Monopsony, Markdowns, and Minimum Wages

Ester Faia
Goethe University Frankfurt and CEPR

Benjamin Lochner
IAB, FAU Erlangen-Nürnberg, and IZA

Benjamin Schoefer
UC Berkeley

December 22, 2025

Preliminary.
Please do not circulate.

Abstract

This paper presents the first *direct* test of two interlinked predictions at the core of the monopsony theory of the labor market: (i) that firms exploit wage-setting power by marking down wages below the marginal revenue product of labor, and (ii) that exogenous wage constraints, if binding, eliminate markdowns. Our research design revisits the 2015 introduction of a high minimum wage in Germany. Drawing on a monopsony model, we derive an empirically tractable difference-in-differences specification that provides a quantitative benchmark for the firm-level markdown response. Our main result is that empirical markdowns respond only 0–25% as much as the monopsony model would have predicted. Hence, at least for the labor market segment we study, (i) markdowns largely reflect other distortions than monopsony, (ii) markdowns are mismeasured, (iii) minimum wages induce widespread labor shortages, or (iv) the standard monopsony model does not provide a full, realistic account of the labor market.

We thank Ihsaan Bassier, Attila Lindner, and Matthias Mertens for useful comments, as well as the audience at the CEPR Macroeconomics and Growth Annual Symposium 2024.

1 Introduction

There is a growing recognition that firms hold some degree of monopsony power over their employees (e.g., Manning, 2013; Azar and Marinescu, 2024; Kline, 2025). Basic economic theory implies that profit-maximizing firms exploit their wage-setting power by marking down wages below the marginal revenue product of labor (e.g., Robinson, 1933; Card, Cardoso, Heining and Kline, 2018; Berger, Herkenhoff and Mongey, 2022; Lamadon, Mogstad and Setzler, 2022; Yeh, Macaluso and Hershbein, 2022). Hence, exogenous wage constraints can eliminate wage-setting power and the markdowns stemming from it. Wage setters become wage takers. This influential view has shaped the interpretation of empirical studies of minimum wages and wage floors in particular regarding their effects on employment and misallocation.[1] Yet, there exists no direct empirical estimate of how wage constraints, such as minimum wages, actually affect wage markdowns.

Our paper aims to provide the first *direct* and dedicated test of how wage constraints affect firms' wage markdowns. We study the 2015 introduction of a high minimum wage in Germany—of 8.50 EUR then.[2] We structurally ground a difference-in-differences (DiD) research design in a model of monopsonistic firms. We show that the firm's predicted change in log markdowns is given by the negative product of (i) the fraction of its pre-reform wage bill subject to the minimum wage and (ii) its pre-reform log markdown. This predicted markdown change is our treatment intensity variable. We then implement this DiD design, akin to a pass-through regression, in IAB establishment panel data before and after the 2015 reform. We merge on hourly wage data to measure a firm's share of the wage bill exposed to the minimum wage. To construct markdowns, we use the proxies based on first-order conditions (see, e.g., De Loecker and Warzynski, 2012; Dobbelaere and Mairesse, 2013; Yeh, Macaluso and Hershbein, 2022; Rubens, Wu and Xu, 2025) and estimate output elasticities following Levinsohn and Petrin (2003) with the Ackerberg, Caves and Frazer (2015) adjustment (and consider robustness to other methods).

We find strikingly limited evidence for the core prediction of monopsony models, that firms'

[1] For examples of the many studies emphasizing the role of monopsony markdowns in mediating minimum wage effects indirectly, see Dickens, Machin and Manning (1999); Manning (2021); Dube and Lindner (2024); Bassier and Budlender (2025); Berger, Herkenhoff and Mongey (2025b). For empirical studies on minimum wages and wage floors, see, e.g., Card and Krueger (2000); Neumark and Wascher (2000); Harasztosi and Lindner (2019); Dube, Lester and Reich (2010); Cengiz, Dube, Lindner and Zipperer (2019); Bailey, DiNardo and Stuart (2021); Dustmann, Lindner, Schönberg, Umkehrer and Vom Berge (2022); Engbom and Moser (2022); Azar, Huet-Vaughn, Marinescu, Taska and Von Wachter (2024); Faia and Pezone (2024); Dustmann, Giannetto, Incoronato, Lacava, Pezone, Saggio and Schoefer (2025). For the role of heterogeneous, e.g., monopsony, markdowns in misallocation and the (at least implicit potential) role of exogenous wages, see Hsieh and Klenow (2009); Berger, Herkenhoff and Mongey (2022); Bachmann, Bayer, Stüber and Wellschmied (2025); Berger, Herkenhoff and Mongey (2025b).

[2] The current levels are 12.82 and 13.90 EUR as of January 1, 2025 and 2026, respectively. Crucially, up until its introduction in 2015, Germany did not have a minimum wage, but relied on increasingly eroding collective bargaining wage floors, which firms seeking to pay low wages were able to evade by not acceding to optional employer associations, with modest coverage shares, particularly in low-wage firms (see, e.g., Jäger, Noy and Schoefer, 2022). The previous evaluations of the minimum wage introduction have focused on employment and wage inequality (Caliendo, Fedorets, Preuss, Schröder and Wittbrodt, 2017; Bosch, 2018; Caliendo, Schröder and Wittbrodt, 2019; Garloff, 2019; Bossler and Gerner, 2020; Burauel, Caliendo, Grabka, Obst, Preuss, Schröder and Shupe, 2020; Dustmann, Lindner, Schönberg, Umkehrer and Vom Berge, 2022; Bossler and Schank, 2023). Bossler, Gürtzgen, Lochner, Betzl and Feist (2020); Haelbig, Mertens and Müller (2024) study the effects on firm-level outcomes but not on markdowns.

markdowns disappear for workers whose wage becomes exogenous due to the minimum wage. This prediction would have yielded a coefficient on treatment intensity of one in our DiD regression. Quantitatively, empirical markdowns move just about 0.25 (SE 0.051) as much as predicted by standard monopsony models, and just 0.047-0.087 (SEs 0.104 and 0.110) as much for our preferred sample of initially monopsonistic firms (with log markdowns initially below zero, i.e., with wages below the marginal revenue product pre-reform) and after a basic pre-trend adjustment. An alternative and equally plausible specification that classifies firms by their 2011 rather than 2013 markdowns reveals substantial room for mean reversion to bias upward the baseline estimate; this earlier sorting year yields point estimates even closer to zero and even slightly negative.

The overall finding—that wage floors hardly eliminate measured markdowns—-holds throughout our event horizon and across multiple robustness checks such as varying the production function estimation method. It also is robust to zooming in on firms with below-median or negative employment growth into the reform. This subsample is informative as these firms are less subject to the concern that labor supply may ration their employment and hence could prop up their markdowns even though the minimum wage binds for them.

There are several potential interpretations of this result. First, markdowns may reflect economic frictions such as adjustment costs, besides monopsony. Second, our markdown variables may be mismeasured. For instance, production function estimation may yield misleading output elasticities, or intermediate inputs may not be competitively supplied and thereby confound our markdown identification. Third, markdowns may be propped up by widespread labor shortages induced by the minimum wage, although we find similar results among shrinking firms. Fourth, the standard monopsony model may not hold up in the data for the labor market segment we study for other reasons, which we—or perhaps follow-up research—have yet to identify. We assess these rationalizations and discuss their implications in the conclusion.

The paper is organized as follows. In Section 2, we present a simple model to explain our research design. Section 3 then develops the full model, which will also guide our empirical analysis. In Section 4, we describe the introduction of the minimum wage. Section 5 describes the research design, data, sample, and variables. Section 6 presents our results. Section 7 concludes.

2 The simple economics of our research design

We first present a simple but rich formal model of monopsony with and without a minimum wage. In this step, we explain markdowns and their predicted disappearance in response to a binding minimum wage, with wage-setting firms becoming wage takers. We also briefly touch on labor demand rationing and its implication for markdowns, which our full model below will sidestep. The model abstracts from the richer features we introduce in the full model in Section 3, such as richly heterogeneous labor and firms, non-labor inputs, etc., all in the shadow of a general equilibrium model.

Monopsony and markdowns without a minimum wage. Figure 1 illustrates the intuitions in a standard graphical representation of labor market monopsony, without the minimum wage (Panel (a)) and with (Panel (b)). It shows the *firm-specific* equilibrium. Panel (a) shows three key curves as a function of firm-level employment n, and is understood in logs (denoted by hats): namely, the (diminishing) marginal revenue product of labor $\widehat{MRPN}$, the upward-sloping *firm-specific* labor supply giving the wage $\hat{w}$, and the marginal cost of labor $\widehat{mcn}$. The marginal cost of labor is an upward shift of the wage curve given by labor supply, and captures the source of monopsony: firms internalize their labor demand effect on all their workers' wages.

Formally, these curves result from the following simple model. Firms have productivity (TFP) z and hire one type of labor to produce $y = zn^{\alpha}$ (normalizing output prices to one). The marginal revenue product of labor is, in logs, $\widehat{MRPN} = \hat{z} + \hat{\alpha} - (1-\alpha)\hat{n}$. Labor supply at the firm is $n^s = w^{\theta} \Leftrightarrow w^s = n^{\frac{1}{\theta}}$ (we could add a multiplicative disutility shifter that would control the intercept in logs), so that $\hat{w}^s = \frac{1}{\theta}\hat{n}$, giving the labor-supply-compliant (hence cost-minimizing) wage as a function of employment. The marginal cost of labor is, for firms on their labor supply curve, $MCN = \frac{\partial(wn)}{\partial n} = (1+\frac{1}{\theta})(n)^{\frac{1}{\theta}} = (1+\frac{1}{\theta})w$, and, in logs, $\widehat{MCN} = \ln(1+\frac{1}{\theta}) + \frac{1}{\theta}\hat{n} = \ln(1+\frac{1}{\theta}) + \hat{w}$. Hence, in Panel (a), markdown $\ln(1+\frac{1}{\theta})$ is the vertical gap between the marginal cost and the wage for any employment. This exposition reiterates the basic micro-foundation of the canonical monopsony markdown through a finite firm-specific labor supply elasticity θ.

The *monopsony equilibrium* is at the cross of $\widehat{MRPN}$ and $\widehat{MCN}$ and denoted by superscript m. The monopsony wage is given by the labor supply curve. The markdown is the *vertical distance* between the marginal product and the wage—equal to $\ln(1+\frac{1}{\theta})$ as discussed above. This *monopsony markdown*—the firm-side gap between the marginal product and the wage—is the focus of our paper.

Importantly, this firm-side markdown is distinct from an alternative gap with the firm-level analog of the *competitive* equilibrium (intersection of labor demand and labor supply, rather than the marginal cost of labo), denoted by ce in Panel (a). (This equilibrium is a hypothetical alternative equilibrium in which firms "forget" about inframarginal wage effects or do not act on their wage-setting power.) Indeed, there, too, log markdowns are zero. However, log markdowns will turn to zero (or considerably shrink) also outside of the competitive equilibrium: specifically, when firms face exogenous wages—even if that exogenous wage does not induce the competitive benchmark (while potentially leaving labor supply rationed). This case and this prediction are of interest to our paper and the monopsony literature. Below, we focus on the monopsony wage markdown as the vertical gap in the figure, namely the firm-side gap between the marginal product and the wage, for which monopsony makes strong predictions.

The setting with a minimum wage. In Panel (b), we now consider the case with a minimum wage, $\underline{w}$. We also consider firms, indexed by i, with heterogeneous productivity z_i and hence $\widehat{MRPN}_i$. Foreshadowing our research design, we will partition firms by their response to the minimum wage regarding employment and wages—and, the central innovation of our paper, markdowns.

Markdowns have always featured prominently in *theoretical* discussions of minimum wages in the conceptual background to guide the interpretation of realized effects on employment.[3] We view our paper as the first direct test of how actual empirical markdowns move in response to minimum wages, and, moreover, relative to a model prediction.

Illustrated in Panel (b), the key manifestation of the minimum wage is a kinked wage curve resulting from $\max\{\underline{w}, w^s(n)\}$—so that, obviously, the wages traced out by the labor supply curve are not feasible below the minimum wage. For monopsony, the key implication is a kinked marginal cost curve: $\max\{\underline{w}, \mathbb{1}(w^m(n) > \underline{w})MCN(n)\}$. The flat portion of this curve foreshadows that the latently monopsonistic firm becomes a *wage-taker* when the minimum wage frustrates its preferred monopsonistic wage.

Effects of the minimum wage, with a focus on markdowns. The minimum wage affects labor demand, our focus, and also labor supply. The higher minimum wage will, at least in the partial, firm-level equilibrium, raise labor supply to the firm. We start by defining point A as the (desired) labor supply at the minimum wage. Formally, it is given by $\underline{n}^s = \underline{w}^\theta$. Analogously, it is also useful to classify firms by their marginal product if hiring at their desired labor demand at the minimum wage, $\underline{n}^d$: $\underline{MRPN}_i = z_i\alpha(\underline{n}^d)^{1-\alpha} = \underline{w}$. This concept will help partition firms with respect to various points we discuss below, regarding their responses in employment, wage, and the markdown.

The first clear-cut group are firms that pay the minimum wage and see their markdown shrink to zero. These are our clean treatment group with the sharpest prediction and our leading case. These are all firms to the left of point A with respect to their $\underline{MRPN}_i$. These firms become wage takers with sufficient labor supply, i.e., $\underline{n}^d \leq \underline{n}^S$.

Moreover, this leading treatment group that fully eliminates markdowns can be further partitioned along point B, into firms that shrink (to the left) and grow (to the right)—but *both* groups set their log markdown to zero. That cutoff point B is the cross of the original, unconstrained marginal cost curve and the minimum wage (i.e., $\ln(1+\frac{1}{\theta}) + \frac{1}{\theta}\hat{n}_i^s = \underline{\hat{w}}$), and captures whether firms perceive labor costs to rise or fall comparing the minimum wage with the previously marked-down wage.

The second clear-cut group are control firms: firms not bound by the minimum wage. Their monopsonistic wage remains feasible, and so their markdowns do not change. In the figure, these are firms whose marginal product $\underline{MRPN}_i$ just crosses point C. We reiterate that our setup is a partial equilibrium one.

Finally, for the residual group—firms with $\widehat{MRPN}_i$ between A and C on the vertical segment of the new MCN curve—markdowns shrink but do not necessarily completely disappear. These firms can pay the minimum wage and set employment at the level of the desired labor supply—which leads to an involuntary markdown that is distinct from the monopsonistic one, yet still *smaller*. To

[3] In this Section 2, our figure, conceptual framework, and discussion build on the "simple models" and graphical analyses in Dickens, Machin and Manning (1999) and Bassier and Budlender (2025), studies that do not directly study markdowns in the data but instead focus on empirical effects on employment and wage dynamics. We gratefully credit specifically Figure 2 and surrounding discussion in Bassier and Budlender (2025) for the representation of the setting along a "kinked" marginal cost curve following the minimum wage and for our useful correspondence related to the question of the choices of firms rationed by labor supply. See also our references in the introduction.

see this, consider the firm just to the right of point A, which without a minimum wage has the standard log markdown of $\ln(1+\frac{1}{\theta})$, but has an essentially zero one following the minimum wage introduction. Firms more productive are to the right of/above A and work their way up the vertical segment of the new MCN curve, with markdowns widening up to their original, monopsonistic level for firms just to the left/below point C. Intuitively, these firms' employment is rationed by labor supply ($\underline{n}^d > \underline{n}^S$) because labor demand increases by more than labor supply following the minimum wage. However, we view it as an open question, and potentially a result of functional forms or parameters, whether some of those firms may prefer the competitive equilibrium—i.e., preferring larger scale at higher wages to a scale distorted by labor supply—which would extend the zero-log-markdown prediction to these firms, too.[4] We conclude that this intermediate firm case may only partially shrink markdowns, and discuss empirical implications below.

Empirical predictions. To conclude, we will sort firms into exposed and unexposed firms (respecting within-firm worker heterogeneity as described below), and estimate the impact of a minimum wage introduction on markdowns among exposed firms. In our full model, we will focus on the clear-cut prediction of log markdowns moving to zero for exposed worker groups, and derive a research design with that prediction.

That is, our full model setup below will sidestep the possibility of the weakly-responding exposed group, where labor supply constraints may entail residual, involuntary markdowns even in exposed firms. However, in our empirical implementation, we will include an empirical check on this issue. Specifically, we will tag one unambiguous subset of exposed firms that will, in the setup above, definitively not face labor-supply constraints: exposed firms that shrink or have below-median employment growth—those to the left of point B in Panel (b). The other exposed firms (to the right of B) grow—and either drop their log markdown to exactly zero, too (firms between point B and A), or else lower markdowns but fall short of totally eliminating them (firms between points A and C). We will pick up this discussion again at the end of Section 3.2 in the full model, and the end of the empirical results in Section 6. The bottom line is that this consideration does not appear to quantitatively affect our empirical findings, as both the clear-cut subgroup and full exposed group (that may contain rationed firms) exhibit the same puzzling departure from the core prediction of the monopsony model for markdowns—so that we can develop our account with the leading clear-cut case below.

[4]Indeed, which of these choices prevails depends on whether the scale gains from the competitive wage exceed the higher wage costs compared to the minimum. Intuitively, the profit gain from scale is given by the triangle area situated to the right of A and between the competitive wage and $\widehat{MRPN}$, while the additional wage cost is given by the trapezoid whose height is the difference between the competitive and the minimum wages and reaches into the competitive equilibrium point. We again thank Ishaan Bassier for a useful discussion, and the clarification that Bassier and Budlender (2025) and in particular Figure Panel (b) is based on the assumption of these firms being restricted to paying the minimum wage, with it potentially being an open question whether (but plausible that) this is indeed the profit-maximizing route, compared to, e.g., the competitive equilibrium alternative.

3 Full model and research design

We now present the full model with a more general production function with multiple labor types and other inputs, heterogeneity, and the possibility that only some of the firm's workforce is exposed to the minimum wage by virtue of their wage levels. Our full model will also serve as the structural basis to derive the DiD regression and for the identification of markdowns. We emphasize that our firm-level focus and DiD design ultimately result in us deriving a partial-equilibrium, firm-level prediction (and are hence subject to the macro-micro critique in Berger, Herkenhoff and Mongey, 2025a). However, we do not see an obvious general equilibrium rationalization for our main finding of the minimal micro effect of the minimum wage on the markdowns of exposed firms.

3.1 Full model: setup

We now set up the full model.

The firm's problem. Firms are atomistic and act as monopolistic competitors in both the product market and the labor market. Each firm i produces a variety y_t^i using a Cobb-Douglas production function with capital k_t^i, intermediate inputs x_t^i, and arbitrary labor types $o \in O$ ("occupation" or skill), $n_t^{i,o}$—which permits us to flexibly classify workers exposed to the minimum wage at a given firm. z_t^i is total factor productivity. For simplicity, we assume that the firm has perfect foresight about the evolution of its productivity and aggregate variables.

The firm chooses the product price and inputs to maximize the following value function:

$$J_t\left(k_{t-1}^i\right) = \max_{p_t^i, k_{t+1}^i, x_t^i, \{n_t^{o,i}\}_o} \left\{ \frac{p_t^i}{p_t} y_t^i - \sum_o w_t^{o,i} n_t^{o,i} - i_t - p_t^x x_t^i + \frac{J_{t+1}(k_t^i)}{1+r_{t+1}} \right\} \tag{1}$$

$$\text{s.t. } y_t^i = z_t^i (k_t^i)^{\alpha^i} (x_t^i)^{\beta^i} \Pi_o (n_t^{o,i})^{\gamma^{o,i}} \tag{2}$$

$$k_{t+1}^i = k_t^i (1-\delta^i) + i_t^i \tag{3}$$

$$y_t^i = \left(\frac{p_t^i}{p_t^{\mathbb{p}(i)}} \right)^{-\eta^{\mathbb{p}(i)}} y_t^{\mathbb{p}(i)} \tag{4}$$

$$w_t^{o,i} = \left(\frac{n_t^{o,i}}{n_t^{\mathbb{m}(i),o}} \right)^{\frac{1}{\theta^{\mathbb{m}(i),o}}} w_t^{\mathbb{m}(i),o} \quad \forall o \tag{5}$$

$$w_t^{o,i} \geq \underline{w}_t \quad \forall o, \tag{6}$$

where δ^i is the capital depreciation rate and i_t is investment, and $p_t^{\mathbb{p}(i)}$ is the price index in product market $\mathbb{p}$ (and p_t will be the aggregate price index that gives the real values used throughout). Equations (4) and (5) reflect monopolistic competition in the product market and monopsonistic competition in the labor market, respectively, with standard representative household microfoun-

dations for constant-elasticity-of-substitution consumption utility and labor supply disutility.[5] $\theta^{\mathbb{m}(i),o}$ denotes the firm-specific labor supply elasticity firm i that operates in labor market $\mathbb{m}(i)$ faces for labor type o. This flexible market definition serves as a pragmatic shortcut for heterogeneous supply curves and hence worker-type-*and*-firm-specific markdowns. Finally, we introduce minimum wages through a simple constraint in Equation (6).

Upon defining mc_t^i as the Lagrange multiplier on the constraint in Equation (2), the first order condition with respect to prices is:

$$\frac{p_t^i}{p_t} = \mu^i mc_t^i, \tag{7}$$

where the markup is $\mu^i = \frac{\eta^{\mathbb{p}(i)}}{\eta^{\mathbb{p}(i)}-1}$ (which we directly index by i for simplicity), where $mc_t^i = \frac{MC_t^i}{p_t}$, and where MC_t^i is the nominal marginal cost. Capital and intermediate input demands are given by:

$$mc_{t+1}^i \alpha^i z_{t+1}^i (k_{t+1}^i)^{\alpha^i-1} (x_{t+1}^i)^{\beta^i} \Pi_o (n_{t+1}^{o,i})^{\gamma^{o,i}} = r_{t+1} - \delta \tag{8}$$

$$mc_t^i \beta^i z_t^i (k_t^i)^{\alpha^i} (x_t^i)^{\beta^i-1} \Pi_o (n_t^{o,i})^{\gamma^{o,i}} = p_t^x. \tag{9}$$

3.2 Monopsonistic labor demand, markdowns, and the minimum wage

Labor demand. When choosing labor demand, the monopsonistic firm takes into account the impact on the wage, as given by Equation (5)—where θ is the labor supply elasticity to the firm. It also operates under the minimum wage constraint in Equation (6). In this environment, demand for labor type o is:

$$\underbrace{mc_t^i \gamma^{o,i} z_t^i (k_t^i)^{\alpha^i} (x_t^i)^{\beta^i} (n_t^{o,i})^{-1} \Pi_{o'} (n_t^{o',i})^{\gamma^{o',i}}}_{\frac{1}{\mu^i} MRPN_t^{o,i}} = w_t^{o,i} + \frac{\partial w_t^{o,i}}{\partial n_t^{o,i}} n_t^{o,i} - \psi_t^{o,i} \frac{\partial w_t^{o,i}}{\partial n_t^{o,i}} \tag{10}$$

$$\Leftrightarrow w_t^{o,i} = \frac{\theta^{\mathbb{m}(i),o}}{1 + \theta^{\mathbb{m}(i),o} - \psi_t^{o,i}/n_t^{o,i}} \frac{1}{\mu^i} MRPN_t^{o,i}, \tag{11}$$

where $\psi_t^{o,i}$ is the Lagrange multiplier on the wage constraint in Equation (6).

When $w_t^{o,i} > \underline{w}_t$ and hence $\psi_t^{o,i} = 0$, we obtain the standard monopsonistic markdown expres-

[5] In each product market, $\mathbb{p}$, firms produce different varieties, i, which are aggregated according to the CES aggregator, $y_t^{\mathbb{p}} = [\int_{i=0}^{1} (y_t^{\mathbb{p}(i)})^{1-\eta^{\mathbb{p}(i)}}]^{\frac{1}{1-\eta^{\mathbb{p}(i)}}}$, with $\eta^{\mathbb{p}(i)}$ being the elasticity of substitution across varieties, and $p_t^{\mathbb{p}}$ being the corresponding price aggregator. Within each product market, firms are atomistic and compete monopolistically. The various $y_t^{\mathbb{p}}$ are CES-aggregated into an aggregate output y_t , with p_t being the corresponding price aggregator. Since the firm is atomistic, it ignores the impact of its price on the aggregate and product-market-level price indices p_t and $p_t^{\mathbb{p}}$. Similarly, on the labor supply side, the representative household considers the CES aggregator $n_t^o = \left[\sum_o (n_t^{i,o})^{\frac{\theta^{\mathbb{m}(i),o}+1}{\theta^{\mathbb{m}(i),o}}}\right]^{\frac{\theta^{\mathbb{m}(i),o}}{\theta^{\mathbb{m}(i),o}+1}}$, where we can omit specifying potential between-occupation effects.

sion:

$$\Leftrightarrow w_t^{o,i} = \frac{\theta^{\mathbb{m}(i),o}}{1+\theta^{\mathbb{m}(i),o}} \frac{1}{\mu^i} MRPN_t^{o,i}. \tag{12}$$

When the constraint binds ($w_t^{o,i} = \underline{w}_t$), we have $\psi_t^{o,i} = n_t^{o,i} > 0$ (also reflecting Hotelling's lemma) and obtain the expression for competitive labor demand—the firm becomes a wage taker:

$$\underbrace{w_t^{o,i}}_{=\underline{w}_t} = \frac{1}{\mu^i} MRPN_t^{o,i}. \tag{13}$$

Markdowns. That is, whenever the constraint does not bind, the standard markdown of $\frac{\theta^{\mathbb{m}(i),o}}{1+\theta^{\mathbb{m}(i),o}}$ emerges, and otherwise there is no wage markdown. Hence, we can define markdowns as follows:

$$\xi_t^{o,i} = \frac{w_t^{o,i}}{\frac{1}{\mu^i} MRPN_t^{o,i}} \tag{14}$$

$$= \begin{cases} 1 & \text{if } w_t^{o,i} = \underline{w}_t \\ \frac{\theta^{\mathbb{m}(i),o}}{1+\theta^{\mathbb{m}(i),o}} \leq 1 & \text{if } w_t^{o,i} > \underline{w}_t. \end{cases} \tag{15}$$

Rationing. As noted before, for our research design below, we focus on the case in which labor demand is not rationed by labor supply at the minimum wage. That is, our firm problem does not feature an explicit rationing constraint whereby labor demand be weakly below labor supply, a condition by definition fulfilled with equality in the unconstrained monopsony equilibrium but that may, but need not, break with wage floors. As discussed above in Section 2, such rationing can lead to involuntary markdowns due to insufficient labor supply. While this possibility is featured in some theoretical equilibrium models (see Azkarate-Askasua and Zerecero, 2025; Berger, Herkenhoff and Mongey, 2025b, for such constraints following minimum wages and collective bargaining wages), this issue has, to our knowledge, not been tagged in the existing empirical literature. To account for this possibility, we will conduct a robustness check for firms with negative or below-median employment growth to look for differential markdown responses, as this group is plausibly unconstrained by the rationing concern (see also the end of Section 2).

3.3 From model to data: deriving the research design

A reform in the model. We now consider two periods as a stepping stone towards our research design. In "pre-period" or "baseline" period 1, there is no wage constraint (equivalently, it does not bind). In the "post-reform" period 2, the minimum wage constraint is introduced. Hence in the first period, the minimum wage is $\underline{w} = 0$, and it never binds, and in the next period, it is $\underline{w}' > 0$ and may bind. We denote by x' the post-reform counterpart of pre-reform variable x. We assume that between those two periods, wage-relevant variables are stable within the firm, so that, absent

the minimum wage binding, we have $w^{o,i'} = w^{o,i}$.

Firm-level markdown changes. Firm i's log markdowns for workers of type o change as follows:

$$\ln \xi^{o,i'} - \ln \xi^{o,i} = \begin{cases} \ln \frac{\theta^{\mathrm{m}(i),o}}{1+\theta^{\mathrm{m}(i),o}} - \ln \frac{\theta^{\mathrm{m}(i),o}}{1+\theta^{\mathrm{m}(i),o}} = 0 & \text{if } w^{o,i'} > \underline{w}' \Leftrightarrow w^{o,i} > \underline{w}' \\ \underbrace{\ln 1 - \ln \frac{\theta^{\mathrm{m}(i),o}}{1+\theta^{\mathrm{m}(i),o}}}_{\text{"Gap from unit markdown"}} = -\ln \frac{\theta^{\mathrm{m}(i),o}}{1+\theta^{\mathrm{m}(i),o}} \geq 0 & \text{if } w^{o,i'} = \underline{w}' \Leftrightarrow w^{o,i} \leq \underline{w}' \,. \end{cases} \tag{16}$$

That is, markdowns ξ become one (i.e., competitive) for worker-type-firm cells for which the minimum wage binds—so that the change in the markdown is equal to the "gap from one" pre-reform, i.e., the distance the markdown has to cover to become competitive. By contrast, for cells where the minimum wage does not bind, no change is expected, also given our above assumption of stability in the markdown-relevant parameters.

Empirical goal. Our paper takes the logic of Equation (16) to the empirical context. The design is straightforward with two extreme firm types: a *control* group composed of firms with 0% share of exposed workers and a treatment group with 100%. Markdowns should drop to zero in the treated firms (this can be seen from the bottom line of Equation (16)), and stay stable for the firms in the control group (top line), again assuming stable markdown determinants and no spillover/equilibrium effects. For instance, a "pass-through" DiD regression with a continuous treatment intensity variable equal to the negative of a firm's pre-reform markdown (again, see the bottom case of Equation (16)) for the firms with 100% of workers exposed to the minimum wage (and zero for the control firms) should yield a coefficient of one.

However, the actual economy is richer, and our data will be generated by firms that are continuously heterogeneous in the share of exposed workers (rather than having either 0% or 100% of its workers exposed with wages bound by the minimum wage) and in baseline markdowns. We next derive the general DiD setup for this realistic context.

Overcoming the challenge of heterogeneity: "composite" labor. We now solve three related challenges. First, in the data, the share of worker types within a firm for whom the minimum wage binds varies continuously between 0% and 100%. Second, rather than a single true firm-level markdown, each firm/worker-type cell will have its own markdown. Third, existing—and our—empirical strategies for estimating markdowns typically lump together all labor types (or consider limited ex ante worker types too coarse for our design, as the same worker type may earn above or below the minimum wage depending on the firm).

To make progress, we define "composite labor" concepts, starting with the composite mark-

down ξ^i:

$$\overline{w}^i(\xi^i)^{-1} = \frac{1}{\mu^i}\overline{MRPN}^i, \tag{17}$$

where the mean wage is $\overline{w}^i = \sum_o \kappa^{o,i} w^{o,i}$ (where $\kappa^{o,i} = \frac{n^{o,i}}{\sum_{o'} n^{o',i}}$) and where $\overline{MRPN}^i = \sum_o \kappa^{o,i} MRPN^{o,i}$ is a "composite labor marginal revenue product," i.e., the employment-weighted average of type-specific MRPNs.[6] (For simplicity, we do not denote ξ^i as $\overline{\xi}^i$.) Then, using the underlying structural equation from the firm-specific wage-setting condition $\frac{1}{\mu^i} MRPN^{o,i} = w^{o,i}(\xi^{o,i})^{-1}$, we have:

$$(\xi^i)^{-1} = \sum_o \phi^{o,i}(\xi^{o,i})^{-1}, \tag{18}$$

where $\phi^{o,i} = \frac{w^{o,i} n^{o,i}}{\overline{w}^i \sum_{o'} n^{o',i}}$ is the wage-bill weight of occupation o in firm i. In words, the composite markdown ξ^i can be viewed as the weighted average of the occupational-level "structural" worker-type-level markdowns, weighted by the types' wage bill shares.[7]

Towards a difference-in-differences design. With the composite labor concepts in hand, we can now consider *intensive-margin* variation in the share of worker types exposed. Partitioning worker types into the subsets of exposed (E^i) workers—i.e., for whom the wage constraint binds— and nonexposed ones (N^i), we can write the composite (inverse) markdown as a weighted average:

$$(\xi^i)^{-1} = \Phi^{E,i}(\xi^{E,i})^{-1} + \Phi^{N,i}(\xi^{N,i})^{-1}, \tag{19}$$

where $\Phi^{E,i} = \sum_{o \in E^i} \phi^{o,i}$ is the payroll share of exposed workers, and $(\xi^{E,i})^{-1} = \sum_{o \in E^i} \frac{\phi^{o,i}}{\Phi^{E,i}} (\xi^{o,i})^{-1}$ is the average markdown for that group, again payroll-weighted, and analogously for N.

Next, we compare the environment post- and pre-reform to guide our identification strategy and formally derive predictions. We again apply the prediction that markdowns move to one for

[6] While mean wages are simple empirical objects, the composite labor marginal product $\overline{MRPN}$ is fictional and linking it to underlying marginal products of the actual labor inputs L_o requires a thought experiment. The bridge we employ is to define it as $\overline{MRPN}^i = \sum_o \kappa^{o,i} MRPN^{o,i}$: the change in output if labor were increased across the board in a way that leaves the skill structure constant. This definition also appears consistent with production function estimation methods. An alternative route may be to log-linearize marginal products.

[7] We had not been aware of other attempts to wrestle with the aggregation of heterogeneous underlying markdowns into a composite markdown such as the commonly estimated ones at the firm level. We thank Matthias Mertens for making us aware of a related discussion in Mertens (2023) (Section 4.3 and specifically Equation (6) therein) for two labor types as the exception. Our derivation applies to multiple arbitrary labor types, alludes to consistency with production function estimation and marginal products of composite labor, and establishes aggregation in terms of payroll shares, which unlike type-specific output elasticities can be easily observed for arbitrary firm-and-worker-type heterogeneity and be flexibly aggregated into subgroups. For instance, below in Equation (19), we will use our results to collapse all workers (without specific type labels) with pre-reform wages below the minimum wage into a firm-level exposed group, and express a general weighted average that uses firm-specific weights that correspond to the share of its own exposed workers in its overall payroll.

exposed types, so that $\xi_o'^{-1} = 1 \; \forall o \in E$ and hence $\xi^{E'^{-1}} = 1$. Hence, the post-reform markdown is:

$$\xi'^{-1} = \Phi^{E'} \cdot 1 + \Phi^{N'} \xi^{N'^{-1}}. \tag{20}$$

The treatment effect emerges by adding and subtracting terms and reformulating as follows:

$$\Delta\xi^{-1} = \xi'^{-1} - \xi^{-1} = \underbrace{\Phi^{E}(1 - \xi^{E^{-1}})}_{\text{Direct markdown effect}} - \underbrace{(\Phi^{N'}(1 - \xi^{N'^{-1}}) - \Phi^{N}(1 - \xi^{N^{-1}}))}_{\text{Indirect spillovers (composition + markdowns)}}. \tag{21}$$

Sharpening the prediction. We highlight two assumptions that deliver a particularly sharp version of Equation (21). We will support those assumptions empirically later. First, we assume that payroll shares and markdowns for nonexposed do not change, which results in:

$$\Rightarrow \Delta\xi^{-1} = \Phi^{E}(1 - \xi^{E^{-1}}). \tag{22}$$

Second, we assume similar *group-level average* baseline markdowns within firms across the exposed group and the nonexposed group—i.e., $\xi^E = \xi^N = \xi$.[8] Then, we can reformulate Equation (22) into an intuitive and, as we show next, particularly tractable percent (and log) terms:

$$\frac{\Delta\xi}{\xi} = \Phi^{E} \cdot (1 - \xi) \tag{23}$$

$$\Rightarrow \Delta \ln \xi \approx -\Phi^{E} \cdot \ln \xi, \tag{24}$$

where the second line uses the fact that $\frac{\Delta\xi}{\xi} \approx \Delta \ln \xi$, and $\ln \xi \approx \xi - 1$ (for ξ slightly below 1).

That is, the markdown will, in percentage terms, rise towards one proportionately to the negative product of (i) the payroll share of exposed workers Φ^E and (ii) the "competitive gap" $1 - \xi \approx \ln \xi$, i.e., how much smaller pre-reform ξ is than 1. Hence, on their own, neither the payroll share of exposed workers nor the initial markdown would sufficiently characterize the predicted effects on markdowns: if the pre-reform markdown is already one, even a considerable payroll exposure will leave no markdown to move, and markdowns will not move if no worker is exposed even though the firm's baseline markdown is far below one.

DiD regression analog of model prediction. Equation (24) can be estimated in a simple difference-in-differences (DiD) regression that relates the firm-level change post- vs. pre-reform in log markdowns to the continuous independent variable that is the negative product of the share of the wage bill below the minimum wage times the firm's pre-reform log markdown, as i Equation (24). A basic example is:

$$\Delta \ln \xi^{i} = \rho D^{i} + v^{i}, \tag{25}$$

[8] Otherwise, with heterogeneous average markdowns between the groups, Equation (23) is: $\frac{\Delta\xi}{\xi} = \Phi^{E} \cdot \left(\frac{\xi}{\xi^{E}} - \xi\right)$.

where we introduce "treatment intensity" following Equation (24):

$$D^i = -\Phi^{i,E} \cdot \ln \xi^i. \tag{26}$$

In practice, we will estimate the DiD specification in a panel regression with fixed effects (see Equation (31) below).

The prediction encoded in Equation (24) for our prototype DiD regression in Equation (25) is straightforward and strong: $\rho = 1$. This benchmark formalizes and quantifies the core prediction of monopsony spelled out in our introduction and in Equations (15) and (16): that firms have wage-setting power, that they take advantage of it by marking down wages below their marginal product, and that wage constraints eliminate these noncompetitive features and hence move markdowns to one.

Stress-testing the prediction. The prediction rests on three assumptions: the correct identification of treatment intensity, sufficient persistence therein, and the specification of the correct model.

Most importantly, measurement error might interfere with the prediction of $\rho = 1$ due to attenuation bias. Consider measurement error in markdowns (e.g., $\ln \widehat{\xi} = \ln \xi + \varepsilon$). Indeed, we will find that a nonnegligible share of wage markdowns are above one, perhaps reflecting this confounder. On the LHS ($\widehat{\Delta \ln \xi} = \Delta \ln \xi + \Delta\varepsilon$), mean-zero error does not interfere with our prediction. However, attenuation bias arises on the RHS ($\widehat{\Phi^E \cdot \ln \xi} = \Phi^E \cdot \ln \xi + \Phi^E \varepsilon$). More precisely, the actually estimated coefficient in a regression of actual markdown changes on predicted ones (based on mismeasured markdowns) will be a weighted average of the structural prediction (one) and zero (due to noise), with the bias in the coefficient (from the unit prediction) given by the share of variance in the regressor that is noise.[9]

Regarding the assumption of stability in markdowns and the payroll share, we will check for pre-period persistence in treatment intensity.

We will also check for the relevance of the assumption that baseline markdowns are on average similar between exposed and nonexposed workers within the firm. For instance, if low-wage exposed workers do not feature markdowns, we would stack the cards against the monopsony prediction (see also Footnote 8 above). As a check for symptoms of markdown heterogeneity within the firm that may bias our estimates, we will compare the pre-reform cross-sectional relationship between share of payroll of exposed workers and the firm-level average markdown, doing so in Section 5.5. We will find little room for this concern, but note that our check remains tentative.

[9] We note one specific source of mismeasurement in markdowns to which our first-order approach to markup identification (detailed in Section 5.2) is subject. That strategy assumes that intermediate inputs do not have a markdown/have fixed prices. Away from that assumption, our wage markdown measure will capture labor's *relative* markdown compared to intermediates, i.e., $\widehat{\ln \xi^i} = \ln \xi^i - \ln \xi^{x,i}$ (see, e.g. Mertens and Schoefer, 2024; Treuren, 2025). Moreover, adjustment costs might generate measured markups or markdowns unrelated to monopsony. For instance, situations of labor hoarding may lead workers to be "overpaid" compared to temporarily low productivity. We find that the markdown measures are relatively stable, which perhaps points away from this concern.

Another possibility is that our model does not accurately describe the labor market structure. Perhaps instead of firms being wage setters, wages might be collectively bargained (at firm or industry levels). In right-to-manage bargaining models specifically, and in efficient bargaining models more generally, there are no firm-side markdowns, so those cells will be in our control group.[10] As noted in the introduction of this section, we do sidestep general equilibrium shifts that may affect our prediction through the labor supply the firm faces or also product market adjustments.

An important consideration we have mentioned at the end of Section 3.2 is the possibility of a subset of firms being rationed by labor supply, which our full model sidesteps but our simple model in Section 2 highlighted. While inherently challenging to proxy for, we, in our empirical analyses below, do include a robustness check that zooms into firms plausibly unlikely to be rationed.

Similarly, we assume that firms set wages occupation by occupation, where our notion of occupation is intentionally flexible (and can in principle denote firm-specific classifications), yielding group-specific markdowns. But even in the presence of frictions in differentiating wages between groups and horizontal equity constraints within the firm (e.g., Saez, Schoefer and Seim, 2019), the minimum wage introduction would, by inducing spillovers through wage increases within the firm that our model's flexible wage setting does not feature, if anything, amplify rather than attenuate the predicted effects on markdowns. This is because the markdown predictions would then partially additionally extend to worker groups our design deems nonexposed (essentially boosting an effective notion of treatment intensity D_i by raising the effective payroll share of exposed workers Φ^E, for all firms with some mass of directly exposed workers (i.e., $\Phi^E > 0$), while leaving the control firms with $\Phi^E = 0$ at the original value).

4 The 2015 introduction of a high minimum wage in Germany

Our empirical setting is the 2015 *introduction* of a minimum wage in Germany.

Background. In 2015, Germany passed a statutory introduction of the federal minimum wage with a bipartisan consensus. There was no general minimum wage before, with the German government having traditionally left wage setting to employer associations and trade unions, or individual firms.[11] The German collective bargaining system, at the sector-regional level, left considerable flexibility for employers' idiosyncratic wage setting due to coverage being optional and up to employers (by choosing to join employer associations), on top of opportunities to opt out (see, e.g., Jäger, Noy and Schoefer, 2022). Employers selecting into coverage often pay premia above

[10] Models with holdup (fixed employment or capital inputs) generate profits some of which workers may appropriate via bargaining, so that wages may be above the revenue marginal product of labor.

[11] More precisely, up to 1996, there were no minimum wages in Germany. Starting in 1996, minimum wages were introduced in the construction sector only, and as an extension of collective bargaining outcomes. More generally, collective bargaining agreements were occasionally extended to apply to an entire sector, but this share had declined in the years leading up to 2015. Indeed, Germany's industrial relations system has applied extensions relatively rarely in international comparison (Jäger, Naidu and Schoefer, 2025).

sectoral wage floors. Consistent with flexible scope for firm-level wage setting, the German wage structure featured considerable wage dispersion around this time (in particular between firms, as documented in Card, Heining and Kline, 2013, for the period leading up to 2015).

The background for policy reform was that in the 2013 coalition negotiations, the Social Democratic Party, a center-left party, was able to negotiate a statutory minimum wage in the coalition agreement as the junior partner of the center-right Christian Democrats. In part, the law was motivated by the decline of the collective bargaining coverage, the rise in inequality, and wage stagnation. In mid 2014, the minimum wage law (*Gesetz zur Stärkung der Tarifautonomie*) passed the two chambers of the German parliament and became effective as of August 2014. It legislated the introduction of a minimum wage of 8.50 EUR from January 1, 2015, and extended the possibility for sectoral minimum wages.

Identification considerations. Four aspects make this variation particularly suitable for our research question. First, the reform was an *introduction* of a minimum wage rather than an adjustment. Second, the minimum wage then remained relatively stable through the end of our analysis window in 2018 (rising only slightly to 8.50 EUR), with the next wave of significant changes not occurring until the 2020s (see Figure 2). These subsequent adjustments are smaller and empirically elusive as our IAB data lack hourly-wage information after 2014 (see below in Section 5). Third, the law applied *erga omnes*, with no avoidance routes (except for workers younger than 18 and apprentices), to every job category and sector. Fourth, the bite of was extensive, as discussed below in Section 5.3—and the existing evaluations of the reform focusing on employment and wage outcomes, but not markdowns, which we cited in Footnote 2 of the introduction.

5 Implementation: data, variables, sample, and research design

Our baseline specification for the event analysis is a difference-in-differences (DiD) design that implements the theoretical prediction from Equation (24). In contrast to the existing minimum wage literature, our main outcome variable is the markdown.

5.1 Data

Our empirical analysis draws on multiple datasets provided by the Institute for Employment Research (IAB) of the German Federal Employment Agency.

Establishment data: production and inputs. Our core dataset is the IAB Establishment Panel (EP, Bellmann, Brunner, Ellguth, Grunau, Hohendanner, Kohaut, Leber, Möller, Schwengler, Stegmaier and Umkehrer, 2022), an annual survey representing establishments across Germany, which we complement with administrative establishment-level information from the Establishment History Panel (BHP, Ganzer, Schmucker, Stegmaier and Wolter, 2023). The EP covers establishments with at least one employee subject to social security contributions as of the reference date, June 30 of

the year preceding the survey. In a given year, the EP interviews (i) establishments with a valid interview in the previous year, (ii) establishments with a valid interview in the year before the previous year, and (iii) newly drawn establishments. The latter are included to replace panel drop-outs and to ensure the required number for representativeness.

From the EP data, we extract all variables needed to estimate production functions (as output elasticities will be required to estimate markdowns, see below in Section 5.2). Specifically, we collect information on establishments' revenues, labor inputs (number of employees), intermediate inputs, capital, wage bill, and industry.

Matched employer-employee data and hours information. To construct hourly wages, we use matched employer-employee data: the Integrated Employment Biographies (IEB, IAB, 2019), which encompass the complete record of employment and unemployment spells (spells cover at most one calendar year and their precision is daily) covered by the German social security system. The IEB includes detailed information on the start and end dates of employment spells, total earnings by spell, occupation and industry codes, and individual worker characteristics such as gender, age, and education.

To translate earnings into hourly wages, we follow Dustmann, Lindner, Schönberg, Umkehrer and Vom Berge (2022) and merge on spell-level (worker-establishment) hours worked data from the Social Accident Insurance, available between 2010 and 2014.

Sample restrictions. Our analysis sample covers the period 2011–2018. We require firms to have at least five employees and exclude those in the financial sector, insurance, and the public sector. This selection procedure yields 22,140 establishment-year observations.

In our baseline specification, we use a strictly balanced panel, i.e., establishments must have data for all analysis years (2011-2018) without gaps. This balanced panel contains 7,232 establishment-year observations. In robustness checks, we relax this restriction and also analyze the results using an unbalanced panel.

5.2 Variable construction

Wage markdowns. We identify wage markdowns following the standard first-order approach (see, e.g., De Loecker and Warzynski, 2012; Dobbelaere and Mairesse, 2013; Yeh, Macaluso and Hershbein, 2022; Rubens, Wu and Xu, 2025). An advantage of the approach is that it is independent of assumptions about market structures and can be easily implemented with firm-level production panel data. A disadvantage is that it relies on assumptions about firm conduct and the standard issues regarding production function estimation. Indeed, part of the motivation of our paper is to put to a reality check these measures in our specific context. That is, the joint hypothesis of our paper inherently combines the prediction of the monopsony model and the empirical performance of this popular and influential route to identifying wage markdowns.

The intuition is that price markups enter both intermediates and labor first-order conditions, while the intermediates are assumed to have exogenous prices , so that only labor may feature a wage markdown.[12] Hence, we can eliminate the price markup and isolate the wage markdown by taking the ratio of the conditions, so that the wage markdown of firm i for labor type o is given by:[13]

$$\xi_t^{o,i} = \frac{\beta^i}{p_t^x x_t^i / y_t^i} / \frac{\gamma^{o,i}}{w_t^{o,i} n_t^{o,i} / y_t^i}, \tag{27}$$

where β and γ correspond to the output elasticities to labor and intermediate inputs we identify through production function estimation (Appendix A).

However, in practice, output elasticities, wages, and markdowns cannot be estimated by sufficiently granular and consistently defined worker type, let alone that, for our purpose, would moreover perfectly capture exposure to the minimum wage. Instead, markdowns are estimated at the firm level—which correspond to the *composite* labor markdown ξ_t^i we derived above in Section 3.3, which hence we identify as follows:

$$\xi_t^i = \frac{\beta^i}{p_t^x x_t^i / y_t^i} / \frac{\overline{\gamma}^i}{\overline{w}_t^i n_t^i / y_t^i}. \tag{28}$$

We obtain cost shares from the establishment survey. We estimate output elasticities in firm-level accounting data from the EP establishment survey using the method of Levinsohn and Petrin (2003), complemented with the Ackerberg, Caves and Frazer (2015) correction, for firms in 2-digit industries. The method uses intermediate inputs as a proxy variable, a choice considered superior to using investment, which has a lumpy nature. More details on the construction of those measures and the description of the sample used are in Appendix A. We perform robustness checks with a different sector classification at the 1-digit level, with a translog production function (to alternatively construct markdowns that permit richer substitution elasticities between intermediates and labor can depart from Cobb-Douglas, see Mertens and Schoefer, 2024), and with and without controls for market share. The latter alleviates concerns that the intermediate input elasticities are affected by oligopsony wedges in this market. All checks are reported in Appendix A. We note that with firm fixed effects and fixed output elasticities, the dependent-variable variation will largely amount to cost-ratio dynamics, although baseline markdowns also enter treatment intensity and we also consider a more flexible translog production function, as noted.

[12] We note the limitations of this approach that assumes the absence of wedges or monopsony in intermediate input markets, and reiterate Footnote 9 regarding measurement error implications for our setting. See also, e.g., Treuren (2025) for this issue. Away from this assumption, our wage markdown measures the relative markdown between labor and intermediates (as embraced in, e.g., Mertens and Schoefer, 2024). See also our robustness check that includes mark shares as controls, mentioned below. We also note the method by Hashemi, Kirov and Traina (2022) as an interesting alternative route, which does not require exogenous intermediate input prices.

[13] To see this, recall that $mc_t^i = \frac{p_t^i}{p_t} \frac{1}{\mu^i}$ and $w_t^{o,i} = \frac{1}{\mu_t^i} \xi_t^{o,i} MRPN_t^{o,i}$. By dividing both sides of the latter equation by $\frac{n_t^{o,i}}{y_t^i}$ and noting that the price markup can also be obtained as $\mu_t^i = \frac{\beta^i}{p_t^x x_t^i / y_t^i}$, we obtain the expression in Equation (27).

Wages, earnings, and hours. To compute hourly wages—and hence construct our measure of payroll share exposed to the reform—, we proceed in two steps. First, we compute gross daily wages by dividing a worker's total earnings with the employer by the total duration of her employment spells with the employer.[14] We follow the usual solution to impute censored wages about the social security contribution limit (Dustmann, Ludsteck and Schönberg, 2009; Card, Heining and Kline, 2013). Second, we divide spell-level earnings by the spell-level hours worked variable described above (as in Dustmann, Lindner, Schönberg, Umkehrer and Vom Berge, 2022). We use the correction method proposed by vom Berge, Umkehrer and Wanger (2023) to homogenize the hours information.

Baseline period. We choose 2013 as our baseline pre-reform period when we calculate the treatment intensity (negative of the product of payroll share exposed and log markdown), and experiment with alternative baseline years.

Payroll share of exposed workers $\Phi^{E,i}$. We construct the pre-reform (2013) payroll share of exposed workers as the sum of workers' earnings y^{ι}_{2013} (computed in the spell data as days worked times average daily wages) paid to workers (indexed by ι) in firm i's workforce (set I^i) whose pre-introduction (spell-specific) hourly wage h^{ι}_{2013} is below the subsequent minimum (8.50 EUR), divided by the overall wage sum:

$$\Phi^{E,i}_{2013} = \frac{\sum_{\iota \in I^i_{2013}} \mathbb{1}(h^{\iota}_{2013} \leq 8.50\text{EUR}) y^{\iota}_{2013}}{\sum_{\iota \in I^i_{2013}} y^{\iota}_{2013}}. \tag{29}$$

We note that this variable intentionally differs from standard exposure variables building on workers' wage gap as a distance. Instead, in our setting, the binary weight reflects the markdown-relevant classification into competitively vs. monopsonistically purchased labor. We do not attempt to correct for wage inflation but note that in our post-reform sample period, inflation was below 2% but with real wage growth picking up (see, e.g., Bundesbank, 2018, for German wage growth through 2017 and its sources).

Treatment intensity D_i. We construct treatment intensity D_i based on pre-reform variables as:

$$D_i = -\Phi^{E,i}_{2013} \cdot \ln \xi^i_{2013}. \tag{30}$$

Winsorization. We winsorize markdowns at the symmetric 2% level, we do not winsorize the payroll share of exposed workers, and generate treatment intensity as the negative product of winsorized log markdowns and unadjusted payroll shares.

[14] We restrict our sample to employment spells subject to social security, i.e., we abstract from marginal employment/mini jobs.

Weighting. We do not weight the establishment-year data.[15]

5.3 Descriptive facts

Table 1 and Figure 3 present summary statistics and the distributions of our key variables, respectively. We reiterate that our dataset is at the establishment level and contains small firms (compared to, e.g., Compustat). All variables are with respect to the baseline year 2013. Besides markdowns, our primary variable is treatment intensity (the negative product of the wage bill share exposed: workers with hourly wages pre-reform in 2013 below the 2015 minimum wage) times the 2013 log markdown. We focus on our analysis sample. Appendix Tables A.1 and A.2 and Appendix Figures A.1 and A.2 replicate these figures for the subsamples of 2013 log markdowns being negative and positive, respectively (where the former is consistent with monopsonistic markdowns, see below).

Summary statistics. Figure 3 Panels (a)-(c) plot the CDFs and densities of the cross-sectional 2013 distribution of treatment intensities, payroll share exposed, and log markdown. The markdown estimates in our analysis sample are broadly in line with estimates in earlier research using the IAB EP survey (see, e.g., Dobbelaere, Hirsch, Müller and Neuschaeffer, 2024). The distribution of the log markdown shows considerable dispersion, with a standard deviation of 0.85. We also note that this wide distribution includes positive values. Monopsony models imply negative values. Positive values, as we discuss when concluding Section 3.3, might reflect measurement error, stem from large wedges on intermediate inputs, or result from different labor market structures than the monopsony models that we take as our point of departure and put to a test. While our baseline specification includes all observations, we will additionally implement an analysis restricted to firms with more plausibly monopsonistic 2013 markdowns (i.e., negative in logs).

Table 1 shows that the share of payroll of exposed workers has a sizable mean of 12.9%, and a median of 3.0%. There is considerable dispersion therein, with a standard deviation of 22.3 ppt, with 75th and 90th percentiles of 11.9% and 45.6%, and 25th and 10th percentiles of 0.9% and 0%.

Finally, the (negative) product of these two factors gives us our treatment intensity variable D_i defined above in Equation (30)—which structurally represents the predicted percent (log) markdown change following the minimum wage introduction. We have considerable power. For instance, while the median predicted change of the log markdown (i.e., treatment intensity) is close to zero at 0.006, the mean prescribes an increase in markdowns (i.e., move towards a competitive benchmark) of 0.056, with a standard deviation of 0.249, and 75th and 90th percentiles of 0.037 and 0.214.

Correlations and decomposition. Underlying the treatment intensity variable is essentially no correlation (0.020 in Table 1) between the two components, with Figure 3 Panel (d) presenting the

[15]Standard practice suggested in Bächmann, Bellmann, Gensicke, Kohaut, Möller, Schwengler, Tschersich and Umkehrer (2023) would use the IAB-provided survey weights for descriptive facts (rather than regressions), and we found similar results for those summary statistics. We prefer to already report unweighted descriptives for consistency with our unweighted regression analysis.

associated binned scatter plot (revealing a regression coefficient of about 0.076).

To further decompose the treatment intensity variable into the two factors, Panel (e) plots treatment intensity on the x-axis. Formally, a variance decomposition of treatment intensity attributes about 83% to the payroll share exposed among firms with a (monopsonistic) negative log markdowns in 2013, and about 48% in the full analysis sample, with the rest onto the markdown variation (with hardly any role for the covariance, as expected given the limited correlation).[16]

Persistence. We also check on persistence in the treatment assignment, plotted in Panel (f), which we discuss again in Section 5.5 below. Treatment intensity persists nearly perfectly across pre-reform years within firms. We cannot check for persistence beyond 2014 not just because of the reform as such but because the hourly wage data end in 2014 (see Section 5.1).

Primary subsample of initially monopsonistic firms: 2013 log markdowns below zero. As mentioned above, we also present the analogous statistics and descriptives for firms whose markdowns appear monopsonistic in 2013 (i.e., feature log markdowns below zero), and report results in Appendix Table A.1 and Appendix Figure A.1. We find broadly similar results, and note that in this sample, if anything, the payroll share exposed to the minimum wage is associated with initially lower (more monopsonistic) markdowns. This pattern will help assuage the concern that we might underestimate the true effect on markdowns, namely by making less likely a scenario that highly exposed firms just did not feature monopsonistic markdowns to begin with. Below, for completeness, we first focus on the pooled sample, and additionally report on the subsample with initially monopsonistic markdowns separately.

5.4 Econometric specifications

We estimate difference-in-differences (DiD) specifications, in different variants, exploiting establishment variation in treatment intensity.

Pooled DiD regression. Our baseline specification is a pooled DiD regression (a fixed effects panel variant of the prototype DiD specification in Equation (25)):

$$Y_{i,t} = \vartheta_t + \zeta_i + \rho D_i \text{Post}_t + \lambda X_{i,t} + v_{i,t}, \tag{31}$$

where the unit of observation is the establishment indexed by i, ϑ_t are time fixed effects, and ζ_i are establishment fixed effects. The post-treatment period indicator, Post_t, is interacted with establishment i's treatment intensity D_i—defined above in Equation (30). The dependent variable $Y_{i,t}$ is, e.g., wage markdown in logs at the establishment level. The time sample for this specification

[16]The calculation relies on an approximate variance decomposition of a product of two random variables, i.e., $\text{Var}(XY) \approx \mu_X^2 \text{Var}(Y) + \mu_Y^2 \text{Var}(X) + 2\mu_X \mu_Y \text{Cov}(X, Y)$. The resulting explained variance, as a fraction of which we report results above, in turn accounts for about a third of the observed variance for the full analysis sample and two thirds for the preferred monopsonistic (log markdown below zero) sample.

is 2011-2018. Standard errors are clustered at the establishment level. As mentioned in Section 5.1, we weight firms equally and consider a balanced panel.

Intuitively, the specification echoes pass-through or tax incidence regressions—here of the minimum wage into markdowns. The theoretical underpinnings for this specification are provided in Section 3.3, in particular in Equation (24). To recap, the model prediction is that $\rho = 1$.

Year-specific DiD effects. We complement this pooled DiD specification with one with year-specific treatment effects, i.e., $\sum_{k=2011,2012,2014,\ldots,2018} \rho_t D_i \mathbb{1}(k = t)$, where we exclude baseline year 2013 (for which the firm-level treatment D_i is absorbed in the firm fixed effect). This perspective reveals the dynamic effects and permits us to gauge pre-trend violations.

Pooled post- and pre-period effects against 2013. Bridging the year-specific and pooled post-effect specifications, we also estimate pooled post- and pre-reform treatment effects in a specification with 2013 as the base period. This specification permits a simple check on pre-trend violations. (Since the resulting post effect is now estimated against 2013, we report this alternative post effect as a complement to the main specification's pooled post effect estimated against the full pre-period through 2014.)

Pre-trend adjustments. In some specifications, we detect minor pre-trend violations—which, if anything, we would actually bias our estimates towards the unit benchmark consistent with monopsony. Albeit for other variables and a specification at the labor market level, such pre-trend violations were also found in one recent existing analysis of the German minimum wage introduction in Dustmann, Lindner, Schönberg, Umkehrer and Vom Berge (2022), which focuses on employment and reallocation outcomes rather than markdowns, and features a simple pre-trend correction by extrapolating the pre-trend differential into the treated years. We follow their pre-trend correction in supplementary specifications throughout, and bootstrap standard errors (see exhibit notes). We note that an alternative approach would be to apply the full-fledged method proposed in Rambachan and Roth (2023).

Mean reversion and placebos. To account for mean reversion that may bias up our estimate of ρ due to spurious markdown variation (as baseline year markdowns enter D_i), we also consider a version in which we consider 2011 rather than 2013 the baseline level at which we compute treatment intensity. The years 2012 to 2014 then serve as a pre-reform placebo period.[17]

Additional outcome variables. To paint a broader picture, we consider additional outcome variables from the existing literature: wages and employment.

[17] A potential complementary specification would instrument our 2013 baseline treatment intensity with lagged, 2011 treamtent intensity (as in the measurement error correction as in, e.g., Drenik, Jäger, Plotkin and Schoefer, 2023; Schoefer and Ziv, 2024, to isolate common wage or productivity factors between groups rather than periods, respectively). This specification would cleanse our findings of mean revision and attenuation bias from noise. Our results below instead present the simpler specifications that, we believe, paint a consistent picture in reduced form.

5.5 Preparatory validation checks

We now support some of the identification assumptions underlying our research design, and then present results in Section 6.

Pre-reform persistence of treatment intensity. Our DiD design relies on persistence in baseline treatment intensity (absent the minimum wage). Idiosyncratic wage evolution, skill structure changes, firm-level shocks and wage growth overall will naturally weaken the persistence, increasingly so as the horizon grows. Imperfect persistence of treatment intensity would lead us to underestimate the true ρ.[18] Figure 3 Panel (f) plots 2014, 2012, and 2011 versions of treatment intensity against our 2013 baseline version. There is considerable persistence in this variable, limiting this source of attenuation.

Markdown heterogeneity within the firm. Our quantitative benchmark also assumes that baseline markdowns between exposed and nonexposed groups of worker types are, within the firm, on average similar, while supply elasticities might be heterogeneous (e.g., Bachmann, Demir and Frings, 2022; Volpe, 2024). If, instead, low-wage exposed workers had log markdowns closer to zero, our unit benchmark would be exaggerated. (See also Footnote 8 and the end of Section 5.3.) While we, of course, do not have firm- and worker-type-specific markdown estimates, we present a tentative check. We compare the pre-reform cross-sectional relationship between the payroll share of exposed workers and the firm-level composite markdown (which necessarily pools all worker types). We had already reported these results in Figure 3 Panel (d) for the full analysis sample, and in Appendix Figure A.1 Panel (d) for the plausibly monopsonistic firms with negative 2013 log markdowns. For these two samples, we find no or, if anything, a negative relationship.[19] Hence, the unit benchmark for the treatment effect on markdowns is unlikely to be confounded by such a compositional effect, although we note the inherent limitations of our indirect check.

[18] To see this, we informally appeal to an instrumental variables (IV) framework, where the first stage would regress future treatment intensity on baseline-year pre-reform treatment intensity as the instrument. While we, of course, cannot actually study effects on treatment intensity post-reform (as, e.g., markdowns are our outcome and, as mentioned before, hourly wage data ends in 2014), the reader may consider an informal adjustment by dividing the ρ estimates by the pre-reform persistence effects for a given similarly spaced horizon depicted in Figure 3 Panel (f). Below, we proceed without such an explicit adjustment.

[19] In our reading, there is no empirical consensus on heterogeneity patterns in measured markdowns or firm-specific labor supply elasticities by worker group. Our finding—no clear relationship between firm-level markdowns and a proxy for low wages, here the share of workers for whom the minimum wage binds—echoes that in Yeh, Macaluso and Hershbein (2022) for the United States, who find no clear evidence for heterogeneous markdowns between production and nonproduction workers on average or by industry (see Table IV and associated discussion). Mertens (2023) documents between-firm markdown heterogeneity within the German manufacturing sector, finding that lower-wage firms appear to have less monopsonistic markdowns (in our notation, closer to one in levels), albeit in a high-wage sector whereas our variation also stems from lower-wage industries and between-industry variation; we will also include industry-year fixed effects in a robustness check.

6 Results

We now present the results on the effect of the minimum wage on firm-level markdowns. Overall, we find only mixed and at best quantitatively moderate support for the core prediction of monopsony theories, that markdowns become competitive when the minimum wage constrains firms' wage-setting power. While some specifications point to effects in that direction, those effects remain quantitatively modest, with point estimates between zero and about 0.25 rather than the benchmark for a coefficient of 1.00 that we derived in the monopsony model. We focus on the full sample and the monopsonistic firms' subsample (and for completeness, also report results for firms with positive rather than negative baseline log markdowns).

Figure 4 collects various pooled DiD effects, including robustness checks below. For our main specification, Table 2 reports the full regression results for the pooled DiD specification. The corresponding year-specific DiD effects are plotted in Figure 5, replicated in Appendix Figures A.7-A.17 for the robustness checks. We always report results with and without a pre-trend adjustment. Additionally, Appendix Figure A.3 reports all pooled post effects for 2013 as the comparison period (rather than comparing to the pooled pre-period). Since this specification also permits a test for pre-trends in the form of a pooled 2011-2013 pre-period effect, we report those in Appendix Figure A.4.

Pooled DiD effect. Table 2 Panel A reports the pooled post effects compared to the full pre-period. Column (1) does so for the entire sample, Column (2) does so for the sample with baseline period log markdowns below 0 (i.e., establishments with log 2013 markdowns below zero, consistent with the monopsonistic range), and Column (3) does so for the smaller residual group.[20]

For the full sample, we estimate an effect of 0.250 (SE 0.0506). Interpreted through the lens of the monopsony model's prediction, which would have predicted an effect of 1.00, markdowns in the data move only about a quarter of the way towards the competitive, wage-taking benchmark. The effect is quite precisely estimated, with confidence intervals ruling out effects larger than 0.35 and smaller than 0.15. Comparing Columns (2) and (3), the effect is exhibited exactly by the firms that have, before the introduction of the minimum wage, monopsonistic markdowns. By contrast, the remaining group (Column (3)) exhibit no meaningful effect around zero albeit imprecisely estimated (-0.041, SE 0.166). This heterogeneity check qualitatively supports the idea of a true causal effect of the reform on monopsony power, although results remain strikingly small for even that group.

Panel B reports the pooled post effects relative to 2013 (rather than the full pre-period). Effects remain economically similar, and quantitatively essentially identical for the primary sample of initially monopsonistic firms (0.270, SE 0.061).

[20] Measurement error in markdowns might lead to attenuation bias, and such measurement error is most plausible for firms whose log markdown exceeds zero, or such firms evidently might not act as monopsonists or are subject to other input market distortions. See Section 3.3 and the end of Section 5.3.

Pre-trends and pre-trend adjustment. This specification with the 2013 baseline year also permits us to estimate a *pre-period* pooled (covering 2011-2012 and 2014) effect against 2013. We find pre-trend violations for the full sample that includes the firms with initially positive log markdowns in 2013, a firm characteristic that would be hard to square with monopsony. Before turning to inspecting dynamic year-specific DiD effects, we note here already that in the secondary panels of the table (Panels A.2 and B.2), we include the pre-trend adjustment suggested and implemented by Dustmann, Lindner, Schönberg, Umkehrer and Vom Berge (2022) in their study of the same minimum wage introduction (for employment/wage outcomes at the local labor market level). With that adjustment, the coefficients return to their original value for the full sample (now 0.266, from 0.341) and even to zero for the firms with plausibly monopsonistic markdowns (now 0.087, from 0.270).[21] Panel (b) of Figure 4 also shows estimated pooled post effects for the full set of robustness checks. This figure makes clear that across all specifications (which we detail below), markdown effects shrink with pre-trend adjustment, and for the monopsonistic firms, estimates are strikingly close to zero.

Year-specific DiD effects. Complementing the pooled estimate, Figure 5 plots the year-specific DiD estimates, obtained by replacing the pooled post dummy in Equation (31) with year-specific effects for the full set of years. The coefficients represent the treatment effects relative to the excluded year of 2013. We start with the baseline effects in Panel (a), and plot pre-trend-adjusted effects in Panel (b).

Consistent with the pooled DiD effect, we estimate relatively stable year-specific DiD effects after the 2015 introduction of the minimum wage, with point estimates tightly centered around 0.25, driven by the monopsonistic subsample. That is, there is little evidence that a single year comes close to the unit benchmark of the model, and there is no increasing pattern (although slow capital adjustment as in Sorkin, 2015; Hurst, Kehoe, Pastorino and Winberry, 2022, need not affect our markdown estimates). (We note, however, that the DiD design naturally loses power over time due to deteriorating first stages, so that a stable reduced form effect plotted here might mask increasing underlying IV-like effects.)

The dynamic perspective also permits us to gauge pre-trends, and to test for parallel trends. In the full sample, we find moderate pre-trend violations, with treatment intensity being associated with markdowns, if anything, increasing (moving towards the competitive benchmark if initially below it). This placebo pattern suggests that the treatment effect is, if anything, overestimated in favor of the monopsony benchmark. We note that when we zoom in on the firms with initially monopsonistic markdowns, we cannot reject the absence of statistically significantly different pre-trends for any individual year, with however a visual inspection revealing somewhat of a positive pre-trend path.

Again, to account for pre-trends and to gauge sensitivity of our estimates, we also apply the detrending procedure proposed by Dustmann, Lindner, Schönberg, Umkehrer and Vom Berge

[21] We bootstrap standard errors for the pre-trend correction with replacement and 100 repetitions, which points to larger standard errors.

(2022), whereby we estimate a linear time trend based on the pre-event period and then extrapolate it to the post-period—which, in particular for the monopsonistic sample, is expected to push down rather than up the estimates. Panel (b) shows the associated year-specific DiD estimates. We find a reduction in estimated effects, particularly for the monopsonistic sample, as discussed above, with the resulting pooled post-reform effect of now 0.087.

Mean reversion, placebos: alternative baseline years. We also check robustness to different baseline years, rather than 2013, namely the first and last, 2011 and 2014. We both construct treatment intensity for these years, and make them the omitted year for the DiD regression. Of particular interest are the year-specific DiD effects, which we report in Figure 6 (2011) and Appendix Figure A.9 (2014), which also report the pooled post effects.

Most importantly, the 2011 specification also permits an interpretation as *placebo* DiD design when focusing on the years before the true reform in 2015. The assumptions underlying our research design so far enlisted to interpret the specification with the 2013 baseline year, would require the coefficients to be stable around zero up until 2015, and then ascend to a moderate treatment effect between about 0.05 and 0.25, then mirroring the baseline 2013 results. By contrast, mean reversion (because treatment intensity includes baseline markdowns, which might naturally decay over time and move towards one) would predict a smoothly increasing emergence of treatment effects already starting in 2012, *already well before—and hence causally unrelated to—the reform*.

Figure 6 Panel (a) reveals exactly such a pattern consistent with mean reversion in markdowns. That is, from 2012 until 2014—before the minimum wage is active—high-treatment-intensity firms already exhibit a movement towards competitive markdowns—with no acceleration in that pattern whatsoever following the actual reform in 2015. These placebo effects, right before the reform, and the absence of an additional compelling treatment effect following the actual reform, leave substantial room for the measured treatment effect with the 2013 base year to simply reflect mean reversion, rather than a true treatment effect.

Again, the positive slope of the pre-trend effects from the 2011 baseline year motivate a pre-trend adjustment, reported in Panel (b). The pre-trend-adjusted coefficients are centered around zero for all groups. After that correction, we find a zero effect for initially monopsonistic-seeming firms, and even a slightly negative effect for the full sample. While our initial research plan had aimed to make 2013 our baseline period, we view the 2011 specification as an equally valid starting point, and so report the pooled effects for robustness checks for this specification in Figure 7 (with the additional results from estimating pooled effects against the baseline year rather than pooled pre-period in Appendix Figures A.5 and A.6)—all of which contain robustness checks we discuss next for both 2013 and 2011 base years.

As noted above in Footnote 17, a complementary one-stop approach would be a DID-IV specification in which we instrument 2013 treatment intensity with its lagged 2011 counterpart, to strip out transitory variation or noise, and hence also shut off mean revision.

Robustness checks. We also conduct a long series of robustness checks, with Figure 4 showing the pooled post effects, and Appendix Figures A.7-A.17 reporting the year-specific DiD effects. In addition to those results based on the 2013 baseline year, Figure 7 does so for the 2011 baseline year, where we favor the pre-trend adjusted versions in Panel (b) as discussed above. The bottom line is that all estimates across those robustness checks confirm the large distance from the unit benchmark that characterizes the monopsony model.

To alleviate concerns that our estimates may be driven by industry composition, we re-estimate the baseline specifications while additionally including industry (2-digit)-year fixed effects, and find no difference in the results. This implies that industry-year shocks or trends do not affect our estimates, and that the results hold up when considering only within-industry markdown and treatment intensity variation. (We note that it would be interesting to zoom into local labor markets as well, but we have shied away from this for sample size reasons.)

We also show robustness to markdown measures using output elasticity estimates from alternative production function estimation methods. To recap, our baseline method is that in Levinsohn and Petrin (2003) with an Ackerberg, Caves and Frazer (2015) correction and estimated on firms with a 2-digit industry classification. To check for robustness, we add firm-level market shares of intermediate inputs as controls (to alleviate concerns that the intermediate input elasticities are affected by oligopsony wedges in this market), drop the Ackerberg, Caves and Frazer (2015) correction, estimate a translog functional form (hence constructing markdowns permitting richer substitution patterns between intermediates and labor that can depart from Cobb-Douglas, see Mertens and Schoefer, 2024), and coarsen to 1-digit industries.

Finally, we reiterate that we focus on a balanced panel to shut off composition effects from, e.g., selective attrition. As a robustness check, we also consider an unbalanced panel.[22] Results remain similar, with the monopsonstic firms' effect remaining below 0.3 throughout (and an only moderately higher effect for the full sample) for the 2013 base year, and zero or even slightly negative for the specification with the 2011 baseline after the pre-trend adjustment (see Figure 7).

Rationing from labor supply as a source of markdowns: an indirect test. As noted in Section 2 and at the end of Section 3.2, our monopsony prediction stems from a condition in which labor demand is not rationed by labor supply at the minimum wage. By contrast, the incipient disappearance of markdowns may trigger a subset of firms to grow so much as to induce "labor shortages." That is, if labor demand exceeds labor supply at the minimum wage, involuntary markdowns emerge—even though firms are wage takers but they cannot find enough workers to push exposed workers' marginal products to the minimum wage. Importantly, those involuntary, labor-supply-constrained markdowns will still be smaller than the original monopsonistic ones. (By contrast, all other exposed firms shrink or grow moderately and eliminate their exposed workers' wage markdowns completely, our clear-cut benchmark and focus so far.)

[22] We stress that the establishment survey does not solely feature sample turnover due to entry/exit but oftentimes due to employers exiting the survey only, so that this check will remain tentative and only suggestively speaks to potential reallocation effects (Dustmann, Lindner, Schönberg, Umkehrer and Vom Berge, 2022).

In the data, we find that this concern is unlikely to account for our findings. Specifically, we split up firms by their employment growth into the reform (log employment in 2016 minus that in 2013). In Appendix Figures A.14-A.17, we replicate the results for firms with below-median and negative employment growth, respectively, for specifications with 2013 and 2011 as baseline years. Intuitively and in the simple model in Section 2, these firms are less subject to the rationing concern. Yet, even for those firms, we find no evidence for strong markdown movement. Instead, the results reported in Appendix Figures A.14 and A.15 are essentially the same as in our baseline specification in Figure 5, and, in particular, fall to values close to or around zero once zooming into the sample with initial log markdowns below 0 and when applying the pre-trend adjustment. Similarly, for the 2011 base year (Appendix Figures A.16 andA.17), we essentially find identical results as in Figure 6 above. We caveat that this test is inherently tentative, as it assumes that the employment growth is not caused by negative shifts in the labor supply curve facing the firm. Moreover, to our knowledge, there exists no direct empirical engagement with this issue in the context of markdowns.

Reduced form specification: interaction effects and base effects. Our main specification is a structurally derived equation, with the dependent variable, treatment intensity, being the product of payroll share exposed and the firm's baseline log markdown. In a suggestive check and for completeness (as in a reduced form regression sense, our treatment intensity can be viewed as an interaction effect), Appendix Table A.4 reports on a reduced form specification that additionally includes the constituent factors of the treatment intensity as regression base effects (and does not multiply their product/interaction by minus one, importantly). The interaction effect's coefficients remain below 0.12 in magnitude, and we stress that the effect on treatment intensity can then no longer be interpreted with a unit benchmark and loses its structural interpretation.

Ancillary outcome variables: employment and wages. The literature has focused on the impact of the minimum wage primarily on wages and employment. To approximate that type of design, we study the effect of payroll share exposed to the minimum wage as the treatment variable. Appendix Tables A.5 and A.6 report effects on log employment and log wages (the log of the mean wage), both taken from the BHP establishment panel. We find positive wage effects (serving as an "intervention check") and moderately negative employment effects. We note that our payroll share exposed measure, as discussed in Section 5.1 surrounding Equation (29), is not directly comparable to existing designs because it does not calculate a "wage gap" but instead counts any worker even slightly below the minimum wage as fully treated. Interestingly, the negative employment effects appear to be driven by firms that are initially not monopsonistic (as indicated by their 2013 markdowns).

7 Conclusion

At the core of the labor market monopsony is that firms exploit wage-setting power by marking down wages below the marginal revenue product of labor. Hence, exogenous wage constraints should eliminate markdowns by eliminating wage-setting power. We have presented the first direct test of these interlinked predictions. Revisiting the 2015 minimum wage introduction in Germany with a focus on how measured markdowns move at the firm level, we find little evidence supporting the monopsony model's prediction, with the empirical markdown responses just 0–25% of the predicted magnitude.

On the one hand, perhaps our test has taken the model and markdown identifications to too strict a quantitative test. In that regard, we view our paper as part of the small but, we would like to believe, growing tradition of putting specific influential labor market models to sharp micro tests of key model conditions, by means of tailored quasi-experiments or other informative micro moments.[23] Moreover, our design delivers a quantitative rather than binary test. On the other hand, one may reasonably argue that our test is at the very nexus where the rubber meets the road for monopsony markdowns and their empirical counterparts, and the interpretation of minimum wage effects. Either way, it seems constructive to inspect specific links in the chain between model prediction and empirical counterpart.

Four potential explanations may account for our findings: (i) markdowns largely reflect sources other than monopsony, (ii) markdowns are mismeasured, (iii) the minimum wage leads to widespread labor shortages, or (iv) the standard monopsony model does not provide a realistic account of the low-wage labor market. For (i) and (ii), our preferred effect estimates, between 0 and 0.25 compared to a unit benchmark, would require 75% to 100% of the variation in measured markdowns to reflect sources other than monopsony (such as noise or adjustment costs[24]). Regarding (iii), our robustness check among the subset of firms with negative or below-median employment growth into the reform points away from rationed labor demand as the mechanism propping up markdowns despite firms acting as wage takers. Regarding (iv), we note the surprising absence of direct evidence on whether and how firms exploit their potential wage-setting power. We hope that future research may contribute additional pieces along all these lines to the mosaic.

We close with a comment on a macro implication less highlighted in the existing empirical literature—and indeed one that we had initially attempted to also tackle. That is, our paper raises the question of how minimum wages affect misallocation (Hsieh and Klenow, 2009), of whose wedge dispersion heterogeneous monopsony markdowns are believed to be an important source (see, e.g., Berger, Herkenhoff and Mongey, 2022, 2025b; Bachmann, Bayer, Stüber and Wellschmied,

[23] For examples following this template of credible identification of key properties forming tests of canonical labor market models and in part finding puzzling results as well, see, e.g., Jäger, Schoefer, Young and Zweimüller (2020); Di Addario, Kline, Saggio and Sølvsten (2023); Jäger, Roth, Roussille and Schoefer (2024); Bíró, Lindner, Prinz, Branyiczki and Márk (2025); Derenoncourt and Weil (2025).

[24] This example calculation assumes that measurement error applies to markdowns only while payroll shares exposed are stable and precisely measured. See the end of Section 3.3 for our back-of-the-envelope calculation gauging attenuation bias.

2025). Our finding that the incipient mechanism appears greatly attenuated or even absent, namely that binding minimum wages straighten out markdowns, has frustrated this more macro empirical analysis plan. But on the bright side, it has led us to dedicate our paper fully to the micro puzzle.[25]

References

Ackerberg, Daniel A, Kevin Caves, and Garth Frazer, "Identification properties of recent production function estimators," *Econometrica*, 2015, *83* (6), 2411–2451.

Addario, Sabrina Di, Patrick Kline, Raffaele Saggio, and Mikkel Sølvsten, "It ain't where you're from, it's where you're at: hiring origins, firm heterogeneity, and wages," *Journal of Econometrics*, 2023, *233* (2), 340–374.

Azar, José, Emiliano Huet-Vaughn, Ioana Marinescu, Bledi Taska, and Till Von Wachter, "Minimum wage employment effects and labour market concentration," *Review of Economic Studies*, 2024, *91* (4), 1843–1883.

Azar, José and Ioana Marinescu, "Monopsony power in the labor market," *Handbook of Labor Economics*, 2024, *102* (5), 761–827.

Azkarate-Askasua, Miren and Miguel Zerecero, "Union and firm labor market power," *Working Paper*, 2025.

Bachmann, Ronald, Gökay Demir, and Hanna Frings, "Labor market polarization, job tasks, and monopsony power," *Journal of Human Resources*, 2022, *57* (S), S11–S49.

Bachmann, Rüdiger, Christian Bayer, Heiko Stüber, and Felix Wellschmied, "Monopsony makes firms not only small but also unproductive: Why East Germany has not converged," *CEPR Discussion Paper*, 2025.

Bailey, Martha J, John DiNardo, and Bryan A Stuart, "The economic impact of a high national minimum wage: Evidence from the 1966 fair labor standards act," *Journal of Labor Economics*, 2021, *39* (S2), S329–S367.

Bassier, Ihsaan and Joshua Budlender, "When do employers share? Rent sharing, monopsony and minimum wages," *Working paper*, 2025.

Bau, Natalie and Adrien Matray, "Misallocation and capital market integration: Evidence from India," *Econometrica*, 2023, *91* (1), 67–106.

Bellmann, Lisa, Laura Brunner, Peter Ellguth, Philipp Grunau, Christian Hohendanner, Susanne Kohaut, Ute Leber, Iris Möller, Barbara Schwengler, Jens Stegmaier, and Matthias Umkehrer, "IAB Establishment Panel (IAB BP) – Version 9321 v1," 2022.

Berger, David, Kyle Herkenhoff, and Simon Mongey, "Labor market power," *American Economic Review*, 2022, *112* (4), 1147–1193.

_ , _ , and _ , "The Macro Approach to Labor Market Power," *Working Paper*, 2025.

_ , _ , and _ , "Minimum Wages, Efficiency, and Welfare," *Econometrica*, 2025, *93* (1), 265–301.

Bíró, Anikó, Attila Lindner, Daniel Prinz, Réka Branyiczki, and Lili Márk, "Firm heterogeneity and the impact of payroll taxes," *Working Paper*, 2025.

Bosch, Gerhard, "The making of the German minimum wage: a case study of institutional change," *Industrial Relations Journal*, 2018, *49* (1), 19–33.

Bossler, Mario and Hans-Dieter Gerner, "Employment effects of the new German minimum wage: Evidence from establishment-level microdata," *ILR Review*, 2020, *73* (5), 1070–1094.

[25] More precisely, we had attempted an analysis of the minimum wage reform on markdown, markup, and overall wedge and productivity dispersion (e.g., a reform-based analysis in the spirit of capital market applications as in Bau and Matray, 2023; Sraer and Thesmar, 2023)—which would have required the micro stepping stone of markdowns to move more closely to the benchmark in exposed firms.

__ and Thorsten Schank, "Wage inequality in Germany after the minimum wage introduction," *Journal of Labor Economics*, 2023, *41* (3), 813–857.

__ , Nicole Gürtzgen, Benjamin Lochner, Ute Betzl, and Lisa Feist, "The German minimum wage: effects on productivity, profitability, and investments," *Jahrbücher für Nationalökonomie und Statistik*, 2020, *240* (2-3), 321–350.

Bundesbank, Deutsche, "Wage growth in Germany: assessment and determinants of recent developments," *Monthly Report*, 2018, *70* (4), 13–27.

Burauel, Patrick, Marco Caliendo, Markus M Grabka, Cosima Obst, Malte Preuss, Carsten Schröder, and Cortnie Shupe, "The impact of the German minimum wage on individual wages and monthly earnings," *Jahrbücher für Nationalökonomie und Statistik*, 2020, *240* (2-3), 201–231.

Bächmann, Ann-Christin, Lisa Bellmann, Miriam Gensicke, Susanne Kohaut, Iris Möller, Barbara Schwengler, Nikolai Tschersich, and Matthias Umkehrer, "IAB Establishment Panel (IAB-BP) 1993–2022," Data report, Institut für Arbeitsmarkt- und Berufsforschung (IAB) 2023.

Caliendo, Marco, Alexandra Fedorets, Malte Preuss, Carsten Schröder, and Linda Wittbrodt, "The short-term distributional effects of the German minimum wage reform," *IZA Discussion Paper*, 2017.

__ , Carsten Schröder, and Linda Wittbrodt, "The causal effects of the minimum wage introduction in Germany–an overview," *German Economic Review*, 2019, *20* (3), 257–292.

Card, David, Ana Rute Cardoso, Joerg Heining, and Patrick Kline, "Firms and labor market inequality: Evidence and some theory," *Journal of Labor Economics*, 2018, *36* (S1), S13–S70.

__ and Alan B Krueger, "Minimum wages and employment: a case study of the fast-food industry in New Jersey and Pennsylvania: reply," *American Economic Review*, 2000, *90* (5), 1397–1420.

__ , Jörg Heining, and Patrick Kline, "Workplace heterogeneity and the rise of West German wage inequality," *Quarterly Journal of Economics*, 2013, *128* (3), 967–1015.

Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer, "The effect of minimum wages on low-wage jobs," *Quarterly Journal of Economics*, 2019, *134* (3), 1405–1454.

Derenoncourt, Ellora and David Weil, "Voluntary minimum wages: The local labor market effects of national retailer policies," *Quarterly Journal of Economics*, 2025, p. qjaf016.

Dickens, Richard, Stephen Machin, and Alan Manning, "The effects of minimum wages on employment: Theory and evidence from Britain," *Journal of Labor Economics*, 1999, *17* (1), 1–22.

Dobbelaere, Sabien and Jacques Mairesse, "Panel data estimates of the production function and product and labor market imperfections," *Journal of Applied Econometrics*, 2013, *28* (1), 1–46.

__ , Boris Hirsch, Steffen Müller, and Georg Neuschaeffer, "Organized labor, labor market imperfections, and employer wage premia," *ILR Review*, 2024, *77* (3), 396–427.

Drenik, Andres, Simon Jäger, Pascuel Plotkin, and Benjamin Schoefer, "Paying outsourced labor: Direct evidence from linked temp agency-worker-client data," *Review of Economics and Statistics*, 2023, *105* (1), 206–216.

Dube, Arindrajit and Attila Lindner, "Minimum wages in the 21st century," *Handbook of Labor Economics*, 2024, *5*, 261–383.

__ , William Lester, and Michael Reich, "Minimum wage effects across state borders: Estimates using contiguous counties," *The Review of Economics and Statistics*, 2010, *92* (4), 945–964.

Dustmann, Christian, Attila Lindner, Uta Schönberg, Matthias Umkehrer, and Philipp Vom Berge, "Reallocation effects of the minimum wage," *The Quarterly Journal of Economics*, 2022, *137* (1), 267–328.

__ , Chiara Giannetto, Lorenzo Incoronato, Chiara Lacava, Vincenzo Pezone, Raffaele Saggio, and Benjamin Schoefer, "Opting Out of Centralized Collective Bargaining: Evidence from Italy," *NBER Working Paper*,

2025.

_ , Johannes Ludsteck, and Uta Schönberg, "Revisiting the German wage structure," *The Quarterly Journal of Economics*, 2009, *124* (2), 843–881.

Engbom, Niklas and Christian Moser, "Earnings inequality and the minimum wage: Evidence from Brazil," *American Economic Review*, 2022, *112* (12), 3803–3847.

Faia, Ester and Vincenzo Pezone, "The cost of wage rigidity," *Review of Economic Studies*, 2024, *91* (1), 301–339.

Ganzer, Andreas, Alexandra Schmucker, Jens Stegmaier, and Stefanie Wolter, "Establishment History Panel 1975–2022 (BHP 7522 v1)," 2023.

Garloff, Alfred, "Did the German minimum wage reform influence (un) employment growth in 2015? Evidence from regional data," *German Economic Review*, 2019, *20* (3), 356–381.

Haelbig, Mirja, Matthias Mertens, and Steffen Müller, "Minimum wages, productivity, and reallocation," *IWH Working Paper*, 2024.

Harasztosi, Péter and Attila Lindner, "Who pays for the minimum wage?," *American Economic Review*, 2019, *109* (8), 2693–2727.

Hashemi, Arshia, Ivan Kirov, and James Traina, "The production approach to markup estimation often measures input distortions," *Economics Letters*, 2022, *217*, 110673.

Hsieh, Chang-Tai and Peter J Klenow, "Misallocation and manufacturing TFP in China and India," *The Quarterly journal of economics*, 2009, *124* (4), 1403–1448.

Hurst, Erik, Patrick Kehoe, Elena Pastorino, and Thomas Winberry, "The distributional impact of the minimum wage in the short and long run," *NBER Working Paper*, 2022.

IAB, "Integrated Employment Biographies," V13.01.01-190111, Nürnberg 2019.

Jäger, Simon, Benjamin Schoefer, Samuel Young, and Josef Zweimüller, "Wages and the Value of Nonemployment," *Quarterly Journal of Economics*, 2020, *135* (4), 1905–1963.

_ , Christopher Roth, Nina Roussille, and Benjamin Schoefer, "Worker beliefs about outside options," *The Quarterly Journal of Economics*, 2024, *139* (3), 1505–1556.

Jäger, Simon, Shakked Noy, and Benjamin Schoefer, "The German model of industrial relations: Balancing flexibility and collective action," *Journal of Economic Perspectives*, 2022, *36* (4), 53–80.

Jäger, Simon, Suresh Naidu, and Benjamin Schoefer, "Collective bargaining, unions, and the wage structure: An international perspective," in "Handbook of Labor Economics," Vol. 6, Elsevier, 2025, pp. 229–372.

Kline, Patrick, "Labor Market Monopsony: Fundamentals and Frontiers," *Handbook of Labor Economics*, 2025, *102* (6), 655–728.

Lamadon, Thibaut, Magne Mogstad, and Bradley Setzler, "Imperfect competition, compensating differentials, and rent sharing in the US labor market," *American Economic Review*, 2022, *112* (1), 169–212.

Levinsohn, James and Amil Petrin, "Estimating production functions using inputs to control for unobservables," *Review of Economic Studies*, 2003, *70* (2), 317–341.

Loecker, Jan De and Frederic Warzynski, "Markups and firm-level export status," *American Economic Review*, 2012, *102* (6), 2437–71.

Manning, Alan, *Monopsony in motion: Imperfect competition in labor markets*, Princeton University Press, 2013.

_ , "The elusive employment effect of the minimum wage," *Journal of Economic Perspectives*, 2021, *35* (1), 3–26.

Mertens, Matthias, "Labor market power and between-firm wage (in) equality," *International Journal of Industrial Organization*, 2023, *91*, 103005.

_ and Benjamin Schoefer, "From labor to intermediates: Firm growth, input substitution, and monopsony," *NBER Working Paper*, 2024.

Müller, Steffen, "Capital Stock Approximation Using Firm Level Panel Data," *Jahrbücher für Nationalökonomie und Statistik*, 2008, *224* (4), 357–371.

Neumark, David and William Wascher, "Minimum wages and employment: A case study of the fast-food industry in New Jersey and Pennsylvania: Comment," *American Economic Review*, 2000, *90* (5), 1362–1396.

Olley, G Steven and Ariel Pakes, "The Dynamics of Productivity in the Telecommunications Equipment Industry," *Econometrica*, 1996, *64* (6), 1263–97.

Rambachan, Ashesh and Jonathan Roth, "A more credible approach to parallel trends," *Review of Economic Studies*, 2023, *90* (5), 2555–2591.

Robinson, Joan, *The Economics of Imperfect Competition*, Macmillan, 1933.

Rubens, Michael, Yingjie Wu, and Mingzhi Xu, "Estimating factor price markdowns using production models," *International Journal of Industrial Organization*, 2025, *102*, 103177.

Saez, Emmanuel, Benjamin Schoefer, and David Seim, "Payroll taxes, firm behavior, and rent sharing: Evidence from a young workers' tax cut in Sweden," *American Economic Review*, 2019, *109* (5), 1717–1763.

Schoefer, Benjamin and Oren Ziv, "Productivity, place, and plants," *Review of Economics and Statistics*, 2024, *106* (5), 1167–1186.

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

Sraer, David and David Thesmar, "How to use natural experiments to estimate misallocation," *American Economic Review*, 2023, *113* (4), 906–938.

Treuren, Leonard, "Wage markups and buyer power in intermediate input markets," *FEB Research Report Department of Economics*, 2025.

Volpe, Oscar, "Job Preferences, Labor Market Power, and Inequality," *Working Paper*, 2024.

vom Berge, Philipp, Matthias Umkehrer, and Susanne Wanger, "A correction procedure for the working hours variable in the IAB employee history," *Journal of Labour Market Research*, 2023, *57*, 10.

Yeh, Chen, Claudia Macaluso, and Brad Hershbein, "Monopsony in the US labor market," *American Economic Review*, 2022, *112* (7), 2099–2138.

Tables

Table 1: Summary statistics and correlations, analysis sample in 2013

	Log mark-down	Payroll share exposed	Treatm. intensity	Log sales	Log employ-ment	Log capital	Log interm. inputs
Mean	-0.465	0.129	0.056	15.215	3.501	15.585	14.402
SD	0.851	0.223	0.249	1.710	1.313	1.817	1.915
Min	-2.705	0.000	-1.191	11.913	1.386	11.553	10.540
p10	-1.462	0.000	-0.025	13.079	1.872	13.246	11.992
p25	-1.007	0.009	0.000	13.886	2.442	14.089	12.963
p50	-0.485	0.030	0.006	15.072	3.434	15.562	14.267
p75	0.066	0.119	0.037	16.396	4.337	16.735	15.712
p90	0.597	0.456	0.214	17.637	5.321	18.065	17.048
Max	1.755	1.000	2.237	19.500	6.895	19.686	18.998
Log markdown	1.000	0.020	-0.480	-0.293	-0.019	-0.246	-0.421
Payroll share exp.	-	1.000	0.363	-0.358	-0.276	-0.241	-0.310
Treatment int.	-	-	1.000	-0.038	-0.148	-0.131	0.033
Log sales	-	-	-	1.000	0.896	0.828	0.971
Log employment	-	-	-	-	1.000	0.778	0.830
Log capital	-	-	-	-	-	1.000	0.785
Log interm. inputs	-	-	-	-	-	-	1.000
Autocorr. 2012-13	0.914	1.000	1.000	0.996	0.996	0.998	0.988
Autocorr. 2013-14	0.923	1.000	1.000	0.995	0.992	0.999	0.987

Notes: The table reports summary statistics and correlations of the main variables for our analysis sample in 2013.

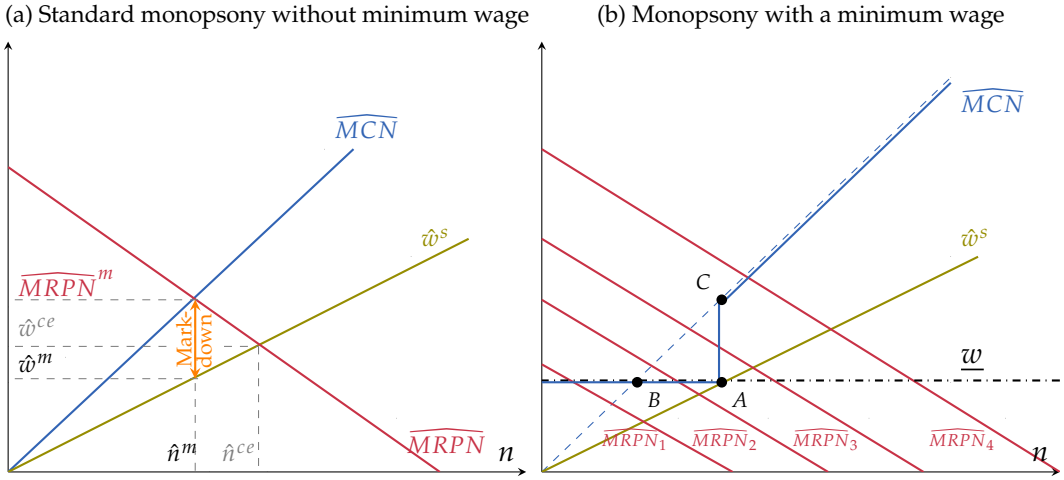
Table 2: Regression results: effect of treatment intensity on log markdowns

	Log markdown (2013)		
A.1: (Rel. to 2011-2014)	Basic DiD specification		
	any	≤ 0	> 0
Post*Treatment intensity	0.250***	0.257***	-0.041
	(0.051)	(0.056)	(0.166)
Constant	-0.441***	-0.792***	0.464***
	(0.001)	(0.003)	(0.007)
Observations	7,232	5,232	2,000
R^2	0.882	0.827	0.714
A.2: (Rel. to 2011-2014)	With pre-trend correction		
Post*Treatment intensity	0.164***	0.047	-0.147
	(0.074)	(0.104)	(0.244)
B.1: (Rel. to 2013)	Basic DiD specification		
	any	≤ 0	> 0
Post*Treatment intensity	0.341***	0.270***	0.075
	(0.055)	(0.061)	(0.171)
Pre*Treatment intensity	0.121***	0.018	0.154
	(0.039)	(0.039)	(0.108)
Constant	-0.446***	-0.794***	0.474***
	(0.002)	(0.004)	(0.009)
Observations	7,232	5,232	2,000
R^2	0.882	0.827	0.714
B.2: (Rel. to 2013)	With pre-trend correction		
Post*Treatment intensity	0.266***	0.087	-0.018
	(0.084)	(0.110)	(0.243)
Pre*Treatment intensity	0.102***	0.040	0.129
	(0.029)	(0.025)	(0.084)

Notes: The table reports results of the DiD establishment-level regression (the specification in Equation (31)) of log markdowns on treatment intensity, reporting the pooled post effect on treatment intensity, for 2013 as the baseline year for treatment intensity and the omitted year. Treatment intensity is defined as the negative product of 2013 log markdown and payroll share exposed to the reform. The DiD regression is motivated structurally in our monopsony model, which would predict a coefficient of one (whereby exposed workers' markdowns move towards the competitive, wage-taking benchmark). Across the columns, the table reports these effects separately for the full analysis sample (Column (1)), as well as for the firms with 2013 log markdowns below zero (Column (2), plausibly monopsonistic) and above (Column (3), the residual category). In Panel A, the estimated effects are relative to the pooled pre-reform period. In Panel B, the estimated effects are relative to 2013, which additionally permits the estimation of a pooled pre-period effect. Each second subpanel (A.2 and B.2) additionally reports pooled post-period effects in a specification that applies a pre-trend correction following Dustmann, Lindner, Schönberg, Umkehrer and Vom Berge (2022). Standard errors are clustered on the establishment level. Standard errors for the specification with the pre-trend correction are bootstrapped (with replacement and 100 repetitions).

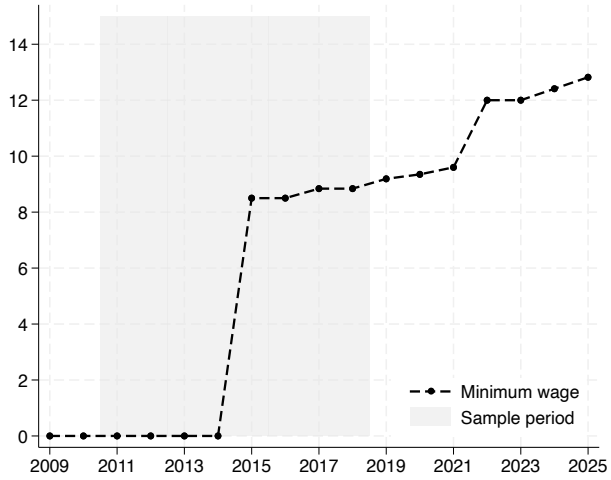
Figures

Figure 1: The simple economics of monopsony and markdowns with and without a minimum wage



Notes: The figure illustrates the simple economics of monopsony with and without a minimum wage, visualizing the simple formal model described in Section 2. Both panels are to be understood in logs (hence hatted variables), and plot the labor supply (inverse wage), marginal cost of labor, and marginal revenue product of labor curves. Panel (a) depicts the solution without the minimum wage for one specific firm (characterized by a marginal product curve). Panel (b) does so with the minimum wage, as well as for different potential firms differentiated by productivity (marginal product curves). The minimum wage introduces a kink and flat segment into the marginal cost of labor curve in the area to the left of point A, where it binds. Firms to the left of point A become wage takers and have their markdown disappear. Point B separates these firms into shrinking (left) and growing (right) firms. Control firms (with productivity situating them to the right of point C) are not exposed. A residual group, on the vertical segment between points B and C, see their markdowns shrink but not necessarily disappear due to their employment being rationed by labor supply.

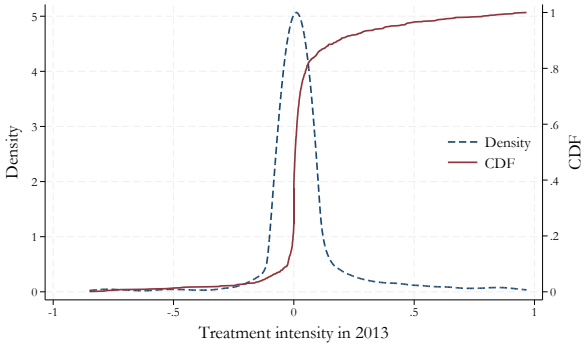
Figure 2: The 2015 introduction of the minimum wage and its evolution



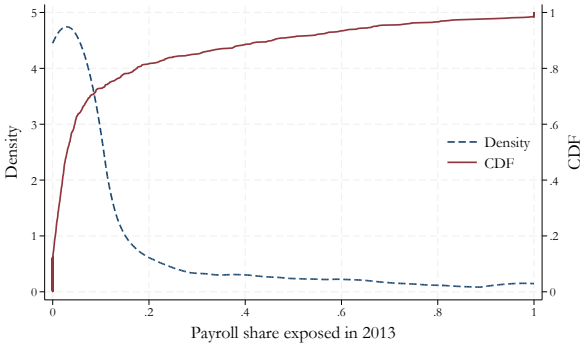
Notes: The figure plots the evolution of statutory minimum wages in Germany. The values are effective as of January 1 of each year, except in 2021 and 2022, when multiple increases occurred. In 2021, the minimum wage rose to 9.50 EUR in January and to 9.60 EUR in July. In 2022, it increased to 9.82 EUR in January, to 10.45 EUR in July, and to 12.00 EUR in October. The shaded area denotes the time window of our DiD analysis. Source: https://www.destatis.de/DE/Themen/Arbeit/Verdienste/Mindestloehne/Tabellen/gesetzlicher-mindestlohn.html.

Figure 3: Descriptives: treatment intensity, share of payroll exposed to minimum wage, and markdowns in 2013.

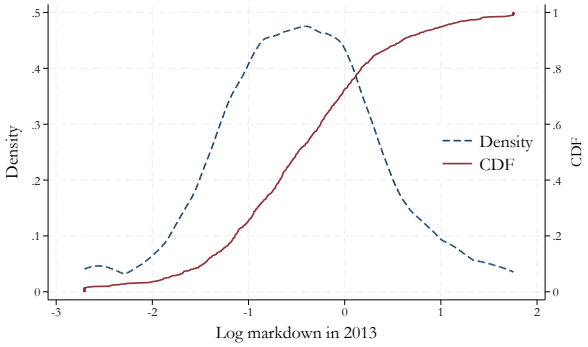
(a) Treatment intensity



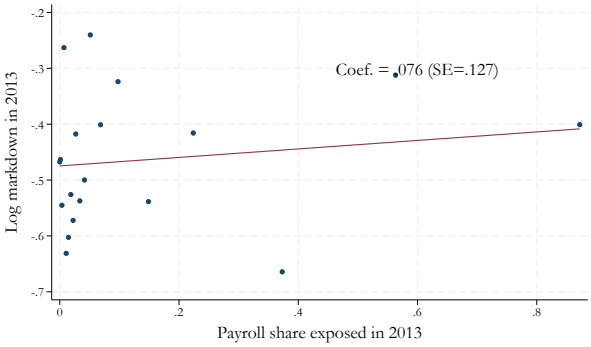
(b) Payroll share exposed



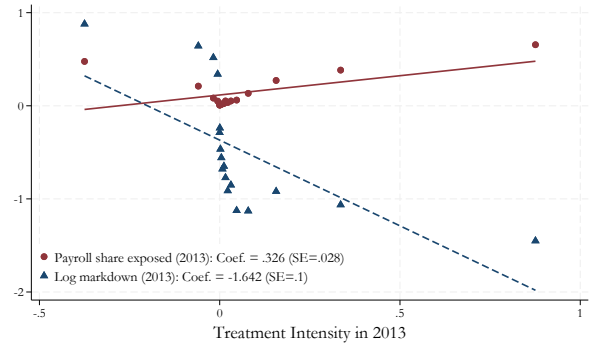
(c) Log markdown



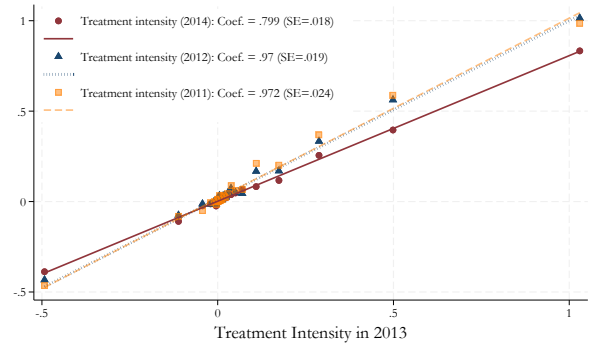
(d) Payroll share exposed and log markdown



(e) Log markdown and payroll share exposed, and treatment intensity



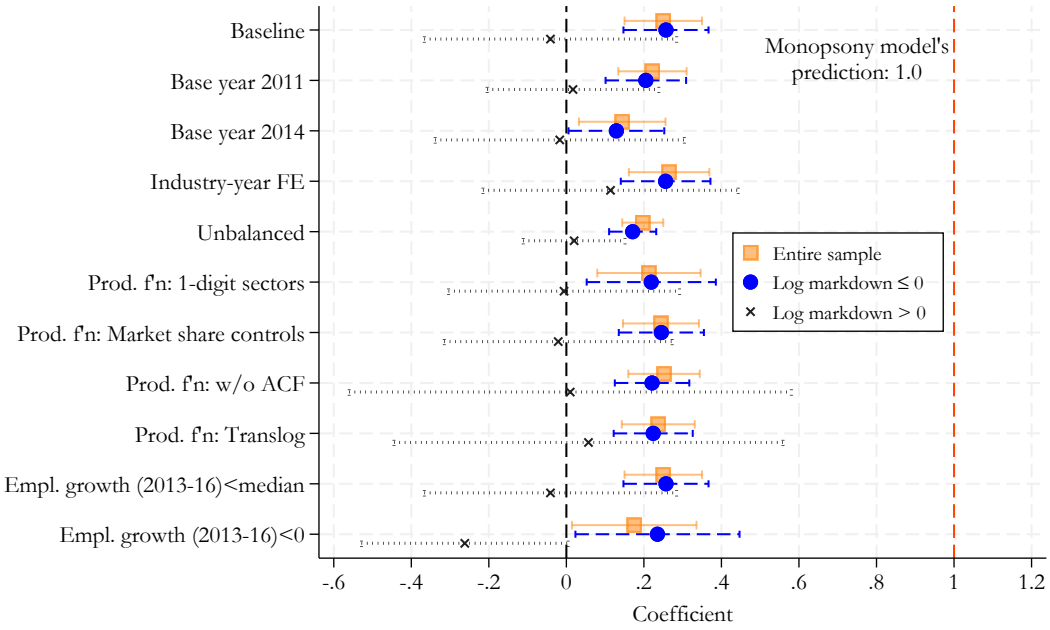
(f) Persistence of treatment intensity



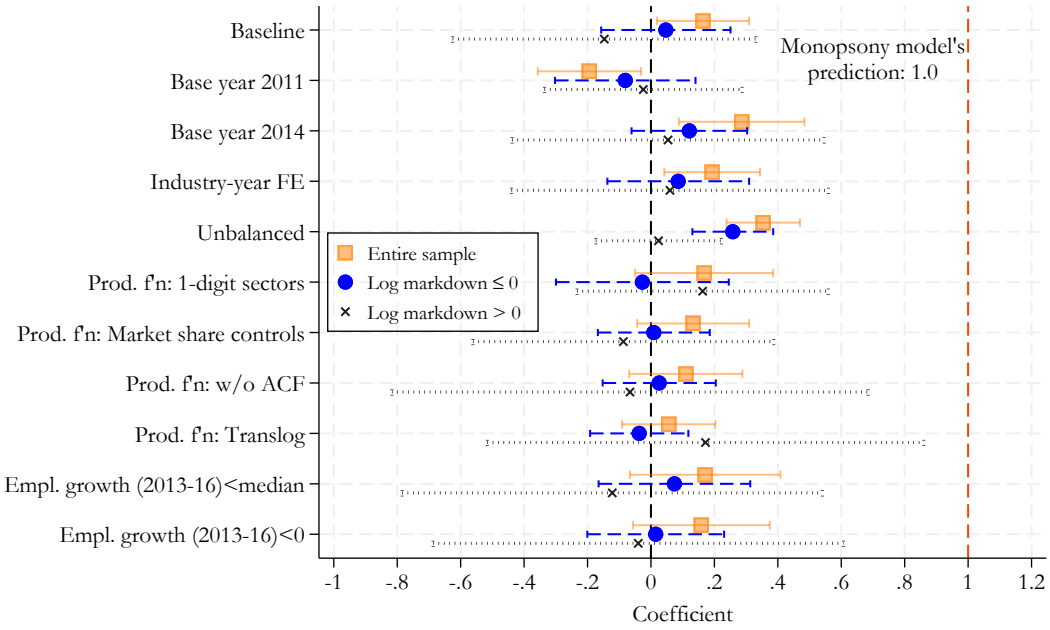
Notes: The figure shows descriptives of our analysis sample of establishments. Panels (a)–(c) show the distributions (kernel density estimates and CDFs) of treatment intensity, the share of payroll exposed to the minimum wage in 2013 (i.e., hourly wages below 8.50 EUR), and the 2013 log markdown; treatment intensity is the product of the latter two (times (-1)). Panel (d) is a binned scatter plot of 2013 log markdowns against 2013 exposed payroll share exposed in 2013. Panel (e) is a binned scatterplot that illustrates the variation underlying treatment intensity by plotting 2013 log markdowns and exposed payroll share against treatment intensity. Panel (f) is a binned scatter plot that reports on the persistence in treatment intensity by plotting 2011, 2012, and 2014 treatment intensity against that in 2013 (and we cannot enter higher years due to the absence of hourly wage data required to compute the payroll share exposed, see Section 5.1). Appendix Figures A.1 and A.2 replicate this figure for the subsamples of firms with negative and positive 2013 log markdowns, respectively.

Figure 4: Overview of treatment effects across specifications (pooled post period compared to pooled pre period)

(a) Baseline specification



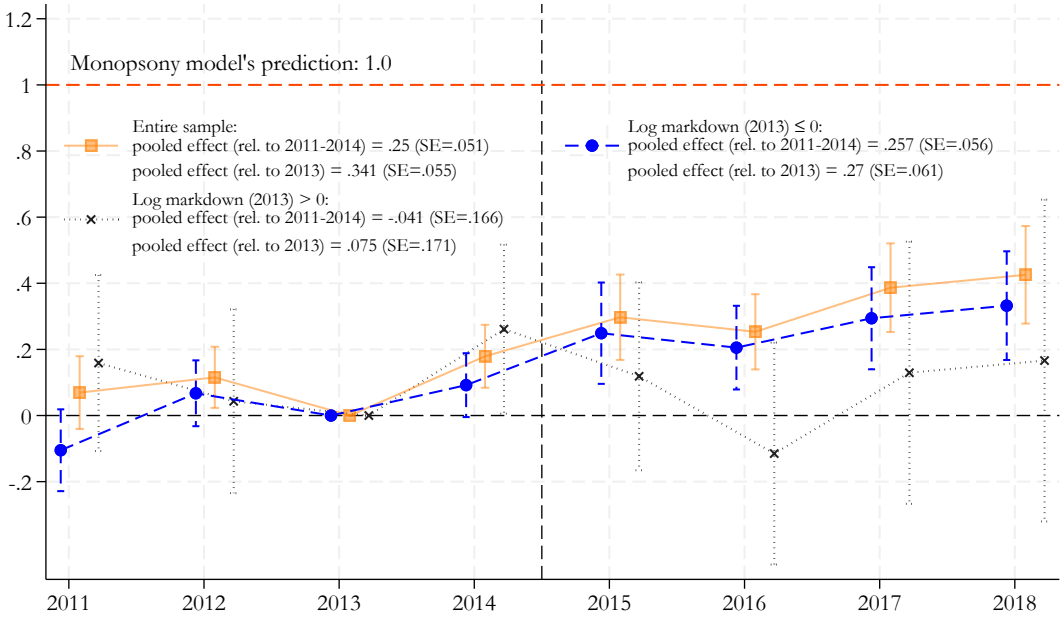
(b) Pre-trend correction



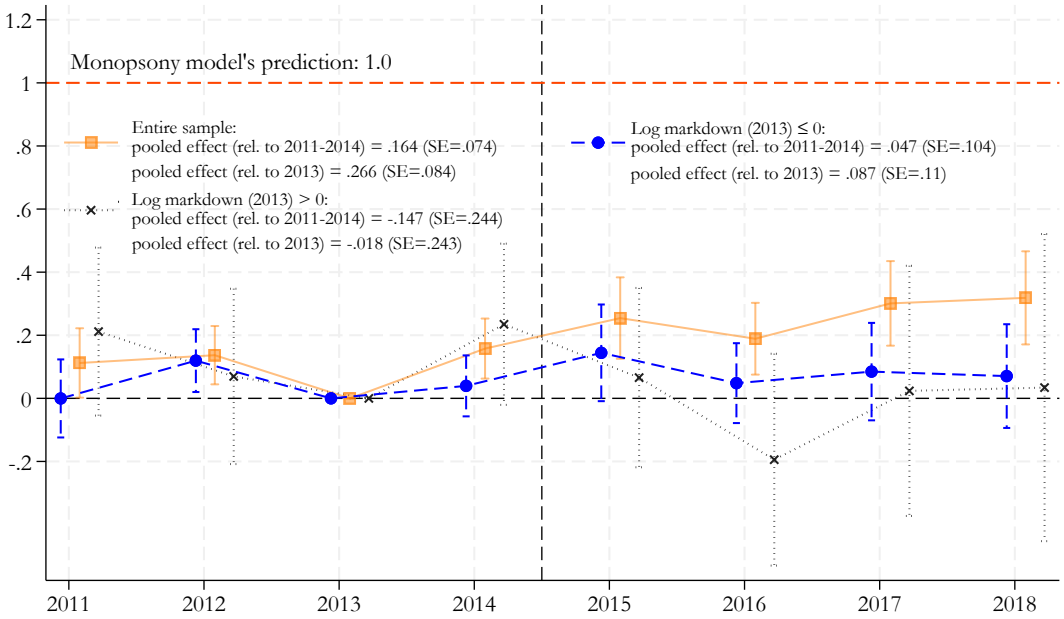
Notes: The figure plots the results of the DiD establishment-level regression (the specification in Equation (31)) of log markdowns on treatment intensity, reporting the pooled post effect on treatment intensity, for 2013 as the baseline year for treatment intensity and the omitted year. Treatment intensity is defined as the negative product of 2013 log markdown and payroll share exposed to the reform. The DiD regression is motivated structurally in our monopsony model, which would predict a coefficient of one (whereby exposed workers' markdowns move towards the competitive, wage-taking benchmark), which we highlight as a unit benchmark. The figure reports these effects separately for the full analysis sample, as well as for the firms with 2013 log markdowns below zero (plausibly monopsonistic) and above (the residual category). Panel (a) reports the effects in our baseline DiD regression and Panel (b) does so including a pre-trend adjustment as in Dustmann, Lindner, Schönberg, Umkehrer and Vom Berge (2022).

Figure 5: Effect of treatment intensity on markdowns: 2013 baseline year

(a) Baseline specification



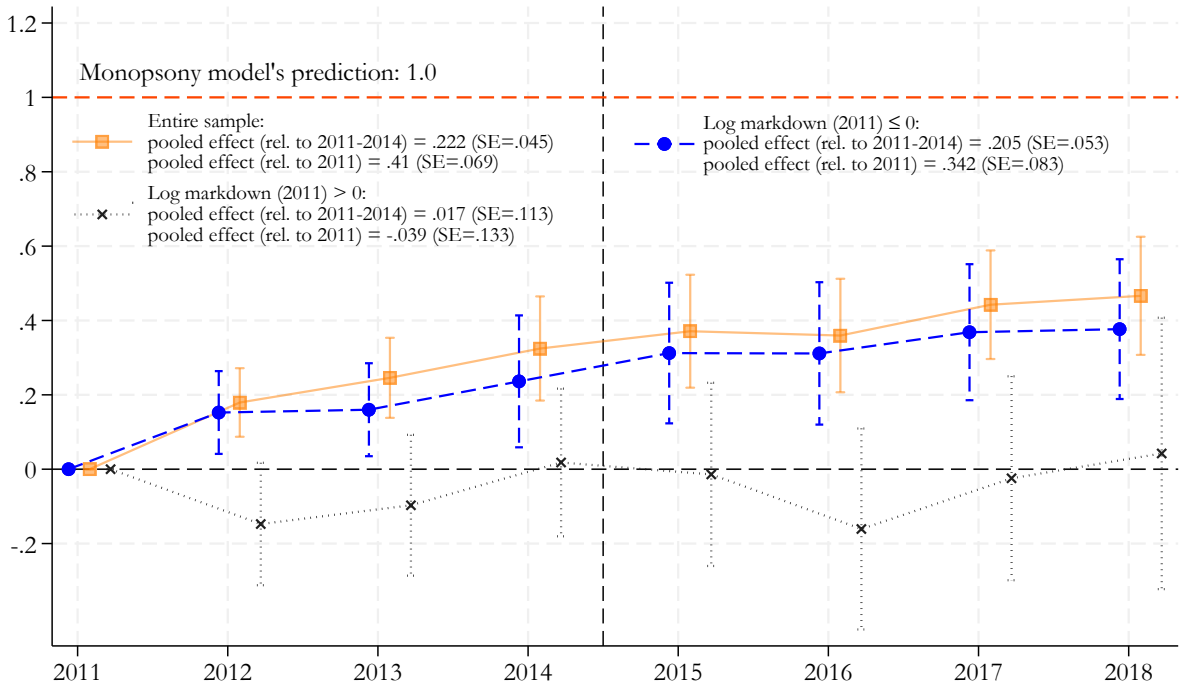
(b) Pre-trend correction



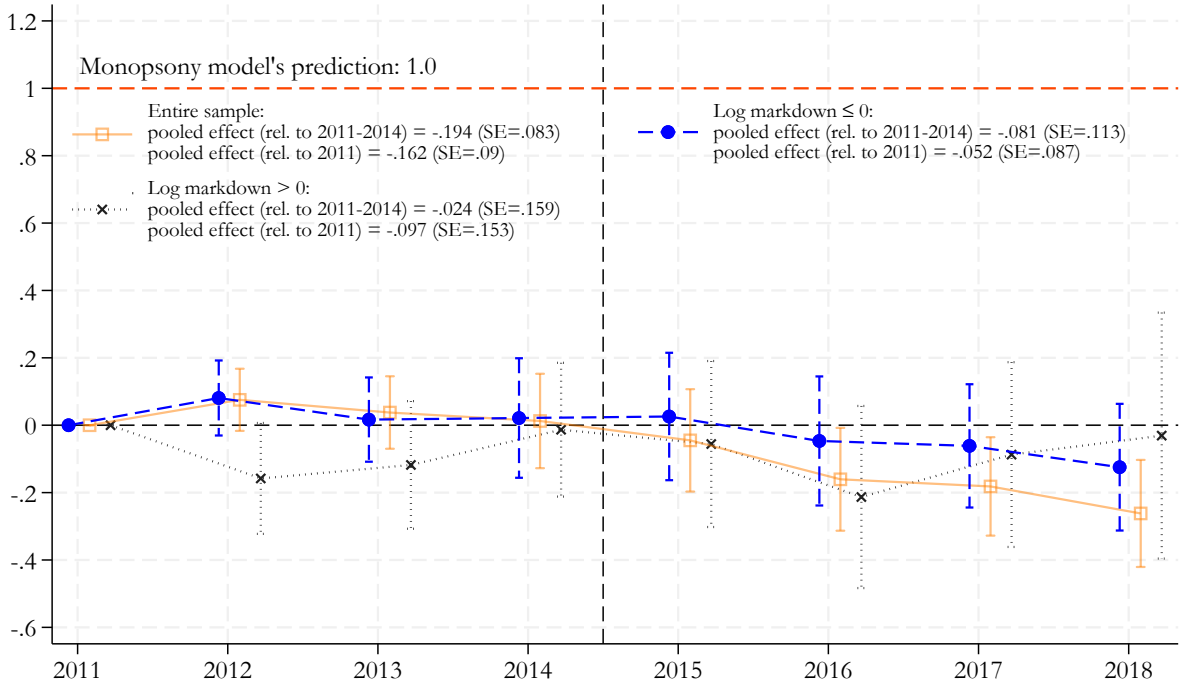
Notes: The figure plots the results of the DiD establishment-level regression (the year-specific version of the specification in Equation (31)) of log markdowns on treatment intensity, reporting the year-specific treatment effects (coefficient on treatment intensity), for 2013 as the baseline year for treatment intensity and the omitted year.. Treatment intensity is defined as the negative product of 2013 log markdown and payroll share exposed to the reform. The DiD regression is motivated structurally in our monopsony model, which would predict a coefficient of one (whereby exposed workers' markdowns move towards the competitive, wage-taking benchmark), which we highlight as a unit benchmark. The figure reports these effects separately for the full analysis sample, as well as for the firms with 2013 log markdowns below zero (plausibly monopsonistic) and above (the residual category). Panel (a) reports the effects in our baseline DiD regression and Panel (b) does so including a pre-trend adjustment as in Dustmann, Lindner, Schönberg, Umkehrer and Vom Berge (2022).

Figure 6: Placebo and robustness check: 2011 base year

(a) Baseline specification



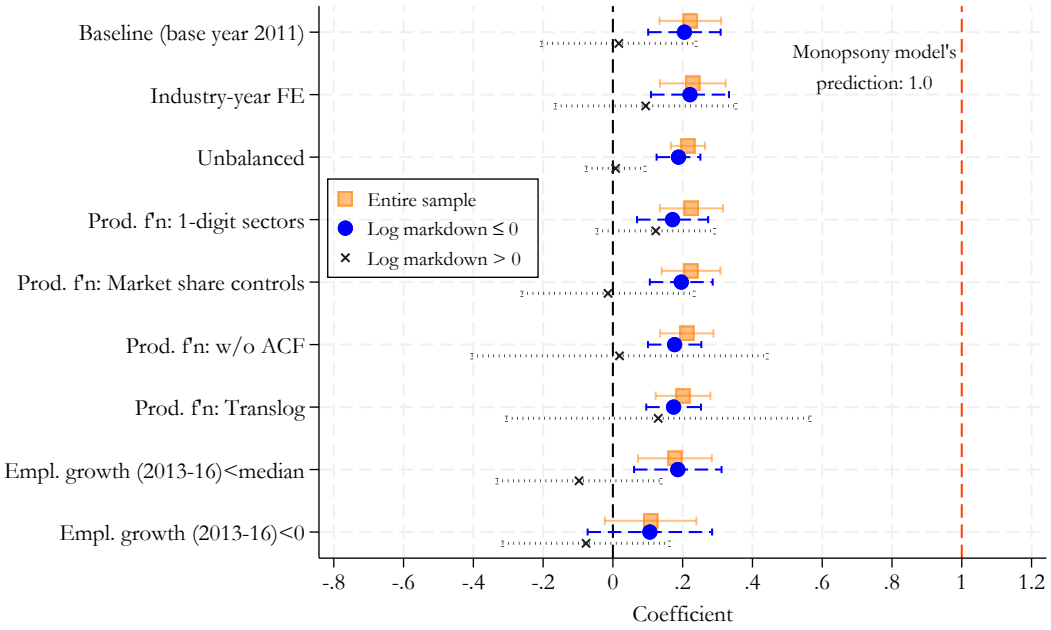
(b) Pre-trend correction



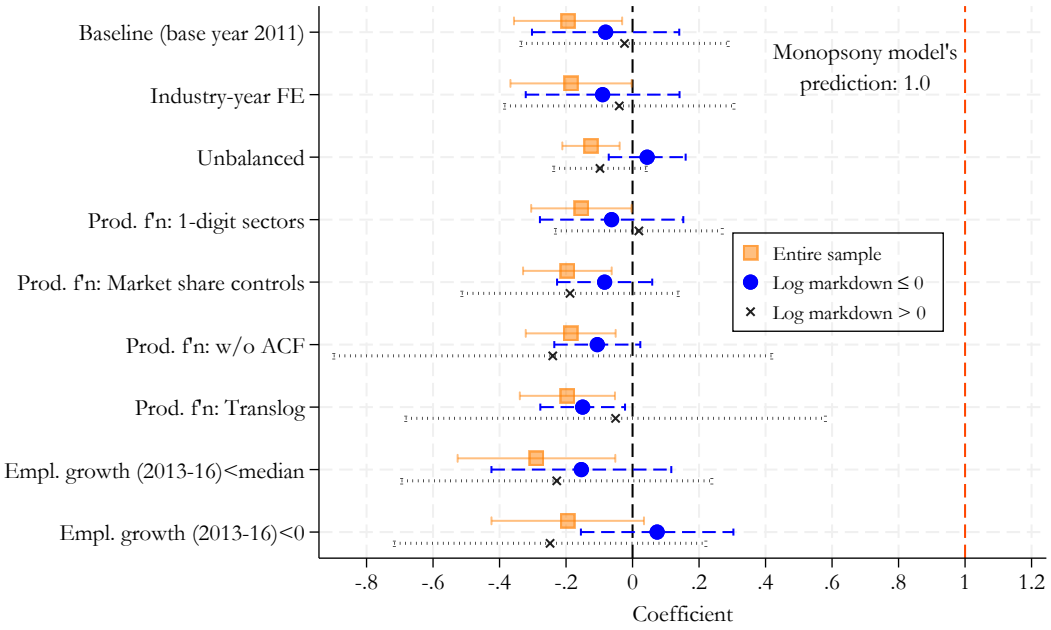
Notes: The figure replicates Figure 5 but uses 2011 rather than 2013 as the baseline year for treatment intensity and the omitted year. While years from 2015 onward remain the years of interest for the treatment effect, the pre-reform effects between 2012 and 2014 can be interpreted as a placebo check for 2015-2017 in the specification with 2013 as the base year, and can detect, e.g., patterns of mean reversion.

Figure 7: Baseline year of 2011: Overview of treatment effects across specifications (pooled post period compared to pooled pre period)

(a) Baseline specification



(b) Pre-trend correction



Notes: The figure replicates Figure 4 but uses 2011 rather than 2013 as the baseline year for treatment intensity and the omitted year.

Appendix

A Methodological appendix: details on production function estimation

The measure of wage markdowns requires the estimation of production function elasticities. We follow the proxy approach based on structural estimation of firms' first order conditions.

Estimating output elasticities. The main idea of the approach is that estimating production function involves a simultaneity problem between the shocks, presumably Hicks neutral productivity, and the choices of the inputs. The problem is spelled out in the original contribution by Olley and Pakes (1996) for a production function with capital and labor. The unbiased estimator then relies on decomposing the shocks in an orthogonal and an anticipated one, which affects the input choices. The proxy method devised in Olley and Pakes (1996) consists in specifying capital as a monotonic function of the anticipated shock: the inverse of this function is then used to substitute out for the anticipated shock. In essence, under certain theoretical and statistical assumptions, one can invert the optimal policy function of a variable input and this allows an econometrician to "observe" unobserved productivity shocks. This is the sense in which capital proxies for the anticipated, or else endogenous, part of the shock. By including the inverse function in the estimation one can control for the unobserved productivity. The whole procedure consists of two stages: in the first stage the proxy variable for anticipated productivity is used as control in a regression that identifies the elasticities to the free variable. The second stage is a regression used to estimate the elasticities to the remaining inputs.

In subsequent work Levinsohn and Petrin (2003) note that investment tends to have a lumpy dynamic with numerous zeros, due to adjustment costs and other non-linearities, and as such it may perform poorly as a proxy variable. Those authors then propose to use intermediate inputs as a proxy variable. We follow their approach and briefly describe it in this appendix. Our estimation also includes the adjustment indicated in Ackerberg, Caves and Frazer (2015), who noted a functional dependence between the production function and the proxy function used to substitute for the anticipated productivity shock. They propose two corrections, which we describe more in detail below.

The underlying conceptual framework is one in which firms have heterogeneous productivity realizations and choose inputs, including variable ones. As in the model of the main text, the firm operates a production function using labor, capital, and intermediate inputs. We assume a Cobb-Douglas production function in logs (where we appeal to homogeneous labor, and hence have switched to the "composite" labor input; moreover, *in this appendix, we depart from the main text and use lower case letters to denote logged variables*):

$$y_t^i = \beta_0 + \alpha k_t^i + \beta x_t^i + \gamma l_t^i + z_t^i. \tag{A.1}$$

The method consists in obtaining structural restrictions to identify the residuals. It assumes that the shock can be decomposed as $z_t^i = \omega_t^i + \epsilon_t^i$, where ω_t^i is an anticipated shock and ϵ_t^i is an ex-post shock. Intermediate inputs, which are observed, are then used as a proxy variable under the assumption of a monotone mapping with ω_t^i. Formally, we can write the demand for intermediate inputs as:

$$x_t^i = f_t^i(l_t, k_t^i, \omega_t^i). \tag{A.2}$$

The function in Equation (A.2) is written conditional on the optimal labor choice following Ackerberg, Caves and Frazer (2015). Assuming strict monotonicity of x_t^i, conditional on capital and labor,

with respect to ω_t^i, Equation (A.2) can be inverted to yield $\omega_t^i = (f_t^i)^{-1}(l_t^i, k_t^i, x_t^i)$. Since intermediate inputs are not chosen based on dynamic considerations the monotonicity assumption is likely to hold. Substituting this into Equation (A.1) gives:

$$y_t^i = \beta_0 + \alpha k_t^i + \gamma l_t^i + \beta x_t^i + (f_t^i)^{-1}(l_t^i, k_t^i, x_t^i) + \epsilon_t^i = \Phi_t^i(l_t^i, k_t^i, x_t^i) + \epsilon_t^i. \tag{A.3}$$

The estimation of the production function, and its input elasticities, is done in two stages. In the first stage, we invert intermediate input demand (Equation (A.2), substitute it in the production function to eliminate ω_t^i and obtain Equation (A.3). We then regress sales (deflated by CPI) on a polynomial approximation of $\Phi_t^i(l_t, k_t^i, x_t^i)$. In other words, the intermediate inputs demand function is used as a proxy to control for unobserved firm-specific productivity shocks. The first stage then attempts to identify the coefficients of the free variable input, such as labor.

Although this step does not identify structural parameters, it delivers an estimate $\hat{\Phi}_t^i(l_t^i, k_t^i, x_t^i)$ and identifies the productivity component ω_t^i from transitory shocks ϵ_t^i.

A second stage regression allows us to compute the elasticity to capital. Following Olley and Pakes (1996) or Levinsohn and Petrin (2003), we assume that firm productivity follows an AR(1) process and that capital does not respond to the error: ξ_t, the innovation in productivity over last period expectation, $\xi_t = \omega_t - E[\omega_t \mid \omega_{t-1}]$. Upon defining, $y_t^{\star i}$ as output net of labor contribution, one obtains:

$$y_t^{\star i} = y_t^i - \gamma l_t^i = \beta_0 + \alpha k_t^i + \beta x_t^i + E[\omega_t \mid \omega_{t-1}] + \epsilon_t^{i\star}, \tag{A.4}$$

where the composite error is $\epsilon_t^{i\star} = \xi_t^i + \epsilon_t^i$. The estimation of the second stage produces consistent estimates of the capital and intermediate input elasticities under two moments conditions: $E[\xi_t + \epsilon_t, k_t] = E[\xi_t k_t] = 0$ and $E[\xi_t + \epsilon_t, x_{t-1}] = 0$.

Finally, note that Ackerberg, Caves and Frazer (2015) demonstrated that a functional dependence (collinearity) problem arises because the same variables (capital, labor, and the proxy) are in both the production function and the (inverted) demand function used as a proxy for productivity. Ackerberg, Caves and Frazer (2015) propose two corrections to deal with this. The first stage is used only to estimate a nonparametric function that captures the combined effect of the capital, the proxy variable, and labor.[26] The coefficients of the variable inputs are identified in the second stage through a Generalized Method of Moments (GMM) approach and it exploits the fact that the lagged values of inputs (chosen prior to the current unpredictable productivity shock) can be valid instruments.

Estimation: sample, summary statistics, and specifications. We conduct the production function estimation in the IAB Establishment Panel, and employ a wider sample from 2001 to 2019. Specifically, we collect information on establishments' revenues, labor inputs (number of employees), intermediate inputs, capital, wage bill, and industry. We measure the labor input in heads. We weight part-time workers by 0.5 relative to full-time workers, as the exact number of working hours is absent in our full data. The capital input is the perpetual inventory method to approximate the establishment-level capital stock, as proposed by Müller (2008), combining net investments and depreciation rates, available from the national accounts, for different types of capital goods. Intermediate inputs is measured as expenditure on intermediate goods and services.

While the initial sample size, from the EP, covers 198,000 establishment-years observations from 2001 to 2019, the requirements of the production function estimation narrows the sample to 36,160 establishment-years observations. We also exclude very small firms and restrict the sample

[26] Olley and Pakes (1996) and Levinsohn and Petrin (2003) employed the first stage to estimate the elasticity of the free input, such as labor.

to firms with at least 5 employees. We also exclude firms in financial, insurance and in the public sector. Sales used in the production function estimation are winsorized at 2%.

Table A.3 reports on the summary statistics of the ingredients in the production function estimation in our estimation sample, and the resulting output elasticities and markdowns in the production function estimation sample.[27]

Our baseline estimation is done on 2-digit industries and taking as baseline a Cobb-Douglas production function. To account for regional heterogeneity, we add state fixed effects. To check for robustness, we add firm-level market shares of intermediate inputs as controls (to alleviate concerns that the intermediate input elasticities are affected by oligopsony wedges in this market), drop the ACF correction, estimate a translog functional form (to permit richer intermediates-labor substitution than Cobb-Douglas), and coarsen to 1-digit industries.

[27]Only two firms in our final analysis sample will exhibit at least one negative estimated output elasticity; we have confirmed that dropping these two firms does not change our results.

Additional tables

Table A.1: Summary statistics and correlations, analysis sample in 2013 with with log markdown (2013) ≤ 0

	Log mark-down	Payroll share exposed	Treatm. intensity	Log sales	Log employ-ment	Log capital	Log interm. inputs
Mean	-0.853	0.119	0.109	15.477	3.551	15.721	14.816
SD	0.607	0.214	0.249	1.736	1.351	1.865	1.878
Min	-2.705	0.000	0.000	11.913	1.386	11.553	10.540
p10	-1.671	0.000	0.000	13.251	1.792	13.271	12.444
p25	-1.174	0.009	0.004	14.167	2.442	14.299	13.391
p50	-0.745	0.027	0.017	15.420	3.450	15.737	14.822
p75	-0.374	0.090	0.070	16.744	4.416	16.996	16.137
p90	-0.151	0.424	0.336	17.729	5.389	18.273	17.234
Max	-0.002	1.000	2.237	19.500	6.895	19.686	18.998
Log markdown	1.000	-0.063	-0.351	-0.189	0.041	-0.232	-0.255
Payroll share exp.	-	1.000	0.796	-0.359	-0.293	-0.290	-0.323
Treatment int.	-	-	1.000	-0.237	-0.257	-0.234	-0.192
Log sales	-	-	-	1.000	0.910	0.827	0.981
Log employm.	-	-	-	-	1.000	0.778	0.869
Log capital	-	-	-	-	-	1.000	0.788
Log interm. inputs	-	-	-	-	-	-	1.000
Autocorr. 2012-13	0.850	1.000	1.000	0.996	0.996	0.999	0.989
Autocorr. 2013-14	0.875	1.000	1.000	0.996	0.991	0.999	0.989

Notes: The table reports summary statistics and correlations of the main variables for our analysis sample in 2013, for firms with log markdown (2013) ≤ 0.

Table A.2: Summary statistics and correlations, analysis sample in 2013 with with log markdown (2013) > 0

	Log mark-down	Payroll share exposed	Treatm. intensity	Log sales	Log employ-ment	Log capital	Log interm. inputs
Mean	0.549	0.154	-0.084	14.530	3.370	15.203	13.319
SD	0.483	0.245	0.187	1.430	1.199	1.625	1.559
Min	0.000	0.000	-1.191	11.913	1.386	11.553	10.540
p10	0.067	0.000	-0.227	12.900	1.872	13.220	11.320
p25	0.162	0.008	-0.070	13.474	2.398	13.970	12.221
p50	0.400	0.039	-0.013	14.472	3.418	15.119	13.240
p75	0.826	0.176	-0.002	15.312	4.025	16.162	14.285
p90	1.297	0.558	0.000	16.436	4.970	17.154	15.399
Max	1.755	1.000	0.000	18.502	6.895	19.686	17.974
Log markdown	1.000	-0.004	-0.360	-0.088	0.030	-0.220	-0.270
Payroll share exp.	-	1.000	-0.723	-0.349	-0.229	-0.053	-0.268
Treatment int.	-	-	1.000	0.287	0.116	0.104	0.283
Log sales	-	-	-	1.000	0.917	0.828	0.947
Log employment	-	-	-	-	1.000	0.778	0.846
Log capital	-	-	-	-	-	1.000	0.824
Log interm. inputs	-	-	-	-	-	-	1.000
Autocorr. 2012-13	0.795	1.000	1.000	0.994	0.995	0.999	0.976
Autocorr. 2013-14	0.772	1.000	1.000	0.993	0.994	0.999	0.976

Notes: The table reports summary statistics and correlations of the main variables for our analysis sample in 2013, for firms with log markdown (2013) > 0.

Table A.3: Production function inputs and estimated wage markdown

Variable	N	Mean	SD	p10	p25	p50	p75	p90
Log sales	36160	14.178	1.482	12.486	13.117	13.886	15.055	16.263
Log labor	36160	2.640	1.081	1.504	1.792	2.351	3.258	4.190
Log capital	36160	13.867	1.538	12.052	12.800	13.600	14.703	15.983
Log interm. inputs	36160	13.268	1.697	11.225	12.082	13.009	14.279	15.645
Elasticity w.r.t. labor	36160	0.445	0.154	0.275	0.336	0.445	0.545	0.644
Elasticity w.r..t. interm. inputs	36160	0.517	0.143	0.338	0.441	0.530	0.620	0.708
Elasticity w.r.t. capital	36160	0.064	0.117	0.013	0.037	0.060	0.093	0.133
Log markdown	36160	-0.540	1.022	-1.768	-1.145	-0.531	0.081	0.708

Notes: The table reports the firm-level variables used in the production function estimation, the resulting output elasticities, and the markdown measure. The sample covers establishments at the 2-digit industry level, excluding finance, insurance, and the public sector. The baseline specification follows Levinsohn and Petrin (2003) with the Ackerberg, Caves and Frazer (2015) correction, using data from 2001–2019. We have checked in our analysis sample for negative output elasticities, and found that only two of our sample firms exhibit at least one negative output elasticity, and confirmed that dropping these firms does not change our results.

Table A.4: Regression results: reduced form DiD regression of log markdowns on 2013 payroll share exposed, 2013 log markdowns, and their interaction

	log markdown (2013)		
	any	≤0	>0
Post*Payroll share (2013)	0.194***	0.156	0.172
	(0.057)	(0.099)	(0.150)
Post*Log markdown (2013)	-0.069***	-0.044*	-0.127**
	(0.017)	(0.024)	(0.060)
Post*Payroll share (2013) *Log markdown (2013)	-0.073	-0.113	-0.004
	(0.066)	(0.106)	(0.267)
Constant	-0.465***	-0.812***	0.488***
	(0.005)	(0.010)	(0.017)
Observations	7,232	5,232	2,000
R^2	0.883	0.827	0.717

Notes: The table reports results of the DiD establishment-level regression (the specification in Equation (31)) of log markdowns on three 2013 independent variables and their interaction with the pooled post effect. The three variabels are the 2013 payroll share exposed to the reform, 2013 log markdown, and the interaction of the two (which is, if multiplied by zero, our treatment intensity variable in the main specificaton). Across the columns, the table reports these effects separately for the full analysis sample (Column (1)), as well as for the firms with 2013 log markdowns below zero (Column (2), plausibly monopsonistic) and above (Column (3), the residual category). Standard errors are clustered on the establishment level.

Table A.5: Regression results: DiD regression of log establishment mean wage (using 2013 payroll share exposed as treatment intensity variable)

	log markdown (2013)		
	any	≤0	>0
VARIABLES	log wage	log wage	log wage
2011 * Payroll share (2013)	-0.037	0.027	–0.045
	(0.0731)	0.046)	(0.032)
2012 * Payroll share (2013)	-0.021	-0.021	-0.021
	(0.0762)	(0.039)	(0.025)
2014 * Payroll share (2013)	-0.015	-0.010	-0.022
	(0.0646)	(0.0194)	(0.0258)
2015 * Payroll share (2013)	0.126***	0.127***	0.136***
	(0.0593)	(0.033)	(0.039)
2016 * Payroll share (2013)	0.134***	0.129***	0.152***
	(0.0549)	(0.033)	(0.040)
2017 * Payroll share (2013)	0.155***	0.136***	0.202***
	(0.0568)	(0.034)	(0.047)
2018 * Payroll share (2013)	0.144***	0.120***	0.199***
	(0.0555)	(0.031)	(0.043)
Constant	-1.690***	-1.662***	-1.765***
	(0.002)	(0.003)	(0.003)
Observations	7,232	5,232	2,000
R^2	0.955	0.960	0.941

Notes: The table reports results of the DiD establishment-level regression (the specification in Equation (31)) but does so for log wages (log mean wage) regressed on the 2013 payroll share exposed as the treatment intensity variable, reporting the year-specific effects, for 2013 as the baseline / omitted year. Across the columns, the table reports these effects separately for the full analysis sample (Column (1)), as well as for the firms with 2013 log markdowns below zero (Column (2), plausibly monopsonistic) and above (Column (3), the residual category). Standard errors are clustered on the establishment level.

Table A.6: Regression results: DiD regression of log employment (using 2013 payroll share exposed as treatment intensity variable)

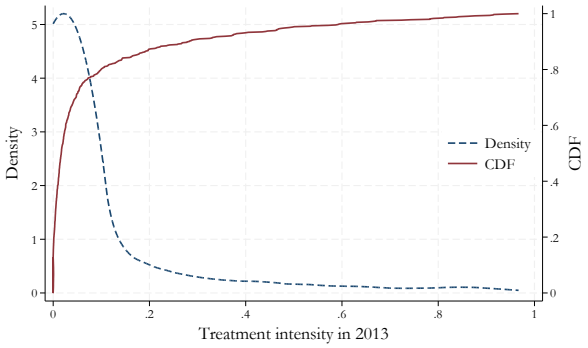
	log markdown (2013)		
	any	≤0	>0
VARIABLES	log employment	log employment	log employment
2011 * Payroll share (2013)	0.051*	0.045	0.065
	(0.027)	(0.041)	(0.036)
2012 * Payroll share (2013)	-0.018	-0.004	-0.043
	(0.0213)	(0.029)	(0.030)
2014 * payroll share (2013)	0.010	-0.001	0.037
	(0.024)	(0.029)	(0.038)
2015 * Payroll share (2013)	-0.0459*	-0.0143	-0.0959*
	(0.027)	(0.030)	(0.051)
2016 * Payroll share (2013)	-0.066**	-0.017	-0.143**
	(0.0309)	(0.0314)	(0.0635)
2017 * Payroll share (2013)	-0.089**	-0.046	-0.151**
	(0.036)	(0.040)	(0.068)
2018 * Payroll share (2013)	-0.116***	-0.048	-0.229**
	(0.044)	(0.045)	(0.092)
Constant	3.518***	3.578***	3.359***
	(0.002)	(0.003)	(0.005)
Observations	7,232	5,232	2,000
R^2	0.989	0.991	0.982

Notes: The table reports results of the DiD establishment-level regression (the specification in Equation (31)) but does so for log employment regressed on the 2013 payroll share exposed as the treatment intensity variable, reporting the year-specific effects, for 2013 as the baseline / omitted year. Across the columns, the table reports these effects separately for the full analysis sample (Column (1)), as well as for the firms with 2013 log markdowns below zero (Column (2), plausibly monopsonistic) and above (Column (3), the residual category). Standard errors are clustered on the establishment level.

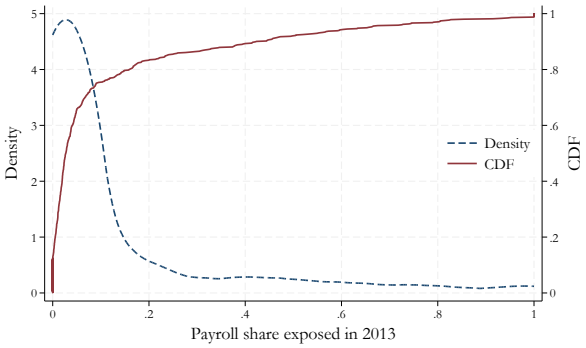
Additional figures

Figure A.1: Descriptives: treatment intensity, share of payroll exposed to minimum wage, and markdowns in 2013. Sample of firms with log markdown (2013) ≤ 0

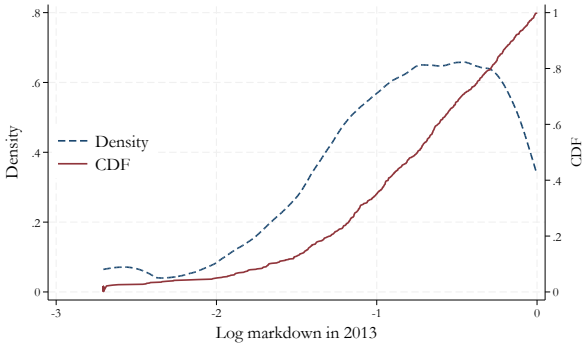
(a) Treatment intensity



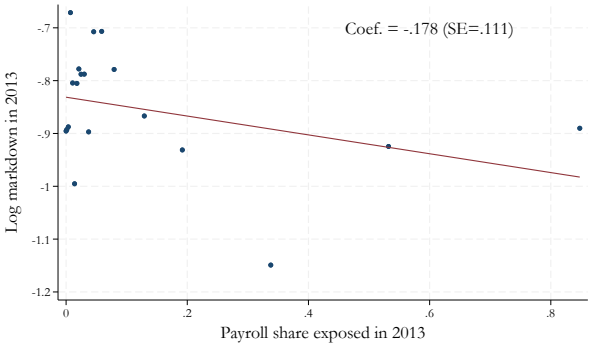
(b) Payroll share exposed



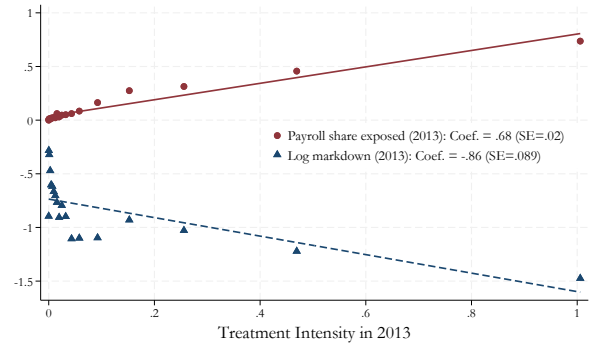
(c) Log markdown



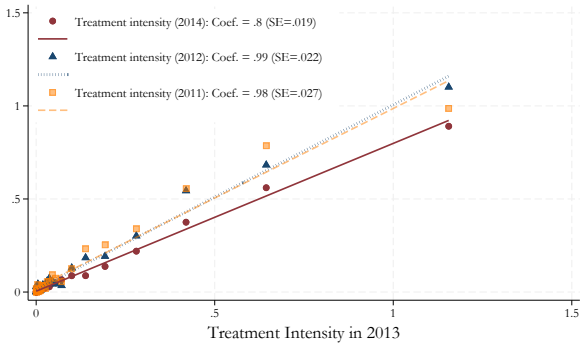
(d) Payroll share exposed and log markdown



(e) Log markdown and payroll share exposed, and treatment intensity



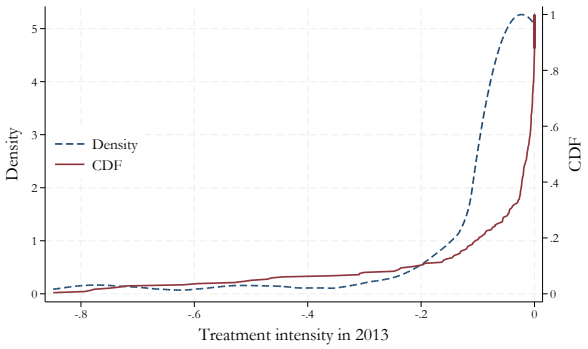
(f) Persistence of treatment intensity



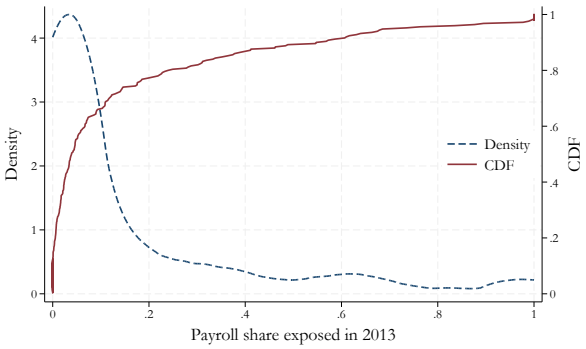
Notes: The figure replicates Figure 3 for the subsample of firms with negative 2013 log markdowns.

Figure A.2: Descriptives: treatment intensity, share of payroll exposed to minimum wage, and markdowns in 2013. Sample of firms with log markdown (2013) > 0

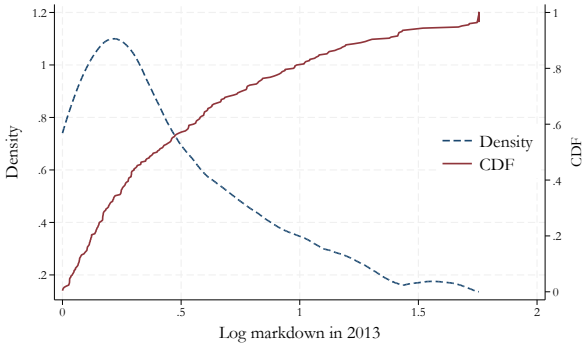
(a) Treatment intensity



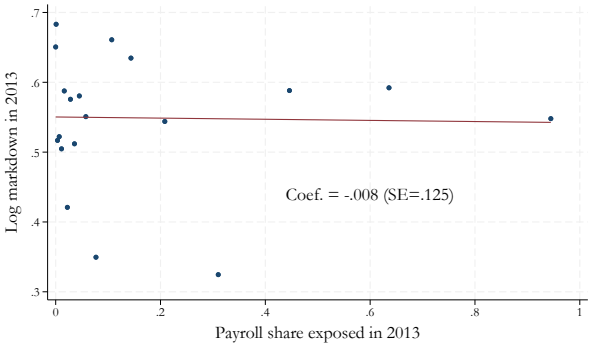
(b) Payroll share exposed



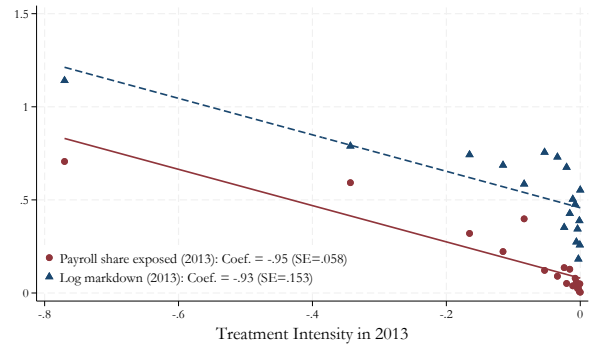
(c) Log markdown



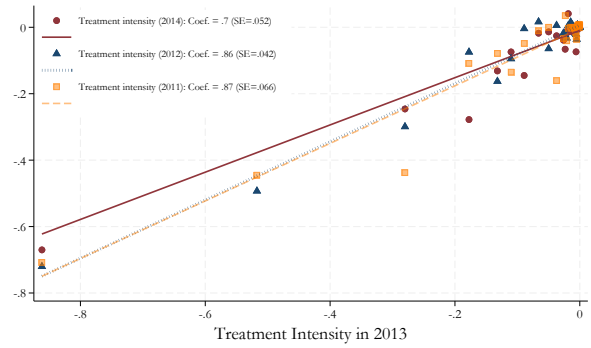
(d) Payroll share exposed and log markdown



(e) Log markdown and payroll share exposed, and treatment intensity



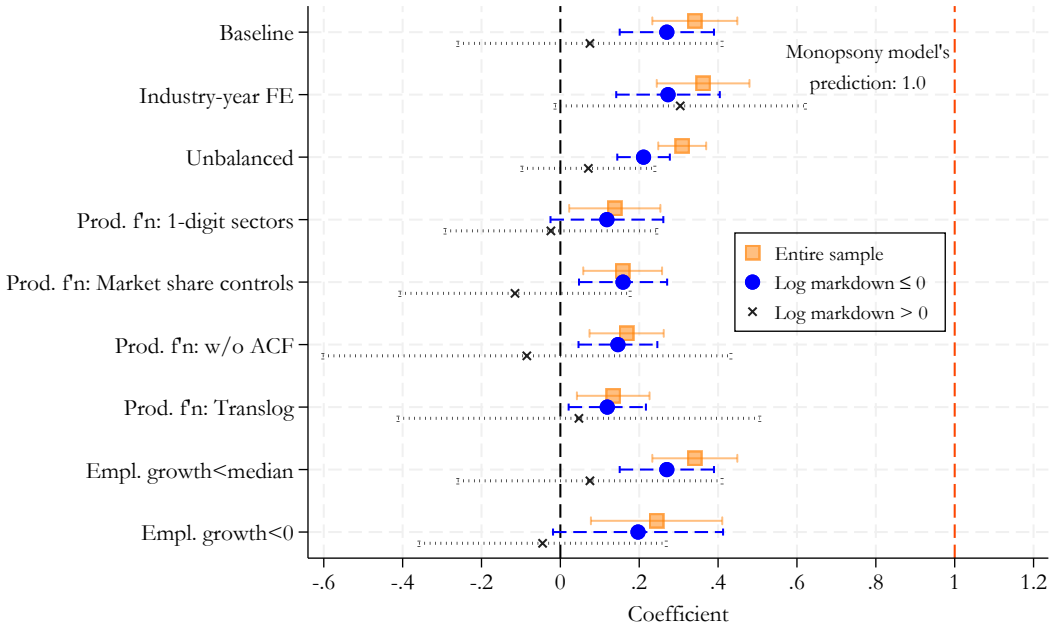
(f) Persistence of treatment intensity



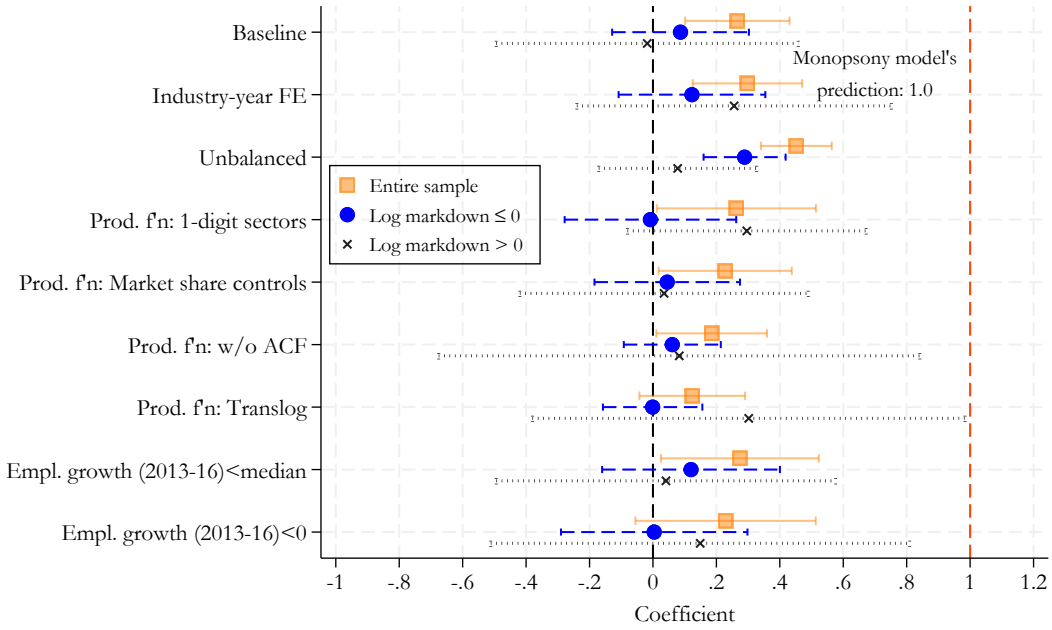
Notes: The figure replicates Figure 3 for the subsample of firms with positive 2013 log markdowns.

Figure A.3: Overview of treatment effects across specifications (pooled **post-**period compared to base year **2013**)

(a) Baseline specification



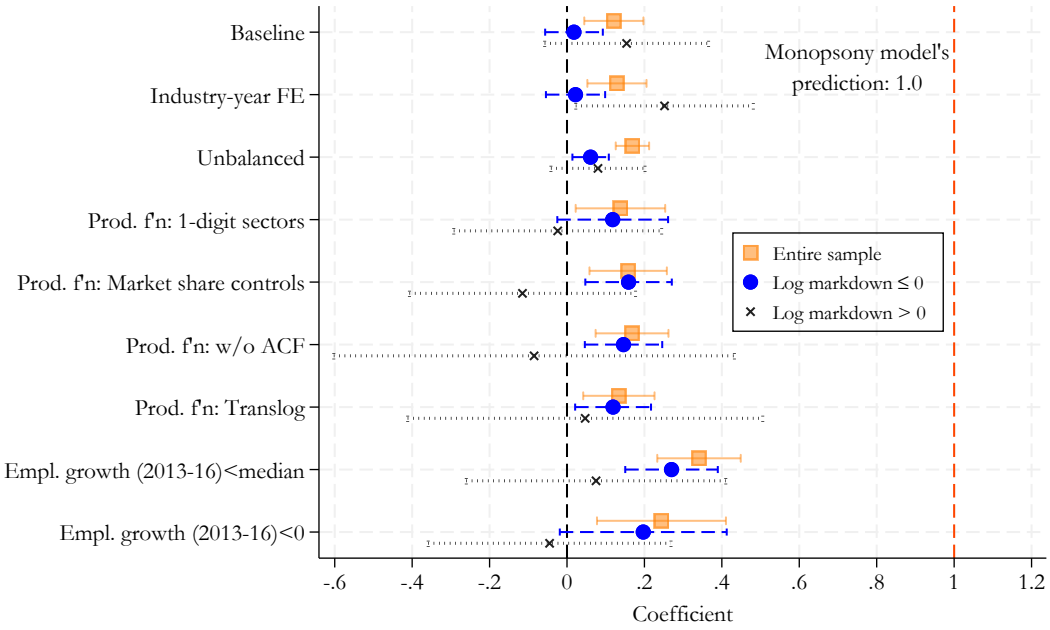
(b) Pre-trend correction



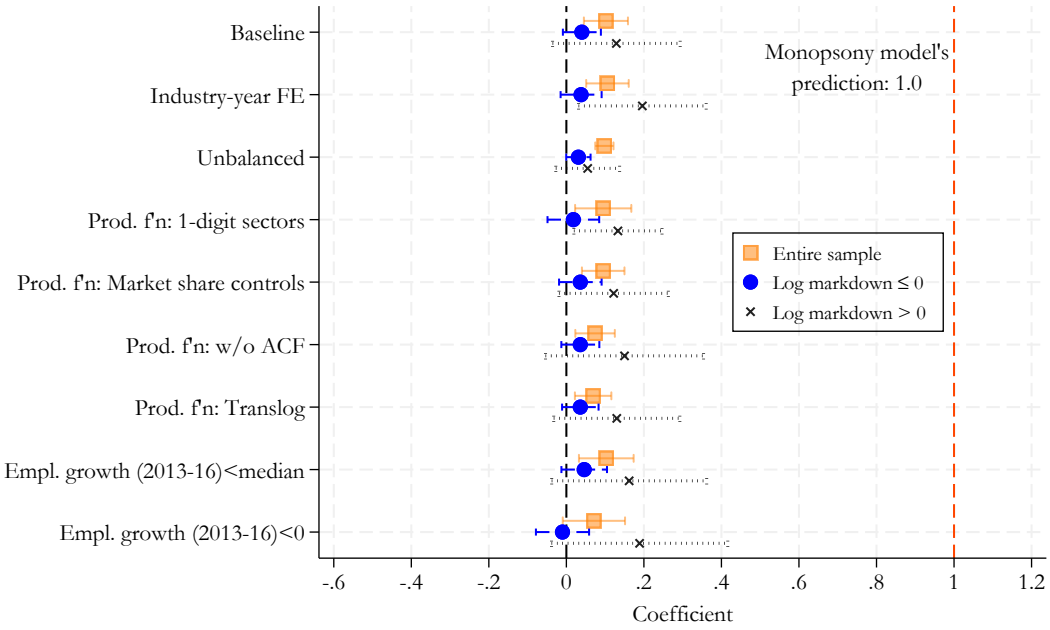
Notes: The figure is organized as Figure 4 but reports the pooled post-period treatment effects compared against 2013 rather than the full pre-period. This specification also permits a pre-period check by estimating the pre-period effects against the 2013 base period, reported in Appendix Figure A.4.

Figure A.4: Overview of pre-period effects across specifications (pooled **pre-**period compared to base year **2013**)

(a) Baseline specification



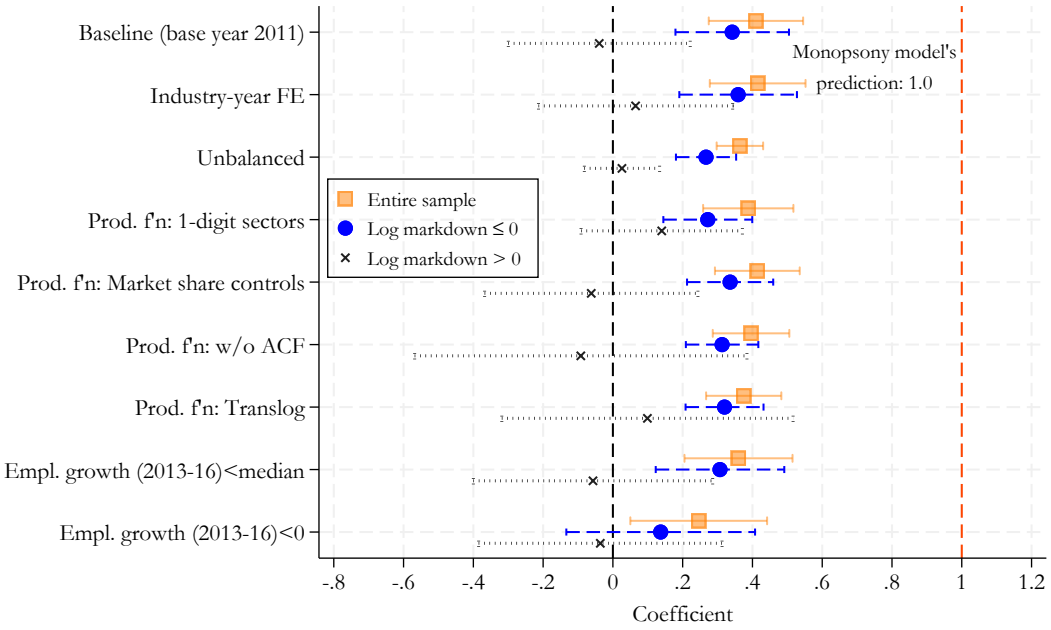
(b) Pre-trend correction



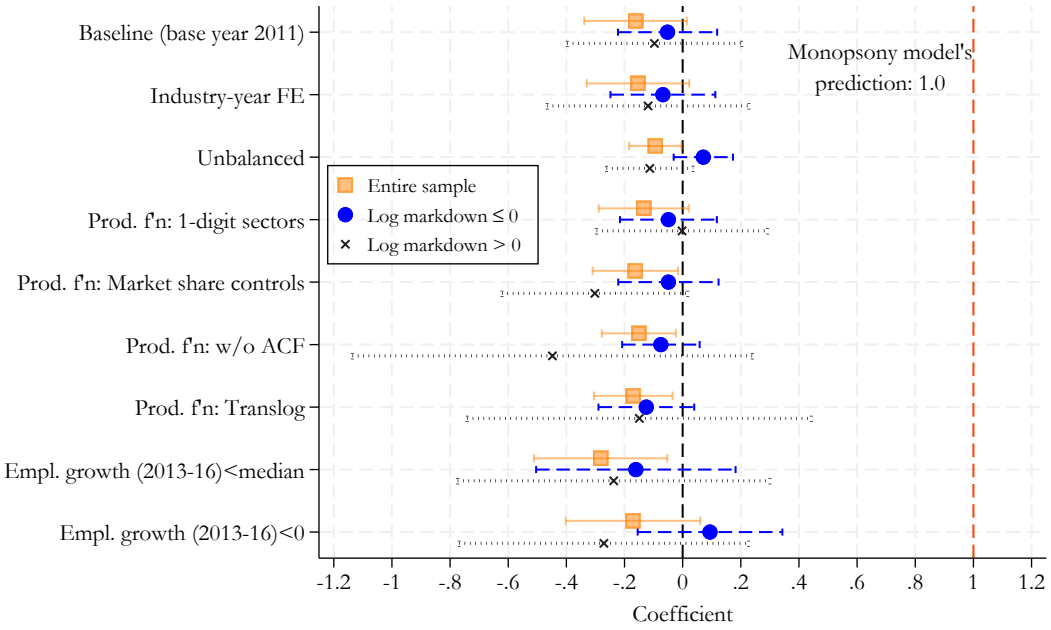
Notes: The figure is organized as Figure 4 but reports the pooled *pre*-period treatment effects compared against 2013. This estimate stems from the DiD specification in which pooled pre- and post-period effects are estimated against the 2013 base period. Appendix Figure A.3 reports on the associated post-period treatment effects.

Figure A.5: Baseline year of 2011: Overview of treatment effects across specifications (pooled **post-**period compared to base year **2011**)

(a) Baseline specification



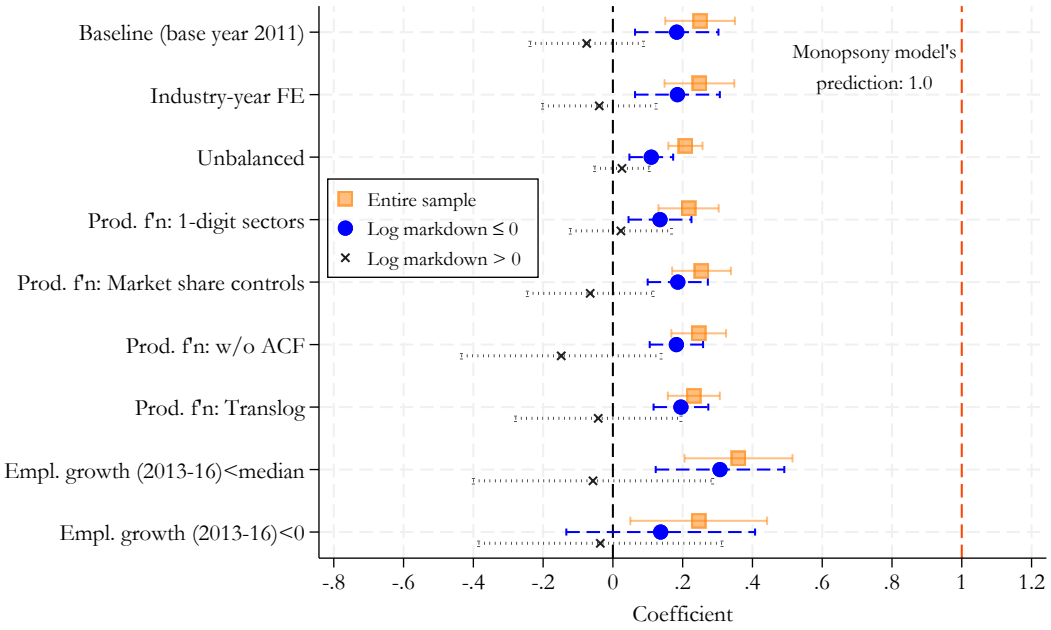
(b) Pre-trend correction



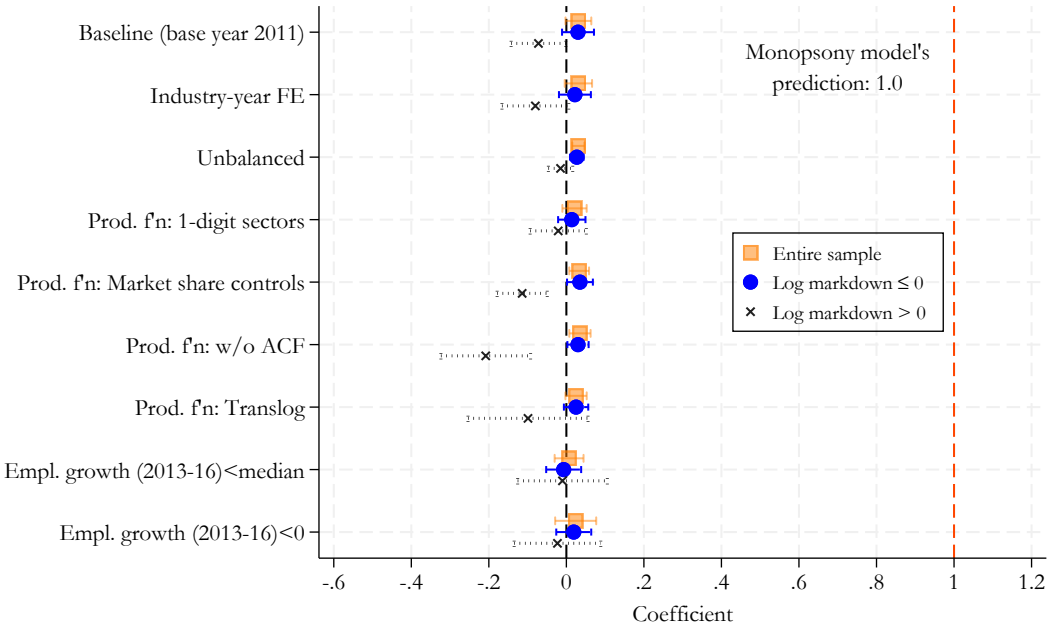
Notes: The figure is organized as Figure 4 but reports the pooled post-period treatment effects compared against 2013 rather than the full pre-period. This specification also permits a pre-period check by estimating the pre-period effects against the 2013 base period, reported in Appendix Figure A.4.

Figure A.6: Baseline year of 2011: Overview of pre-period effects across specifications (pooled **pre-**period compared to base year **2011**)

(a) Baseline specification



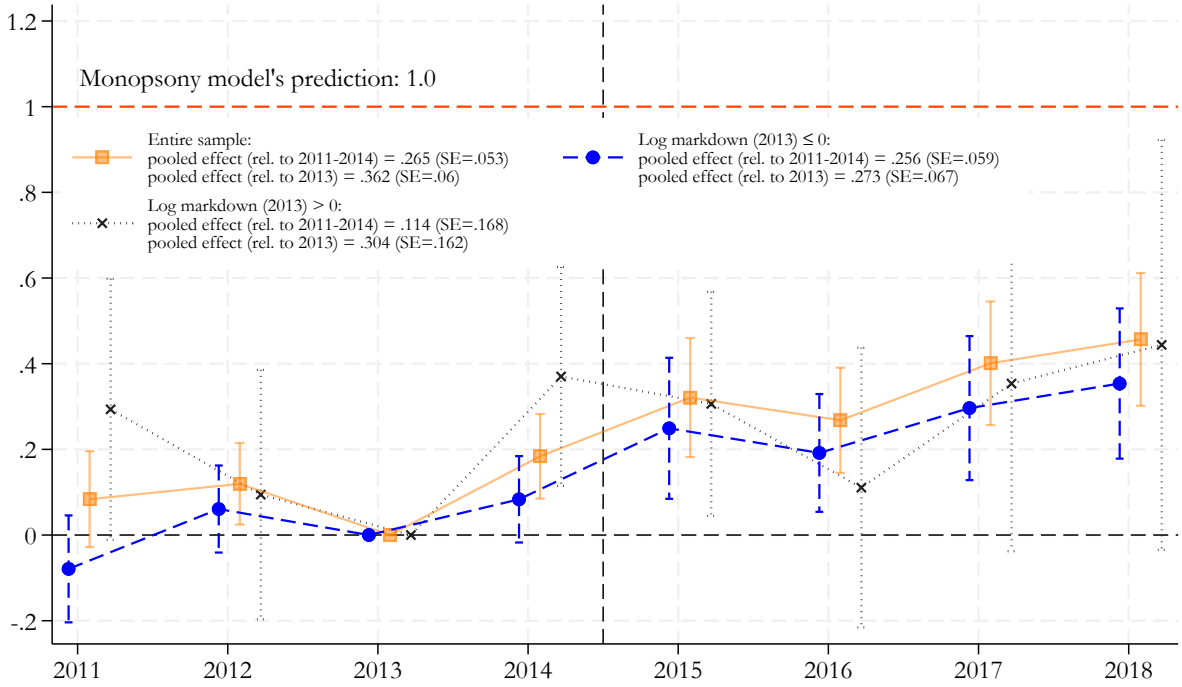
(b) Pre-trend correction



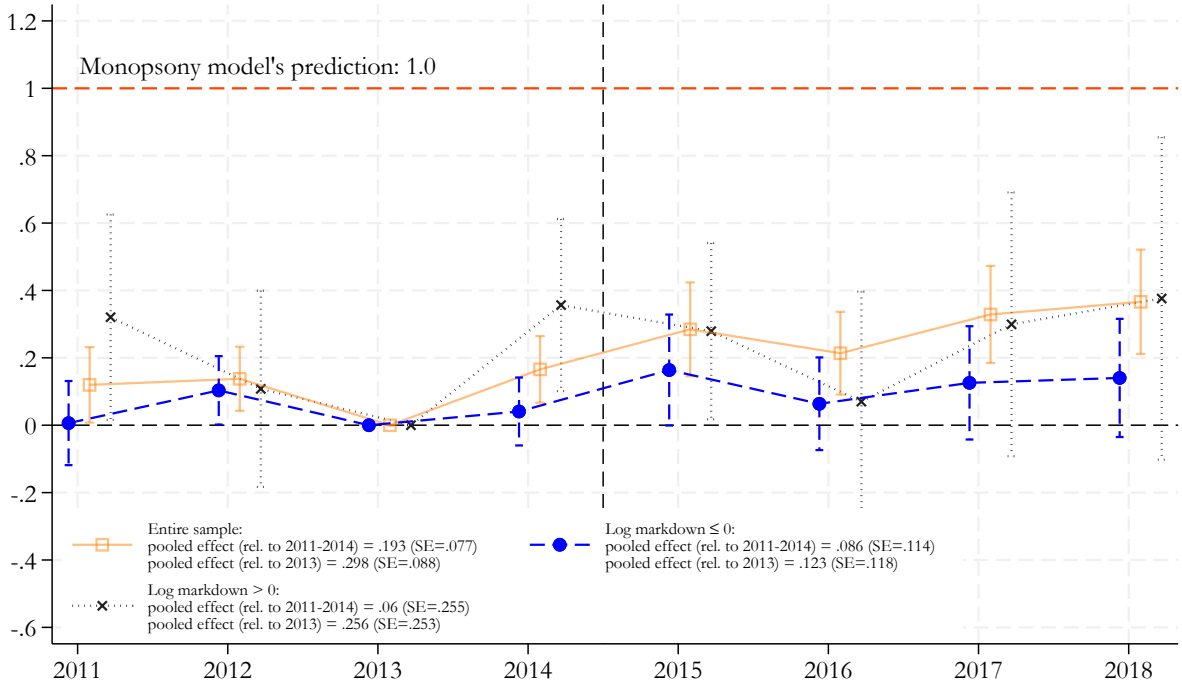
Notes: The figure is organized as Figure 4 but reports the pooled *pre*-period treatment effects compared against 2013. This estimate stems from the DiD specification in which pooled pre- and post-period effects are estimated against the 2013 base period. Appendix Figure A.3 reports on the associated post-period treatment effects.

Figure A.7: Robustness check: controlling for year-industry fixed effects

(a) Baseline specification



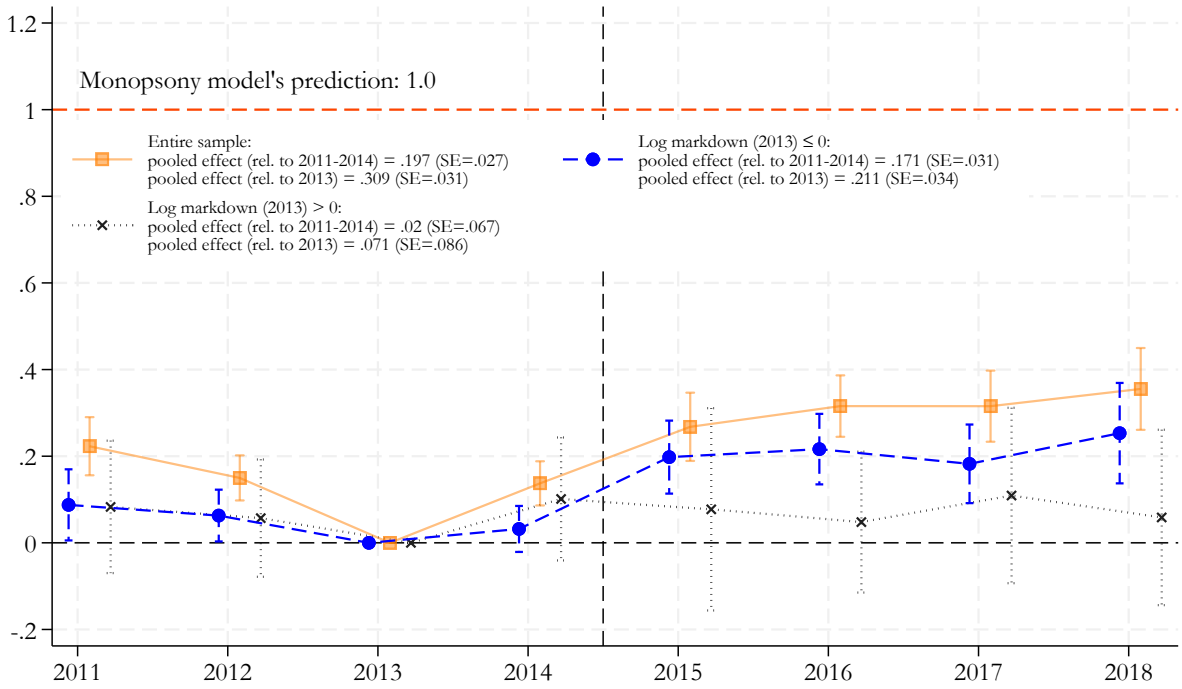
(b) Pre-trend correction



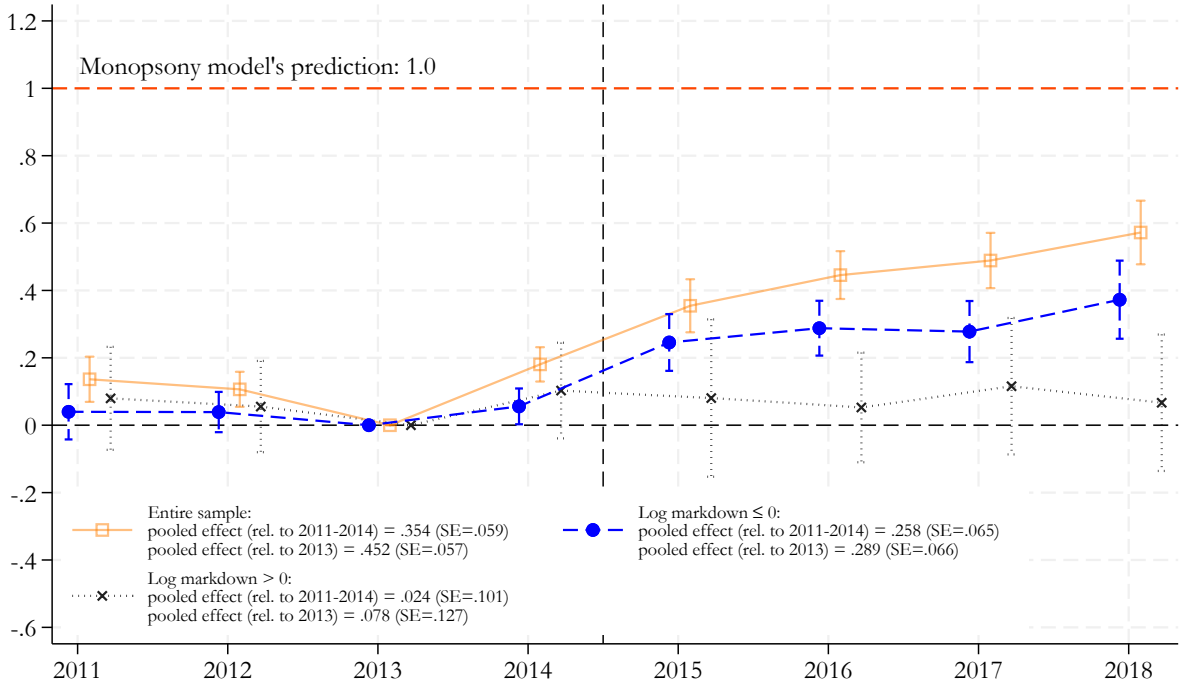
Notes: The figure replicates Figure 3 but reports on the DiD specification with industry-year FEs.

Figure A.8: Robustness check: unbalanced sample

(a) Baseline specification



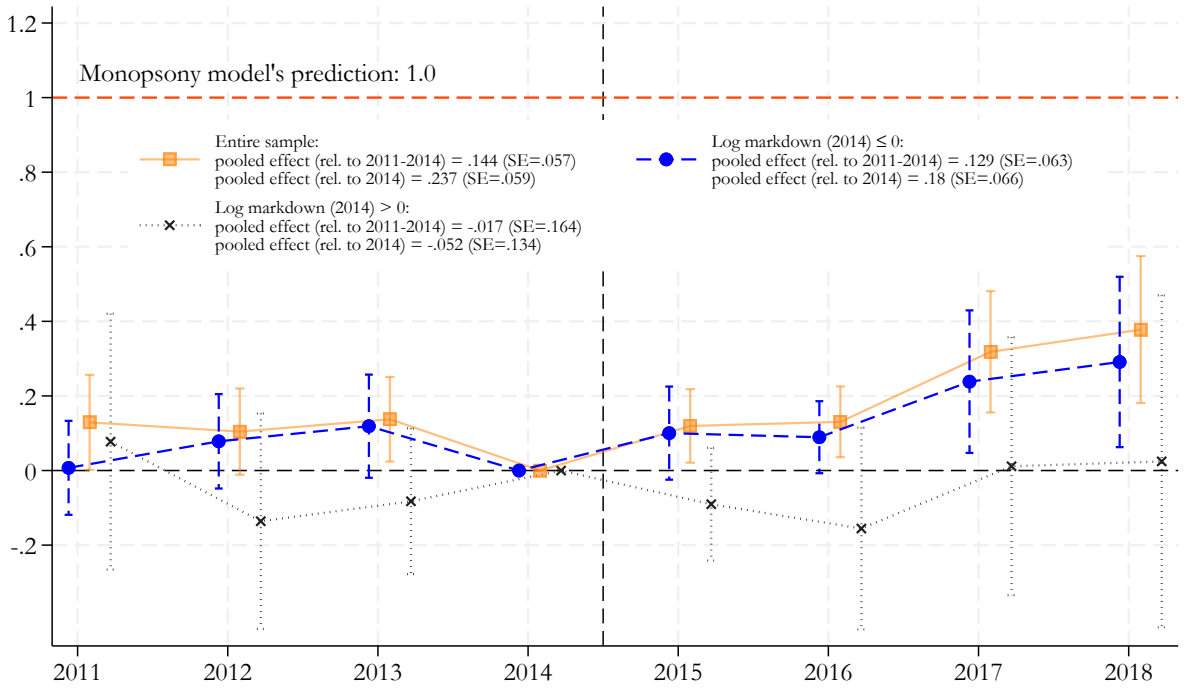
(b) Pre-trend correction



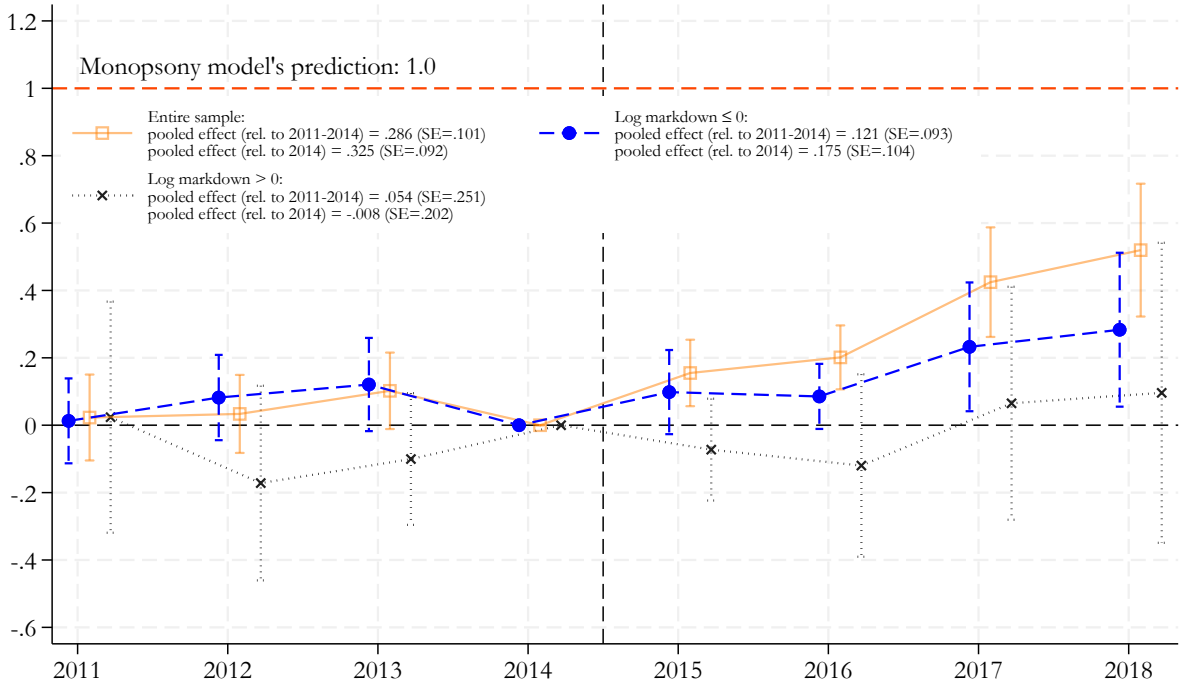
Notes: The figure replicates Figure 3 but reports on the DiD specification with an unbalanced firm panel.

Figure A.9: Robustness check: 2014 base year

(a) Baseline specification



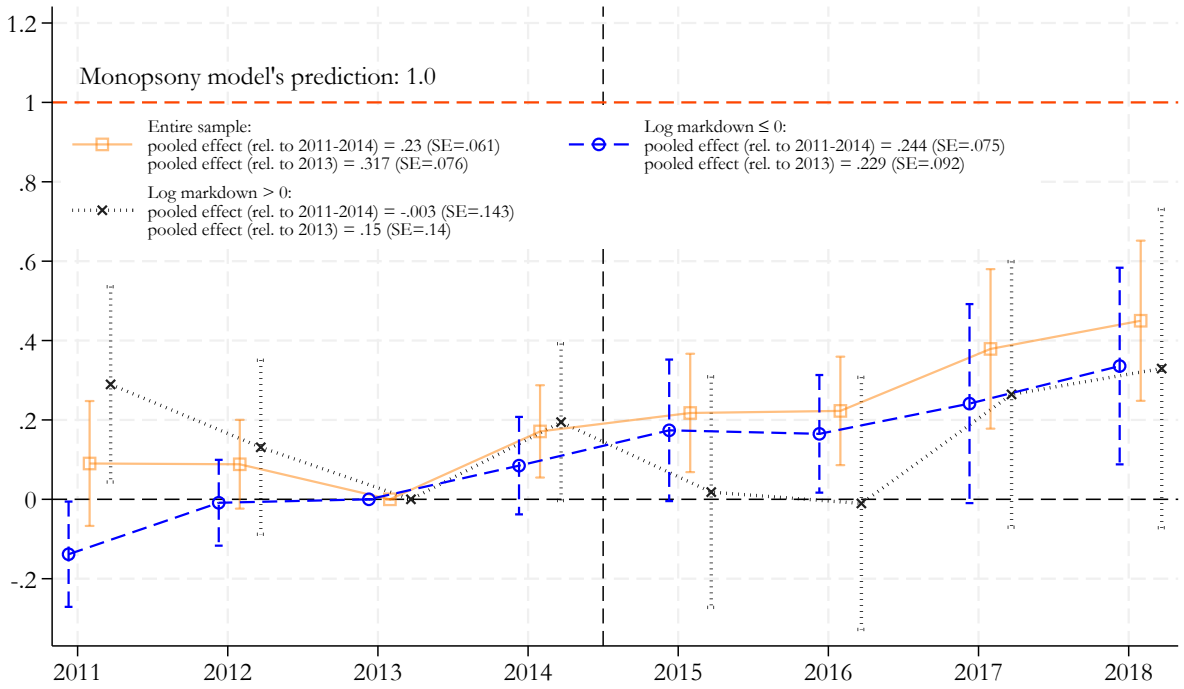
(b) Pre-trend correction



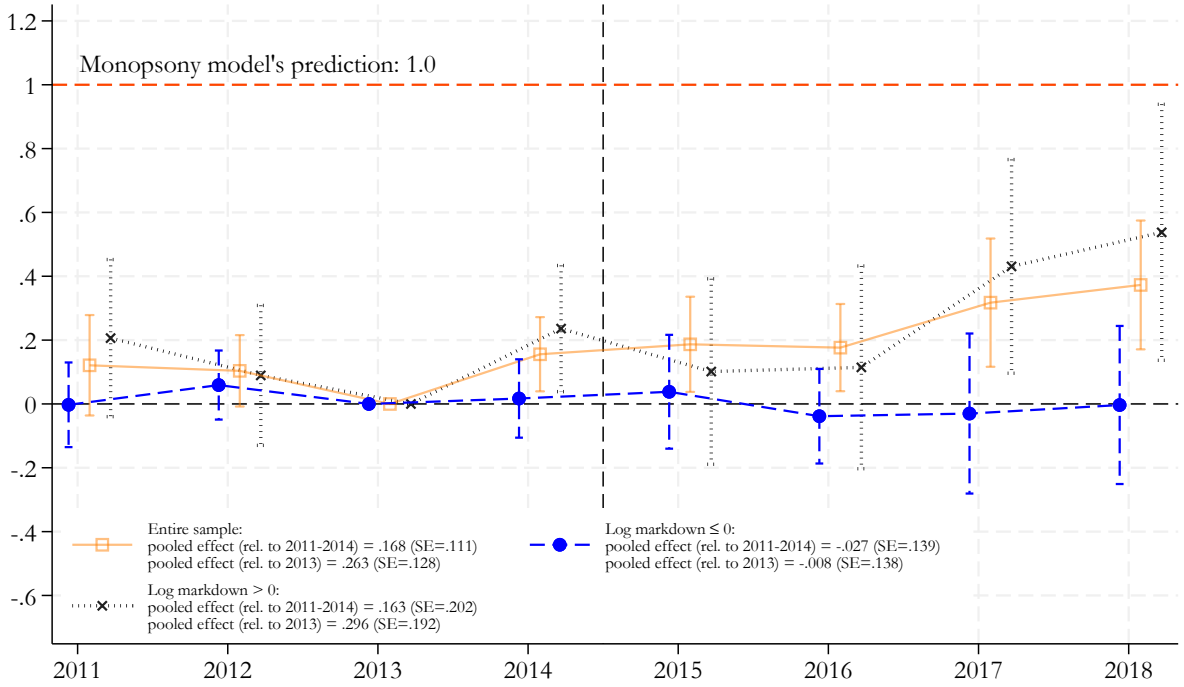
Notes: The figure replicates Figure 3 but reports on the DiD specification using 2014 as the base year, i.e., to compute treatment intensity and as the omitted year.

Figure A.10: Robustness check: production function estimation at 1-digit industry level

(a) Baseline specification



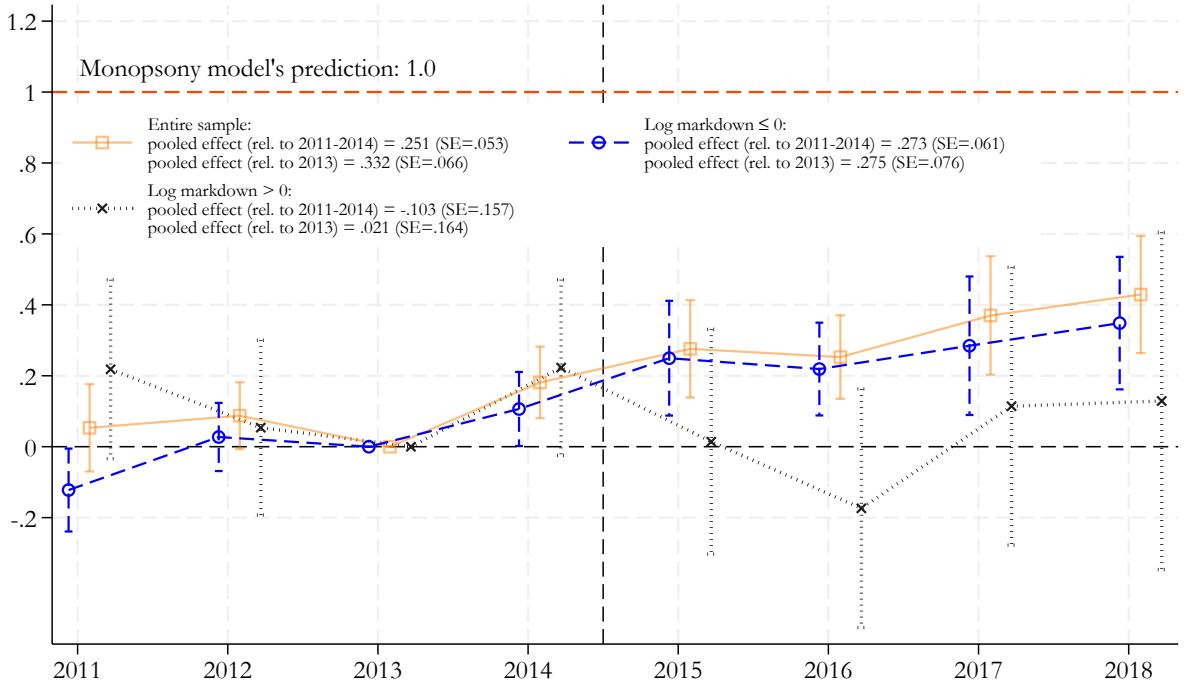
(b) Pre-trend correction



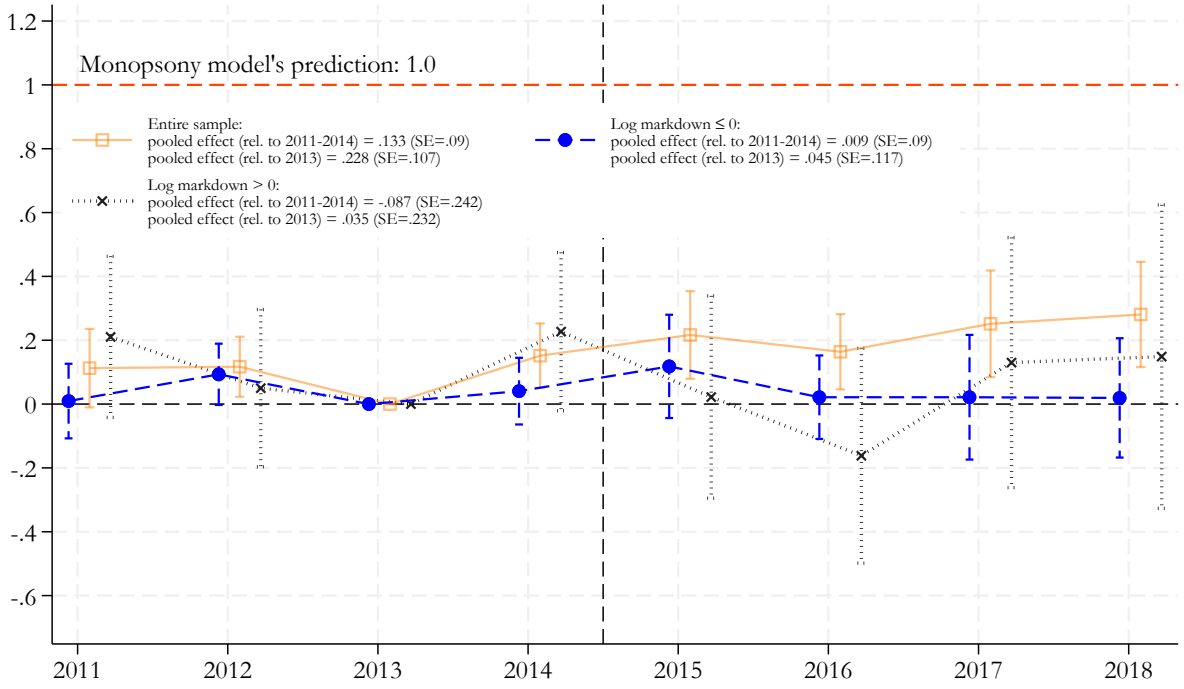
Notes: The figure replicates Figure 3 but reports on the DiD specification using markdown estimates stemming from an alternatively production function estimation method: at the 1-digit industry level rather than 2-digit.

Figure A.11: Robustness check: production function estimation with firm market shares of intermediate inputs as a control

(a) Baseline specification



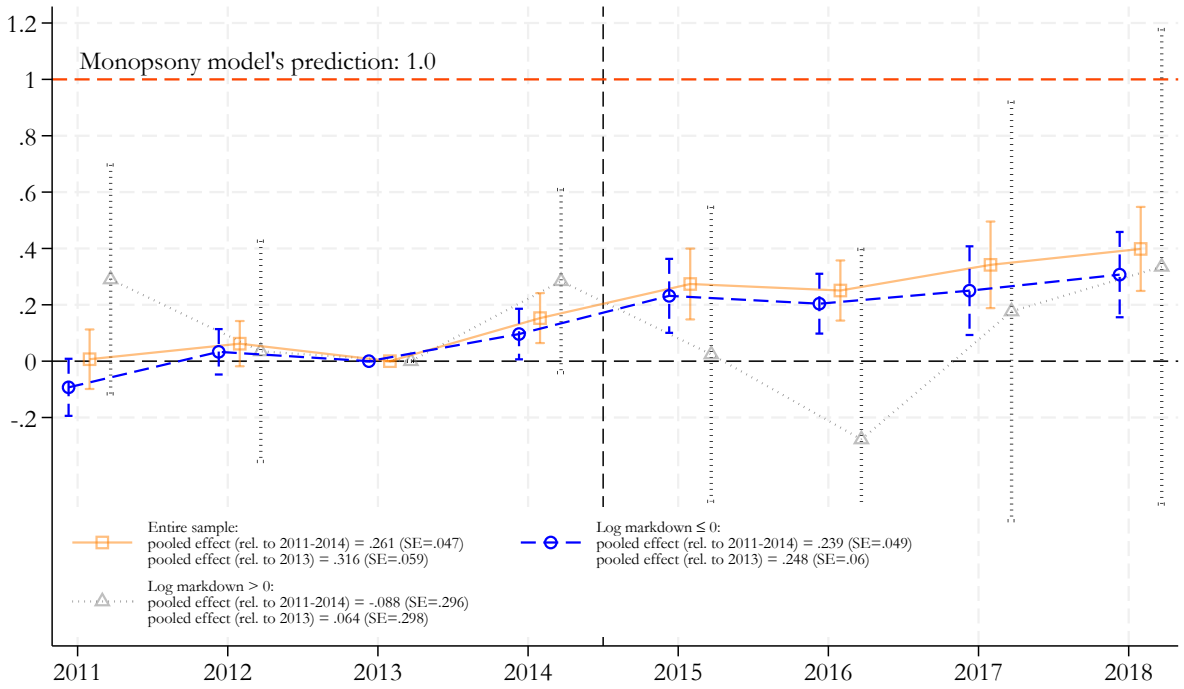
(b) Pre-trend correction



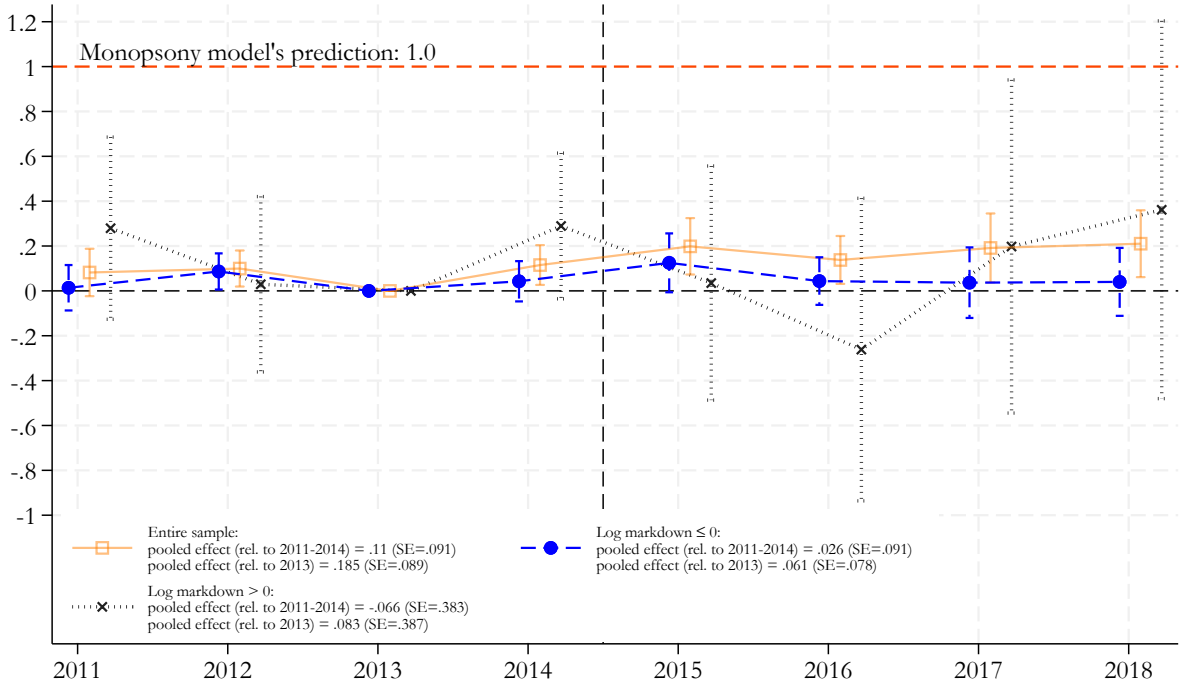
Notes: The figure replicates Figure 3 but reports on the DiD specification using markdown estimates stemming from an alternatively production function estimation method: using mark market shares of intermediate inputs as a control.

Figure A.12: Robustness check: production function estimation dropping ACF correction

(a) Baseline specification



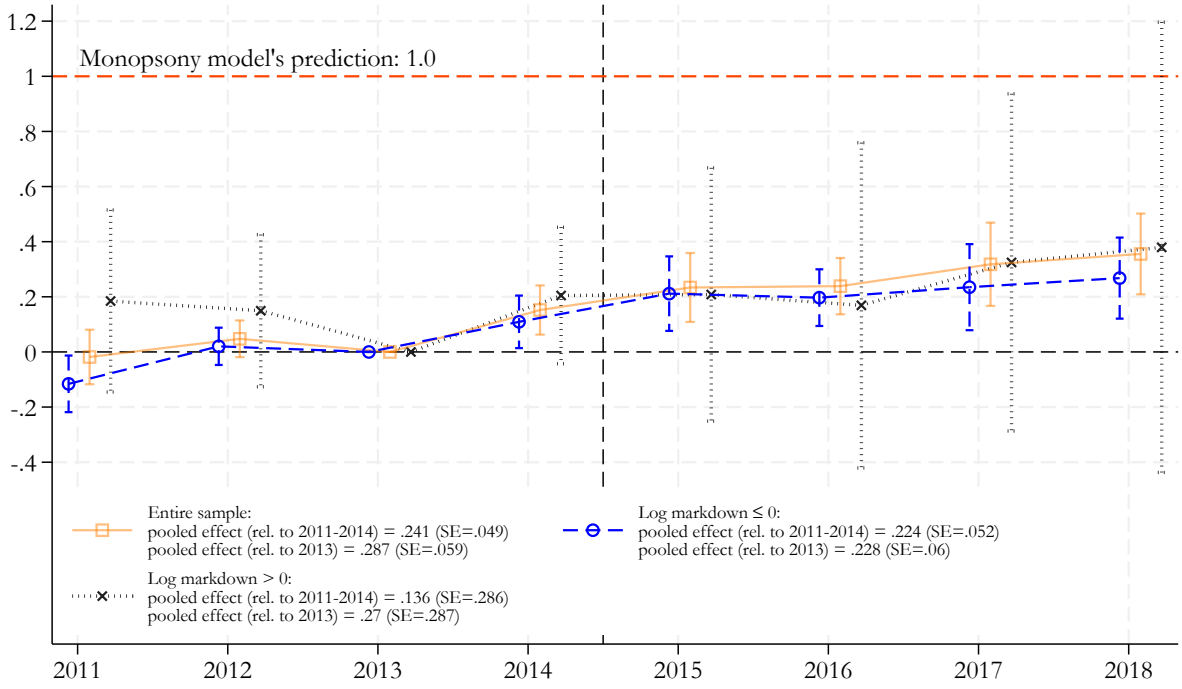
(b) Pre-trend correction



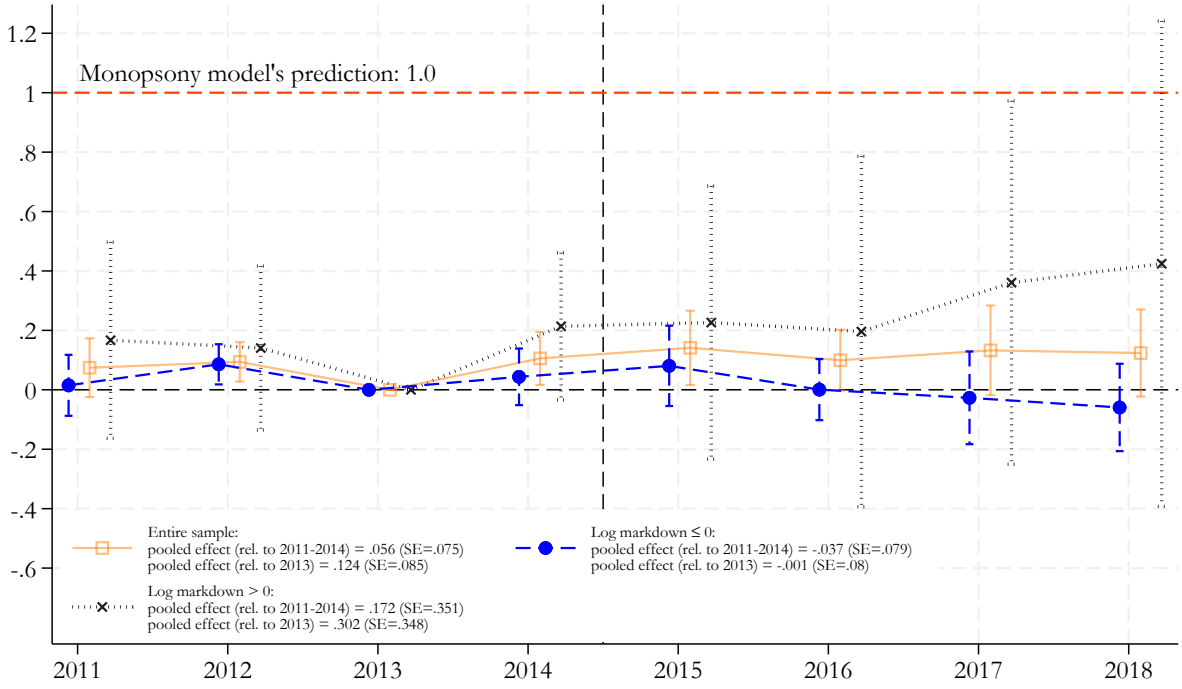
Notes: The figure replicates Figure 3 but reports on the DiD specification using markdown estimates stemming from an alternatively production function estimation method: dropping the ACF (Ackerberg, Caves and Frazer, 2015) correction.

Figure A.13: Robustness check: production function estimation using translog production function, which is more flexible than Cobb Douglas regarding substitution patterns

(a) Baseline specification



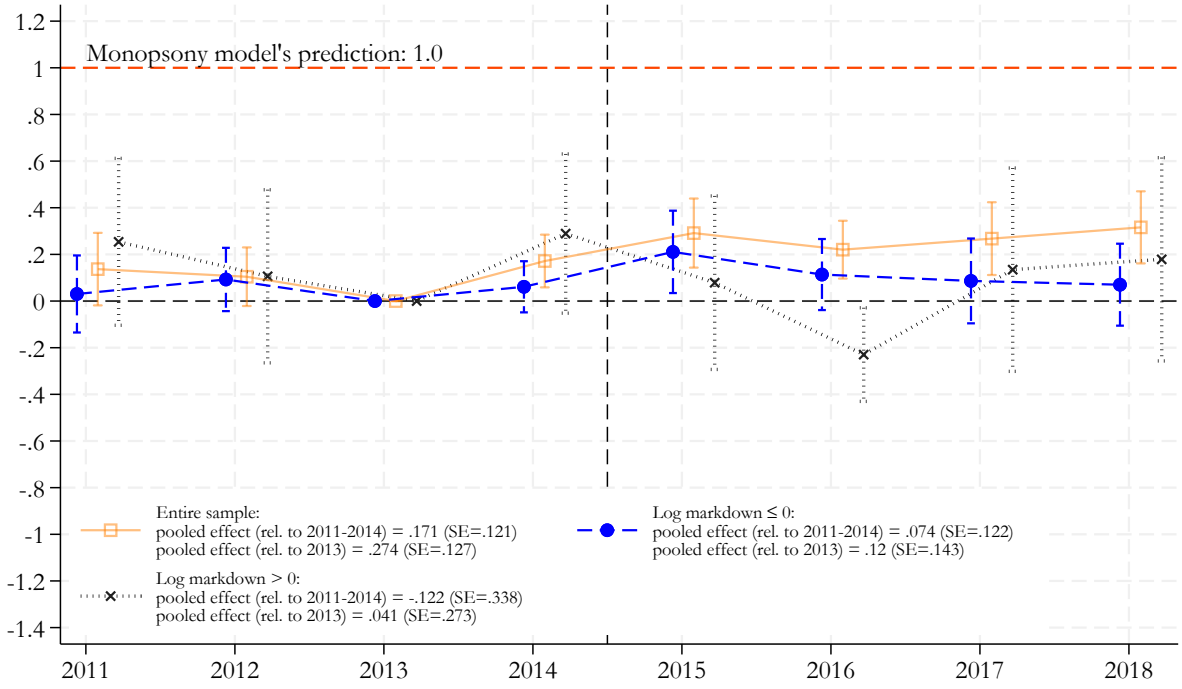
(b) Pre-trend correction



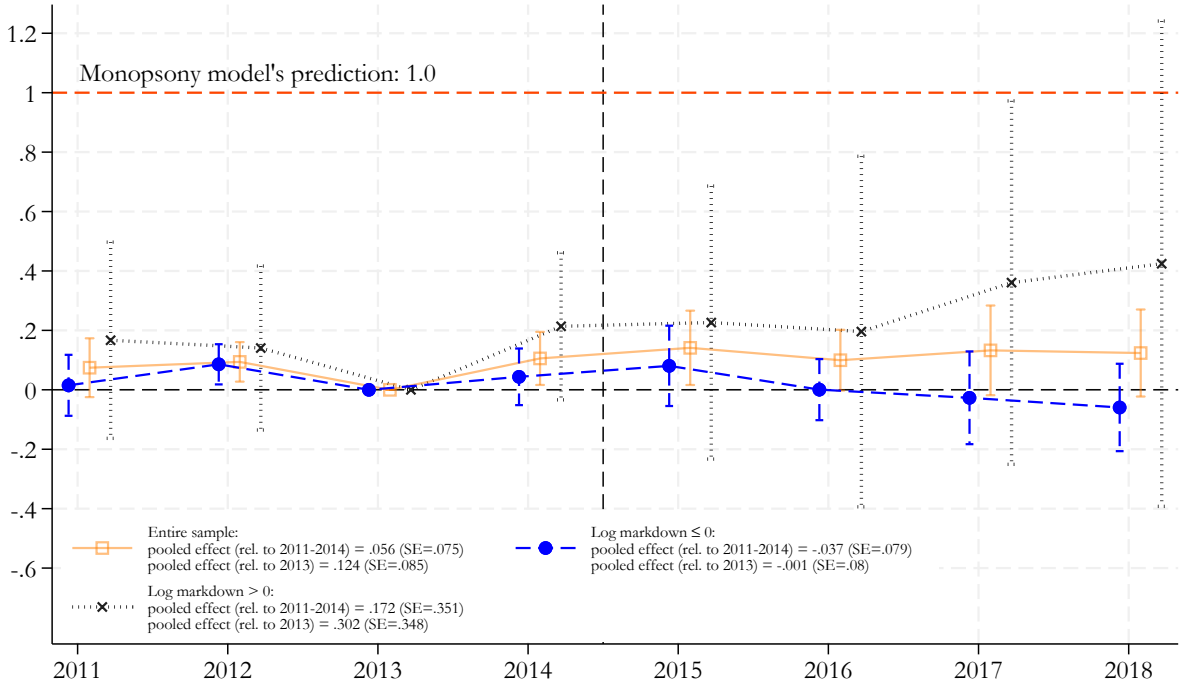
Notes: The figure replicates Figure 3 but reports on the DiD specification using markdown estimates stemming from an alternatively production function estimation method: using a translog production function.

Figure A.14: Robustness check: zooming into firms with **below-median employment growth** from 2013 to 2016, which are plausibly less likely to be rationed by labor supply — **2013** base year

(a) Baseline specification



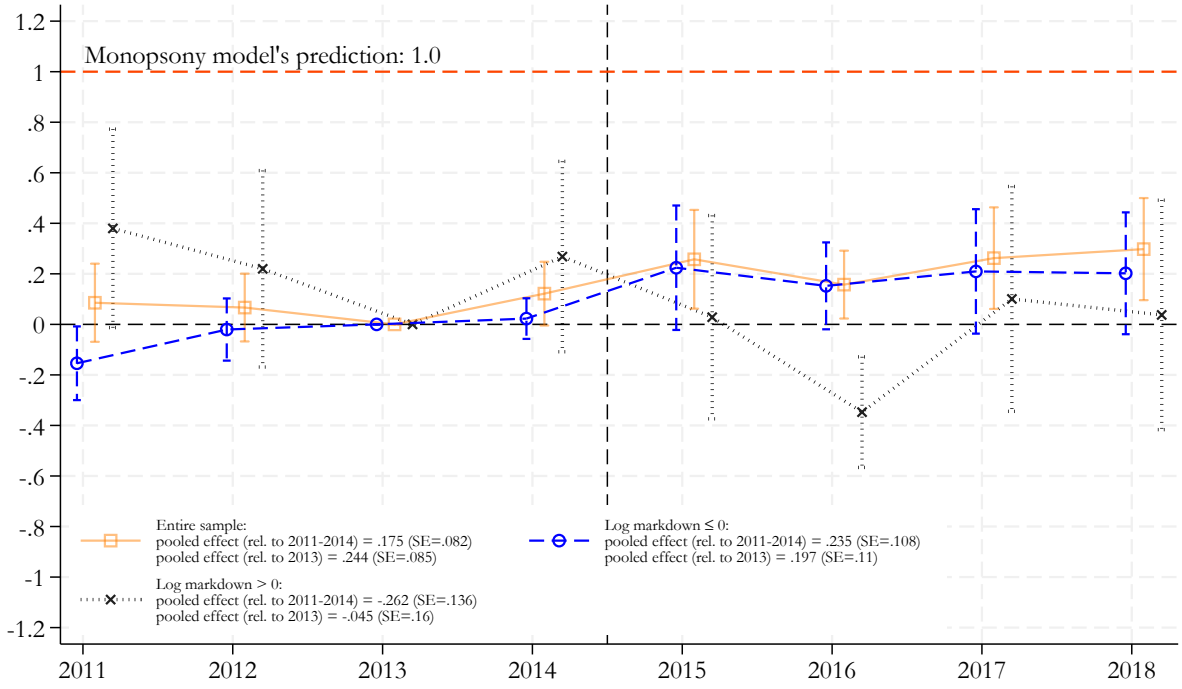
(b) Pre-trend correction



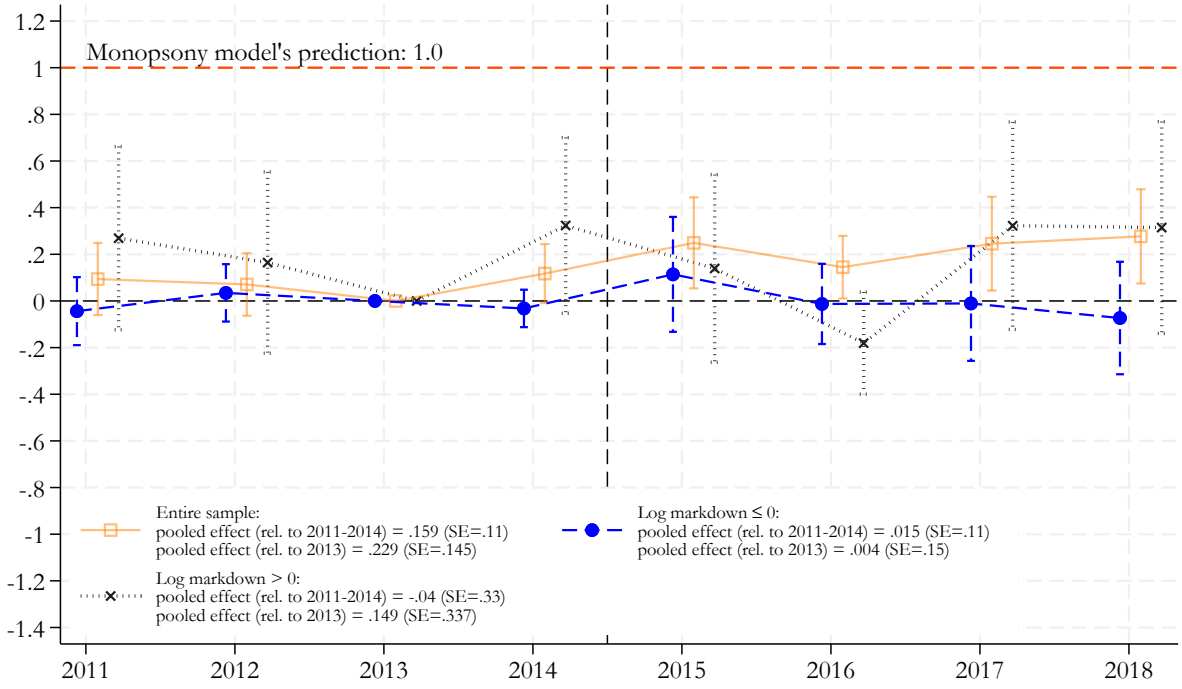
Notes: The figure replicates Figure 3 but reports on the DiD specification drawing on the sample of firms with below-median employment growth from 2013 to 2016.

Figure A.15: Robustness check: zooming into firms with **negative employment growth** from 2013 to 2016, which are plausibly less likely to be rationed by labor supply — **2013** base year

(a) Baseline specification



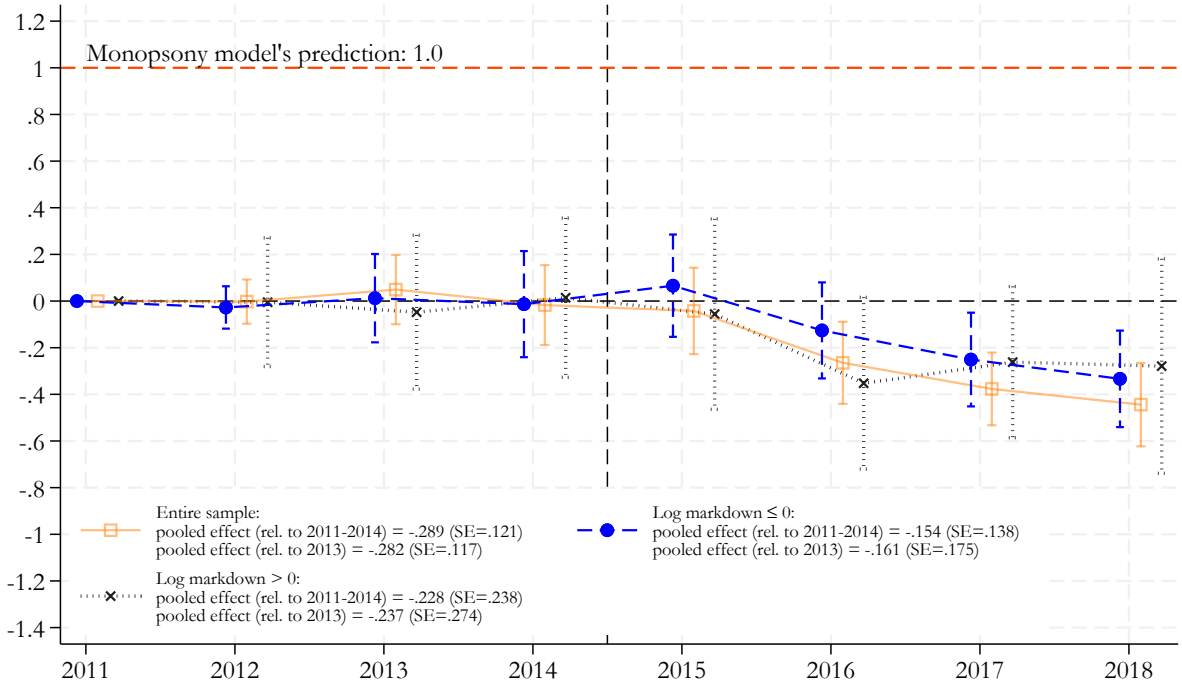
(b) Pre-trend correction



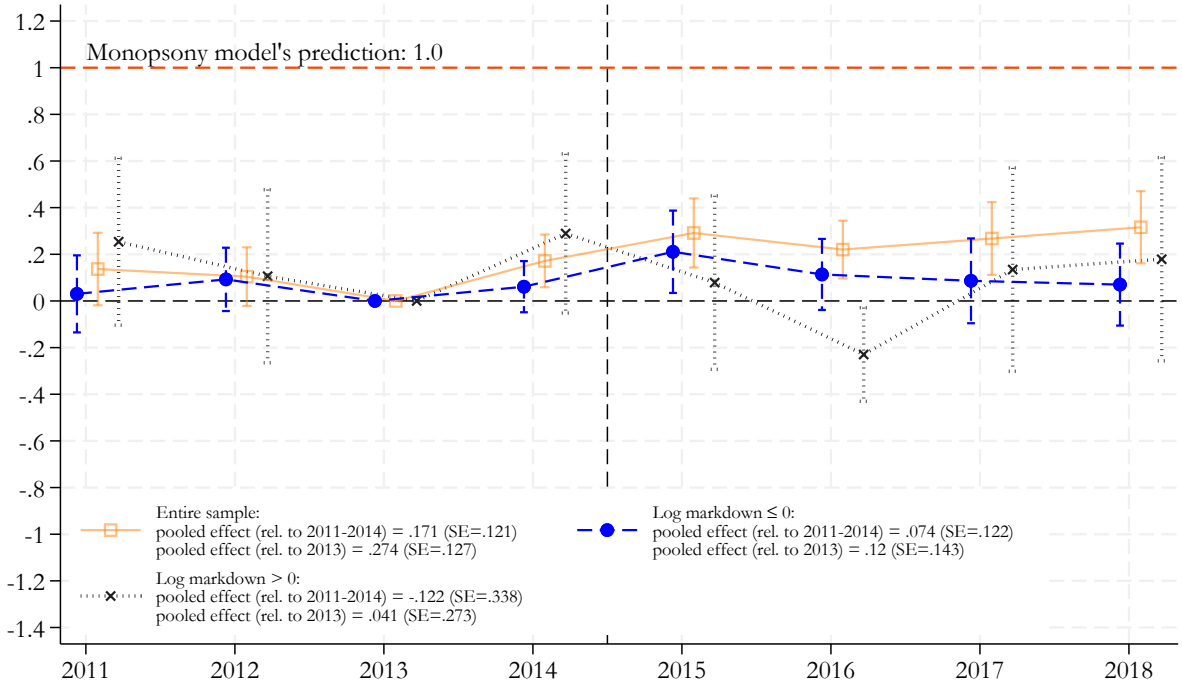
Notes: The figure replicates Figure 3 but reports on the DiD specification drawing on the sample of firms with below-median employment growth from 2013 to 2016.

Figure A.16: Robustness check: zooming into firms with **below-median employment growth** from 2013 to 2016, which are plausibly less likely to be rationed by labor supply — **2011** base year

(a) Baseline specification



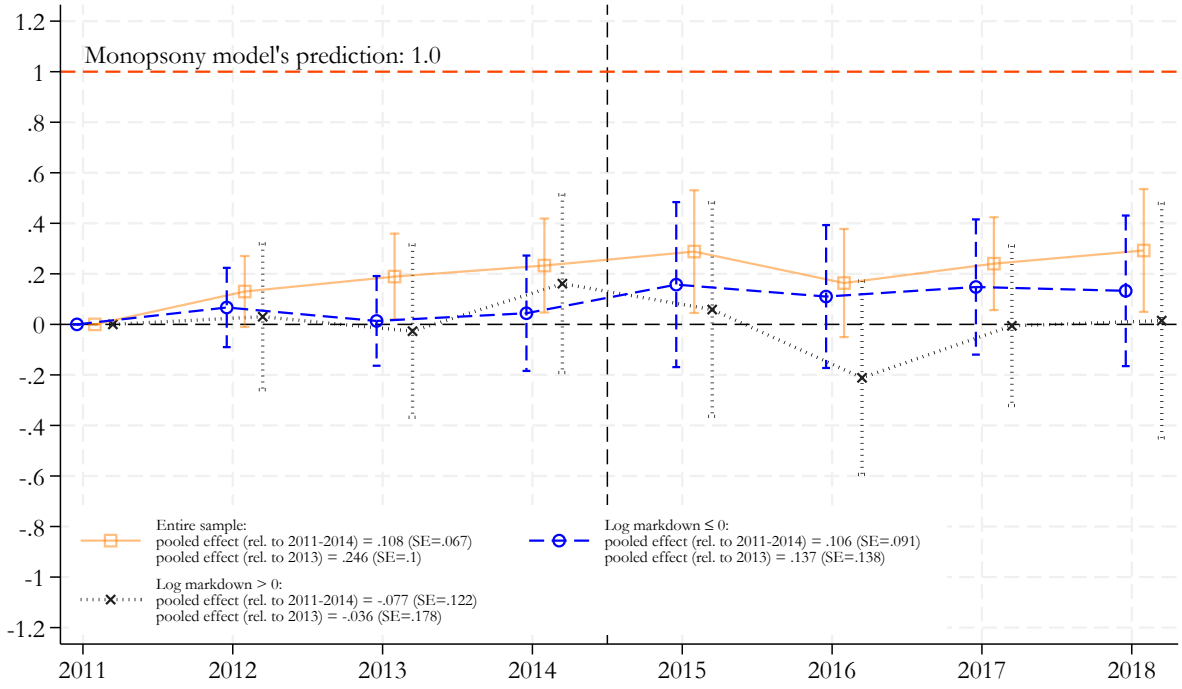
(b) Pre-trend correction



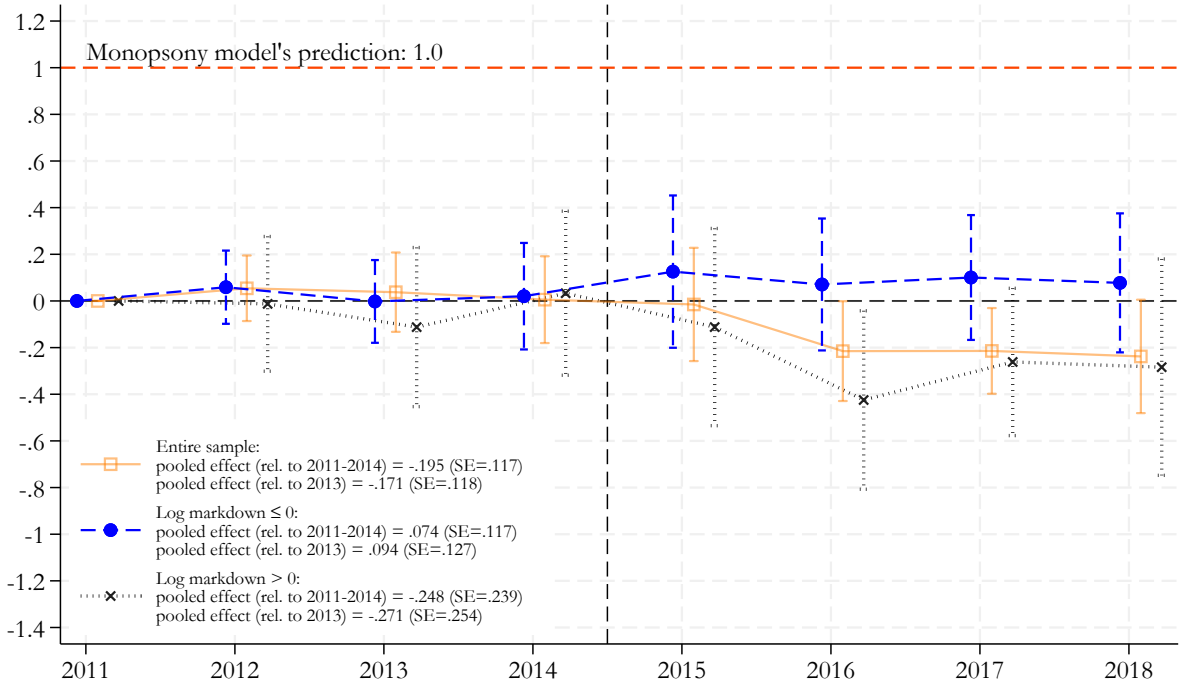
Notes: The figure replicates Figure 3 but reports on the DiD specification drawing on the sample of firms with below-median employment growth from 2013 to 2016, and with 2011 as the baseline year rather than 2013.

Figure A.17: Robustness check: zooming into firms with **negative employment growth** from 2013 to 2016, which are plausibly less likely to be rationed by labor supply — **2011** base year

(a) Baseline specification



(b) Pre-trend correction



Notes: The figure replicates Figure 3 but reports on the DiD specification drawing on the sample of firms with below-median employment growth from 2013 to 2016, and with 2011 as the baseline year rather than 2013.