051	0.009	0.004
	(0.001)	(0.005)	(0.003)	(0.001)	(0.001)

Notes: This table is analogous to Table 6, except that it reports results for a difference-in-differences estimator based on a binary version of treatment. Treated status is based on initial median wages being below some simulation-specific threshold, chosen such that the share of treated units corresponds to the desired level.

Table A6: Difference-in-differences with instrumental variables

	Outcome				
	Emp.	p10	p25	p50	p90
Panel A: Small initial min. wage, truncation/censoring only					
True average causal effect	-0.006	0.016	0.008	0.004	0.002
FA instrumended by GAP	-0.008	0.021	0.010	0.006	0.003
	(0.000)	(0.001)	(0.001)	(0.001)	(0.001)
GAP instrumended by FA	-0.006	0.015	0.007	0.004	0.002
	(0.000)	(0.001)	(0.001)	(0.001)	(0.001)
Panel B: Large initial min. wage, truncation/censoring only					
True average causal effect	-0.031	0.118	0.036	0.020	0.010
FA instrumended by GAP	-0.040	0.182	0.044	0.027	0.013
	(0.000)	(0.010)	(0.001)	(0.001)	(0.001)
GAP instrumended by FA	-0.028	0.131	0.031	0.019	0.009
	(0.000)	(0.008)	(0.001)	(0.001)	(0.001)
Panel C: Small initial min. wage, positive emp. effects					
True average causal effect	0.010	-0.002	-0.012	-0.006	-0.003
FA instrumended by GAP	0.002	0.069	-0.001	-0.001	-0.000
	(0.000)	(0.003)	(0.001)	(0.001)	(0.001)
GAP instrumended by FA	0.001	0.051	-0.001	-0.001	-0.000
	(0.000)	(0.002)	(0.001)	(0.001)	(0.001)
Panel D: Large initial min. wage, positive emp. effects					
True average causal effect	-0.002	0.149	0.039	0.002	0.001
FA instrumended by GAP	-0.039	0.128	0.144	0.027	0.012
	(0.002)	(0.007)	(0.004)	(0.002)	(0.001)
GAP instrumended by FA	-0.027	0.094	0.102	0.019	0.008
	(0.001)	(0.006)	(0.003)	(0.001)	(0.001)

Notes: This table is analogous to Table 6, except that it reports results for a difference-in-differences estimator where the main regressor (the interaction between one treatment intensity variable and an indicator for the post period) is instrumented with an alternative treatment intensity variable interacted with the dummy for the post period.

Table A7: Difference-in-differences with quadratic treatment intensity

	Outcome				
	Emp.	p10	p25	p50	p90
Panel A: Small initial min. wage, truncation/censoring only					
True average causal effect	-0.006	0.016	0.008	0.004	0.002
Quadratic on FA	-0.007	0.017	0.009	0.005	0.002
	(0.000)	(0.002)	(0.002)	(0.002)	(0.002)
Quadratic on GAP	-0.006	0.015	0.008	0.004	0.002
	(0.000)	(0.001)	(0.001)	(0.001)	(0.001)
Panel B: Large initial min. wage, truncation/censoring only					
True average causal effect	-0.031	0.118	0.036	0.020	0.010
Quadratic on FA	-0.034	0.281	0.038	0.023	0.011
	(0.001)	(0.019)	(0.003)	(0.003)	(0.003)
Quadratic on GAP	-0.030	0.211	0.033	0.020	0.010
	(0.000)	(0.008)	(0.002)	(0.002)	(0.002)
Panel C: Small initial min. wage, positive emp. effects					
True average causal effect	0.010	-0.002	-0.012	-0.006	-0.003
Quadratic on FA	0.007	0.029	-0.010	-0.004	-0.002
	(0.000)	(0.007)	(0.002)	(0.002)	(0.002)
Quadratic on GAP	0.004	0.040	-0.006	-0.002	-0.001
	(0.000)	(0.004)	(0.001)	(0.001)	(0.001)
Panel D: Large initial min. wage, positive emp. effects					
True average causal effect	-0.003	0.149	0.039	0.002	0.001
Quadratic on FA	-0.013	0.246	0.126	0.007	0.004
	(0.001)	(0.004)	(0.014)	(0.004)	(0.004)
Quadratic on GAP	-0.019	0.167	0.118	0.012	0.006
	(0.001)	(0.005)	(0.006)	(0.002)	(0.002)

Notes: This table is analogous to Table 6, except that it reports results for a difference-in-differences that allows for treatment effects to vary with the treatment intensity through a quadratic functional form.

Table A8: Sensitivity of the Fraction Affected design

	Outcome				
	Emp.	p10	p25	p50	p90
Panel A: Only permanent differences in location					
True average causal effect	-0.010	0.026	0.012	0.007	0.003
Fraction affected	-0.013	0.036	0.015	0.009	0.004
	(0.000)	(0.003)	(0.001)	(0.001)	(0.001)
Panel B: Adding location shocks, stable distributions					
True average causal effect	-0.010	0.026	0.012	0.007	0.003
Fraction affected	-0.010	0.059	0.042	0.036	0.033
	(0.001)	(0.007)	(0.008)	(0.008)	(0.009)
Panel C: Adding dispersion differences and shocks, stable distributions					
True average causal effect	-0.010	0.026	0.012	0.007	0.003
Fraction affected	-0.008	0.072	0.047	0.031	0.005
	(0.002)	(0.008)	(0.008)	(0.008)	(0.012)
Panel D: Average dispersion falls over time					
True average causal effect	-0.010	0.027	0.013	0.007	0.003
Fraction affected	-0.005	0.065	0.044	0.030	0.005
	(0.001)	(0.008)	(0.008)	(0.008)	(0.012)

Notes: This table is analogous to Table 7, except that it shows results for the Gap design instead of the Fraction Affected design.

Table A9: Sensitivity of the Gap design: placebo

	Outcome				
	Emp.	p10	p25	p50	p90
Panel A: Only permanent differences in location					
True average causal effect	-0.000	0.000	0.000	0.000	0.000
Gap measure	0.000	0.000	0.000	0.000	0.000
	(0.000)	(0.001)	(0.001)	(0.001)	(0.001)
Panel B: Adding location shocks, stable distributions					
True average causal effect	-0.000	0.000	0.000	0.000	0.000
Gap measure	0.001	0.017	0.018	0.019	0.019
	(0.000)	(0.005)	(0.006)	(0.006)	(0.006)
Panel C: Adding dispersion differences and shocks, stable distributions					
True average causal effect	-0.000	0.000	0.000	0.000	0.000
Gap measure	0.002	0.026	0.022	0.016	0.004
	(0.001)	(0.007)	(0.006)	(0.006)	(0.008)
Panel D: Average dispersion falls over time					
True average causal effect	-0.000	0.000	0.000	0.000	0.000
Gap measure	0.004	0.022	0.021	0.016	0.004
	(0.001)	(0.007)	(0.006)	(0.006)	(0.009)

Notes: This table is analogous to Table 7, but it reports a placebo scenario with no increase in the national minimum wage. The Gap measure, however, is calculated as if the national log minimum wage would increase by 0.2 between periods (as is the case in Table 7).

Table A10: Canonical model of labor demand, Fraction Affected and Gap design

	Outcome				
	Emp.	p10	p25	p50	p90
Panel A: Initial minimum wage is low, elast. subs. is 3.0					
True average causal effect	-0.009	0.024	0.012	0.007	0.003
Fraction affected	-0.009	0.022	0.011	0.007	0.004
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Gap measure	-0.008	0.018	0.009	0.005	0.003
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Panel B: Initial minimum wage is low, elast. subs. is 1.4					
True average causal effect	-0.009	0.022	0.011	0.006	0.003
Fraction affected	-0.009	0.020	0.011	0.007	0.004
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Gap measure	-0.007	0.017	0.009	0.005	0.003
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Panel C: Initial minimum wage is high, elast. subs. is 3.0					
True average causal effect	-0.042	0.087	0.050	0.030	0.014
Fraction affected	-0.040	0.058	0.044	0.030	0.017
	(0.000)	(0.001)	(0.000)	(0.000)	(0.000)
Gap measure	-0.031	0.044	0.034	0.023	0.013
	(0.000)	(0.001)	(0.001)	(0.000)	(0.000)
Panel D: Initial minimum wage is high, elast. subs. is 1.4					
True average causal effect	-0.039	0.083	0.048	0.029	0.013
Fraction affected	-0.036	0.057	0.044	0.031	0.017
	(0.000)	(0.001)	(0.001)	(0.000)	(0.000)
Gap measure	-0.028	0.045	0.035	0.025	0.014
	(0.001)	(0.002)	(0.001)	(0.001)	(0.000)
Panel E: Initial minimum wage is very high, elast. subs. is 3.0					
True average causal effect	-0.086	0.138	0.096	0.063	0.031
Fraction affected	-0.068	0.060	0.069	0.061	0.039
	(0.000)	(0.001)	(0.001)	(0.001)	(0.000)
Gap measure	-0.051	0.045	0.052	0.046	0.029
	(0.001)	(0.001)	(0.001)	(0.001)	(0.000)
Panel F: Initial minimum wage is very high, elast. subs. is 1.4					
True average causal effect	-0.081	0.134	0.093	0.061	0.029
Fraction affected	-0.058	0.064	0.073	0.065	0.038
	(0.001)	(0.002)	(0.002)	(0.001)	(0.001)
Gap measure	-0.045	0.049	0.057	0.050	0.029
	(0.001)	(0.002)	(0.002)	(0.001)	(0.000)

Notes: This table is similar in structure to Table 6, but the simulation is based on the Canonical model instead of the Normal-markdown model. Panels differ in the initial level of the minimum wage and the elasticity of substitution between skill levels in the Canonical model. See Appendix A.2 for details.

Table A11: Canonical model of labor demand, Effective Minimum Wage design

	Outcome			
	Emp.	p10 - p50	p25 - p50	p90 - p50
Panel A: Initial minimum wage is low, elast. subs. is 3.0				
True average causal effect	-0.009	0.017	0.005	-0.004
Effective min. wage	0.416	-0.239	-0.103	-0.040
	(0.011)	(0.020)	(0.009)	(0.003)
Effective min. wage, no region FE	-0.009	0.018	0.003	0.043
	(0.000)	(0.001)	(0.001)	(0.002)
Panel B: Initial minimum wage is low, elast. subs. is 1.4				
True average causal effect	-0.009	0.016	0.005	-0.003
Effective min. wage	0.285	-0.184	-0.077	-0.014
	(0.006)	(0.011)	(0.006)	(0.003)
Effective min. wage, no region FE	-0.010	0.002	-0.007	0.077
	(0.000)	(0.002)	(0.001)	(0.002)
Panel C: Initial minimum wage is high, elast. subs. is 3.0				
True average causal effect	-0.042	0.057	0.020	-0.016
Effective min. wage	0.447	0.485	0.110	-0.197
	(0.023)	(0.013)	(0.008)	(0.013)
Effective min. wage, no region FE	-0.043	0.054	0.017	0.032
	(0.000)	(0.001)	(0.001)	(0.002)
Panel D: Initial minimum wage is high, elast. subs. is 1.4				
True average causal effect	-0.039	0.054	0.019	-0.016
Effective min. wage	0.232	0.280	0.068	-0.098
	(0.007)	(0.019)	(0.007)	(0.007)
Effective min. wage, no region FE	-0.045	0.042	0.008	0.066
	(0.001)	(0.001)	(0.001)	(0.002)
Panel E: Initial minimum wage is very high, elast. subs. is 3.0				
True average causal effect	-0.086	0.075	0.033	-0.032
Effective min. wage	0.257	0.313	0.153	-0.201
	(0.016)	(0.001)	(0.000)	(0.019)
Effective min. wage, no region FE	-0.091	0.071	0.029	0.018
	(0.001)	(0.000)	(0.000)	(0.002)
Panel F: Initial minimum wage is very high, elast. subs. is 1.4				
True average causal effect	-0.081	0.074	0.032	-0.032
Effective min. wage	0.114	0.217	0.106	-0.098
	(0.003)	(0.008)	(0.004)	(0.010)
Effective min. wage, no region FE	-0.098	0.064	0.022	0.051
	(0.001)	(0.001)	(0.001)	(0.002)

Notes: This table is similar in structure to Table 1, but the simulation is based on the Canonical model instead of the Normal-markdown model. Panels differ in the initial level of the minimum wage and the elasticity of substitution between skill levels in the Canonical model. See Appendix A.2 for details.

Table A12: Quality of fit of the task-based, monopsonistic model

	Data		Model		
	1998	2012	1998	2012	R2
Moments	(1)	(2)	(3)	(4)	(5)
Log wage gaps between educational groups					
Secondary / less than secondary	0.498	0.168	0.486	0.15	0.77
Tertiary / secondary	0.965	1.038	0.995	0.932	0.131
Variances of log wages within educational groups					
Less than secondary	0.41	0.241	0.387	0.225	0.575
Secondary	0.684	0.355	0.647	0.335	0.831
Tertiary	0.702	0.624	0.69	0.644	0.051
Two-way fixed effects decomposition					
Variance establishment effects	0.116	0.056	0.117	0.057	0.652
Covariance worker, estab. effects	0.049	0.048	0.058	0.048	0.421
Formal employment rates by educational group					
Less than secondary	0.266	0.337	0.266	0.336	0.951
Secondary	0.435	0.508	0.435	0.508	1.0
Tertiary	0.539	0.629	0.539	0.631	0.878
Minimum wage bindingness					
Log min. wage - mean log wage	-1.418	-0.922	-1.418	-0.922	1.0
Share < log min. wage + 0.05	0.031	0.053	0.03	0.074	0.696
Share < log min. wage + 0.30	0.086	0.212	0.099	0.218	0.892

Notes: This table is adapted from Haanwinckel (2023). "Data" corresponds to statistics calculated at the *microregion* level using Brazilian data. "Model" corresponds to the model fit using that data as input. Columns (1) through (4) report averages for all regions for each of the two years, using region weights based on total formal employment. Column (5) reports the usual R2 metric $r_e^2 = 1 - \left[\sum_{t=1}^{2}\sum_{r=1}^{151} s_r \left(Y_{e,r,t} - \hat{Y}_{e,r,t}\right)^2\right] / \left[\sum_{t=1}^{2}\sum_{r=1}^{151} s_r \left(Y_{e,r,t} - \bar{Y}_e\right)^2\right]$, where e indexes the specific target moment, $\hat{Y}_{e,r,t}$ is the model prediction, and $\bar{Y}_e$ is the sample average using the region weights employed in the estimation of the model. See Haanwinckel (2023) for details.