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

Daniel Haanwinckel *

April 2025

Abstract

This paper evaluates the validity of estimators that exploit regional wage differences to study the effects of a national minimum wage. It shows that variations of the "fraction affected" and "effective minimum wage" designs are vulnerable to bias from measurement error and functional form misspecification, even when standard identification assumptions hold, and that small deviations from these assumptions can substantially amplify the biases. Using simulation exercises and a case study of Brazil's minimum wage increase, the paper illustrates the practical relevance of these issues and assesses the performance of potential solutions and diagnostic tools.

Keywords: minimum wage, employment, wage distribution, Brazil, Kaitz index, fraction affected

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

1 Introduction

One approach to measuring the labor market effects of minimum wages is to exploit between-region differences in wage levels within a country. For instance, Mississippi and South Carolina should be more strongly affected by an increase in the US federal minimum wage than Texas or Georgia because wages in the latter states are higher. Based on this idea, the econometrician estimates panel regressions where the treatment intensity variable is a function of region-specific wage distributions. Classic examples are the "fraction affected" design of Card (1992) and the "effective minimum wage" design of Lee (1999).

Both Card (1992) and Lee (1999) used data from the US, where much of the identifying variation came from state-specific minimum wage laws. Since then, researchers have developed more modern estimators that directly exploit that source of variation (e.g. Cengiz et al., 2019; Ghanem, Kédagni and Mourifié, 2023). However, these newer estimators are infeasible in contexts without regional minimum wages, explaining why variations of the effective minimum wage and fraction affected designs are still in use. Recent applications published by leading 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), Brazil (Engbom and Moser, 2022), and the UK (Giupponi et al., 2024).

This paper examines the validity of the fraction affected and effective minimum wage designs in such contexts. Specifically, it uses a combination of economic theory, simulation exercises, and original data analysis to investigate whether these techniques can accurately capture the causal effects of a national minimum wage on regional employment levels and wage distributions.

I begin by discussing the sources of identification for each estimator and documenting identification threats that have not been recognized in the existing literature. For the effective minimum wage design, I show that even when the assumptions emphasized by Lee (1999) are satisfied, the design suffers from correlated measurement error, which can lead to economically significant biases. In addition, I demonstrate that even small deviations from the central identifying assumption in Lee (1999)—namely, the absence of correlation between the mean and standard deviation of latent log wage distributions across regions—can generate large biases. This finding is particularly important because there are economic reasons to expect such correlations, and empirical evidence from both the U.S. and Brazil suggests that they are not negligible.

The fraction affected design suffers from analogous issues. Even when the estimator is expected to perform well under the "parallel trends" assumption, functional form misspecification can lead to substantial biases. These biases are more likely to arise when the minimum wage has nontrivial employment effects—whether positive or negative. I also show that the parallel trends assumption can be violated by seemingly benign factors that affect all regions, such as changes in the dispersion of latent log wages driven by skill-biased technical change.

After presenting the econometric issues, I evaluate potential diagnostics tests and alternative regression specifications proposed in academic literature. For the effective minimum wage design, I discuss 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, trends, and controls included in the regression. The leading diagnostics test for the fraction-affected design is the well-known check for parallel pre-trends. As for alternative specifications, I consider other treatment intensity measures (such as the Gap measure or binary indicators), instrumental variables designs, and the estimator of Giupponi et al. (2024).

I employ two approaches to quantify the practical relevance of these issues and the effectiveness of potential solutions. The first is a battery of simulation exercises using different economic models as data-generating processes. They include scenarios where the minimum wage causes disemployment in the lower tail of the wage distribution; others where the minimum wage causes no employment effects but significantly increases wages in the lower tail, perhaps due to the reallocation of workers from low-wage to high-wage firms; others where both wages and employment increase, corresponding to scenarios with significant monopsony power; and still others where minimum wages can affect the returns to skill in a "canonical" model in the style of Katz and Murphy (1992).

The second approach is a detailed analysis of the rise of the Brazilian federal minimum wage starting in 1995. Several published and working papers use the effective minimum wage and fraction affected designs to estimate minimum wage effects in this context. I start by reviewing that literature and showing that it contains significant discrepancies in estimation results. Next, I conduct original data analysis to investigate the sensitivity of those estimators to specification choices. I also benchmark those estimates against the minimum wage effects predicted by a structural model that is estimated on the same data but does not suffer from the same econometric problems.

Both exercises conclude that when those estimators are used in contexts where the minimum wage is set at the national level, economically significant biases may be the norm rather than the exception. The effective minimum wage design requires the existence of shocks that change the bindingness of minimum wages in particular regions but are otherwise unrelated to both the shape of the wage distribution and employment levels. The only plausible shock of that kind explicitly mentioned in the literature is region-specific minimum wage regulations. Even if those shocks exist, misspecification biases still apply. In those cases, an instrumental variables approach in the style of Autor, Manning and Smith (2016) can perform better than the ordinary least squares estimator.

Out of the two estimators, the fraction affected design is more appealing, especially because tests based on pre-treatment data can detect many of the issues I discuss in the paper (if properly implemented). However, those tests are only feasible when a sufficiently long pre-treatment period with constant minimum wages exists. Even with parallel pre-trends, practitioners should investigate the possibility of functional form misspecification. Longer-term studies should also consider regression to the mean and regional convergence issues.

The paper is structured as follows. Section 2 clarifies the scope of my analysis and relates it to existing literature. Sections 3 and 4 analyze the effective minimum wage and the fraction affected estimators, respectively. Section 5 presents additional simulation results to assess the generality of previously discussed issues. Section 6 analyzes the rise of the Brazilian minimum wage. The final section concludes with a summary of recommendations for researchers.

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 but are independent across regions.

I assume that $mw_1 > mw_0$. There is no loss of generality in the DGP from Equation (1), as outcomes depend only on current values of mw_t and $\boldsymbol{\theta}_{r,t}$. There would be loss of generality in models with nominal rigidities or other dynamic concerns, but I abstract from such issues in this paper.

There are two natural ways to define the ceteris paribus causal effects of the rise in the national minimum wage:

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

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

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 $\theta_{r,t}$ distribution is time-invariant. I use the average of these two definitions as the object of interest to be recovered by the econometric designs:

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

Does this definition really correspond to the central object of interest in papers using the effective minimum wage and fraction affected designs? I claim that the answer is *yes*. Appendix A substantiates this point by showing that all of the recently published papers cited early in the introduction seek to identify an average (or aggregate) effect of this type.

I 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. 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 minimum wage effects taking some time to materialize. Using Canadian data, Baker, Benjamin and Stanger (1999) document that employment responds differently to low- or high-frequency minimum wage variation. Sorkin (2015) argues that short- and long-run effects of the minimum wage can be different, and that it may be impossible to identify long-run effects with regressions in some contexts. Meer and West (2016) discusses econometric challenges that arise if minimum wages change employment growth rates instead of employment levels. Vogel (2023) documents that minimum wage effects "trickle up" the wage distribution over a 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 in this paper are also different from the mechanical bias 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 if the data includes 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 (Cadena, 2014; Huang, 2019).

The issues I discuss complement recent work on the econometrics of difference-in-differences models. Because the analysis is restricted to a two-period setting, concerns related to staggered treatment timing do not apply; see de Chaisemartin and D'Haultfœuille (2020) and Roth et al. (2023). However, two other issues from this literature are closely related to my results: violations of the "strong parallel trends" assumption in models with continuous treatment intensity (as studied by Callaway, Goodman-Bacon and Sant'Anna, 2024), and biases arising from heterogeneous treatment effects (see, e.g., de Chaisemartin and D'HaultfŒuille 2018; and Dube and Lindner 2024 for a discussion of how the issues raised in Sun and Shapiro 2022 apply in the minimum wage context).

Relative to these papers, my contribution is to show that such problems arise naturally from economic theory in minimum wage studies; that the resulting biases can be economically meaningful; and that general-purpose solutions proposed in the broader literature may not be effective in this setting. I elaborate further on these connections after describing my results.

Finally, my paper focuses on designs where regions are the unit of analysis. For this reason, I do not evaluate "fraction affected" designs where the unit of analysis is workers (Dustmann et al., 2021) or firms (Harasztosi and Lindner, 2019). These estimators are not designed to capture labor market-level treatment effects, as defined in Equation (2), but rather how outcomes for different groups of firms or workers evolve in response to the minimum wage. Even if the control group in such designs is not directly affected by minimum wage regulations, the resulting estimands generally do not correspond to market-level effects. For example, the reallocation effects documented in

Dustmann et al. (2021) imply that high-wage (control) firms may still be indirectly affected by an inflow of low-skilled labor, which could alter internal task allocations, as in Haanwinckel (2023). Alternatively, minimum wages may affect firm entry and exit (Aaronson et al., 2018) or worker search effort (Adams, Meer and Sloan, 2022), in which case individual-level designs would not capture the impact of compositional changes on the aggregate labor market. Because the estimands differ in important ways, a detailed econometric analysis of these methods does not fit neatly within the framework of this paper and is left to future work.

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:

$$y_{i,r,t} = \alpha_{i,r} + \delta_{i,t} + \beta_i \left[mw_t - w_{0.5,r,t}\right] + \gamma_i \left[mw_t - w_{0.5,r,t}\right]^2 + \varepsilon_{i,r,t}, \tag{3}$$

where i indexes the 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, a measure of the bindingness of the national minimum wage in a particular region-time. This baseline specification includes region and time fixed effects. I discuss alternative specifications later.

To calculate the predicted treatment effects of the minimum wage increase in each region, 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 an estimate of average treatment effects as defined in Equation 2:

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

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.[1]

[1] 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

In order to discuss identifying assumptions, it is helpful to introduce the semiparametric conceptual framework of Lee (1999). 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. The cumulative distribution function for those latent log wages is:

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

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 (3)—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 and controls included in the regression.

The central message of the analysis below is that, even when those assumptions are satisfied, the effective minimum wage estimator may not perform well, and that minor deviations from the second assumption can introduce large biases.

3.2 Issue #1: Correlated measurement errors

The effective minimum wage design is predicated upon minimum wage effects being stronger where it bites more into the *latent* wage distribution. 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 the minimum wage introduces deviations between $w_{0.5,r,t}$ and $\mu_{r,t}$ that can cause significant biases even if they are small in magnitude, because these deviations correlate with minimum wage effects on the outcomes of interest.

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

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.

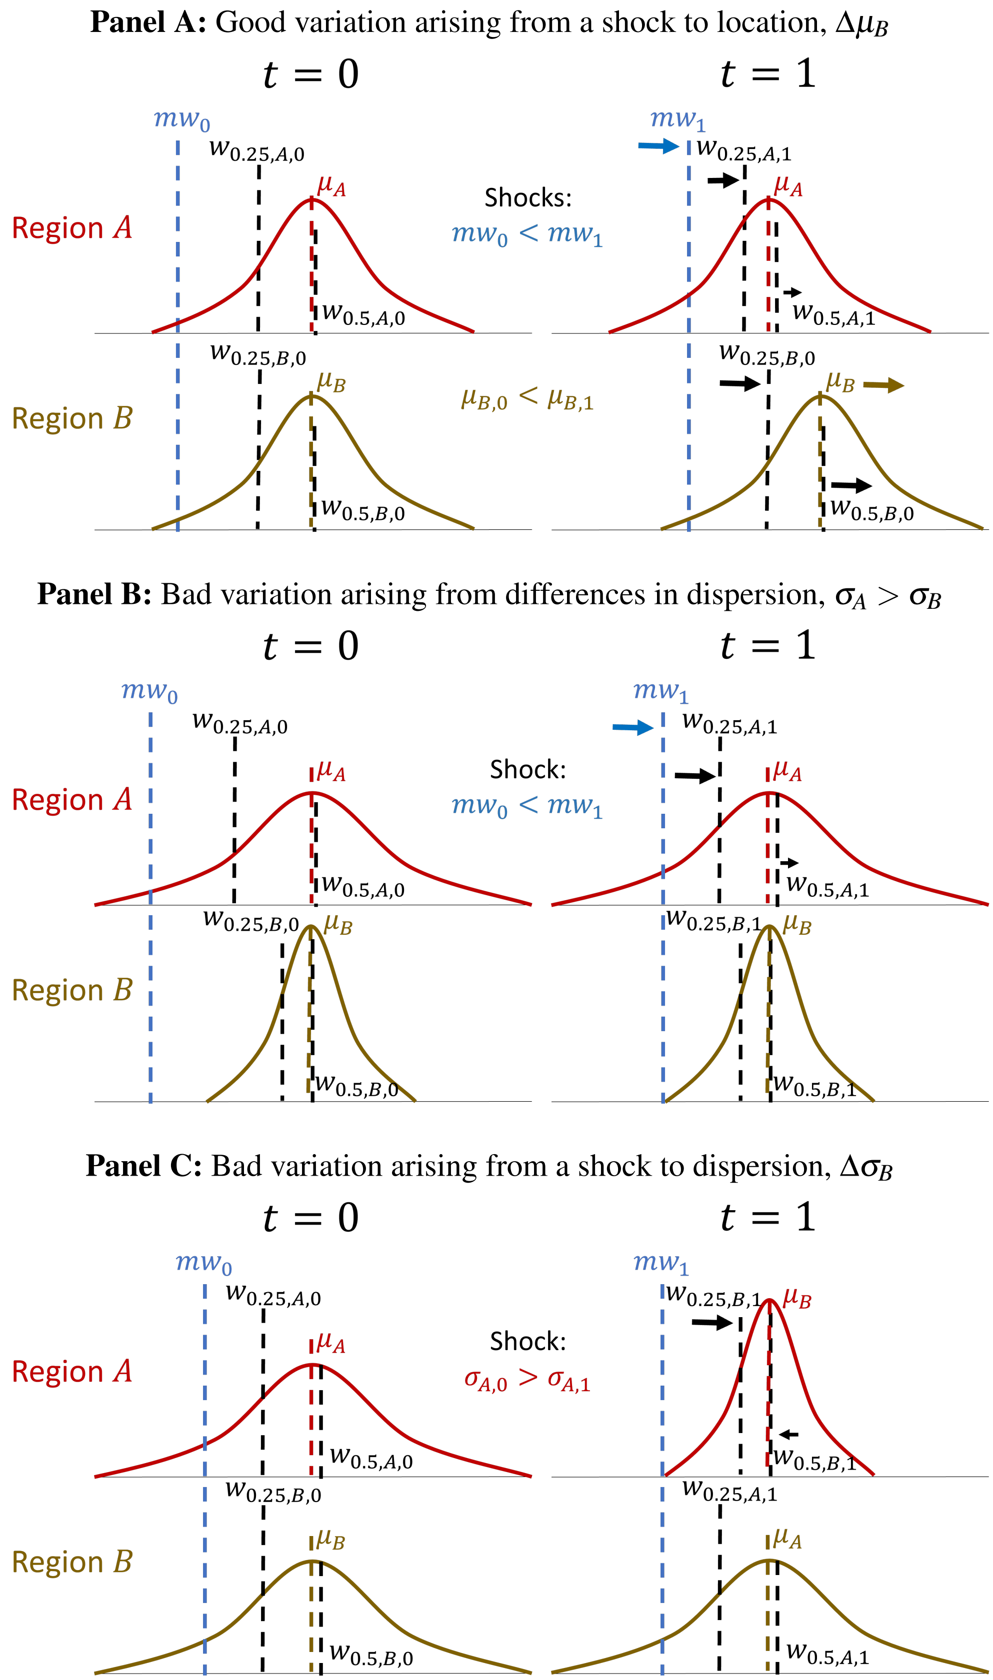

Figure 1: Good and Bad Variation in the Effective Minimum Wage Design

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 $\mu_{r,t}$.[2]

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 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 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.

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

[2]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 Subsection 3.1 that after estimating Equation (3), the predicted treatment effects are calculated by multiplying the coefficients and changes in the effective minimum wage, not changes in the national minimum wage.

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, one needs stronger assumptions about the DGP to infer the sign of the total bias. 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 impose $m = 0.7$.

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 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 $t = 1$).[3] As explained below, each particular simulation exercise makes different

[3]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

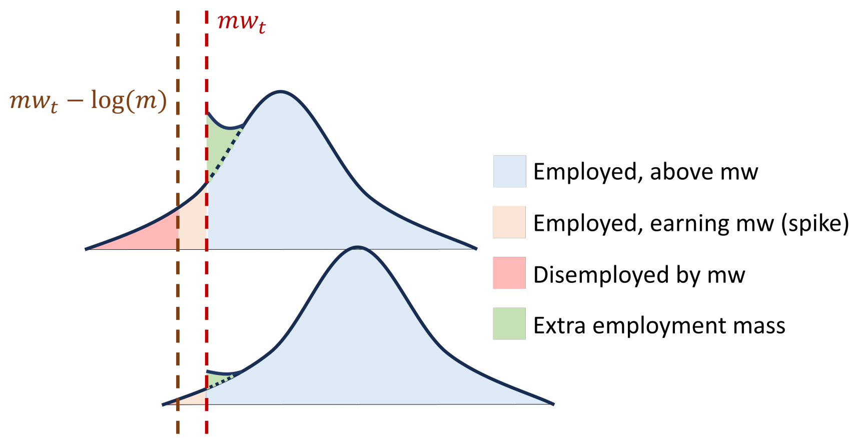

Figure 2: Minimum wage effects in the Normal-markdown simulation model

assumptions about that meta-distribution of parameters across regions. In all exercises shown, the data contains 200 regions. See Appendix B 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.

The first panel in Table 1 shows a case where all regions have the same dispersion parameters $\sigma_{r,t} = \sigma_t$, 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 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 B shows that their correlation is 0.999. Still, the estimator displays biases. They arise because, though the differences between $w_{0.5,r,t}$ and $\mu_{r,t}$ constitute a tiny share of the variation,

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. Each panel corresponds to different assumptions on the data-generating process. In **Panel A**, regions differ only in the location parameter $\mu_{r,t}$. **Panel B** includes differences in the dispersion parameter $\sigma_{r,t}$. **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 B for details.

they are systematically correlated with the outcomes of interest.[4]

Panel C shows that biases are larger 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). 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 $\mu_{r,t}$.

Panel D highlights how differences in the dispersion of latent log wages between regions constitute the source of the bias. Increasing those differences by 50% is enough to 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

[4]The direction of the biases suggests that the issue illustrated in Panel C of Figure 1 is more serious than the one in Panel B. I confirm this intuition by verifying that a model that maintains cross-sectional dispersion in $\sigma_{r,t}$ but makes it invariant over time is almost unbiased.

for individuals with latent wages just above it, corresponding to the green areas in Figure 2 (see Appendix B.3 for details). Such effects could emerge from, e.g., increased search effort by workers (Adams, Meer and Sloan, 2022; Piqueras, 2024). The results (available upon request) show that biases exist even when employment effects are zero or positive on average. That's because the correlated measurement error issue arises from heterogeneity in employment effects across regions, not from its average effects. In Section 5, I will come back to this issue by quantifying biases a large number of alternative DGPs, many of which have zero or positive employment effects of the minimum wage.

Connection to previous econometric literature: The dispersion in shape parameters $\sigma_{r,t}$, combined with the non-linearity of minimum wage effects in standard economic models, implies that the effective minimum wage design may be subject to econometric concerns related to heterogeneous treatment effects in two-way fixed effects settings (Blundell and Dias, 2009; de Chaisemartin and D'HaultfŒuille, 2018; Sun and Shapiro, 2022; Callaway, Goodman-Bacon and Sant'Anna, 2024). As a result, these regressions may yield biased estimates of the average treatment effect of the national minimum wage, even if the econometrician observes the latent centrality parameter $\mu_{r,t}$. Accordingly, part of the bias observed in my simulation exercises stems from this treatment effect heterogeneity.

That said, the econometric concerns discussed in this section go beyond those highlighted in the existing literature. In particular, prior work does not consider the endogeneity of the treatment intensity variable in the effective minimum wage design, which is the central focus of my analysis.

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

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)

Notes: 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.

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.

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.

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. Another economic factor that could generate this correlation would be regional differences in endowments that could affect industrial composition, which in turn would lead to non-random sorting of workers and firms tho different regions.

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. For example, a persistent shock reducing labor demand over a long period would introduce a negative correlation between the effective minimum wage and employment. Alternatively, regions may differ in the shape of their distribution of worker skills, leading to different baseline levels for quantile wage gaps. Such concerns make specifications with fixed effects overwhelmingly more popular in published 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.

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 argument on 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,

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	0.034	0.015	-0.023
	(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: 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.

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). In Section 6, I discuss those issues in the Brazilian context.

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

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?

Sometimes, the econometrician may have a 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

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	0.220	0.177	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.

when the observed log wage distribution is identical to the latent wage distribution.[5] 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.

As with any test, one should consider the possibility of false positives and false negatives. False positives 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

[5] See Lee's discussion around Equation (5) in page 996.

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.

False negatives may emerge from two reasons. The first are cases where the true model includes spillovers in the upper tail. The second possibility is that the estimator may be biased for lower tail outcomes, or for employment, while being unbiased for spillovers 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).

In Section B.5, I will come back to the effectiveness of this diagnostics test when examining biases in a large number of DGPs.

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.

In Appendix C, I investigate these scenarios. I show that, the more variation coming from regional minimum wages, the smaller the biases in the effective minimum wage design. In addition, I show that the remaining biases can be removed by using an instrumental variables approach inspired by Autor, Manning and Smith (2016).

3.8 Taking stock

The main takeaway from the previous discussion is that the econometrician should have a clear sense of what constitutes the identifying variation behind the effective minimum wage design. Exogenous changes in state-level minimum wages are the clearest example of such variation. If the data includes little or no variation in state-level minimum wage laws, then identification relies on the existence of a latent residual shock that shifts the location of latent log wage distributions but has no independent effects on their shape, nor on 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, as labor supply and demand shocks are typically biased towards some types of workers and affect employment levels in addition to wages.

Even if a plausible shifter of regional minimum wage bindingness exists, the estimator may still have economically significant biases arising from correlated measurement error. These biases can be more severe if the specification includes many controls, as they may diminish the residual variance of the "good" source of identification. If the bindingness shifter is observable, the econometrician may reduce misspecification biases by using an instrumental variables approach in the spirit of Autor, Manning and Smith (2016).

If one must use this regression design anyway, my analysis confirms the advice from Lee (1999). First, the effective minimum wage should be constructed using the median rather than other quantiles of the wage distribution. Second, time fixed effects should always be used. As for whether to include region fixed effects, regional trends, or other regional controls, econometric theory provides little guidance.

Some of the issues discussed in this section arise from the fact that the treatment intensity variable used in the effective minimum wage design may be endogenous to the minimum wage itself. One potential strategy to address this concern is to construct the treatment variable using *pre-treatment* information. In that case, identification would no longer rely on *shocks* to the location parameters $\mu_{r,t}$, but rather on their initial levels. This idea serves as the central motivation for the estimators introduced in the following section.

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:

$$y_{i,r,t} = \alpha_{i,r} + \delta_{i,t} + \beta_i FA_r \cdot \mathbf{1}\{t = 1\} + \varepsilon_{i,r,t} \tag{4}$$

where the subscript i indexes a specific equilibrium outcome, such that different i 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.[6] The regressions include region and

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

time fixed effects. The estimated average treatment effect ATE_i is obtained by multiplying the average of FA_r by the $\hat{\beta}_i$ estimate.

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 emphasized that much of the identifying variation in his application comes from state-level minimum wage laws passed in the 1980s. 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 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. The Gap measure corresponds to the increase in regional average wage if all low-wage workers in that region 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

As in any regression, unbiasedness requires correct specification of the conditional expectation function. That is, the effects of the national minimum wage should be approximately linear in the continuous treatment intensity variable. If that sensitivity measure is misspecified, estimates of average treatment effects based on that estimator may be biased.

In this section, I show that the fraction affected estimator and its variants are sensitive to this issue, at least when the DGP is of the Normal-markdown class I introduced previously on the paper. Later, I show that the same is true when the DGP comes from the "canonical" model of labor demand from Katz and Murphy (1992) (and in that case, functional form misspecification issues are relevant for the effective minimum wage design as well).

To focus on the role of functional form assumptions, I design the simulations of the Normal-markdown model 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 μ_r, the location parameter of their latent log wage distributions. 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 5 reports results for four different DGPs. They differ in two dimensions. First, the starting minimum wage may be lower or higher, in which case the minimum wage increase has stronger effects (as the minimum bites more into the latent distribution). Second, in some DGPs, I allow for positive employment effects, as illustrated in the green areas of Figure 2.

I find 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; 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 when it causes positive employment effects. For example, in Panel C, 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 these particular DGPs, 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.

Table 5: 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	0.186	0.044	0.026	0.013
	(0.000)	(0.009)	(0.001)	(0.001)	(0.001)
Gap measure	-0.028	0.128	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	-0.002	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	-0.027	0.091	0.101	0.019	0.008
	(0.001)	(0.006)	(0.003)	(0.001)	(0.001)

Notes: See the notes below Table 1 for an explanation of the structure of the table, and Appendix B for a full description of the simulation model. The national minimum wage increases by 20 log points in all panels. In Panels B and C, the initial minimum wage is higher, such that the minimum wage bites more into the latent wage distribution. In Panels C and D, positive employment effects are allowed, as illustrated in the green region of Figure 2.

Relationship to previous literature and alternative specifications. Misspecification biases arise because the effect of a marginal increase in the treatment intensity variable is not constant, either within or between regions. In this sense, the issue can be viewed as a deviation from the "strong parallel trends" assumption formalized in Callaway, Goodman-Bacon and Sant'Anna (2024), a paper that studies difference-in-differences designs with continuous treatment variables. It is also related to unit-level heterogeneity in treatment effects emphasized by Sun and Shapiro (2022) and recently discussed in the minimum wage context by Dube and Lindner (2024).[7]

My contribution relative to those methodological papers is threefold. First, the previous results

[7]Previous work on similar econometric issues include Blundell and Dias (2009); de Chaisemartin and D'HaultfŒuille (2018); de Chaisemartin and D'Haultfœuille (2020); de Chaisemartin et al. (2024).

illustrate how these problems arise naturally in the minimum wage context from the basic mechanics of standard labor economics models. Second, I show that these misspecification biases can be large in general, by exploring a large range of possible data-generating processes (see Section 5). Third, as I discuss below, I show that possible solutions discussed in those methodological papers may not be effective in this context.

Callaway, Goodman-Bacon and Sant'Anna (2024) argue that, if there is a control group composed of entirely untreated units, then a binary difference-in-differences specification can recover an average treatment effect on treated units. However, in the minimum wage context, it may be hard to argue that any single region is entirely untreated by a national minimum wage without imposing strong structural assumptions on the data-generating process. Still, some applied papers explore this idea by creating a binary version of treatment intensity. Recent examples in the Brazilian case, which will be further discussed in Section 6, are Derenoncourt et al. (2021) and Parente (2024).

In Appendix Table A7, I show that rather than solving the problem, a binary treatment specification can suffer from more substantial biases than the baseline models. In this exercise, 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 loss of precision, and significant biases remains.

Section 3.2.4 in Dube and Lindner (2024) discusses the fraction affected and related estimators. They emphasize the issue that treatment effect heterogeneity may lead to biases, as shown by Sun and Shapiro (2022). These papers argue that, if treatment effects are constant within within fine observable groups (say, industries), it is relatively straightforward to adjust the econometric model to solve the misspecification bias. However, in the discussion above, I show that the problematic heterogeneity in treatment effects is not of the form discussed in Sun and Shapiro (2022) and Dube and Lindner (2024). Rather, the fundamental problem is that with monopsony power or minimum wage-induce reallocation effects, minimum wage effects are likely to be nonlinear within fine groups. In that case, re-parameterizing the econometric model is much less straighforward, and can only be done if the econometrician is willing to make parametric structural assumptions on the channels through which minimum wages affect local labor markets.

One may wonder whether adding a quadratic term to the regression specification could effectively capture that kind of treatment effect heterogeneity. Appendix Table A9 reports results for those models. Biases become smaller for some outcomes but larger for others. Thus, this approach is not a practical solution for the misspecification problems either.

One alternative specification discussed by Dube and Lindner (2024) is the bin-level estimator of Giupponi et al. (2024). It differs from the baseline fraction affected design in three ways. The first is that it uses a binary treatment intensity measure, with the regions in the top decile of wages being the control group and the other nine deciles being the treated group. Second, this split is based not on the averaged wages but rather location wage premiums estimated in a Mincerian regression, netting out the influence of worker's gender, education, occupation, industry, and even person effects in some specifications. Third, it measures impacts at the level of "skill" bins (with skill defined as wage purged from the location wage premium). That is, one can view their strategy as separate binary fraction affected designs, one for each bin.

Those alterations are interesting but do not fundamentally alter the conclusions from the previous specifications; the estimator is subject to the same concerns. It is unclear whether functional form misspecification biases should be larger or smaller using this estimator. As discussed above, binary versions of the fraction affected design are more likely to perform poorly. On the other hand, defining treatment based on regional wage premiums rather than raw wages, and estimating impacts at the skill bin level, could reduce omitted variable bias in the presence of skill-specific shocks correlated with minimum wage bindingness. Finally, the bin-specific approach relies strongly on the assumption that treatment and control groups differ only on the location of the latent log wage distributions, not their shape. Thus, like the effective minimum wage design, this estimator may be very sensitive to correlations between location and dispersion parameters across regions. The baseline effective minimum wage design does not share that sensitivity.

A final potential solution 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 A8 reports the results of using such strategies. They have no impact on the estimates compared to the basic OLS estimator.

4.3 Issue #2: Trends in the dispersion of latent wages

As with any difference-in-differences design, the validity of the fraction affected and gap estimators depends critically on the parallel trends assumption. In Appendix E, I discuss two potential threats to this assumption that have been raised in the prior literature. The first is regression to the mean: if regional wage levels contain a transitory component, then defining treatment intensity based on wages from a single period (or a small number of periods) can introduce a dynamic form of bias. The second, closely related concern is regional convergence driven by other factors, which can also result in violations of the parallel trends assumption. One key takeaway from that dis-

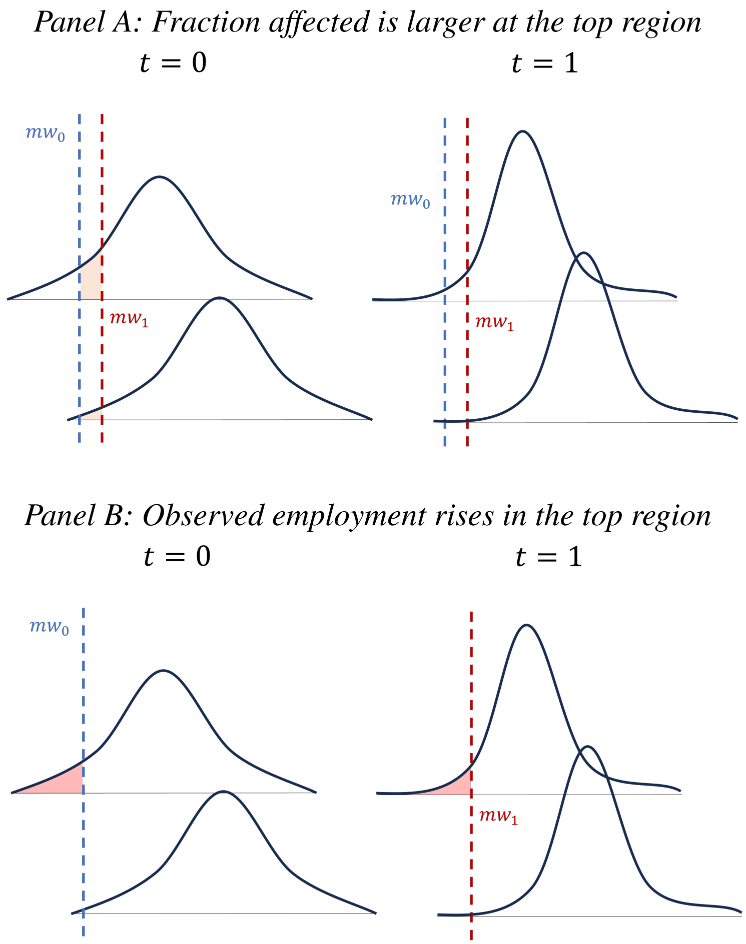

Figure 3: A truncation model with a secular decline in latent log wage dispersion

cussion is that, although such problems can often be detected using standard tests for pre-trends, these tests are only informative when the treatment intensity variable is constructed using a single pre-treatment period, rather than being averaged across multiple periods.

In this section, I highlight another potential threat to identification that, to my knowledge, has not been discussed in the existing literature. Suppose that, across all regions, there is a simultaneous change in the dispersion of latent log wages that coincides with an increase in the national minimum wage. Since this structural shock is common to all regions, one might expect the time fixed effects included in the regression to account for it. However, as I show below, this is not necessarily the case when the minimum wage is already binding in the initial period.

Consider the model illustrated in Figure 3. Regions differ in a location parameter μ_r that is constant over time. The shape parameter σ_t is common across regions and decreases over time. All else equal, regions with lower μ_r have a higher fraction affected and a larger (more negative) causal effect of the minimum wage. However, as σ_t declines, these same regions may experience a relative *increase* in employment. This is because the truncation effects of the minimum wage may weaken 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, even though the causal effects are negative for all observations in all periods.

The comparison between Panels C and D in Appendix Table A10 illustrates this issue. Panel C, discussed in Appendix E, features idiosyncratic shocks to both the location and dispersion parameters, but the distributions of these parameters remain stable over time. Panel D differs only in that the average dispersion parameter $\sigma_{r,t}$ declines from 0.54 in $t = 0$ to 0.51 in $t = 1$.[8] Even this small change is enough to meaningfully increase the bias in the estimated employment effects.

In Appendix E, I also discuss how tests for parallel pre-trends may effectively detect biases arising from regression to the mean and regional convergence. Similarly, such tests may be able to detect the issue discussed in this section, provided that the trend in dispersion begins prior to the minimum wage change and has a relatively stable effect on employment over time. This is illustrated in the placebo regressions reported in Appendix Table A11. An additional suggestion for practitioners is to check for the existence of trends in the variance of log wages before the implementation of the national minimum wage.

4.4 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 may be particularly sensitive to functional form misspecification. To assess the possibility of misspecification biases, practitioners are encouraged to report predicted employment and wage effects from both the Fraction Affected and Gap designs, and perhaps from quadratic specifications as well. In addition, I showed that binarized versions of

[8] These values are based on the calibration discussed earlier, using state-level statistics from CPS data for 1989 and 2004. In the U.S. context, the decline in dispersion could partly reflect the increased prevalence of state-level minimum wages in 2004 compared to 1989, and should not necessarily be interpreted as a change in *latent* log wages. As noted before, the purpose of this exercise is to illustrate an econometric issue using empirically grounded parameters, not to evaluate the effects of the national minimum wage in the U.S.

those estimators should be avoided as they display more significant misspecification biases.

5 Quantifying biases under alternative DGPs

The previous sections presented a series of econometric issues that may invalidate the use of the effective minimum wage and fraction affected designs. The remainder of the paper is devoted to gauging their practical relevance. In this section, I use simulation exercises to estimate the likelihood of significant biases arising under a wide range of economic scenarios. Later in the paper, I explore the same question in a particular empirical setting.

5.1 Additional simulations of the Normal-markdown model

In the main exercise of this section, I quantify biases for 500 different DGPs in the Normal-markdown class introduced in Subsection 3.2. Each DGP corresponds to different assumptions on distribution of region-specific parameters $[\mu_{r,0}, \sigma_{r,0}, \mu_{r,1}, \sigma_{r,1}]$, markdown parameter m, positive employment effect parameters P_{base}, P_{height}, and levels of the national minimum wage mw_0, mw_1 (measured relative to the center of the latent log wage distribution). They run the gamut from models with significant disemployment effects, to positive employment effects, to no employment effects but significant wage spillovers. The specific parameter values are randomly drawn, with most distributions centered around the corresponding values observed for US states (used for previous simulations in the paper).

I make conservative choices regarding DGP elements that have been shown to be problematic. For example, the contemporaneous correlation between location and dispersion parameters is allowed to be positive, but its value is never larger than 0.20. That number is smaller than the correlation between regional mean log wages and regional standard deviation of log wages observed for the US in 2004 (0.26) or Brazil in 1998 (0.34). See Appendix B.5 for details.

For each DGP, I simulate 1,000 samples, each with 200 regions and two periods. For each sample, I estimate the causal effects of the national minimum wage using the effective minimum wage and fraction affected designs. Then, I average those estimates over the 1,000 samples to obtain the expected value of those estimators under that particular DGP, assuming a sample size of 200. I classify an estimator as being biased for a specific outcome, under that specific GDP, if the expected estimated effect is significantly different than the expected causal effect of the minimum wage on that outcome under that DGP.

What does "significantly different" mean in this context? For employment effects, the estimator is considered biased if the gap between expected estimate and expected causal effect is at least 0.005

Table 6: Quantifying biases under 500 alternative DGPs

Outcome	Emp	p10-p50	p25-p50	p90-p50
Panel A: 128 DGPs with positive employment effects				
Average causal effects	0.013	0.153	0.034	0.000
Eff. min. wage: share with positive bias	0.16	0.55	0.46	0.00
· : share with negative bias	0.05	0.00	0.01	0.80
Fraction affected: share with positive bias	0.00	0.44	0.43	0.02
· : share with negative bias	0.82	0.20	0.00	0.23
Panel B: 241 DGPs with negative employment effects				
Average causal effects	-0.018	0.151	0.039	-0.009
Eff. min. wage: share with positive bias	0.86	0.45	0.50	0.00
· : share with negative bias	0.00	0.00	0.00	0.88
Fraction affected: share with positive bias	0.19	0.39	0.39	0.00
· : share with negative bias	0.56	0.25	0.00	0.27
Panel C: 131 DGPs with small emp. effects but large wage effects				
Average causal effects	0.000	0.203	0.051	-0.002
Eff. min. wage: share with positive bias	0.22	0.36	0.50	0.02
· : share with negative bias	0.00	0.01	0.00	0.74
Fraction affected: share with positive bias	0.00	0.46	0.55	0.01
· : share with negative bias	0.40	0.32	0.02	0.30

Notes: See the text for an explanation of the criterion used to classify an estimator as biased under a specific DGP, and Appendix B.5 for a description of the DGPs used in this exercise.

(half a percentage point) *and* at least 25% of the true causal effect. For quantile gaps, the estimator is biased if the gap is at least 0.05 (five log points) *and* at least 25% of the true causal effect. For example, an estimator is considered biased for the p10-p50 quantile gap if the expected estimate is 0.04 while the expected causal effect is 0.09. However, it is not classified as biased if the expected estimate is 0.48 and the expected causal effect is 0.40.

Table 6 reports the results of that exercise. To provide additional clarity on when each estimator is expected to perform better or worse, I split DGPs in three groups: positive employment effects (0.005 or more), negative employment effects (-0.005 or less), and small employment effects but significant wage effects.[9] All numbers with two decimal places in the table correspond to the share of DGPs in a given group for which one of the estimators (rows) is positively or negatively biased for a given outcome (columns).

The broad message of the table is that, in all three categories of DGPs, both estimators are more

[9]To be included in the analysis, the DGP must be such that the minimum wage causes effects of magnitude at least half a percentage point on the employment to population ratio or 5 log points on at least one quantile gap. This choice rules "trivial" cases with non-binding minimum wages. I also exclude DGPs where the causal effects of the minimum wage on the employment to population ratio is of four percentage points or more, either positive or negative. This criterion ensures that the biases are not arising from atypical scenarios such as a large minimum wage hike in a country where latent log wages have low dispersion.

likely than not to be biased for at least one outcome, and often for two or more. A careful examination reveals some interesting patterns. The effectively minimum wage design is more likely to be positively- than negatively-biased for employment effects, especially when the true effects are negative. For the fraction affected design, the opposite is true: it is more likely to be negatively biased for employment effects, especially when the true effects are positive. The effective minimum wage design is also more likely to overestimate, rather than underestimate, the degree to which minimum wages reduce inequality in the lower tail of the wage distribution.

Appendix Table A12 repeats this exercise using a more stringent requirement for classifying an estimator as biased. Specifically, it requires gaps between expected estimates and expected causal effects that are twice as large as in the baseline exercise. It finds that large biases are common for all combinations of estimator and DGP group. Taken together, these results suggest that the reliability of the effective minimum wage and fraction affected designs should not be taken for granted if the data does not include regional variation in minimum wage laws.

5.2 Upper-tail effects as a diagnostics tool?

In Subsection 3.6, I discussed using estimates of wage effects in the upper tail of the wage distribution for model selection. Appendix Table A13 replicates the exercise in Table 6 restricting attention to DGPs where the effective minimum wage estimate of the p90-p50 quantile gap effect is less than 0.05 in magnitude. This procedure approximates a two-step process where the econometrician first measures the upper-tail spillovers and abandons the project if they find large effects.

The first takeaway from that table is that there are few DGPs without upper-tail spillovers: only 77 out of 500. This result is consistent with the fact that papers using the effective minimum wage design commonly detect upper-tail spillovers. Using US data, both Lee (1999) and Autor, Manning and Smith (2016) find evidence of such spillovers in some specifications and argue that they should be used for model validation. In the Brazilian context, Hinojosa (2019) and Engbom and Moser (2022) also detect upper-tail spillovers but interpret them as causal effects, rather than evidence of misspecification.

The second takeaway is that, while the incidence of biases is indeed reduced in that subsample, they are not completely eliminated. For example, there remains a 16% chance of positive biases on the p10-p50 effect when the true effect on employment is positive, and a 40% chance of positive bias on the employment effect when the true employment effect is negative.

In Appendix Table A14, I implement a similar exercise for the fraction affected design: estimates are considered invalid if the estimator detects positive effects on the 90th quantile of the wage distribution. This procedure is less stringent than the one for the effective minimum wage, leaving

345 out of 500 DGPs. However, economically significant biases remain common. Thus, a lack of wage effects on the upper tail is no guarantee of unbiasedness for the fraction affected design.

5.3 The canonical model of labor demand

All of the simulations performed so far used the Normal-markdown class of simulation models. That class is very flexible and can approximate a broad range of economic models. However, it imposes symmetry of the latent log wage distribution and implicitly assumes that the marginal returns to labor are invariant to the minimum wage. To address those limitations, I perform an additional simulation exercise based on an equilibrium labor market model very similar to that of Katz and Murphy (1992), with parameters calibrated to reflect US states as of 1989.

Appendix D describes that exercise and presents the results. They mirror previous results for the fraction affected estimator regarding functional form issues, except that they now extend to the effective minimum wage design as well. Specifically, model misspecification introduces economically significant biases even in cases that, conceptually, should be ideal applications of those estimators. These results suggest that simulation results based on the Normal-markdown model may be conservative in the sense that misspecification biases can be larger in cases where minimum-wage-induced employment effects or reallocation introduce changes in marginal returns to labor.

6 The rise of the Brazilian federal minimum wage

This section discusses the limitations of the effective minimum wage and fraction affected estimators in a specific context: the rise of the Brazilian federal minimum wage beginning in 1995, illustrated in Figure 4. Brazil underwent a series of macroeconomic reforms in 1993 and 1994 which successfully stabilized the economy after years of hyperinflation. Over the 22 years following stabilization, three different presidents promoted yearly adjustments that, in total, increased the minimum wage by 167 percent in real terms. Productivity also grew, but there is no question that the minimum wage became more binding. The red dash-dot line in Figure 4 shows that the minimum wage grew substantially relative to median wages. Because the Brazilian workforce became more skilled over the same period, the relative change is more stark if one conditions on a specific worker group, such as young men with high school degrees (the blue dashed line).[10][11]

This setting is ideal for analyzing the performance of the effective minimum wage and fraction

[10]Young, high school-educed men are a positively selected subsample earlier in the studied period because men earned significantly more than women, and because most workers had less than complete high school.

[11]Figure A1 in the appendix displays alternative bindingness measures based on shares of workers earning close to the minimum wage. They reinforce the view that the "bite" of the minimum wage grew substantially before 2007.

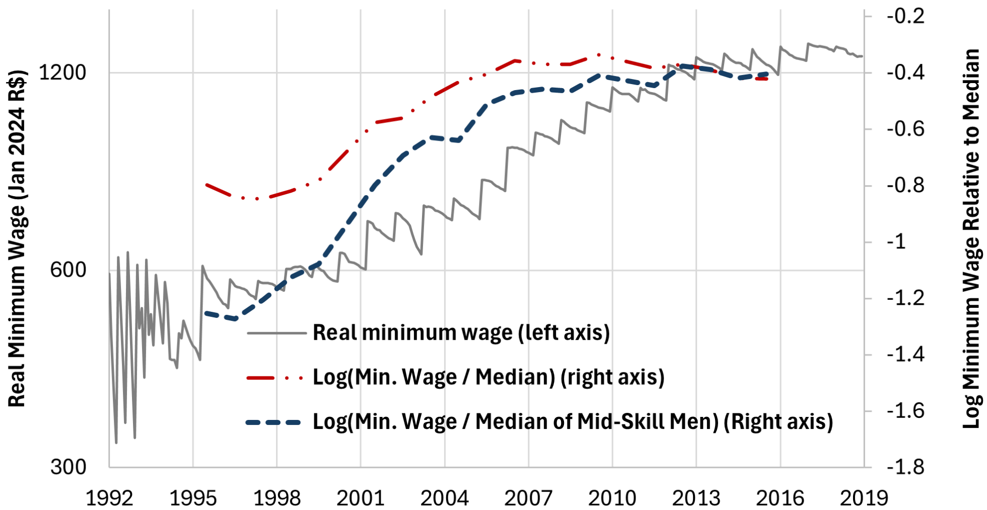

Figure 4: Evolution of the federal minimum wage in Brazil (logarithmic scale)

Notes: The monthly real minimum wage series is constructed based on IPEA Data time series for nominal minimum wage and IPCA price index. The jagged pattern in earlier years corresponds to frequent nominal minimum wage adjustments during a hyperinflationary period. The other two series are constructed based on yearly PNAD survey data from IBGE (processed using the DataZoom tool from PUC-Rio), using the September minimum wage value for each year. They show the minimum wage—legally defined as a minimum monthly earnings level—relative to median monthly earnings in workers' main jobs. The sample includes salaried and self-employed workers between ages 18 and 54. "Mid-Skill Men" are those up to 30 years of age with exactly 11 years of schooling. The relative series start in 1995 because there was no PNAD survey in 1994, and before that, macroeconomic instability makes it difficult to calculate bindingness measures.

affected estimators for four reasons. First, Brazil is a large country with substantial variation across regional labor markets. Second, there is no regional variation in minimum wage laws, making those estimators more appealing.[12] Third, there are several published papers and recent working papers using those estimators in the Brazilian context, allowing me to discuss how those strategies are implemented in practice and to evaluate their results. Fourth, in previous work (Haanwinckel, 2023), I developed a structural model that predicts minimum wage impacts at the regional level and estimated it using a panel of Brazilian regions. This model yields region- and time-specific estimates of minimum wage effects that are directly comparable to those coming from the reduced-form designs, but that are not subject to the econometric issues discussed in this paper. Thus, they can be used as benchmark values for minimum wage effects in this context.

In the following, I first show that existing papers often come to conflicting conclusions about its impacts on employment and wages, and that econometric theory provides little guidance on which of them—if any—is correct. Next, I conduct an original data analysis based on Brazilian data to

[12]Brazilian law allows for state- and occupation-specific wage floors. However, they exist in few states, apply to few occupations, and are imperfectly enforced (Moura and Neri, 2008; Corseuil, Foguel and Hecksher, 2015).

investigate the sensitivity of those estimators to arbitrary specification choices, and to benchmark the estimates I find against those coming from the structural model.

6.1 Estimates from previous literature: which one is the right one?

I focus on the six papers listed in Table 7, which were selected based on a search on the IDEAS database.[13] The table summarizes the data sources for each paper, the preferred econometric specification, and the main results.

Even though all of those papers use regional variation in wage levels to identify the effects of the same national minimum wage increase, they often reach different conclusions. Parente (2024) estimates an increase in informal employment relative for formal employment, while Derenoncourt et al. (2021) and Engbom and Moser (2022) estimate zero effects on that margin. Parente (2024) also finds that wage dispersion rises for all workers (formal and informal), while Hinojosa (2019) finds the opposite. Neumark, Cunningham and Siga (2006) finds negative employment effects, while most other papers looking at that outcome do not.[14] Estimates of spillover effects are different in Hinojosa (2019) relative to Engbom and Moser (2022).

An even clearer illustration comes from the sensitivity analysis from Engbom and Moser (2022). The central component of their masterful study of the Brazilian minimum wage is a structural labor market model with search frictions and two-sided heterogeneity. But in addition, their motivating descriptive analysis is thorough and exceptionally transparent. Figure B.14 in their Online Appendix, replicated as Appendix Figure A2 here for the reader's convenience, shows that their effective minimum wage estimates vary dramatically depending on specification choices. Based on previous discussions in Subsections 3.4 and 3.5, it may be warranted to restrict attention to specifications that include time fixed effects and that use the median wage as the deflator, as opposed to the 90th percentile. Even then, estimates still differ substantially depending on the inclusion of region fixed effects or region-specific time trends.

In many cases, disagreements between papers studying the same empirical context can be informative. For example, estimates of returns to education may differ depending on whether an instrument is used or not, or which specific instrument is used. Researchers can rationalize those differences

[13] The search was executed on 10/29/2024. I used the keywords "Brazil minimum wage" and included published papers from 2000 or later, or working papers from 2019 or later, that (i) used the methods studied in this paper; and (ii) used data from the post-estabilization period (i.e., after 1994). Derenoncourt et al. (2021) was not present in the IDEAS data set, but was included because it was known to me through participation in conferences.

[14] Using a different design (still based on regional variation in wage levels), Hinojosa (2019) also finds negative and statistically significant effects on employment, though the magnitude of those effects is smaller than in Neumark, Cunningham and Siga (2006). Studying the period from 1982 to 2000 with PNAD data, Lemos (2009) also reports that estimates of employment effects are sensitive to the choice of estimator, with some specifications leading to negative employment effects and others yielding zero effects.

Table 7: Summary of existing literature studying the Brazilian minimum wage

Paper	Data	Preferred Specification	Results
Hinojosa (2019)	PME, 2002-2016; PNAD, 2002-2016	Eff. min. wage based on p60; region and time FEs; region-specific trends; some specifications instrument with lagged features of wage distribution.	Significant wage compression both for formal workers and all workers; spillovers up to the median.
Saltiel and Urzúa (2022)	RAIS, 2003-2012	Eff. min. wage based on p50; region and time FEs; region-specific trends; control for regional GDP and adult population.	No effects on employment.
Engbom and Moser (2022)	RAIS, 1996-2018	Eff. min. wage based on p90; region FEs; region-specific trends.	Significant wage compression for formal workers; spillovers up to p80 or more.
	PNAD, 1996-2012	As above, but based on p50.	No employment effects.
	PME, 2002-2012	As above, but based on p50.	No effects on transitions between formal and informal sectors.
Neumark, Cunningham and Siga (2006)	PME, 1996-2001	"Rolling" fraction affected based on wage distribution from previous quarter; region and time FEs; includes lags.	Large, significant reduction in employment of head of household; reduction in household income for families in the bottom 30%.
Derenoncourt, Gérard, Lagos, and Montialoux (2021)	PNAD, 1995-2015	Binary version of fraction-affected-style diff-in-diff based on minimum-to-median as of 1999; region and year FEs; controls for regional GDP and demographic composition.	No effect on informality rates.
Parente (2024)	PNAD 1996-2012	Binary version of fraction-affected with multiple treatments, grouping states according to "share under" as of 1999; region and year FEs; controls for demographic composition.	Large, significant increases in informality; increase in wage dispersion for all workers.

based on transparent discussions of identification assumptions: the presence or not of omitted variable bias, instrument validity, and who is the complier population for each instrument.

However, in the case of the papers in Table 7, it is difficult to argue which specification should be preferred over the others because the source of identification is not as transparent. In the first three papers in that table (using effective minimum wage designs), identification relies on the existence an unobserved shock that shifts the location of the latent wage distribution, but is orthogonal to determinants of its shape and of latent employment levels. None of those papers specify what economic factor could play that role. The most relevant shocks affecting Brazilian labor markets between 1995 and 2018 are changes in educational achievement, trade liberalization, the commodities boom, import competition from China in the manufacturing sector, increased enforcement of labor regulations, the expansion of technologies such as broadband, computers, and numerical control machines, and social programs like conditional cash transfers (Firpo and Portella, 2019). All of those factors are biased towards either low-, mid-, or high-skilled workers (meaning that they are not simply a locational shift of the log wage distribution), and all are likely to have independent effects on employment.

Similar arguments can be made for latter three papers in Table 7, which use difference-in-differences estimators analogous to the fraction affected design. The parallel trends assumption cannot be tested given the lack of a pre-treatment period with a stable minimum wage. In addition, the long-run estimates from the last two papers can be subtantially biased if there are other structural factors causing regional convergence (see Manzi et al., 2023; Abreha, Ornelas and Zaourak, 2024, for papers documenting regional convergence in Brazil in the 2000s).

One may wonder whether controlling for all of the shocks known to affect the Brazilian economy in this period would be sufficient to ensure unbiased estimation. Indeed, this is one motivation behind the use of region-specific trends in the effective minimum wage papers.[15] But using controls only tells us what is *not* being used as the source of identification. The estimators still rely on the existence of unobserved regional parameters satisfying some properties—except that now these parameters are defined as residuals, netting out the contribution of the included controls. Without knowing what that residual latent parameter is, it is impossible to evaluate the identification assumptions. In addition, the inclusion of controls can magnify biases arising from misspecification if they reduce the variance of identifying variation.

[15] Whether state-specific trends adequately capture the impact of those factors depends on whether their joint effect in each state is linear, or at least more linear than the impact of the minimum wage. I discuss this assumption in Appendix D.8.4 of Haanwinckel (2023).

6.2 Sensitivity analysis and comparison with a structural model

The previous discussion has two limitations. First, some of the conflicting results in existing literature may be coming from using different samples, different treatment of dynamic issues, or different ways to display the results (though Figure B.14 in Engbom and Moser 2022 goes a long way in addressing this limitation). Second, one would ideally want to benchmark the potentially problematic numbers using estimates that are not subject to the biases discussed in this paper.

In this final analysis, I address those limitations with quantitative exercises based on Brazilian data. To keep the analysis comparable to the remainder of the paper, the data includes exactly two periods: 1998 and 2012, the same years used in Haanwinckel (2023). This exercise aims at the long-run effects of the minimum wage, starting from a period a few years after the stabilization reforms and ending at a period where the economy has had a few years to adjust to the higher minimum wage levels (from Figures 4 and A1, bindingness measures are mostly stable following 2007). The outcomes of interest are formal employment and formal wages. They are measured using an administrative matched employer-employee data set (*"Relação Anual de Informações Sociais,"* henceforth RAIS).

Each spatial unit is a "microregion" from the Brazilian Statistical Bureau (IBGE). These units are analogous to commuting zones in the United States in terms of size, population, and in that less than 5% of workers live in one microregion and work in another. They are commonly used to determine the boundaries of labor markets in studies using Brazilian data (see e.g., Costa, Garred and Pessoa, 2016; Dix-Carneiro and Kovak, 2017; Ponczek and Ulyssea, 2021).[16]

The reference estimates of minimum wage effects come from Haanwinckel (2023). In that paper, I develop an imperfectly competitive labor market model with worker and firm heterogeneity to study the impact of rising schooling achievement, labor demand shocks, and the federal minimum wage on Brazilian labor markets. The predicted impacts of the federal minimum wage incorporate several channels discussed in previous literature: disemployment of very low productivity workers (i.e., truncation of the latent log wage distribution); mechanical wage increases for some workers (i.e., censoring); increase in employment for some worker types due to the presence of monopsony power; reallocation of workers to firms with higher wage premiums (as in Dustmann et al. 2021); changes in returns to skill, both in the aggregate economy and differential changes within firms; effects on the composition of firms operating in the economy (as in Aaronson et al. 2018); and pass-through effects to consumer prices (as in Harasztosi and Lindner 2019).

[16] Microregions are grouped to ensure constant boundaries throughout the studied period. Following the sample selection criterion in Haanwinckel (2023), I select regions with at least 15,000 workers in the RAIS data set in each year, and at least 1,000 in each educational group (less than high school, high school, and complete college or more). Those criterion select 151 out of 486 microregions that, together, account for 73% of the adult population in Brazil.

The magnitudes of those mechanisms are regulated by parameters which, in turn, are estimated using a simultaneous equation nonlinear least squares procedure. The estimator targets several endogenous outcomes: employment rates, measures of wage inequality between and within educational groups, the contribution of firm wage premiums to inequality, and a measure of sorting of high-wage workers to high-wage firms (the latter two being measured using the methodology developed by Kline, Saggio and Sølvsten, 2018). To account for potential confounders on the labor supply side, the model directly incorporates the role of education in affecting labor market outcomes such as the distribution of, and returns to, skill. The model also includes three types of unobserved labor demand shifters at the regional level, which are partly determined by initial sectoral shares in agriculture and manufacturing. These labor demand shocks absorb the influence of important transformations mentioned above, such as trade liberalization and technical change.[17]

Table 8 reports the average causal effects of the minimum wage predicted by that structural model, along with estimates of causal effects using different versions of the effective minimum wage and fraction affected estimators based on the same data. The structural model predicts sizable disemployment effects: a reduction of 3.3 percentage points in the share of the adult population (18 to 54 years old) employed in the formal sector. Both the baseline fraction affected design and the baseline effective minimum wage design also predict reductions in formal employment that are economically and statistically significant. However, the magnitudes differ a lot. For the fraction affected, the decline is smaller at 1.9 percentage point, while for the effective minimum wage it is larger at 6.6 percentage points.

The structural model predicts sizable compression of wages in the lower tail of the wage distribution, but no inequality-reducing effects on the upper tail.[18] The effective minimum wage estimate is very close for the p10-p50 quantile gap, but overestimates the compression of the p25-p50 gap relative to the structural model. It also predicts an *increase* in inequality in the upper tail. The baseline fraction affected estimate predicts large wage effects that are mostly uniform over the wage distribution. To the extent that there is wage compression, it happens in the upper tail.

Table 8 also reports estimates for alternative estimators. They encompass using binary versions of the fraction affected design, changing the fixed effects included in the effective minimum wage design, using an instrumental variables version of this estimator, or constructing the effective minimum wage using the 90th percentile as the deflator, instead of the median. Those specifications mirror the variety of implementations in existing literature, summarized in Table 7. Each of those

[17] Appendix Table A15 shows that the estimated model fits the original data well. See Haanwinckel (2023) for additional information on the model and identification.

[18] Wage levels rise by about nine log points in that part of the distribution because the disemployment of low-skilled workers means that the same quantile of the wage distribution corresponds to more skilled workers following the minimum wage hike.

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

Panel A: Fraction Affected and Gap Designs

	Outcome				
	Emp.	p10	p25	p50	p90
Structural model	-0.033	0.297	0.159	0.090	0.095
Fraction affected	-0.019	0.248	0.235	0.239	0.183
	(0.007)	(0.013)	(0.014)	(0.015)	(0.020)
Binary measure, 50% treated	-0.006	0.149	0.138	0.149	0.090
	(0.006)	(0.011)	(0.012)	(0.014)	(0.017)
Binary measure, 90% treated	-0.017	0.185	0.193	0.219	0.137
	(0.018)	(0.033)	(0.037)	(0.040)	(0.032)

Panel B: Effective Minimum Wage Designs

	Outcome			
	Emp.	p10 - p50	p25 - p50	p90 - p50
Structural model	-0.033	0.208	0.070	0.005
Effective min. wage	-0.061	0.188	0.127	0.076
	(0.022)	(0.020)	(0.015)	(0.033)
Effective min. wage, no region FE	-0.090	0.158	0.091	-0.059
	(0.009)	(0.014)	(0.009)	(0.022)
Effective min. wage, no time FE	0.101	0.200	0.103	-0.127
	(0.004)	(0.006)	(0.005)	(0.013)
AMS, no time FE	0.118	0.201	0.101	-0.149
	(0.006)	(0.006)	(0.005)	(0.013)

Panel C: Effective Minimum Wage based on percentile 90

	Outcome			
	Emp.	p10 - p90	p25 - p90	p50 - p90
Structural model	-0.033	0.202	0.064	0.000
Effective min. wage, p90	0.005	0.357	0.341	0.000
	(0.022)	(0.033)	(0.037)	(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. It is based on Brazilian data from the RAIS data set, using a sample of 151 microregions and two time periods: 1998 and 2012. Estimated standard errors (clustered at the microregion level) are shown in parentheses.

estimators finds different numbers for the causal effects of interest, helping explain the diversity of results in existing literature. If one takes the structural estimates as the correct benchmark, the baseline specifications of the effective minimum wage and fraction affected estimators seem to perform better than the alternative ones.

All in all, the analysis of the Brazilian case shows that when researchers try use regional variation in wage levels to identify the effects of a national minimum wage, the estimates may differ substantially based on specification choices for which economic theory provides little to no guidance. Based on these results, practioners are urged to pinpoint the source of identification in their empirical context, discuss threats to identification, and report results from several specifications to gauge the sensitivity to arbitrary choices.

7 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.

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) to construct the treatment intensity variable using only one pre-treatment year rather than averaging over years; (ii) to check for differential pre-trends; (iii) to assess the possibility of regression to the mean/regional convergence and, if needed, account for that in the estimation model; (iv) to check for trends in the dispersion of log wages in the pre-treatment period, which may cause biases as discussed in Section 4.3; and (v) *not* to use a binary version of the treatment intensity variable. Even if the econometrician takes such precautions, the estimator may be subject to economically significant misspecification biases arising from functional form assumptions, especially when the minimum wage increase is significant. Econometricians may consider reporting results of both the "fraction affected" and "gap" designs, and also a version of those estimators that included a quadratic term, 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, it may be tempting to 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 ones include time fixed effects and construct 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 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, as explained in Appendix C.

In contexts that are not well-suited to either estimator, the econometrician may consider two alternative strategies. One is to employ 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. As discussed at the end of Section 2, such designs do not directly identify the policy counterfactual analyzed in this paper; recovering those effects from the estimated coefficients requires additional structural assumptions. Moreover, the econometrician should be aware that some of the concerns discussed earlier—such as regression to the mean—may still apply. See Dustmann et al. (2021) for an example of how this issue can be addressed in a worker-level analysis.

The second alternative strategy is to estimate a parametric structural 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 for 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 that adequately controls for the issues studied in this paper.

References

Aaronson, Daniel, Eric French, Isaac Sorkin, and Ted To. 2018. "Industry Dynamics and the Minimum Wage: a Putty-Clay Approach." *International Economic Review*, 59(1): 51–84.

Abreha, Kaleb, Rafael Amaral Ornelas, and Gabriel Roberto Zaourak. 2024. "Regional Convergence in Brazil." The World Bank Policy Research Working Paper Series 10904.

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.

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.

Blundell, Richard, and Monica Costa Dias. 2009. "Alternative Approaches to Evaluation in Empirical Microeconomics." *Journal of Human Resources*, 44(3): 565–640.

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

Cadena, Brian C. 2014. "Recent immigrants as labor market arbitrageurs: Evidence from the minimum wage." *Journal of Urban Economics*, 80: 1–12.

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.

Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer. 2019. "The Effect of Minimum Wages on Low-Wage Jobs." *The Quarterly Journal of Economics*, 134(3): 1405–1454.

Corseuil, Carlos Henrique, Miguel Foguel, and Marcos Hecksher. 2015. "EFEITOS DOS PISOS SALARIAIS ESTADUAIS SOBRE O MERCADO DE TRABALHO: UMA NOVA ABORDAGEM EMPÍRICA." *Economia Aplicada*, 19(1): 131–169.

Costa, Francisco, Jason Garred, and João Paulo Pessoa. 2016. "Winners and losers from a commodities-for-manufactures trade boom." *Journal of International Economics*, 102: 50–69.

de Chaisemartin, C, and X D'HaultfŒuille. 2018. "Fuzzy Differences-in-Differences." *The Review of Economic Studies*, 85(2): 999–1028.

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

de Chaisemartin, Clément, Diego Ciccia, Xavier D'Haultfœuille, and Felix Knau. 2024. "Two-way Fixed Effects and Differences-in-Differences Estimators in Heterogeneous Adoption Designs."

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.

Dix-Carneiro, Rafael, and Brian K. Kovak. 2017. "Trade Liberalization and Regional Dynamics." *American Economic Review*, 107(10): 2908–46.

Dube, Arindrajit, and Attila Lindner. 2024. "Chapter 4 - Minimum wages in the 21st century." In . Vol. 5 of *Handbook of Labor Economics*, , ed. Christian Dustmann and Thomas Lemieux, 261–383. Elsevier.

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.

Firpo, Sergio, and Alysson Portella. 2019. "Decline in Wage Inequality in Brazil: A Survey." World Bank Policy Research Working Paper 9096.

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

Ghanem, Dalia, Désiré Kédagni, and Ismael Mourifié. 2023. "Evaluating the Impact of Regulatory Policies on Social Welfare in Difference-in-Difference Settings."

Giupponi, Giulia, Robert Joyce, Attila Lindner, Tom Waters, Thomas Wernham, and Xiaowei Xu. 2024. "The Employment and Distributional Impacts of Nationwide Minimum Wage Changes." *Journal of Labor Economics*, 42(S1): S293–S333.

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

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

Hinojosa, Paul Garcia. 2019. “The Role of the Minimum Wage on the Declining Wage Inequality in Latin America: Evidence from Brazil.” *Review of Economic Analysis*, 11(1): 39–82.

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.

Lemos, Sara. 2009. “Comparing employment estimates using different minimum wage variables: the case of Brazil.” *International Review of Applied Economics*, 23(4): 405–425.

Manzi, Giancarlo, Ahmed Alsayed, Humberto Martins, and Giorgio Saibene. 2023. “Regional economic convergence in federation contexts: a comparative analysis of Brazil and the European Union.” *Regional Studies*, 57(11): 2269–2284.

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

Moura, Rodrigo Leandro de, and Marcelo Côrtes Neri. 2008. “Efetividade do "salário mínimo estadual": uma análise via regressões quantílicas para dados longitudinais.” *Economia Aplicada*, 12(2): 239–273.

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.

Neumark, David, Wendy Cunningham, and Lucas Siga. 2006. “The effects of the minimum wage in Brazil on the distribution of family incomes: 1996-2001.” *Journal of Development Economics*, 80(1): 136–159.

Parente, Rafael Machado. 2024. “Minimum Wages, Inequality, and the Informal Sector.” International Monetary Fund IMF Working Papers 2024/159.

Piqueras, Jon. 2024. “Search Effort and the Minimum Wage.” Unpublished manuscript.

Ponczek, Vladimir, and Gabriel Ulyssea. 2021. “Enforcement of Labour Regulation and the Labour Market Effects of Trade: Evidence from Brazil.” *The Economic Journal*, 132(641): 361–390.

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.

Saltiel, Fernando, and Sergio Urzúa. 2022. "Does an Increasing Minimum Wage Reduce Formal Sector Employment? Evidence from Brazil." *Economic Development and Cultural Change*, 70(4): 1403–1437.

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

Sun, Liyang, and Jesse M. Shapiro. 2022. "A Linear Panel Model with Heterogeneous Coefficients and Variation in Exposure." *Journal of Economic Perspectives*, 36(4): 193–204.

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

Online Appendix

A What do empirical minimum wage papers identify?

In this appendix, I argue that the central goal in all of the recently published papers mentioned early in the introduction is to measure treatment effects in a way that is equivalent to the definition used in this paper (Equation 2).

Bosch and Manacorda (2010): Their goal is to explain average changes in quantile gaps, where averages are taken over local labor markets—exactly in the same way as in Equation 2. Those effects are estimated based on coefficients obtained from the effective minimum wage, in the same way that is described in Subsection 3.1. They write: *"Our analysis reveals that a substantial part of the growth in inequality between 1989 and 2001, and essentially all the growth in inequality in the bottom end of the distribution, is due to the steep decline in the real value of the minimum wage."*

Dinkelman and Ranchhod (2012): The paper uses an identification strategy similar to the Fraction Affected/Gap designs, but with the continuous treatment intensity variable being the initial effective minimum wage from Lee (1999). The authors do not explicitly specify how they infer aggregate effects from the regression coefficients, as that requires a normalization to deal with the "missing intercept" issue. That said, the language in the paper suggests that it attempts to measure aggregate effects similar the definition in Equation 2. Specifically, the paper begins with the question *"What happens to wages and employment in the informal sector after the introduction of a minimum wage in that sector?"*, indicating that the object of interest is the aggregate impact of a ceteris paribus introduction of a minimum wage. The introduction summarizes their strategy and results as follows: *"We evaluate the effects of South Africa's 2002 minimum wage law for domestic workers by exploiting time-series variation in the application of the law and pre-existing cross-sectional variation related to the intensity of the law to identify wage and employment effects. (...) We complement the before-after analysis with a difference-in-differences strategy that adopts the methods in Lee (1999) to statistically examine the effects of the law. (...) We find that wages increase by a statistically significant 13–15% in the wake of the law."*

Dustmann et al. (2021): The authors write in the introduction: *"We investigate the wage, employment, and reallocation effects of the introduction of a nationwide minimum wage in Germany that affected 15% of all employees. Based on identification designs that exploit variation in exposure across individuals and local areas, we find that the minimum wage raised wages but did not* *lower employment."* In the introduction, speaking specifically about their Gap design exploiting pre-existing variation in regional wage levels, they write: *"Findings from an analysis that exploits variation in the exposure to the minimum wage across 401 local areas (similar to Card 1992) corroborate our findings from the individual-level analysis: the minimum wage boosted wages but did not reduce employment in heavily affected areas relative to less affected. Our findings therefore do not confirm the fears of many economists that the minimum wage would cause substantial job losses. Rather, our findings support the idea that the minimum wage helped reduce wage inequality without reducing employment across individuals and across local areas."* Given this language, it is reasonable to argue that the authors are measuring aggregate, ceteris paribus effects of the introduction of the national minimum wage in Germany.

Bailey, DiNardo and Stuart (2021): Their abstract reads: *"This paper examines the short- and longer-term economic effects of the 1966 Fair Labor Standards Act (FLSA), which increased the national minimumw age to its highest level of the twentieth century and extended coverage to an additional 9.1 million workers. Exploiting differences in the "bite" of the minimum wage owing to regional variation in the standard of living and industry composition, this paper finds that the 1966 FLSA increased wages dramatically but reduced aggregate employment only modestly. However, some evidence shows that disemployment effects were significantly larger among African American men, 40% of whom earned below the new minimum wage."* A passage in the introduction makes the aggregate-effects, ceteris-paribus interpretation even more explicit: *"Across the United States, our estimates suggest that average wages increased by 6.5% because of the 1966 FLSA."*

Engbom and Moser (2022): The central results of the paper come from a counterfactual simulations of a ceteris paribus minimum wage increase, based on an estimated structural model. Thus, their findings correspond to the definition in Equation 2 in a setting with a single "national" region in the model (such that no averaging across locations is necessary) and with a static model (such that $\mathbf{ATE}_0$ and $\mathbf{ATE}_1$ are identical).

Giupponi et al. (2024): From their abstract: *"We assess the impact of nationwide minimum wages on employment throughout the whole wage distribution by exploiting geographical variation in the level of wages. We find a substantial increase in wages at the bottom of the wage distribution, while we detect a small, statistically insignificant negative effect on employment."* The authors specify in detail how they can measure aggregate effects, as opposed to only relative effects, using their methodology: *"Since in practice no region is entirely unaffected by a national minimum wage policy, our approach shares the characteristic of other regional variation approaches of identifying the relative effect of the minimum wage on employment in lower-wage areas compared with higher-*

wage ones. Retrieving an absolute effect requires additional assumptions that we describe below."

B Simulation details for the Normal-markdown model

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 where otherwise noted.

B.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 $m = 0.7$.

B.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.

B.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 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

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	$\mu_{r,t}, w_{0.5,r,t}$
Table 1																
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
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 B 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 and A4																
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 A5																
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 A6																
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 B for details.

of minimum wage on employment will become negative.

B.4 The Normal-markdown model with state-level minimum wages

For the exercise shown in Table A3, 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 A3. 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.

B.5 Generation of 500 DGPs for the exercise in Section 5

Each of the 500 DGPs corresponds to the parameters determining the joint distribution of the vector of region-specific parameters $[\mu_{r,0}, \sigma_{r,0}, \mu_{r,1}, \sigma_{r,1}]$, coupled with a markdown parameter m, the positive employment parameters P_{base}, P_{height}, and the levels of the national minimum wage, mw_0, mw_1. I randomly draw these features of the DGP as specified below:

Average location $\mu_{r,t}$: Set to zero in both periods in all DGPs.

Standard deviation of location $\mu_{r,t}$: That parameter is uniformly drawn in the interval (0.056, 0.185) for each of the two periods (that is, the variability of location parameters between regions may change over time). That interval corresponds to 0.5 times the corresponding parameter for the US in 2004 and 1.5 times for the US in 1989, respectively.

Correlation between initial and final location, $Corr(\mu_{r,1}, \mu_{r,2})$: Uniformly drawn in the interval (0.838, 0.95). The number observed for the US, 0.894, corresponds to the center of the interval. The maximum was set to 0.95 to avoid DGPs where the initial and final location parameters are too similar, which means that the variance of location *shocks* would be too small. In that scenario, the effective minimum wage design performs poorly. Alternatively, if the correlation is too low, the lack of persistence across time handicaps the fraction affected design. With this narrow range, all DGPs have persistence levels similar to what can be inferred from the US between 1989 and 2004.

Average dispersion $\sigma_{r,t}$: For the first period, that average is uniformly drawn in the interval (0.255, 0.813). That interval corresponds to 0.5 times the corresponding parameter for the US in 2004 and 1.5 times for the US in 1989, respectively. Thus, the simulated DGPs are substantially heterogeneous in the within-region dispersion of latent log wages, corresponding to different empirical applications (e.g., Germany vs. Brazil).

The *change* in average dispersion from the first to the second period is uniformly drawn in the interval (-0.047, 0.047). The extreme values correspond to 1.5 times the corresponding statistic observed in the US. Thus, trends in the dispersion of latent log wages—which I showed to be problematic for the fraction affected design—can never be much more significant than what can be inferred from US data and are typically smaller than that.

Standard deviation of dispersion $\sigma_{r,t}$: Uniformly drawn in the interval (0.012, 0.074) for each of the two periods. That interval corresponds to 0.5 times the corresponding parameter for the US in 1989 and 1.5 times for the US in 2004, respectively.

Correlation between initial and final dispersion, $Corr(\sigma_{r,1},\sigma_{r,2})$: Uniformly drawn in the interval (0.228, 0.684). That interval is the corresponding parameter for the US multiplied by 0.5 and 1.5, respectively.

Contemporaneous correlation between location and dispersion, $Corr(\mu_{r,t},\sigma_{r,t})$: Uniformly drawn in the interval (0.0, 0.2). Thus, while I allow for positive correlations between location and dispersion, the magnitudes are small compared to the data from the US and Brazil. The corresponding numbers for the US are 0.076 in 1989 and 0.264 in 2004. Using Brazilian data from 1998 (see Section 6 of the paper for details), I find a correlation of 0.339 between the mean log wage and standard deviation of log wage in the 151 Brazilian microregions included in my sample. One may be worried that the minimum wage itself generates this correlation, such that it does not reveal information about the *latent* distribution of log wages. To test this hypothesis, I calculate the correlation using only 88 microregions where the share of workers earning at most the minimum wage plus 30 log points is 10% or less. Rather than decreasing, the correlation between mean log wage and standard deviation of log wage increases to 0.528 in this subsample, suggesting that the correlation between location and dispersion of latent log wages in Brazil is significantly higher than 0.2. Thus, the simulated DGPs are conservative in that regard.

Markdown parameter m: Uniformly drawn in the interval (0.1, 0.9). The lower half of this interval may seem implausible under partial-equilibrium models of monopsonistic firms, as they would imply labor supply elasticities to the firm below one. I include such numbers to allow for the possibility that the minimum wage causes significant wage increases in the lower tail of the distribution without causing any employment effects. One could interpret the very-low-markdown model as an approximation for a richer model where the minimum wage reallocates workers from firms that would be paying less than the minimum wage to other firms paying exactly the minimum wage.

Positive employment effect parameters: P_{base} is uniformly drawn in the interval (0.01, 0.5), implying that the minimum wage may attract workers earning up to 50 log points above the minimum wage in the highest-draw DGP. P_{height} is uniformly drawn in the interval (0.01, 0.5) as well, implying an extra mass of workers of up to 50% of the density of log wages implied by the latent log wage distribution (right above the minimum wage). See Subsection B.3 above for details.

Minimum wage levels: The initial log minimum wage is uniformly drawn in the interval (-1.5, -0.5), where zero corresponds to the center of the latent log wage distribution. The increase in the log minimum wage—which should be interpreted as an increase relative to average TFP growth—is uniformly drawn in the interval (0.15, 0.5). Thus, the simulations include minimum wage increases that are relatively small, at 15 log points, up to more significant increases similar in magnitude to that observed in Brazil between 1998 and 2012.

Simulation procedure: Sometimes, the combination of meta-parameters may be invalid because it leads to a covariance matrix for the $[\mu_{r,0}, \sigma_{r,0}, \mu_{r,1}, \sigma_{r,1}]$ vector that is not positive definite. When that happens, that draw is discarded and a substitute is used.

After obtaining a valid DGP candidate, I perform 1,000 simulations, with each simulation corresponding to a sample of 200 regions and two periods (as in the rest of the paper). Using this sample, I calculate average treatment effects as defined in Section 2. The DGP is included in the analysis if the minimum wage effects on employment are at least 0.005 (half a percentage point), or if any of the spillovers relative to the median wage is at least 0.05 (five log points). If that is not the case—typically because of a low starting minimum wage coupled with a slight minimum wage increase—then that DGP is discarded, and I randomly draw another one to substitute for it. I also discard and find a replacement if minimum wage effects on employment are more than 0.04. The process is completed when there are 500 DGPs with non-trivial, but not excessively large, minimum wage effects.

C Regional minimum wages and instrumental variables

If the data includes region-level minimum wage changes, 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 A3 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 B.4. Panel A showcases the baseline model, where there is a slight correlation

Table A3: 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	0.050	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	0.041	0.008	-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	0.053	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)	-0.021	0.053	0.010	-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 B.4 for details.

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 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.

D Simulations of the canonical model of labor demand

D.1 Model description

Overview. 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.

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.

The critical difference between this model and the Normal-markdown model used in other simulations in the paper 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.

Mathematical description. 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:

$$\begin{aligned} p_{i,r,t} &= \log F_i\left(s_r E(p_{1,r,t}, mw_t), (1-s_r)E(p_{2,r,t}, mw_t)\right) \qquad i \in \{1,2\} \\ \text{where } E(p, mw) &= \int_{mw_t - p}^{\infty} \exp(e)\phi\left(\frac{e}{D}\right) de \\ \text{and } F_i(L_1, L_2) &= \frac{dF(L_1, L_2)}{dL_i} \end{aligned}$$

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.

D.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.

D.3 Results

Table A4 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 5 above, illustrating ideal scenarios for the Fraction Affected and Gap designs. Also, as specified above, 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 5. 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.

Table A5 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.

Table A4: 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	0.056	0.050	0.029
	(0.001)	(0.002)	(0.002)	(0.001)	(0.000)

Notes: This table is similar in structure to Table 5, 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 D for details.

Table A5: 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	-0.238	-0.102	-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)	(0.012)	(0.006)	(0.003)
Effective min. wage, no region FE	-0.010	0.002	-0.007	0.078
	(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	0.449	0.484	0.110	-0.198
	(0.023)	(0.013)	(0.008)	(0.012)
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	0.279	0.068	-0.097
	(0.007)	(0.018)	(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	0.258	0.313	0.153	-0.202
	(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	0.215	0.104	-0.100
	(0.003)	(0.009)	(0.004)	(0.010)
Effective min. wage, no region FE	-0.097	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 D for details.

E Regression to the mean and regional convergence in the fraction affected design

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 not always investigated 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 A6 illustrates this issue in the context of the Normal-markdown simulation model. It reports results for the fraction affected design; Table A10 in Appendix F shows similar results for the Gap design. In Panel A, regions only differ in a time-invariant location parameter μ_r. Panel B introduces 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 explores 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 param-

Table A6: 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	0.058	0.041	0.036	0.032
	(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	0.013	0.007	0.003
Fraction affected	-0.008	0.072	0.047	0.031	0.006
	(0.001)	(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	0.064	0.044	0.030	0.005
	(0.001)	(0.008)	(0.008)	(0.008)	(0.012)

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.

eters 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 A11 illustrates this concept using a placebo exercise that parallels Table A6. The fraction affected measure is calculated as if the national minimum wage would increase by 0.2, but in reality, the minimum wage remains stable. 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. One way to do so is to mirror the individual-level design in Dustmann et al. (2021). Alternatively, one may formally estimate the time series properties of the error term using pre-treatment data.

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.[19]

[19] 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.

F Additional Tables and Figures

Table A7: 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	0.117	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	0.016	-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	-0.002	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 5, 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 A8: 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	0.181	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
	(0.001)	(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 5, 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 A9: 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	0.021	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	0.030	-0.010	-0.004	-0.002
	(0.000)	(0.006)	(0.002)	(0.002)	(0.002)
Quadratic on GAP	0.004	0.041	-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	-0.002	0.149	0.039	0.002	0.001
Quadratic on FA	-0.013	0.246	0.125	0.006	0.004
	(0.001)	(0.004)	(0.014)	(0.004)	(0.004)
Quadratic on GAP	-0.019	0.167	0.117	0.012	0.006
	(0.001)	(0.005)	(0.006)	(0.002)	(0.002)

Notes: This table is analogous to Table 5, 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 A10: 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	0.024
	(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	0.023	0.008
	(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	0.045	0.031	0.023	0.009
	(0.001)	(0.006)	(0.006)	(0.006)	(0.009)

Notes: This table is analogous to Table A6, except that it shows results for the Gap design instead of the Gap design.

Table A11: Sensitivity of the Fraction Affected 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
Fraction affected	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
Fraction affected	0.001	0.027	0.028	0.029	0.029
	(0.000)	(0.008)	(0.009)	(0.009)	(0.009)
Panel C: Adding dispersion differences and shocks, stable distributions					
True average causal effect	-0.000	0.000	0.000	0.000	0.000
Fraction affected	0.002	0.041	0.034	0.024	0.002
	(0.001)	(0.010)	(0.008)	(0.008)	(0.012)
Panel D: Average dispersion falls over time					
True average causal effect	-0.000	0.000	0.000	0.000	0.000
Fraction affected	0.005	0.034	0.031	0.023	0.002
	(0.001)	(0.010)	(0.009)	(0.009)	(0.013)

Notes: This table is analogous to Table A6, but it reports a placebo scenario with no increase in the national minimum wage. The fraction affected measure, however, is calculated as if the national log minimum wage would increase by 0.2 between periods (as is the case in Table A6).

Table A12: Quantifying *large* biases under 500 alternative DGPs

Outcome	Emp	p10-p50	p25-p50	p90-p50
Panel A: 128 DGPs with positive employment effects				
Average causal effects	0.013	0.153	0.034	0.000
Eff. min. wage: share with positive bias	0.09	0.32	0.20	0.00
· : share with negative bias	0.03	0.00	0.00	0.58
Fraction affected: share with positive bias	0.00	0.30	0.24	0.00
· : share with negative bias	0.53	0.15	0.00	0.14
Panel B: 241 DGPs with negative employment effects				
Average causal effects	-0.018	0.151	0.039	-0.009
Eff. min. wage: share with positive bias	0.66	0.24	0.19	0.00
· : share with negative bias	0.00	0.00	0.00	0.62
Fraction affected: share with positive bias	0.12	0.26	0.21	0.00
· : share with negative bias	0.41	0.20	0.00	0.11
Panel C: 131 DGPs with small emp. effects but large wage effects				
Average causal effects	0.000	0.203	0.051	-0.002
Eff. min. wage: share with positive bias	0.13	0.21	0.18	0.01
· : share with negative bias	0.00	0.00	0.00	0.50
Fraction affected: share with positive bias	0.00	0.37	0.34	0.00
· : share with negative bias	0.21	0.24	0.02	0.15

Notes: This table is similar to Table 6, except the classification of estimators into the "biased" category is more stringent.

Table A13: Quantifying biases: discarding DGPs with upper-tail spillovers

Outcome	Emp	p10-p50	p25-p50	p90-p50
Panel A: 25 DGPs with positive employment effects				
Average causal effects	0.011	0.117	0.026	0.000
Eff. min. wage: share with positive bias	0.12	0.16	0.00	0.00
· : share with negative bias	0.04	0.00	0.00	0.08
Panel B: 20 DGPs with negative employment effects				
Average causal effects	-0.012	0.095	0.025	-0.006
Eff. min. wage: share with positive bias	0.40	0.00	0.00	0.00
· : share with negative bias	0.00	0.00	0.00	0.00
Panel C: 32 DGPs with small emp. effects but large wage effects				
Average causal effects	0.000	0.198	0.078	-0.004
Eff. min. wage: share with positive bias	0.09	0.09	0.03	0.03
· : share with negative bias	0.00	0.00	0.00	0.00

Notes: This table is similar to Table 6, except that it only considers the subset of DGPs where the effective minimum wage estimate of p90-p50 effects is less than 0.05 in magnitude.

Table A14: Quantifying biases: discarding DGPs with upper-tail wage effects

Outcome	Emp	p10-p50	p25-p50	p90-p50
Panel A: 95 DGPs with positive employment effects				
Average causal effects	0.012	0.139	0.028	0.002
Fraction affected: share with positive bias	0.00	0.34	0.27	0.00
· : share with negative bias	0.82	0.19	0.00	0.01
Panel B: 161 DGPs with negative employment effects				
Average causal effects	-0.018	0.139	0.035	-0.008
Fraction affected: share with positive bias	0.22	0.22	0.22	0.00
· : share with negative bias	0.46	0.29	0.01	0.00
Panel C: 89 DGPs with small emp. effects but large wage effects				
Average causal effects	0.000	0.203	0.044	-0.001
Fraction affected: share with positive bias	0.00	0.37	0.45	0.00
· : share with negative bias	0.38	0.38	0.01	0.00

Notes: This table is similar to Table 6, except that it only considers the subset of DGPs where the fraction affected estimate of wage effects on the 90th quantile is less than 0.05 in magnitude.

Table A15: Quality of fit of the task-based, monopsonistic model

	Data		Model		R2
	1998	2012	1998	2012	Model
Moments	(1)	(2)	(3)	(4)	(5)
Wage inequality measures					
Secondary / less than secondary	0.498	0.168	0.478	0.168	0.755
Tertiary / secondary	0.965	1.038	0.981	0.954	0.127
Within less than secondary	0.41	0.241	0.401	0.233	0.607
Within secondary	0.684	0.355	0.67	0.331	0.848
Within tertiary	0.702	0.624	0.701	0.637	0.139
Two-way fixed effects decomposition					
Variance establishment effects	0.126	0.054	0.126	0.04	0.586
Covariance worker, estab. effects	0.052	0.046	0.056	0.059	0.374
Formal employment rates					
Less than secondary	0.266	0.337	0.266	0.335	0.953
Secondary	0.435	0.508	0.435	0.508	1.0
Tertiary	0.539	0.629	0.539	0.63	0.89
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	0.046	0.084	0.528
Share < log min. wage + 0.30	0.086	0.212	0.107	0.201	0.873

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.

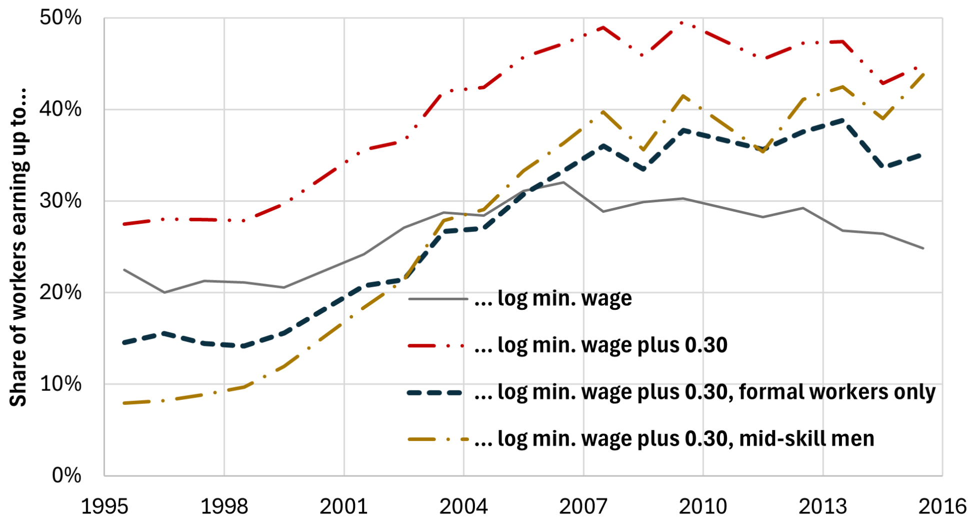

Figure A1: Alternative bindingness measures for the federal minimum wage in Brazil

Notes: All series are constructed based on yearly PNAD survey data from IBGE (processed using the DataZoom tool from PUC-Rio), using the September minimum wage value for each year. They show the share of workers with monthly earnings in their main jobs up to a given threshold based on the minimum wage. The base sample includes all workers and self-employed individuals between ages 18 and 54. The third series conditions on workers being employed in the formal sector. The fourth series conditions on workers being male, up to 30 years old, and with exactly 11 years of schooling (i.e., complete high school).

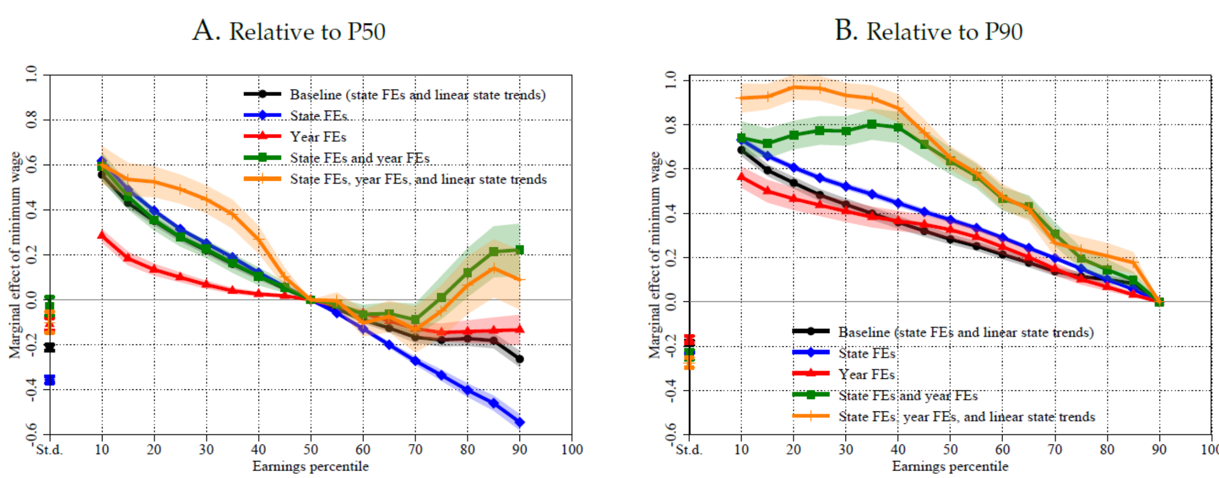

Figure A2: Copy of Figure B.14 in Engbom and Moser (2022)—sensitivity of the effective minimum wage design to specification choices