Does the Minimum Wage Reduce Inequality? Current vs. Lifetime Perspectives

FRANCISCO LIBANO-MONTEIRO*

London School of Economics

Click here for the latest version

Abstract

An increase in the minimum wage can strongly compress the cross-sectional distribution of current earnings, but do *individual* gains to low-wage workers persist over time? Or do they quickly dissipate, or perhaps even come at the expense of lower future earnings? Using Portuguese administrative data covering all private-sector employment, I study a reform that raised the real national minimum wage by 16.2% between 2006 and 2010, an increase sustained through 2014. Exploiting heterogeneous exposure to the reform across the initial 2006 wage distribution in a difference-in-differences design, I find that it generated strong individual wage and earnings gains, but these dissipated relatively quickly on average. The reason is that most workers rapidly grow out of minimum-wage exposure and spend only a small fraction of their working lives near the floor. The data show little evidence of adverse effects on employment, hours worked, or career progression. To extrapolate beyond the eight-year empirical horizon, I estimate a statistical wage process that embeds the minimum wage as a floor, allowing me to consider the impact of a permanently higher floor on lifetime earnings. While the higher minimum wage raises current wages and earnings of the lowest-paid workers by the full 16.2%, it raises lifetime earnings of workers at the bottom of the lifetime earnings distribution by 7.3%. Finally, I show that the smaller lifetime equity effects need not necessarily lower the optimal minimum wage.

Keywords: Minimum wage, labor income dynamics, lifetime inequality.
JEL codes: J31, J38, D31, E24.

*Contact: M.F.LIBANO-MONTEIRO@LSE.AC.UK. I am deeply grateful to Wouter Den Haan, Xavier Jaravel, Ricardo Reis, Alan Manning, and Vernon Henderson for their guidance and support. I also thank Marco Bellifemine, Francesco Caselli, Rui Costa, Adrien Couturier, Maarten de Ridder, Matthias Doepke, Joe Hazell, Stephan Hobler, Steve Machin, Johannes Matt, Guy Michaels, Ben Moll, Carlos Oliveira, Jane Olmstead-Rumsey, Steve Pischke, Johannes Spinnewijn, Christiane Szerman, and Silvana Tenreyro for helpful comments and discussions. I thank participants at the LSE Macroeconomics, Labour Economics, and Public Economics work-in-progress seminars, the Lisbon Micro Group, the GEE/GPEARI seminar at the Portuguese Ministry of Finance, and the SNDE conference at Católica Lisbon for helpful comments. Marina Santos-Feliciano provided excellent research assistance. I gratefully acknowledge financial support from a STICERD Research Grant and from the Fundação para a Ciência e a Tecnologia (FCT). I thank GEP and DGEEC for facilitating access to the Quadros de Pessoal data.

1 Introduction

Over the last two decades, policymakers around the world have raised minimum wages considerably, even relative to median wages (Figure A1). The case for a higher minimum wage partly rests on efficiency arguments, which have been strengthened by growing evidence that labor markets are imperfectly competitive (Manning, 2021; Card, 2022) and that moderate increases do not cause large employment losses (Dube and Lindner, 2024). However, larger increases may become distortionary, limiting how high a minimum wage can be justified on efficiency grounds alone. Above that level, a higher minimum wage can still be justified on equity grounds (Berger et al., 2025).

Evidence on the equity effects of the minimum wage shows that it raises wages and earnings at the bottom of the distribution, compressing lower-tail inequality (Lee, 1999; Cengiz et al., 2019). In settings where spillovers to workers earning above the minimum are strong, it can also reduce broader measures of inequality (Engbom and Moser, 2022). However, this evidence is mostly cross-sectional, as it measures how much the policy reduces inequality at a given point in time. Individual workers, by contrast, move through that distribution over their careers, so those at the bottom when the policy is enacted need not remain there throughout their lifetimes.

This paper takes a longitudinal perspective, tracking individual workers rather than the cross-sectional distribution, to measure the lifetime distributional effects of the minimum wage. I show that the lifetime effects are considerably smaller than cross-sectional evidence suggests. When the policy's current individual gains to low-wage workers fade over time, either because they move up the distribution or because the policy harms their future earnings prospects, the minimum wage compresses lifetime inequality less than current inequality.

I make three contributions. First, using Portuguese administrative data, I show that a permanent minimum-wage increase generates strong earnings gains that dissipate relatively quickly over an eight-year horizon because workers move up the distribution, not because the policy harms their hours, employment, or career prospects. Second, to quantify the policy's lifetime effects, I estimate a wage process with a minimum-wage floor. I find that workers with the lowest lifetime earnings gain less than half as much as those at the bottom of the current distribution. The lifetime gain nevertheless remains sizable because these workers still spend a non-trivial share of their working lives near the floor. Third, I assess the implications of my findings for the optimal minimum wage by extending the model of Berger et al.

(2025) to allow for income mobility. I show that weaker lifetime equity effects do not necessarily imply a lower optimal minimum wage. The optimal minimum wage does not change when individuals cannot smooth consumption, as the minimum wage now provides insurance against low-income periods.

Empirics. My empirical setting is Portugal, which has one of the highest minimum wages relative to the median and where the minimum wage has been shown to strongly compress the wage distribution (Oliveira, 2023; Section 2.3). I study a sizable and sustained reform that raised the real national minimum wage by 16.2% between 2006 and 2010. The floor remained substantially above its 2006 level through 2014. I leverage matched employer-employee administrative data on the universe of private-sector workers, which include high-quality *individual* (hourly) wages and rich firm and worker characteristics. My difference-in-differences design follows incumbent workers and exploits heterogeneous exposure to the reform across the initial 2006 wage distribution. Using the pre-reform period, when the real minimum wage was flat, I construct the counterfactual individual wage and earnings trajectories workers would have followed in the absence of the reform. The identifying assumption is that, absent the reform, relative wage growth across the initial wage distribution would have followed its stable pre-reform shape. The main threat to identification is that the financial and sovereign debt crises, which overlap the post-reform period, may have differentially affected workers across the initial 2006 wage distribution. I argue, however, that this threat is of minor concern when looking at incumbents' individual wage trajectories. I find support for my identifying assumption by looking across other major downturns in my sample period, such as the 1993 recession, the early-2000s slowdown and the COVID-19 recession, suggesting that those post-reform crises are unlikely to drive my results.

My main finding is that the reform generated strong current gains in individual wages and earnings, but these gains dissipated relatively quickly after 2010. For workers earning the minimum wage in 2006, wages and earnings had grown about ten percentage points more than in the counterfactual by 2010. By 2014, the gain had fallen by about sixty percent to around four percentage points, even though the real floor was roughly at its 2010 level. Workers initially above the minimum wage also gained while the floor was rising, with spillovers extending up to around the median of the wage distribution, but by 2014 their individual wages had largely converged to their counterfactual trajectories on average.

This dissipation is unlikely to be explained by adverse effects on employment, hours, or career paths, which could have scarring effects on future earnings trajectories. Instead, there is considerable upward

income mobility, as most workers would have grown out of the 2006 minimum wage even absent the reform. For most of them, the minimum wage is a transitory state, so a lasting increase in the floor need not translate into equally large lifetime earnings gains, even when current gains are high.

Wage process. To extrapolate beyond the eight-year empirical horizon and quantify the lifetime incidence of the policy, I estimate a statistical wage process that embeds the minimum wage as a floor and allows for spillovers to wages further up the distribution. The process includes standard features from the income-dynamics literature (Meghir and Pistaferri, 2011), such as an individual fixed effect, persistent and transitory shocks and deterministic lifecycle growth. Its distinctive feature is the explicit minimum-wage floor, which becomes increasingly important for characterizing wage dynamics at the bottom of the distribution as minimum wages become more binding around the world. I estimate it to match the treatment effects from the empirical analysis together with cross-sectional and individual dynamic moments that help discipline individual wage dynamics. The process also distinguishes why workers earn the minimum wage, whether because of a low permanent type, an early career stage, or a transitory negative shock, distinctions that cross-sectional data alone may not reveal.

I show that the lifetime gains from the minimum wage are considerably smaller than current gains, but remain sizable. I simulate earnings paths from ages 18 to 64 under two permanent regimes: one fixes the real minimum wage at its 2006 level, while the other fixes it at its 16.2% higher 2010 level. Workers earning the minimum wage, who occupy the bottom 2.5 percent of the current-earnings distribution, receive the full 16.2% increase in current earnings. By contrast, mean lifetime earnings rise by 7.3% for workers in the bottom 2.5 percent of the lifetime-earnings distribution. Workers who remained at the minimum wage throughout their working lives would also receive a 16.2% lifetime gain, holding employment and hours fixed. Two factors explain why the lifetime gain remains sizable. First, workers at the bottom of the lifetime-earnings distribution spend a non-trivial share of their working lives near the floor. Second, through spillovers, workers keep benefiting from the higher minimum wage even after moving above it. Without spillovers, the lifetime gain would fall to 3.8%, about half of the lifetime gain with spillovers and about a quarter of the current gain.

Optimal minimum wage revisited. The final part of the paper assesses the implications of these findings for the optimal minimum wage. The optimal minimum wage balances equity gains against

efficiency costs, so a weaker lifetime equity effect could, in principle, call for a lower minimum wage.

To test this, I extend the model of Berger et al. (2025), in which households have permanent types (flat income profiles over time) and their gains from permanent minimum-wage increases are themselves permanent. I introduce idiosyncratic type risk, which generates income mobility, while keeping the cross-sectional distribution of types fixed. Income mobility partially shifts the minimum wage's role from redistribution between permanently rich and poor households toward insurance against low-income periods within their working lives.

Whether a weaker lifetime equity motive lowers the optimal minimum wage depends on households' ability to smooth consumption. When households are hand-to-mouth, the insurance value of the minimum wage fully offsets its weaker lifetime redistribution. These households can neither borrow nor save and therefore face an effectively static problem. Thus, households of a given current type make the same consumption and labor choices as in the permanent-type benchmark. The cross-sectional distribution of consumption, hours and period utility therefore remains unchanged, even though households move through this distribution over time. Since a utilitarian planner evaluates the discounted sum of period utilities across all individuals and over time, social welfare coincides with that of the permanent-type benchmark at each minimum wage. The optimum therefore remains unchanged at 53.6% of the baseline median wage. With complete consumption insurance, the optimal minimum wage falls to 47.6% of the baseline median wage, as the policy's insurance role is reduced when households can privately insure themselves.

Related literature. This paper contributes to four different strands of literature. First, this paper contributes to the literature on the distributional effects of minimum wages. Existing work shows that minimum wages compress lower-tail wage inequality and, in some settings, affect broader parts of the earnings distribution, not just the very bottom (Lee, 1999; Oliveira, 2023; Bossler and Schank, 2023; Giupponi et al., 2024; Engbom and Moser, 2022). Several of these studies document lasting compression of the future cross-sectional distribution.[1] However, their focus remains on the cross-

[1]Related work shows minimum wages can also compress earnings differences across racial, gender, and regional groups (Derenoncourt and Montialoux, 2021; Caliendo and Wittbrodt, 2022; Bossler and Schank, 2023). To the extent that belonging to these groups persistently lowers pay, the minimum wage's lifetime incidence for individuals belonging to these groups may exceed that implied by current wage rank.

sectional distribution, rather than on the lifetime distribution.[2] Hurst et al. (2026) also echo the idea that current gains erode over time, as firms take time to substitute low-wage workers with capital. Because they model workers as permanent types, their model is more informative about the long-run evolution of the bottom of the cross-section. I instead follow the impact of a permanent minimum-wage increase on individual workers' wage and earnings trajectories.

There is also a literature that follows individual workers, but over short horizons (Jardim et al., 2022; Clemens and Wither, 2019), whereas I follow them for up to eight years and extrapolate to the lifetime. Others study temporary differences in the minimum wage faced by treated and comparison workers, either by design, as with the youth subminimums in Cardoso (2019), or because inflation erodes nominal increases, as in the U.S. state studies below, whereas I find fading of individual effects even when the real minimum wage stays permanently higher. Closest to my horizon, Rinz and Voorheis (2018) and Kerr et al. (2025) follow workers for several years after U.S. state minimum-wage increases and reach different conclusions about the persistence of the gains. In those settings, however, inflation and wage growth alter the real bite of the original increase over time, making it difficult to separate fading treatment from fading individual effects. Kerr et al. (2025) compare minimum-wage workers with all workers earning more than a dollar above the minimum, a group that likely receives spillovers itself, which pulls the estimated gain toward zero.[3] I instead study an economy-wide, permanent increase in the minimum wage.

Second, I contribute to the literature on how firms and workers respond to minimum-wage increases on margins beyond wages and earnings directly. For example, minimum-wage increases may lead to adjustments in employment, hours, job promotion, mobility across firms, occupations, and industries, and educational choices. While most of this literature focuses on the contemporaneous effects of the minimum wage on these margins, with important exceptions (Meer and West, 2016; Harasztosi and Lindner, 2019), I offer worker-level evidence up to an eight-year horizon.

Third, I contribute to the income-dynamics literature (Gottschalk and Moffitt, 1994; Meghir and Pistaferri, 2011; Guvenen, 2009; Arellano et al., 2017; De Nardi et al., 2020) by estimating a wage process that explicitly allows for a binding minimum wage and spillovers above it. To my knowledge,

[2] In principle, a policy may permanently reduce cross-sectional inequality without changing lifetime inequality, if the gains individuals receive during their low-income years are offset by losses later in their individual careers.

[3] Cengiz et al. (2019), Autor et al. (2016) and Rinz and Voorheis (2018) find U.S. minimum-wage effects reaching roughly the first quarter of the distribution, up to about $3 above the minimum, whereas Kerr et al. (2025) classify as minimum-wage workers only those earning less than the minimum plus $1.00.

no existing income process embeds the existence of a minimum-wage floor.

Finally, the paper contributes to quantitative work on the optimal minimum wage. I introduce idiosyncratic risk into an otherwise permanent-type household problem, generating income mobility. This individual dynamic perspective, familiar in macroeconomics and the New Dynamic Public Finance literature, contrasts with the "static" assumption in models used to study the optimal minimum wage (Berger et al., 2025; Hurst et al., 2026; Ahlfeldt et al., 2022). I analyze the distinction between the redistribution and insurance roles of redistributive policies, studied for example in taxes and transfers (Farhi and Werning, 2013; Stepner, 2023), in the context of the optimal minimum wage. The minimum wage has been shown to have important insurance value because it limits how far wages can fall after a negative shock (Adamopoulou et al., 2027).

Outline. Section 2 describes the Portuguese minimum wage reform, the data, and the institutional setting. Section 3 presents the empirical strategy. Section 4 presents the results. Section 5 develops the wage model. Section 6 studies the optimal minimum wage. Section 7 concludes.

2 Institutional Context, Data and Descriptive Evidence

This section describes the Portuguese minimum wage reform, the administrative data, and the institutional wage-setting environment.

2.1 The Portuguese Minimum Wage Reform

Portugal's national minimum wage is set by the government, typically adjusted on January 1 of each year.[4] It is defined as a monthly value for full-time normal working hours and applies proportionally to part-time workers; overtime hours must also be compensated at or above the minimum wage rate. The national minimum wage applies to workers of all ages during the analysis period (2002–2014), as the youth subminimum wage was abolished in 1998. The evolution of the real minimum wage has historically depended on the governing party. Under center-right governments, nominal increases roughly matched inflation, leaving the real minimum wage unchanged. Under center-left governments, the real minimum wage has typically risen.

[4] All but one minimum wage adjustments took effect on January 1. The exception is 2014, when the increase from €485 to €505 took effect on October 1.

This paper exploits the contrast between three distinct phases, summarized in Table 1. From 2002 to 2006, Portugal was governed by the center-right, and the real minimum wage remained essentially constant: the nominal minimum wage rose from €348 to €386, approximately matching cumulative CPI inflation. From 2006 to 2010, following the unanticipated resignation of the center-right prime minister and the subsequent snap elections of February 2005, the incoming center-left government implemented a substantial minimum wage reform. In December 2006, a tripartite agreement between the government, labor unions, and the employers' confederation set an explicit target of raising the minimum wage to €500 by 2011. Although the government can set the minimum wage by decree, securing a consensus with social partners lends legitimacy and facilitates compliance.[5]

The nominal minimum wage rose from €386 in 2006 to €403 (2007), €426 (2008), €450 (2009), and €475 (2010), a cumulative nominal increase of 23% over four years. Deflated by the Consumer Price Index published by *Instituto Nacional de Estatística* (INE, Portugal's national statistical office), the real minimum wage rose by 16.2% from 2006 to 2010.

The third phase was triggered by the European sovereign debt crisis. Prime Minister Sócrates resigned in 2011, requesting an international bailout. Under the subsequent austerity program led by the incoming center-right government, the nominal minimum wage was frozen at €485 from January 2011 through September 2014, before a modest increase to €505 in October 2014.

[5]Compliance is high: the wage distribution displays a sharp mass at the statutory floor and little mass below it (Figure 2).

Table 1: Portuguese National Minimum Wage, 2002–2014

Year	Nominal MW (€/month)	Real MW Index	Cum. Real Change
2002	348	100.0	—
2003	357	99.6	−0.4%
2004	366	99.9	−0.1%
2005	375	99.7	−0.3%
2006	386	100.0	0%
2007	403	101.8	+1.8%
2008	426	105.1	+5.1%
2009	450	112.7	+12.7%
2010	475	116.2	+16.2%
2011	485	113.9	+13.9%
2012	485	111.5	+11.5%
2013	485	111.8	+11.8%
2014	505	116.4	+16.4%

Notes: The nominal minimum wage is a monthly value for full-time normal working hours; part-time workers receive the proportional amount. Real MW Index uses 2006 as base year (=100), deflated by the Consumer Price Index (INE). All minimum-wage changes during this period took effect on 1 January, except in 2014, when the increase from €485 to €505 took effect on 1 October.

Figure 1 plots the real hourly minimum wage over time. The three phases are clearly visible: a flat pre-reform period, a steep increase from 2006 to 2010, and a post-reform plateau.

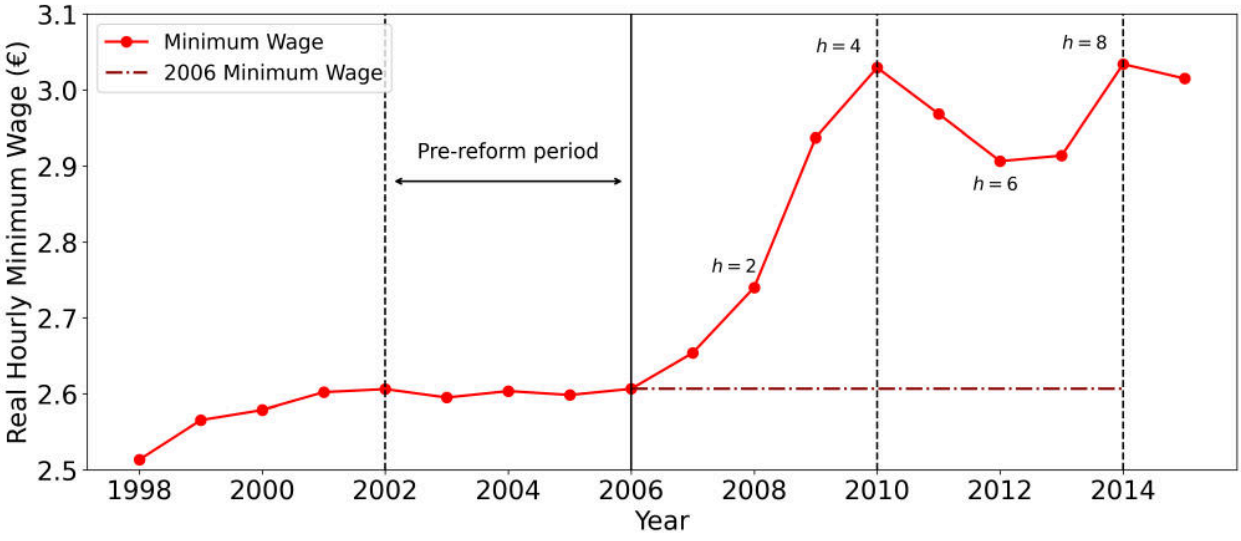


Figure 1: Evolution of the Real Hourly Minimum Wage in Portugal

Notes: Statutory real hourly minimum wage in 2019 euros: the monthly minimum divided by 173 full-time normal working hours per month, deflated by the INE Consumer Price Index. The dash-dotted policy series holds the real minimum wage at its 2006 level. The 2006 vertical marker is solid; the 2002, 2010 and 2014 markers are dashed. The reform phased in over 2006–2010 and was followed by a higher-level plateau, with the interim real fluctuations shown.

The period from 2006 to 2014 saw small negative growth in real GDP per capita and only meager

wage growth throughout: the median real hourly wage rose by 4.1% over these eight years (Figure A2). The reform therefore raised the minimum wage substantially relative to the median wage: the Kaitz index, the ratio of the minimum to the median hourly wage, sat around 55% throughout the pre-reform period and rose by around 12% between 2006 and 2014, reaching 61%. Portugal is part of a broader international trend. Figure A1 documents considerable increases in the Kaitz ratio across countries over time.

2.2 Data: Quadros de Pessoal

I use data from Quadros de Pessoal (QP), an annual matched employer–employee census conducted by the Portuguese Ministry of Employment. Every firm with at least one salaried worker is legally required to report its full roster of employees each October, along with detailed information on worker compensation, hours, demographics, and firm characteristics. To reduce misreporting, firms must also post their employee roster (including names, job titles, and monthly pay) inside their premises. The Ministry distributes an electronic version of the data with longitudinal identifiers for each worker, establishment, and firm, enabling matched employer–employee analysis. No wave was collected in 1990 or 2001.

QP covers the universe of private-sector employment in Portugal. It excludes public-sector employees (approximately 15% of total employment) and the self-employed (approximately 18%), capturing roughly two-thirds of total employment: in 2006, for example, the raw data contain 3.1 million workers, against 4.75 million employed individuals in the Labour Force Survey. The data contain rich information at the worker level (base earnings, overtime pay, regular and irregular bonuses, normal and overtime hours, age, gender, education, four-digit occupation and job title, tenure, and part-time status) as well as at the firm level (region, four-digit industry, firm size, and annual revenue). QP provides high-quality data on *hourly* wages, in contrast to administrative panels like U.K. PAYE and U.S. Social Security earnings records, which also have high-quality earnings measures but limited or no information on hours. This richness at both the worker and firm level is important for the analysis: beyond estimating the effects on wage and earnings growth, it allows me to examine whether the minimum wage increase affected hours, occupational mobility, industry switching, and reallocation across firms of different sizes and productivity levels.

For the analysis, I measure total compensation as the sum of base earnings, overtime pay, and regular

and irregular payments[6], and construct real hourly earnings (i.e., wages) by dividing total compensation by total hours (normal plus overtime), deflated to 2019 euros using the CPI.[7] I also conduct the analysis using regular earnings, defined as base pay plus regular and seniority payments, and regular hourly wages, obtained by dividing these earnings by normal hours worked. I exclude observations with zero or missing reported hours.

2.3 Cross-Sectional Wage Compression

Figure 2 presents suggestive evidence of how strongly the minimum-wage reform compressed the cross-sectional wage distribution. Low wages that otherwise would have been bunched at the 2006 real minimum wage are instead bunched at the 16.2% higher 2010 level. The figure also suggests spillovers to workers earning above the new minimum. Absent spillovers, the mass of wages below the new minimum would all be bunched at the 2010 minimum, but the rest of the distribution would remain the same.[8] Instead, the shift extends further up the distribution and becomes weaker around the median, consistent with Portuguese evidence in Oliveira (2023) and Caires (2025).

Spillovers may have several sources. One source is firms preserving pay hierarchies (Dube et al., 2019). Firms may do so, for example, for fairness reasons or because they fear higher reservation wages further up the distribution (Butcher et al., 2012; Phelan, 2019).[9] Another possible source is the adjustment of collectively bargained wage floors in response to an increase in the national minimum wage. Portugal belongs to the Continental European labor-market model, with extensive sectoral collective bargaining (Vilares and Reis, 2022).[10] These agreements establish occupation-specific wage floors above the national minimum, although the large majority of these floors lie below the median wage (Card and Cardoso, 2022; Martins, 2021). Most covered workers earn above their applicable floor, with a typical wage cushion of about 20%.

Caires (2025) finds that spillovers come mainly from firms preserving pay hierarchies, and much

[6]Irregular bonuses generally account for a larger share of earnings higher up the distribution: in 2006, their average share of total monthly earnings is around 2% among workers in the bottom 20%, compared with around 10% among those in the top 20%.

[7]Throughout the paper, "wages" refer to hourly pay, while "earnings" are wages multiplied by hours worked.

[8]This comparison abstracts from aggregate effects, which were relatively small during this period. When plotted separately, the cross-sectional wage distributions for 2002–2005 lie almost on top of the 2006 distribution, also due to small aggregate effects during the pre-reform period.

[9]Spillovers may also operate in the opposite direction: Gregory and Zierahn (2022) document negative wage effects further up the distribution in Germany's roofing industry.

[10]Evidence from Anglo-Saxon labor markets, such as the United States and the United Kingdom, generally shows smaller spillovers (Autor et al., 2016; Dickens and Manning, 2004; Stewart, 2012; Giupponi et al., 2024).

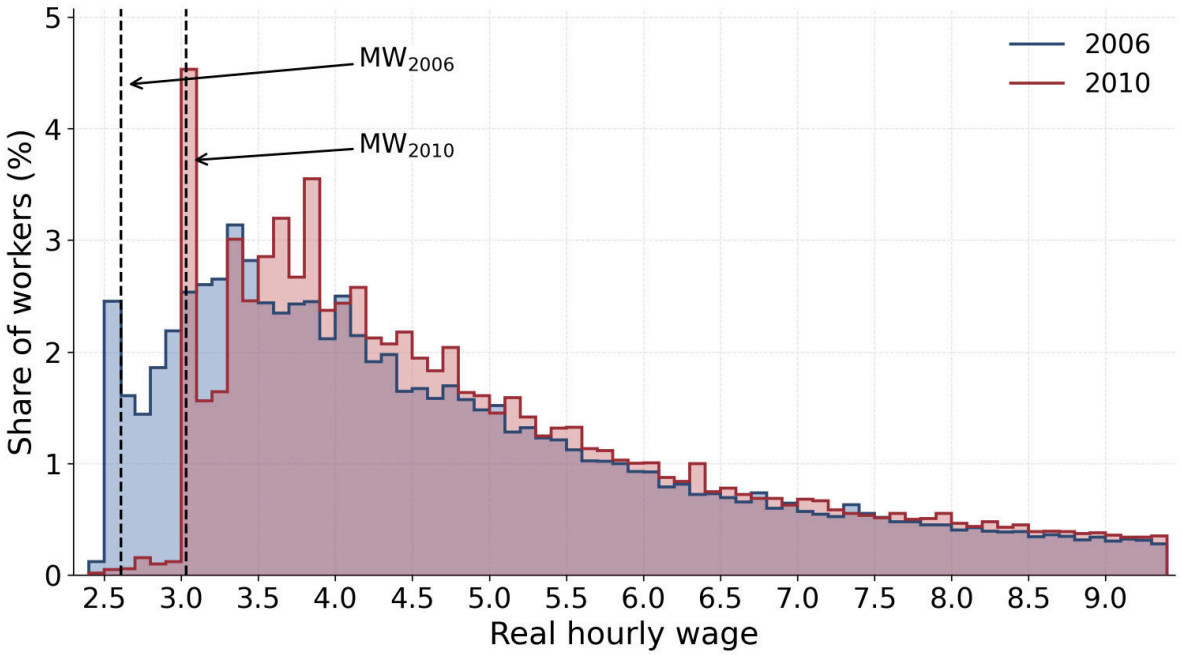


Figure 2: Cross-Sectional Hourly Wage Distribution, 2006 vs 2010

Notes: Distribution of real hourly wages (2019 euros, deflated by the Consumer Price Index) among workers observed in each year with real hourly wages above €2.45, in bins of €0.10/h. Because the 2006 minimum wage of €2.6065/h falls just above a bin threshold, that threshold is moved from €2.60 to €2.61 so that the bin absorbs almost all minimum-wage workers, as in the first wage bin of the empirical analysis. Bars show the share of the full sample in each bin. Dashed vertical lines mark the real hourly statutory minimum wage in each year, M_{2006} = €2.61 and M_{2010} = €3.03. The horizontal axis stops at €9.4/h, approximately the 83rd percentile of the 2006 distribution; the median hourly wage was €4.78 in 2006 and €5.08 in 2010. Sample sizes are 2,044,169 workers in 2006 and 2,112,758 in 2010.

less from adjustments to collectively bargained wage floors.[11] While this paper is agnostic about the sources of spillovers, it emphasizes their importance for the lifetime impact of minimum-wage policy. Even if workers grow out of the minimum wage, they may remain in the part of the wage distribution affected by spillovers and continue to receive partial gains later in life.

The figure suggests that workers at the bottom may experience strong current gains, but the persistence of these gains depends on whether—and how quickly—they escape the bottom of the wage distribution. A permanent rise in the minimum wage permanently treats particular positions in the wage distribution, rather than particular individuals, who may grow out of the minimum wage; I now examine how quickly they do so.

[11] Caires (2025) studies the later minimum-wage increases of 2014 and 2016. Related evidence from Card and Cardoso (2022) shows that bargained wage floors fell relative to the national minimum over 2010–2016, suggesting that any increases in these higher floors were, on average, less than proportional to the increase in the national minimum.

2.4 Growing Out of the Minimum Wage and Fraction of Lifetime Spent at the Minimum Wage

Figure 3 shows the share of workers initially at the minimum wage who earn above the contemporaneous minimum at different horizons: one, two, four, and eight years later. For example, the highlighted 40% indicates that around 40% of workers at the minimum wage in 2004 earn above the minimum wage one year later, while the other 60% still earn the minimum wage.[12] For 2004 minimum-wage workers, this share rises to around half after two years, around two-thirds after four years, and around three-quarters after eight years. These shares are relatively stable across the years. The remaining quarter are still at the minimum wage after eight years and remain fully, directly exposed to the policy, absent employment effects.

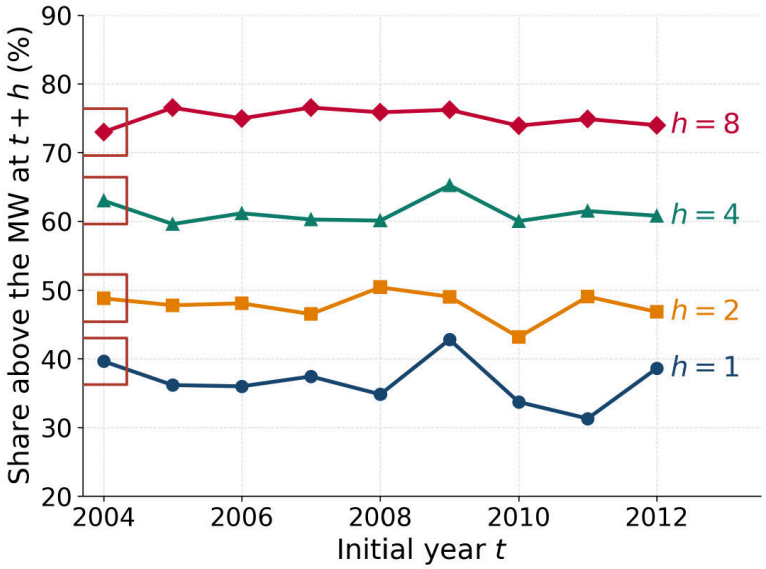


Figure 3: Share of Minimum-Wage Workers Who Grow Out by Horizon h

Notes: Each line shows the share of workers with total hourly wages $w_{i,t} \leq 1.05M_t$ in starting year t who earn $w_{i,t+h} > 1.05M_{t+h}$ at horizon $h \in \{1,2,4,8\}$, conditional on being observed again at $t+h$. The denominator is the number of baseline workers observed at that endpoint; the contemporaneous statutory floor is used at each date. Starting years are $t = 2004,\ldots,2012$; red boxes highlight the four 2004 observations. Where a worker has multiple spells in a year, the highest-paid spell is retained. These are endpoint transitions, rather than exits at any time within the window or permanent exits from the minimum wage.

For the United States, more than half of minimum-wage workers earn a higher wage one year later (Smith and Vavrichek, 1992; Kerr et al., 2025). The comparable figure in the Portuguese data is about half when the 5% band is not applied.

[12] These shares are conditional on workers being observed in employment at the relevant horizon. The figure notes define the band around the statutory minimum.

These are average growing-out rates; Table 2 shows how they vary across basic demographic groups. In particular, women, older workers, and less-educated workers are less likely to grow out of the minimum wage, suggesting greater direct exposure over their remaining working lives. Nevertheless, even for older, less-educated women, around half have grown out of the minimum wage after four years.

Table 2: Probability of Growing Out of the Minimum Wage

$$\mathbf{1}\{\,\mathrm{MW}^{+}_{i,t+h}\,\big|\,\mathrm{MW}_{i,t}\,\} \;=\; \tau_t \;+\; X_{i,t}\,\beta \;+\; \varepsilon_{i,t+h}$$

	Share earning above the minimum wage at $t+h$			
	$h=1$	$h=2$	$h=4$	$h=8$
Base group: Men, Age <30, Higher Education	49.0%	65.9%	80.5%	89.4%
Women	-3.49***	-4.43***	-4.13***	-4.15***
	(0.14)	(0.17)	(0.21)	(0.22)
Age 30–39	-8.21***	-10.16***	-10.21***	-8.35***
	(0.17)	(0.21)	(0.24)	(0.23)
Age 40–49	-11.33***	-14.22***	-15.37***	-12.94***
	(0.19)	(0.23)	(0.27)	(0.29)
Age 50+	-14.97***	-19.01***	-21.20***	-19.12***
	(0.20)	(0.26)	(0.33)	(0.42)
Upper Secondary Education	-1.77***	-2.16***	-2.11***	-2.33***
	(0.49)	(0.59)	(0.67)	(0.65)
Lower Secondary Education	-4.06***	-5.01***	-4.90***	-4.92***
	(0.47)	(0.58)	(0.65)	(0.63)
Primary Education	-5.23***	-6.55***	-7.06***	-6.58***
	(0.47)	(0.57)	(0.65)	(0.63)
Year Fixed Effects	Yes	Yes	Yes	Yes
Mean outcome, all workers	36.5%	47.7%	61.2%	75.1%
Observations	727,999	629,706	524,570	435,299

Notes: $\mathrm{MW}_{i,t} = 1$ if worker i is at the minimum wage in year t, defined as earning within 5% of the statutory real minimum hourly wage ($w_{i,t} \le 1.05\,M_t$); $\mathrm{MW}^{+}_{i,t+h} = 1$ if i earns more than 5% above it in year $t+h$ ($w_{i,t+h} > 1.05\,M_{t+h}$). One regression per horizon h, estimated by OLS on workers with $\mathrm{MW}_{i,t} = 1$ in base year $t \in \{2004,\ldots,2012\}$ and observed again at $t+h$, with the outcome multiplied by 100, so coefficients are in percentage points; τ_t are year fixed effects. Reference group: men aged 18–29 (Age <30) with higher education; Age 50+ spans 50–64. Schooling categories by years completed: higher education ≥14 (ISCED 5–8), upper secondary 12 (*ensino secundário*, ISCED 3), lower secondary 9 (3rd cycle of basic education, ISCED 2), primary ≤6 (1st–2nd cycle of basic education, ISCED 1). A dummy for missing schooling (1.2% of the sample) is included and not reported. Standard errors clustered by worker (a worker can appear in several base years) in parentheses. * $p<0.10$, ** $p<0.05$, *** $p<0.01$.

While workers may grow out of the minimum wage, they may return to it later. I therefore complement these transition rates by examining the fraction of working life spent at the minimum wage. Figure 4 shows the fraction of observed working life spent there. I follow workers born between 1961 and 1966 over the prime-age window from 25 to 55, in line with the convention adopted by the GRID

project (Guvenen et al., 2022b). Complete histories covering my empirical age range from 18 to 64 are not available. This may understate the share of workers who spend a large fraction of their working life at the minimum wage, since workers are more likely to earn the minimum wage when very young or older.[13]

The left panel includes all workers, while the right panel conditions on ever having been at the minimum wage. Around three-quarters of workers are never observed at the minimum wage. Even among workers who earn the minimum wage at some point, the median share is 11.1%, or around three years.[14] Nevertheless, around 0.1% of all workers spend almost their entire observed working life at the minimum wage—around 95%.



Figure 4: Share of Working Life Spent at the Minimum Wage

Notes: Distribution of the fraction of observed prime-age years in which a worker's real base hourly wage lies in the symmetric band $[0.95M_t, 1.05M_t]$. The 1961–1966 birth cohorts are followed over ages 25–55, consistent with the prime-age convention in Guvenen et al. (2022b), and restricted to workers observed in at least 14 years, roughly half of the 29–30 available QP waves over their age-specific window. The number of available waves varies across cohorts because 1990 and 2001 were not surveyed. This roughly-half-history restriction selects workers with stronger labor-market attachment, in the spirit of the lifetime-attachment screen in Guvenen et al. (2022a). Each worker's fraction uses their observed years as the denominator. Left: all eligible workers; right: conditional on ever entering the band. The first bar represents exact zero; subsequent adjoining bars group positive fractions in intervals of width 2/30. Shares use the corresponding full or conditional worker denominator. The share at zero is 75.57%.

Taken together, these patterns suggest that while workers at the minimum wage may see strong gains, minimum-wage status is mostly transitory. Therefore, these gains may fade on average. I now test this more formally.

[13]The sample requires observation in at least 14 of the 29–30 available waves, selecting workers with stronger labor-market attachment. The distributions are therefore suggestive evidence rather than a full-career population distribution.

[14]The percentage uses observed working years as the denominator. Around three years is its equivalent over the full prime-age window, rather than the median number of observed minimum-wage years.

3 Research Strategy

This section presents the empirical strategy used to estimate the treatment effects and discusses the main threats to identification.

3.1 Empirical Design

The empirical strategy exploits the fact that workers were differentially exposed to the real minimum wage rise from 2006 onwards, depending on their position in the 2006 wage distribution. Whereas workers earning at or near the minimum wage at the onset of the reform are highly exposed, as one moves up the wage distribution, workers become plausibly unaffected beyond a certain point and these serve as the control group.[15] The exact cutoff dividing treatment and control groups is discussed below.

The main challenge with worker-level panel data is mean reversion: a worker who happens to earn a low wage in a given year tends to grow faster afterwards, even absent any policy change. This rules out a difference-in-differences (DiD) in wage log-levels, as the expected trend in low-wage workers' wages differs from the trend in high-wage workers' wages. Working with *individual growth rates* can address the problem if aggregate economic conditions shift individual growth rates by approximately the same amount. This approach has precedent in the minimum wage literature: Stewart (2012) and Dustmann et al. (2022) both work with individual wage changes to address mean reversion.[16]

The aim of the research design is to characterize h-year individual wage growth rates, $\log w_{i,t+h} - \log w_{i,t}$, by initial wage, $w_{i,t}$, and evaluate how these growth rates change with the reform. Since a worker's wage trajectory can be described by an initial wage and the sequence of h-year growth rates at each horizon, computing those individual growth rates by initial wage traces out the trajectories themselves. I do so at every annual horizon up until 2014 (eight-year horizon). Each worker's actual trajectory from 2006 onwards is compared to a counterfactual trajectory built from pre-reform (2002–2006) individual growth rates, and the gap between the two is the causal effect of the reform at each horizon.

[15] As formalized below, the design also allows workers at the upper end of the distribution to be affected by general-equilibrium channels of the policy, provided those effects are approximately common across the wage distribution. Examples are provided in Section 3.2.

[16] The concern dates back to Ashenfelter (1978), whose "dip" in trainees' pre-program earnings made mean reversion look like a treatment effect.

Identifying assumption. The identifying assumption is that the shape of individual growth rates along the initial wage distribution (the pattern of mean reversion) would have remained constant absent the reform.[17] The growth-rate profile can still shift up or down with aggregate economic conditions other than the reform, but the shape is assumed to be unaffected by such aggregate effects. Subtracting the average growth-rate difference across the high-wage control group nets out these aggregate movements without distorting the shape. Implementing this adjustment requires drawing the line between treated and high-wage control workers, specified in the **Control adjustment** paragraph.

Specification. The regression below recovers, non-parametrically, the relationship between a worker's position in the initial wage distribution, $w_{i,t}$, and their subsequent h-year wage growth, $\log w_{i,t+h} - \log w_{i,t}$. The baseline specification, run separately for each horizon $h \in \{1, \ldots, 8\}$, is[18]:

$$\begin{aligned}\log w_{i,t+h} - \log w_{i,t} \;=\; & \sum_{k=1}^{K} \gamma_k^{(h)} \mathbf{1}\{b_{k-1} < w_{i,t} \le b_k\} \;+\; \lambda^{(h)} X_{i,t} \\ & +\; \text{Post}_t \cdot \left[\sum_{k=1}^{K} \delta_k^{(h)} \mathbf{1}\{b_{k-1} < w_{i,t} \le b_k\} \;+\; \tilde{\lambda}^{(h)} X_{i,t} \right] + \varepsilon_{i,t+h}. \qquad (1)\end{aligned}$$

The variable $w_{i,t}$ is worker i's real hourly wage at time t; the indicator $\mathbf{1}\{b_{k-1} < w_{i,t} \le b_k\}$ assigns each worker to one of K bins by their wage at the start of the h-year period, with bin edges $b_0 < b_1 < \cdots < b_K$ (Appendix B.1); $X_{i,t}$ is a vector of demographic controls such as age, gender and education; and Post_t equals one if the individual growth rate belongs to the POST period.

Each regression stacks individual growth rates from two h-year periods: one set of growth rates from the PRE period, the other from the POST period. As a running example, the $h = 2$ regression stacks all 2004–2006 (PRE) and 2006–2008 (POST) individual growth rates. The set of coefficients $\gamma_k^{(2)}$ describes the shape of the average 2004–2006 individual growth rates by initial 2004 wage bin, the cross-bin profile. These average growth rates are the counterfactual 2006–2008 growth rates under stable aggregate conditions and absent any minimum-wage change. The coefficients $\delta_k^{(2)}$ are the *excess* 2006–2008 growth rates relative to the 2004–2006 PRE growth rates. Standard errors are clustered at the worker level, allowing for correlation between observations of the same individual in the PRE and POST samples.

[17] By "shape" I mean the entire growth-by-bin profile (steeper at the bottom of the wage distribution and flatter higher up) rather than a single slope, since the relationship is non-linear.

[18] Section 4.2 replaces individual wage growth on the LHS with other outcomes (earnings growth, employment status, and firm retention). The general specification is shown in Appendix B.2.

The identifying assumption states that, under stable aggregate conditions and in the absence of real minimum-wage changes, the 2006–2008 growth rates should equal those of 2004–2006; the $\delta_k^{(2)}$ would therefore all be zero. More generally, the assumption allows for different aggregate conditions in 2006–2008 relative to 2004–2006: if no minimum-wage change had occurred, the $\delta_k^{(2)}$ would be the same across bins. This is what happens when contrasting 2004–2006 against 2003–2005 growth rates: the recovery from the wage slowdown that followed the 2002–2003 crisis makes 2003–2005 growth rates higher than the subsequent 2004–2006 ones, but the shift is approximately parallel across the distribution, as shown in Appendix B.4. When the real minimum wage rises, excess growth in 2006–2008 is higher for lower bins than for higher bins, since low-wage workers in 2006 are more exposed to the policy change than high-wage workers.

Control adjustment. To purge the common aggregate component, I subtract the mean PRE-to-POST shift in individual wage growth among control workers, equivalently the count-weighted average of $\delta_k^{(h)}$ across the high-wage control bins:

$$\beta_k^{(h)} \equiv \delta_k^{(h)} - \delta_C^{(h)}, \qquad \delta_C^{(h)} = \sum_{k \in \mathscr{C}} s_k \, \delta_k^{(h)} = \frac{1}{N_{\mathscr{C}}} \sum_{i : w_{i,t} \geq €5} \delta_{k(i)}^{(h)}, \qquad s_k = \frac{N_k}{\sum_{k' \in \mathscr{C}} N_{k'}}, \tag{2}$$

where $\mathscr{C}$ is the set of control bins, N_k is the bin's POST-period sample count, $N_{\mathscr{C}} = \sum_{k \in \mathscr{C}} N_k$, and $k(i)$ denotes worker i's bin: $\delta_C^{(h)}$ is the average growth-rate difference among control-group workers. The resulting $\beta_k^{(h)}$ is the causal effect of the minimum wage on h-year wage growth for workers starting in bin k. I designate workers earning above €5/h (around the 54th percentile of the 2006 wage distribution) as the control set $\mathscr{C}$.[19] The choice is grounded in the Portuguese minimum-wage spillover literature: Oliveira (2023) estimates spillovers reaching up to the 54th percentile of the wage distribution and Caires (2025) up to the 47th, placing the €5/h threshold just above the upper end of documented spillover reach. It is corroborated by the empirical pattern of the data: in the reform-period regressions reported below, the cross-bin profile of $\delta_k^{(h)}$ flattens around that point and remains flat through the higher bins.

Synthetic pre-reform growth for $h > 4$. The pre-reform period spans only four years, so no PRE growth rate is directly observed for $h > 4$. For those horizons I chain each worker's actual 2002 →

[19] Main results are robust to the two alternative thresholds of €5.4/h and €6.2/h, around the 60th percentile and two-thirds of the distribution, respectively.

2006 growth with a predicted continuation over the remaining $h-4$ years. I leverage the identifying assumption of stable mean reversion to predict the continuation: the growth-by-bin profile is assumed to hold had the real minimum wage stayed constant.[20] The continuation is evaluated at the worker's 2006 position rather than their 2002 one: a worker who started near the minimum wage in 2002 and moved higher in the wage distribution by 2006 is by then four years older and further along their earnings path, and both imply slower growth over the next four years. It is therefore the average pre-reform $(h-4)$-year growth of workers in that 2006 bin, and, in specifications with controls, at the worker's 2006 characteristics. Stacked against the worker's actual $2006 \to 2006+h$ growth in Equation (1), this makes $\beta_k^{(h)}$ the treatment effect on wage growth between 2006 and $2006+h$, defined exactly as at the shorter horizons. Appendix B.3 gives the construction in full.

3.2 Threats to Identification

The validity of the research design rests on parallel shifts in individual wage growth rates between treatment and control groups absent the minimum wage increase. What would violate this assumption is any shock, policy, or trend, other than the minimum wage reform, that *differentially* affects wage growth rates at different points of the 2006 wage distribution during the post-reform period.

Macroeconomic shocks. The reform period overlaps with both the 2008 financial crisis and the 2011–2013 Portuguese sovereign debt crisis, two major events that negatively affected economies throughout the world and Southern Europe in particular. Several arguments mitigate the concern that these shocks could confound the estimated treatment effects.

First, the analysis looks at the growth rates of workers already employed in 2006. If firms prefer to penalize the pay of entrants rather than that of incumbent workers during recessions, consistent with Haefke et al. (2013) and, for Portugal, Carneiro et al. (2012), then it is possible to have asymmetric impacts of the business cycle across the cross-sectional wage distribution (Barrela et al., 2024) but a similar impact on growth rates across that distribution. The cross-section moves for two reasons, because wages change and because who is in it changes (Solon et al., 1994); the growth rates here follow the same workers throughout, so only the first reason is at play.

Second, to evaluate this claim I look at the other recessions in the sample, which runs from 1986 to

[20] The assumption would be violated if the crises in the POST period had differential individual growth effects across the 2006 distribution, discussed in Section 3.2.

2021, and find further support for the identifying assumption. As already shown in Section 4.1.1, after the wage slowdown of the early 2000s the real minimum wage was constant, both in absolute terms and relative to the median, and the growth-rate profiles shifted approximately in parallel. The adjacent comparison of 2000–2002 with 2002–2004 also yields similar growth profiles (Appendix Figures A10 and A11). During COVID, individual wage growth was considerably smaller in 2020–2021 than in 2019–2020, and the profiles were again approximately parallel (Appendix Figures A6 and A7), consistent with the identifying assumption.[21] During the 1993 recession, wage growth in 1992–1993 was considerably smaller than in 1993–1994. Here, however, the real minimum wage was not stable: it fell by 6.1% across 1992–1993 and then rose by 5.8% across 1993–1994, with the Kaitz index dropping from 55.7% to 52.4% before recovering to 53.6%. Accordingly, the 1993–1994 profile displays a positive tilt relative to 1992–1993, in the expected direction (Appendix Figures A8 and A9).

Finally, if crises produce differential effects across the wage distribution by hitting low-wage sectors or occupations harder, industry and occupation fixed effects absorb the variation; results are robust to this specification.

Reverse causality. If the government raised the minimum wage in response to a slowdown trend in low-wage workers' wage growth relative to that of higher-wage workers, the estimated effects could be biased. However, the reform originated from a snap election in 2005, called after the unexpected resignation of the prime minister, which brought to power a center-left government with a longstanding commitment to raising the minimum wage. The placebo results in Section 4.1.1 provide direct evidence against this concern: the gradient of individual wage growth rates across the initial wage distribution is stable throughout the 2002–2006 pre-reform period, indicating that the wage dynamics the identification strategy relies on were not shifting in ways that might have prompted or coincided with the policy. The same stability shows no sign of anticipation effects: pseudo-treatment effects in the pre-reform period are uniformly zero. Anticipation was in any case unlikely, as discussed in Section 2.1.

Confounding policies. Existing time-invariant policies with differential effects across the wage distribution (such as progressive taxation) are not a threat to identification. Such policies contribute to the already-existing pattern of mean reversion in wage growth, but only *changes* in policy or its incidence alter the pattern and pose a threat. During the reform period, I am not aware of policy changes with nontrivial differential effects. One possibility is the *Novas Oportunidades* initiative, which pre-

[21] The minimum wage grew relative to the median in this period, but by about the same amount across the two windows.

dated the reform but was expanded in 2005 and targeted low-skilled adult workers with the aim of improving their qualifications. However, the literature on active labor market policies typically finds skill-enhancing programs to be generally ineffective for adults (Heckman et al., 1999). Furthermore, Costa Dias and Varejão (2012) evaluate Portuguese active labor market policies over this period and find limited effects on participants' employability, consistent with the international evidence. The austerity programme also raised income tax progressivity, increasing the top marginal rate from 46.5% to 48% in 2013 for incomes above €80,000, which falls above the 99th percentile of the 2006 wage distribution, far above the control group.

Differential general equilibrium effects. The minimum wage increase may have generated general equilibrium effects. Such effects are not a concern for identification insofar as they are common across wage bins: the difference-in-differences specification identifies treatment effects relative to the control group, not in absolute terms (the classical missing-intercept problem). This is appropriate for assessing the inequality-reducing effects of the policy, since common general equilibrium effects would shift all workers' wage growth equally and leave the relative position of low- versus high-wage workers unchanged. Harasztosi and Lindner (2019) find that most of the incidence of a large minimum-wage increase falls on consumers through higher prices, while Ashenfelter and Jurajda (2022) document near-full pass-through into restaurant prices. Although there is no Portuguese evidence on differential exposure to this price shock across the income distribution, MaCurdy (2015) for the United States and Harasztosi and Lindner (2019) for Hungary find little evidence of substantial differential exposure, whereas Renkin et al. (2022) find some differential exposure in U.S. grocery spending, diminishing the *real* minimum-wage gains of low-wage workers. Another general-equilibrium channel is an increase in aggregate productivity through the cleansing of low-productivity firms (Alexandre et al., 2022). Lastly, if there were a strong differential impact across the wage distribution, it should also be visible within the control group; there is no evidence of one at either shorter or longer horizons.

4 Empirical Results

This section estimates the causal impact of the large and sustained minimum-wage reform on individual wage trajectories over the eight-year horizon from 2006 to 2014. The reform's 16.2% real minimum-wage increase was phased in between 2006 and 2010, and the estimated treatment effects build up

as the floor is gradually raised, then fade out after 2010. Section 4.2 examines whether the policy affected outcomes that could shape future earnings and career paths, including employment, hours worked, and moves across firms, occupations, and industries, among others. Section 4.3 draws out implications for cross-sectional versus lifetime inequality. The outcome throughout is wage growth; effects on employment and hours are very small (Section 4.2, Appendix B.8), so the causal effects on earnings growth are of a similar magnitude.

4.1 Results on Individual Wage Trajectories

I begin by documenting that individual wage growth rates by initial wage are stable across the pre-reform period, consistent with the identifying assumption.[22] I then present the effects on individual wage trajectories during the reform and after the reform at the permanently higher minimum wage. The main wage growth results display the even horizons ($h \in \{2, 4, 6, 8\}$): two and four years into the reform (2008 and 2010), and two and four years into the freeze at the higher minimum wage (2012 and 2014). The average wage trajectory figures in Section 4.1.4 use all horizons.[23]

4.1.1 Individual Wage Growth During the Pre-Reform Period (2002–2006)

This section characterizes the empirical relationship between a worker's initial wage and their subsequent individual wage growth, using the pre-reform individual wage growth rates (2002–2006), during which the real minimum wage was essentially constant. It provides direct evidence that the gradient of growth rates across the initial wage distribution is stable across the pre-reform period. Stability of this gradient is the identifying assumption (parallel shifts in individual wage growth rates) on which the difference-in-differences specification developed in Section 3 relies. I document this stability for two-year wage growth, the baseline horizon of the main results. Appendix B.5 presents the same evidence for one-year and three-year growth.

Figure 5 compares the two-year wage growth profiles of the three pre-reform windows: 2004–2006 (the baseline), 2003–2005, and 2002–2004. Aggregate conditions differed across these windows, as the recovery from the 2002–2003 crisis lifted wage growth in the two earlier ones (Section 2.1, Figure A2).

[22] For clarity, the figures present estimates without covariate controls. The treatment effects barely change when including a quartic in age, gender, and education, or in a robustness check that additionally includes industry fixed effects.

[23] The first year of the reform saw a small real minimum-wage rise of less than 2% from 2006 to 2007, against a 5.1% cumulative real increase by 2008 (Table 1), so the shortest horizon emphasized in the main results is $h = 2$.

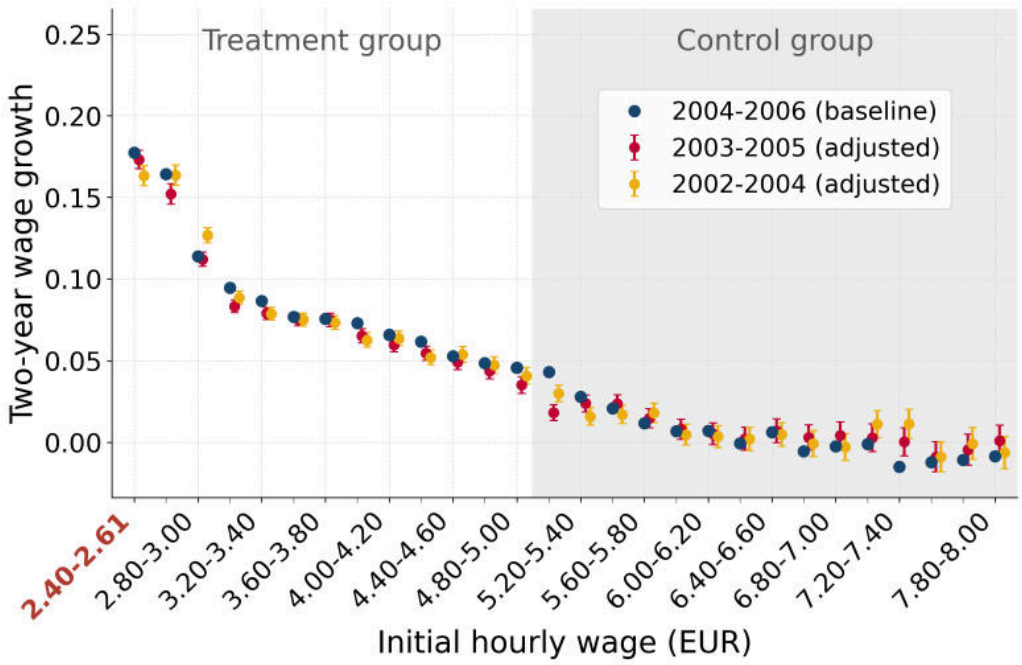


Figure 5: Individual Wage Growth Before the Reform: Stability of Two-Year Growth Profiles

Notes: The 2004–2006 baseline two-year wage-growth profile and the 2003–2005 and 2002–2004 placebo profiles, adjusted for their count-weighted mean difference from the baseline in the control bins. The 99% normal intervals around the adjusted placebo profiles quantify uncertainty in their contrasts with the baseline. Initial total hourly wages above €5/h define the shaded control region. Appendix Figure A4 also reports the corresponding placebo treatment effects.

Each profile is therefore adjusted by subtracting the mean growth rate across the high-wage control bins, as specified in Section 3; Appendix B.4 illustrates the adjustment for the 2003–2005 window. The three adjusted profiles lie essentially on top of each other across the entire distribution. Appendix Figure A4 reports both these profiles and the placebo treatment effects of the two earlier windows relative to the baseline, which are statistically close to zero.

Over the 2002–2006 period, the shape of the growth-by-bin profile thus appears stable, consistent with the identifying assumption developed in Section 3. The shape itself is also informative: the profile slopes downward, reflecting mean reversion, as initially lower-wage workers experience faster wage growth. This mean reversion could be due to lifecycle growth, since many low-wage workers are young and still inexperienced, or to the reversal of a negative shock, such as starting in a low-paying firm.[24] Section 5 explores these sources of wage growth and their relative importance in shaping workers' lifetime trajectories.

[24]Mean reversion remains very strong even conditional on age, due to the existence of such shocks.

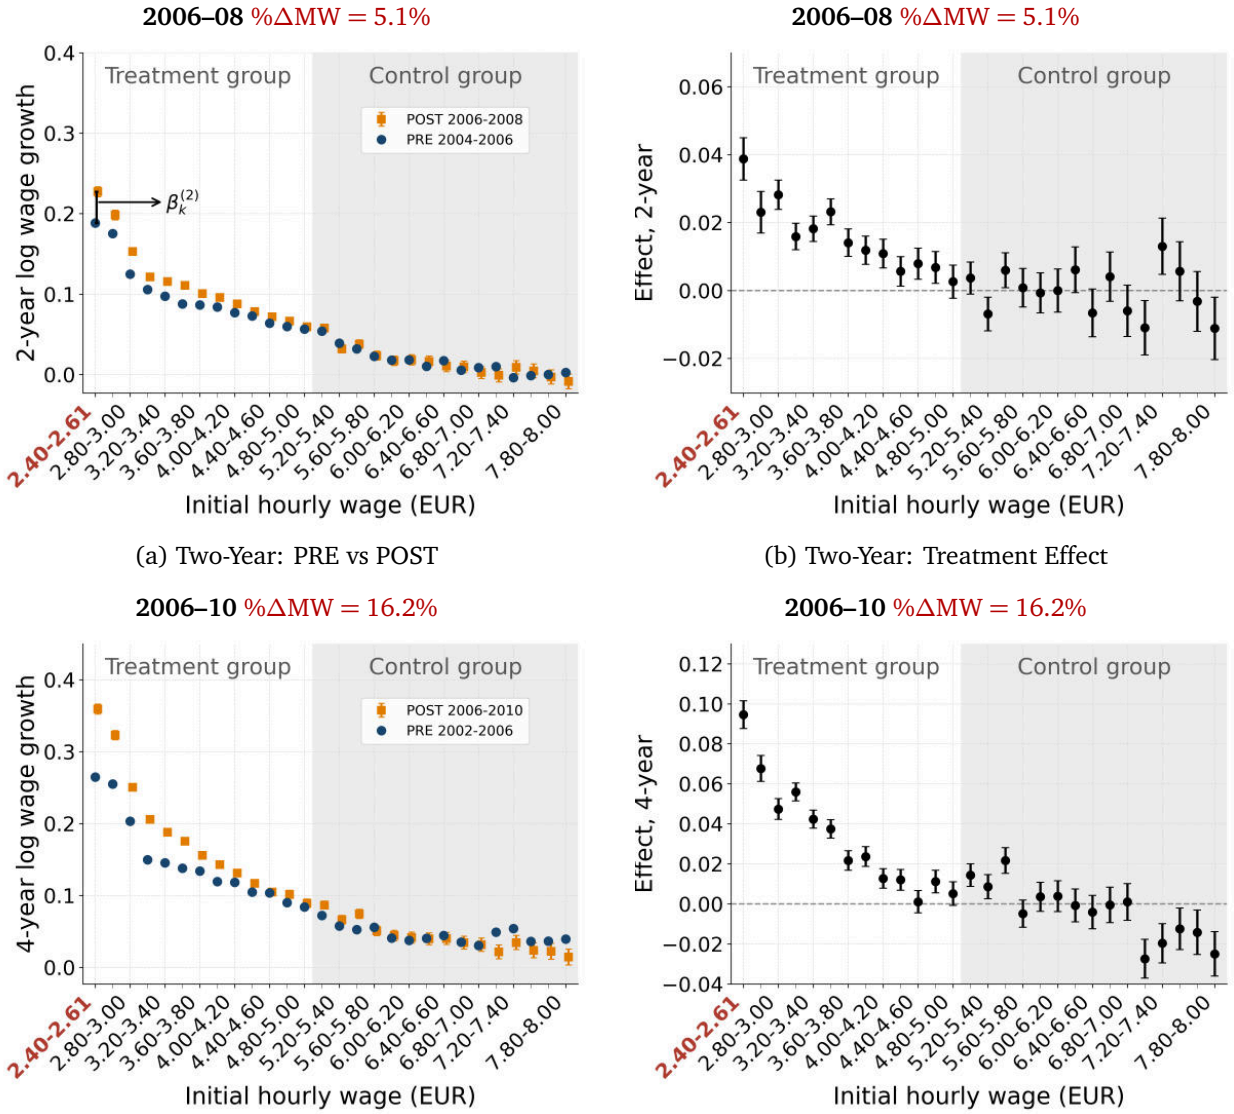
4.1.2 Individual Wage Growth During the Reform (2006–08 & 2006–10)

Figure 6, panel (a), compares actual two-year individual wage growth (2006–2008) with the counterfactual two-year growth the same workers would have experienced had the real minimum wage remained unchanged. These counterfactual wage growths are the adjusted pre-reform growth rates (2004–2006). Panel (b) shows the treatment effects on wage growth $\beta_k^{(2)}$. Panels (c) and (d) plot the same objects at the four-year horizon: individual wage growth between 2006 and 2010, against the counterfactual implied by the 2002–2006 pre-reform growth rates.

In contrast with the pre-reform comparisons of Figure 5, where the adjusted profiles sat on top of each other, the observed post-reform profile now lies clearly above its counterfactual. The gap is widest at the bottom of the 2006 wage distribution and shrinks to zero moving up. Workers earning the minimum wage in 2006 would have grown roughly 19% on average by 2008 had the real minimum wage stayed constant. With the 5.1% real minimum-wage increase over 2006–2008 (Table 1), their average wage growth was instead 23%: an average treatment effect of approximately 4 percentage points. Over the four-year window, the same workers would have grown 26.5% on average; with the cumulative 16.2% increase of 2006–2010, their average growth was instead 36%, 9.5 percentage points higher. Treatment effects shrink moving up the 2006 wage distribution but are still statistically distinguishable from zero up to approximately the median of the 2006 wage distribution.

Interpreting the magnitude of the average treatment effects and the degree of fading. The estimated effects are averages of heterogeneous individual treatment effects. The outcome is wage growth over a fixed window, and workers differ in how long they would have remained near the minimum wage within it. A worker who would have permanently stayed at the minimum wage throughout is exposed to the full increase, with 100% pass-through: absent employment effects (Section 4.2), their wage growth rises by the entire 5.1 percentage points at two years and 16.2 at four.[25] A worker who would have grown modestly stays close to the floor and is still strongly affected, while a worker who would have quickly outgrown the floor is barely affected. The estimated 4 and 9.5 percentage points average across all of these workers, and their shortfall relative to the full increases reflects how many workers outgrow the floor within each window, a channel documented in Section 2.4. While the empirical analysis

[25] With a constant real minimum wage, a worker permanently at the minimum would experience zero real wage growth. Under the actual minimum-wage path, their wage rises one-for-one with the floor: 5.1% by 2008 and 16.2% by 2010. The treatment effect on their wage growth is therefore the full cumulative increase.



(a) Two-Year: PRE vs POST

(b) Two-Year: Treatment Effect

(c) Four-Year: PRE vs POST

(d) Four-Year: Treatment Effect

Figure 6: Individual Wage Growth During the Reform: Two- and Four-Year Horizons

Notes: Panels (a,c) compare observed (POST: $\gamma_k^{(h)} + \delta_k^{(h)}$) and counterfactual (PRE: $\gamma_k^{(h)} + \delta_C^{(h)}$) mean log wage growth by initial 2006 wage bin for 2006–2008 and 2006–2010. Here, $\delta_C^{(h)}$ is the control-group shift. Panels (b,d) show their difference, $\beta_k^{(h)}$. Shading marks controls (initial hourly wage above €5). Intervals are 99% confidence intervals for the contrast, drawn around POST in the profile panels.

identifies average effects, the model of Section 5 captures their heterogeneity.

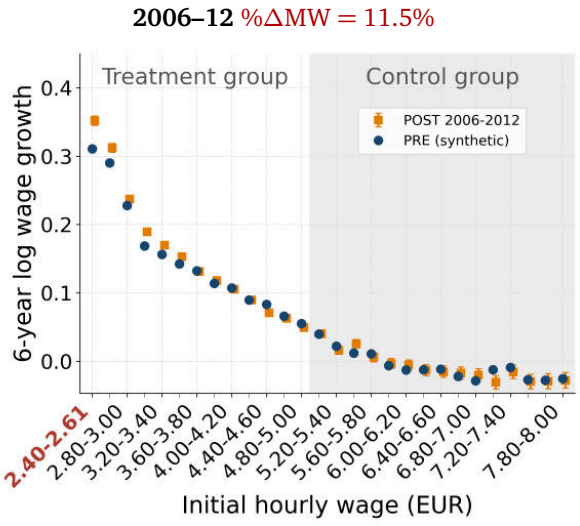
4.1.3 Individual Wage Growth After the Reform (2006–12 & 2006–14)

From 2010 onwards the successive minimum-wage increases stopped, but the real minimum wage remained at a permanently higher level. The cumulative real increase since 2006 was 11.5% by 2012 and 16.4% by 2014, so the 2014 minimum was essentially back at its 2010 peak (Table 1). This plateau allows me to trace whether the individual effects fade once the policy stops rising, and how quickly. Figure 7, panel (a), compares actual six-year individual wage growth (2006–2012) with its synthetic counterfactual, discussed in Section 3, and panel (b) shows the corresponding treatment effects; panels (c) and (d) repeat the exercise for the eight-year window (2006–2014).

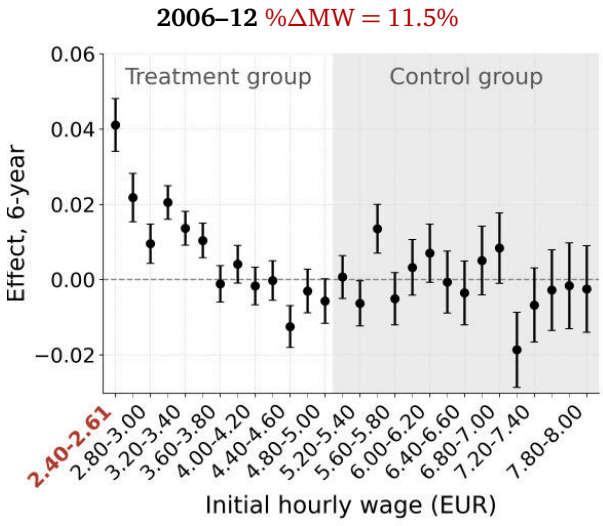
For workers earning the minimum wage in 2006, the treatment effect falls from its 9.5 percentage point peak in 2010 to 4.2 points by 2014: individual wages partially converge back toward their counterfactual paths even though the higher floor remains in place. The 2012 estimate of 4.1 points overstates the speed of this convergence, since it combines fading with the temporary dip of the real minimum during 2011–2013; the cleanest comparison is 2014, when the real minimum was back at its 2010 peak, so the fall from 9.5 to 4.2 points cannot be attributed to a retreat of the policy. To assess fading over time, a useful benchmark is to compare the average treatment effects to the full-passthrough effect, the treatment effect of a worker who remains at the minimum wage at every year, equal to the cumulative real minimum-wage increase at each horizon. For workers initially earning the minimum wage, pass-through declines to 77.5%, 62.9%, 37.7%, and 27.8% at the two, four, six, and eight-year horizons shown above, respectively. This declining pass-through is consistent with more workers growing out of minimum-wage exposure over time. Higher up the distribution, convergence appears complete: from around the 30th percentile of the initial wage distribution onward, the six- and eight-year effects are statistically indistinguishable from zero, against spillovers reaching the median during the active reform.

4.1.4 Average 2006–2014 Wage Trajectories by Initial 2006 Wage

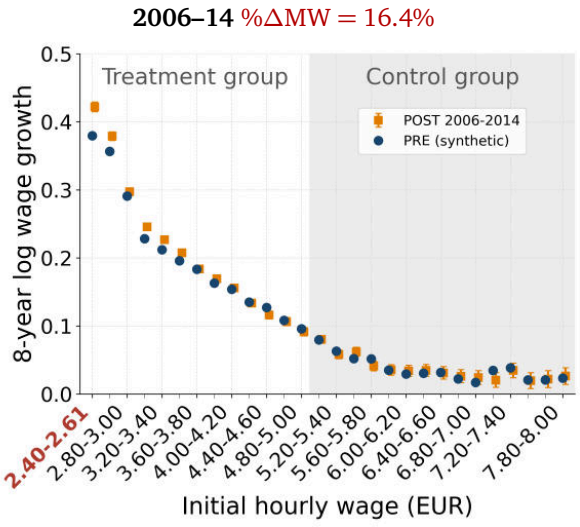
Figure 8 summarizes the results. Panel (a) plots the average log real wage path of workers earning the minimum wage in 2006 over 2006–2014, actual against counterfactual; panel (b) repeats the exercise



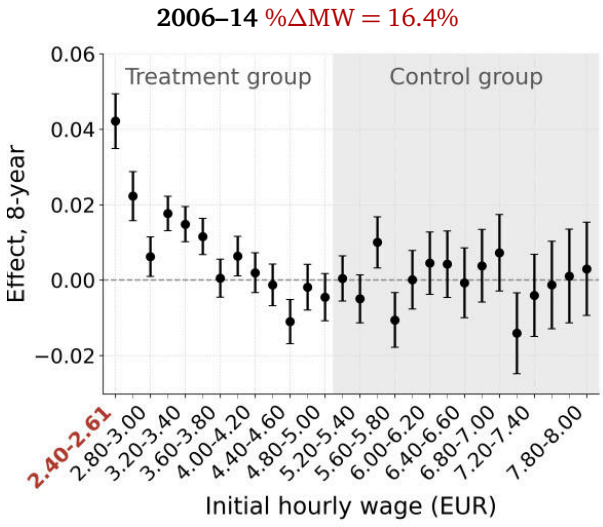
(a) Six-Year: Synthetic PRE vs POST



(b) Six-Year: Treatment Effect



(c) Eight-Year: Synthetic PRE vs POST

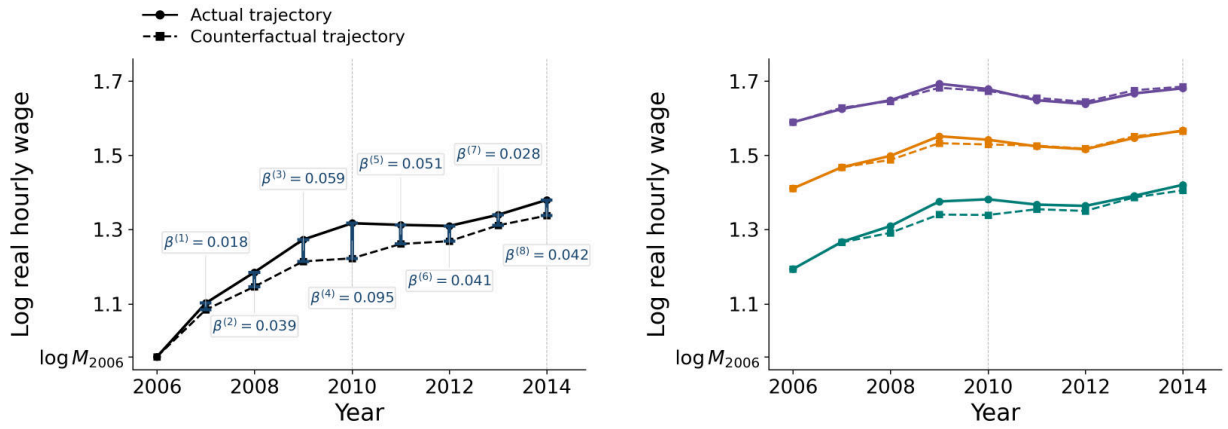


(d) Eight-Year: Treatment Effect

Figure 7: Individual Wage Growth After the Reform: Six- and Eight-Year Horizons

Notes: Panels (a,c) compare observed and synthetic counterfactual mean log wage growth for 2006–2012 and 2006–2014; panels (b,d) show treatment effects. Definitions and estimation conventions follow Figure 6. Shading marks controls (initial hourly wage above €5). Intervals are 99% confidence intervals for the contrast, drawn around POST in the profile panels.

for workers starting higher in the 2006 distribution. The vertical gap between the paths at each horizon is the average treatment effect estimated above: it opens while the minimum wage rises, peaks in 2010, and narrows during the freeze. Treatment effects are smaller for workers starting higher in the 2006 distribution and converge faster. Appendix B.1 gives the exact group definitions and initial wage anchors.



(a) 2006 Minimum-Wage Workers

(b) Higher-Wage Workers

Figure 8: Average 2006–2014 Wage Trajectories by Initial 2006 Wage: Actual vs Counterfactual

Notes: Solid circles show actual mean log real hourly wages; dashed squares show counterfactuals using Figure 6's control adjustment. Panel (a) follows 2006 minimum-wage workers; brackets mark treatment effects. Panel (b) shows separate paths near the 20th percentile in teal, the 40th percentile in orange, and the median in purple. Paths share their 2006 anchors; counterfactual growth is synthetic beyond four years.

4.2 Results on Other Outcomes

The fading of individual wage gains could reflect upward mobility, but it could also reflect adverse effects of the reform on the determinants of future earnings. More separations, fewer hours, or moves to lower-paying employers could weaken workers' subsequent wage trajectories. Conversely, reallocation to more productive firms (Dustmann et al., 2022) or greater labor supply (Coviello et al., 2022) could make the gains more persistent. I therefore examine these margins using the same difference-in-differences design as for wages.[26]

Table 3 and the appendix figures show the two- and eight-year results, while the discussion draws on estimates across all available short- and longer-term horizons and specifications. Short-term effects refer broadly to horizons of up to three years, and longer-term effects to horizons beyond three and up

[26]Changes in unobserved determinants, such as effort, cannot be ruled out.

to eight years. The detailed profiles by initial wage are in Appendix B.8.

Table 3: Short- and Longer-Term Effects of the Minimum Wage on Wages and Their Future Determinants

Outcome	Short-Term Effects	Longer-Term Effects	Appendix
Wages*	++	+	figures
Earnings*	++	+	figures
Hours	∼ 0	∼ 0 (+)	figures
Employment (separations)	∼ 0	∼ 0 (−)	figures
Firm Retention	∼ 0	∼ 0	figures
Revenue per Worker, New Firm	∼ 0	∼ 0	figures
Occupational Retention	∼ 0	∼ 0	figures
Industry Retention	∼ 0	∼ 0	figures
Educational Upgrading	(−)	(−)	figures
Job Promotion	∼ 0 (−)	(+)	figures

Notes: Short term: 2006–2008; longer term: 2006–2014, using the synthetic counterfactual. Wage bins are defined by initial 2006 hourly wages. Symbols summarize economic magnitude: ++ and + denote large and smaller positive wage or earnings gains; ∼ 0 denotes broadly small effects. Signs in parentheses flag departures from zero or estimates interpreted cautiously, rather than significance levels. This qualitative summary does not imply that every bin or horizon has a zero effect; employment and firm retention fall at the longer horizon for initial minimum-wage workers. Educational-upgrading and promotion estimates require the qualifications discussed in the text. *Main results presented above.

Employment and hours. The short-run effects on employment and hours are small and statistically indistinguishable from zero. These findings are consistent with Card and Krueger (1994); Manning (2021); Cengiz et al. (2019); Hirsch et al. (2015). Since the design follows incumbent workers, it measures separations and does not capture reduced hiring of entrants. In the short term, the point estimates tend to be negative for employment and positive for hours, but are generally statistically insignificant. The effects are generally more pronounced at longer horizons, consistent with workers and firms taking time to adjust, but remain relatively small across specifications. A complementary regional-age design suggests negative employment effects, as shown in Appendix B.9, consistent with Gopalan et al. (2021), who find no adverse effects on incumbents' employment retention but reduced hiring of new workers. Centeno et al. (2011) also find negative employment effects when studying the same Portuguese reform. Nevertheless, differential recession exposure makes identification more difficult when including entrants, as discussed in Section 3.2.

Other outcomes. Across specifications, the effects on firm retention, the revenue per worker of destination firms, and occupational and industry retention are negligible. Some specifications, mainly at longer horizons, find that treated workers are less likely to remain with their initial employer, and that

those who change firms move to employers with slightly higher revenue per worker. This pattern is consistent with Dustmann et al. (2022).[27] Because the reform was sizable but moderate, these margins may simply not have moved, as the main specification suggests. Overall, the results provide little evidence that the reform altered workers' career paths, through changes in employer quality or in occupational and industry mobility, in ways that could explain the fading wage gains.

Educational upgrading and job-title promotion are two exceptions, which I interpret with caution. The education estimates are strong, negative, and stable at both short and longer horizons. Nevertheless, I interpret these estimates cautiously, as the education classification changed precisely in 2006.[28] International evidence on educational responses is mixed. Higher minimum wages have been found to reduce school enrollment (Neumark and Wascher, 1995) and community-college enrollment, although with little evidence of reductions in completed certificates or degrees (Schanzenbach et al., 2026).[29]

To my knowledge, this paper is the first to estimate the effects of the minimum wage on job-title promotions. Quadros de Pessoal records each worker's job title as defined in the collective agreement covering the job, and each title carries its own negotiated wage floor. Employers also report the date of each worker's last promotion. I identify a promotion when a worker who remains with the same employer has a reported promotion within the horizon. The estimates are quite oscillatory across specifications and horizons.

4.3 Implications for Inequality: Individual vs Cross-Sectional Dynamics

Although individual-level effects tend to fade over time, the fading refers specifically to the effects on individual wage trajectories. Cross-sectional distributional effects can persist indefinitely, as new cohorts of low-wage workers continuously enter the labor market and become exposed to the higher floor, even when individual-level effects fade.

Consider the inequality measure p90/p10, the ratio of the 90th and 10th percentiles of the cross-

[27] This is a worker-level analysis. It remains possible that the policy drives out low-productivity or low-paying firms, including firms in low-wage industries, with their workers absorbed by more productive employers within the same industry. Related evidence documents firm-exit responses to minimum-wage increases (Mayneris et al., 2018), including in Portugal (Alexandre et al., 2022) and among restaurants (Luca and Luca, 2019). Lower firm retention contrasts, however, with Portugal and Cardoso (2006), who find that a higher youth minimum wage in Portugal reduced separations, and with the broader evidence that minimum wages reduce worker turnover (Dube and Lindner, 2024).

[28] The 2006 questionnaire introduced finer distinctions for non-tertiary post-secondary and postgraduate qualifications. I aggregate these categories to match the earlier classification.

[29] Minimum wages may also reduce on-the-job training (Neumark and Wascher, 2001), although Acemoglu and Pischke (2003) and Arulampalam et al. (2004) find no evidence of a reduction.

sectional wage distribution. When the 10th percentile coincides with the minimum wage (which happens when more than 10% of workers earn the minimum) or sits very close to it, a one-percent permanent rise in the real minimum wage translates roughly one-for-one into a one-percent reduction in the p90/p10 ratio, assuming the p90 is unaffected by the policy. When the increase raises the minimum relative to the median, as in the Portuguese reform, this compression can be large and lasting, as documented in Engbom and Moser (2022).[30] Figure 2 shows the Portuguese version of this compression directly, comparing the full 2006 and 2010 cross-sectional wage distributions.

However, a worker's individual wage trajectory partially converges back toward its counterfactual path after the reform, so the cumulative lifetime wage gain from the policy is much smaller than the cross-sectional gain.

NPV impact on the wage trajectories of 2006 minimum-wage workers. A simple undiscounted NPV of the 2006–2014 average wage trajectory of 2006 minimum wage workers, shown in panel (a) of Figure 8, makes this concrete.[31] Exponentiating the annual mean log-wage paths and summing the resulting wage indices separately for each trajectory gives an undiscounted cumulative gain of approximately 4.4%.[32] This gain is approximately 44% of the roughly 10% cumulative gain for a worker who remains at the minimum wage throughout the period.[33]

The gap between the cross-sectional and lifetime effect would widen further if one extends the horizon beyond the 8-year horizon (2006-2014). If individual exposure to the minimum wage continues to weaken over time, as Figure 3 suggests, the NPV impact for a worker initially at the minimum wage may fall below 4.4%. By contrast, the effect for the permanent-MW worker would tend toward the

[30] This statement refers to a permanent rise relative to the median: if the economy experiences long-term wage growth even a permanent minimum-wage increase may eventually cease to have lasting effects on the lower percentiles of the cross-sectional distribution, as the minimum wage loses its bite on the lower tail. This is a small concern during our period of analysis of meager wage growth.

[31] Hours are normalized to one throughout: the longitudinal profile of mean monthly hours varies by only about 2.4% over ages 18–64 and 0.8% over ages 25–55 (Appendix B.7), and hours did not meaningfully respond to the reform (Table 3).

[32] Let w_t^A and w_t^C denote the actual and counterfactual wage indices obtained by exponentiating each annual mean log-wage value in the figure. Summing the resulting indices separately for each trajectory over 2006–2014 gives the undiscounted cumulative gain:

$$G = \frac{\sum_{t=2006}^{2014} w_t^A}{\sum_{t=2006}^{2014} w_t^C} - 1 \approx 0.044 = 4.4\%. \tag{3}$$

[33] The 2006–2014 effect on the permanent-MW worker is smaller than the eventual 16.2% because the minimum wage rose gradually from 2006 to 2010 rather than in a single jump; as the horizon extends, the average gap between the actual and constant-2006-MW path approaches 16.2%.

16.2% cumulative real minimum-wage increase.

5 A Wage Process with a Minimum Wage Floor

The empirical results show that the effects of the minimum wage on individual wage trajectories tend to fade over time. To better characterize this fading, I estimate a statistical model of individual wage dynamics that embeds an explicit minimum wage floor.

The model serves three main purposes. First, it extends the analysis beyond the eight-year empirical horizon, quantifying what the reform implies for lifetime earnings. Second, it uncovers the heterogeneity of individual trajectories behind the average trajectories the empirics identifies. Characterizing this heterogeneity is essential for getting the treatment effects right: the average effect mixes workers who would have remained at the floor with workers who would have quickly outgrown it, and the model recovers that mix to match the estimated effects and their fading. Third, it disentangles the forces that carry workers out of the minimum wage, such as lifecycle growth and reversal of negative shocks.

On this last point, a worker may be at the minimum wage because they were unlucky to start at a low-paying firm, because they are young and early in their career, or because their earning capacity is permanently low. These reasons are indistinguishable in a single cross-section but imply very different lifetime exposures to the policy; separating them is one of the model's central tasks.

5.1 Setup

The model has two building blocks: a *latent wage process* and a *wage floor function*. The latent wage process is specified in the spirit of Meghir and Pistaferri (2011), including components that the income-dynamics literature has found important for characterizing individual wage dynamics. The distinctive ingredient is the explicit wage floor, increasingly relevant as many countries have spent the past decade pushing minimum wages toward the median wage.

The latent wage, $Z_{i,t}$, is the wage that worker i would be paid in period t in the absence of the minimum wage. The wage floor function, $W(Z, M)$, maps the latent wage Z into the actual wage W, ensuring all workers are paid at or above the national minimum wage M. It allows for spillovers for workers above the new higher minimum wage, lifting pay at and near the bottom of the distribution. When Z is sufficiently high, the actual wage W equals the latent wage Z.

The latent wage dynamics are assumed to be invariant to a moderate rise in the floor.[34] Section 4.2 broadly supports this assumption. Accordingly, the model should not be extrapolated to increases far outside the range of the minimum wage variation.

Hours vary little over the life cycle and were not meaningfully affected by the reform, so I abstract from them and normalize them to one.[35] Accounting for hours would not materially change the results, and the normalization makes wages equal earnings, so the lifetime quantities below are lifetime earnings rather than ill-defined sums of wage rates.

5.1.1 The Latent Wage Process

Time is discrete and one period spans two calendar years, so $t = 0$ corresponds to 2006, $t = 1$ to 2008, and so on, aligning with the two-year empirical horizon and its multiples. Let $Z_{i,t} \geq 0$ denote worker i's latent wage in period t (in levels), and $z_{i,t} = \log Z_{i,t}$ the corresponding log latent wage. The log latent wage is the sum of five components: a permanent fixed effect, an individual-specific lifecycle profile, an aggregate time effect, and persistent and transitory stochastic components.

In the model, there is wage stickiness, so it is useful to separate two latent wages.[36] The *flexible* latent wage $Z_{i,t}^{\text{flex}}$ is the wage that would prevail absent stickiness, when pay is free to move with underlying productivity; it evolves every period, as specified in (5) below. Let $d_{i,t}$ be an update indicator, equal to one when the worker's latent wage resets to the flexible wage and zero when it remains unchanged. Then

$$d_{i,t} \sim \text{Bernoulli}(1-\theta), \qquad z_{i,t} = (1-d_{i,t})z_{i,t-1} + d_{i,t}z_{i,t}^{\text{flex}}, \qquad t \geq 1. \tag{4}$$

[34] The latent wage process is disciplined by wage dynamics observed under the stable pre-reform floor. These dynamics may themselves depend on the existence of a minimum wage. The assumption therefore concerns a moderate increase from the pre-reform floor, rather than a comparison with an economy without a minimum wage.

[35] The model also assumes continuous employment. This assumption is more consequential, since the probability of remaining employed varies with age and with the wage; the reform itself, however, left the separation margin unaffected (Section 4.2).

[36] The inaction probability helps fit the leptokurtic distribution of wage changes: most wages barely move from one period to the next, while a minority move a lot (Grigsby et al., 2021; Guvenen et al., 2021; Hazell and Taska, 2025). Large wage changes typically coincide with a promotion or a job change.

Individual dynamics. Let $\widetilde{a}_{i,t}$ denote normalized age.[37] The log flexible latent wage is the sum of the five components:

$$z_{i,t}^{\text{flex}} = \underbrace{\alpha_i}_{\text{fixed effect}} + \underbrace{g_i(\widetilde{a}_{i,t})}_{\text{lifecycle}} + \underbrace{\mu_t}_{\text{time effect}} + \underbrace{\omega_{i,t}}_{\text{persistent}} + \underbrace{\xi_{i,t}}_{\text{transitory}}. \tag{5}$$

The fixed effect α_i stays constant and normalized age advances deterministically, $\widetilde{a}_{i,t+1} = \widetilde{a}_{i,t} + 2$. The time effect μ_t is a single scalar common to all workers.

Permanent heterogeneity. The individual fixed effect α_i collects everything about a worker that translates into permanently higher or lower pay. Dimensions of heterogeneity the model does not spell out (gender, education, ability) are absorbed into α_i to the extent that they map into permanent pay differences. Workers with a sufficiently low α_i are more likely to be the model's *persistent* minimum wage workers. In the baseline specification, α_i follows a flexible sinh-arcsinh (Jones–Pewsey) distribution,

$$\alpha_i = \alpha_{\text{loc}} + \alpha_{\text{scale}} \cdot \sinh\left(\frac{\operatorname{arcsinh}(\zeta_i) + \alpha_{\text{skew}}}{\alpha_{\text{tail}}}\right), \qquad \zeta_i \sim \mathcal{N}(0,1), \tag{6}$$

whose four parameters control location, dispersion, skewness, and tail weight. This flexibility helps fitting the right-skewed wage distribution with a heavy upper tail, unlike a Gaussian distribution.

Lifecycle growth. I allow for heterogeneous income profiles, in the spirit of Guvenen (2009). Specifically, the individual lifecycle profile $g_i(\widetilde{a})$ consists of a common profile $g(\widetilde{a})$ and a worker-specific component $s_i h_\tau(\widetilde{a})$. The common profile $g(\widetilde{a})$ is modeled flexibly as a degree-nine polynomial expressed in a Hermite basis, with nine free parameters and normalized to $g(0) = 0$, allowing it to fit the shape of the wage–age profile and capture lifecycle wage growth, an important force behind upward mobility. The individual slopes s_i have mean zero and standard deviation σ_s, and are allowed to correlate with α_i, as explained in Appendix C.1. The function $h_\tau(\widetilde{a}) = \tau\left[1 - e^{-\widetilde{a}/\tau}\right]$ allows worker-specific differences in lifecycle growth to fade with age at a rate governed by τ.[38]

Transitory and persistent shocks. The two stochastic components separate transitory luck from persistent luck. Economically, a worker may be temporarily unlucky if they start at a low-paying firm but soon move on, or persistently unlucky if they remain at that firm for several years. The transitory compo-

[37] $\widetilde{a}_{i,t} = a_{i,t} - 18$, so $\widetilde{a}_{i,t} \in [0, 46]$, corresponding to ages 18–64, matching the empirical age range.

[38] In the limiting case $\tau \to 0^+$, $h_\tau(\widetilde{a}) \to 0$, so every worker shares the common mean profile $g(\widetilde{a})$. As $\tau \to \infty$, $h_\tau(\widetilde{a}) \to \widetilde{a}$, so individual profiles diverge linearly from the common profile according to their worker-specific slopes s_i.

nent is drawn independently every period, while the persistent component follows an AR(1). Following Arellano et al. (2017), I allow its persistence to depend nonlinearly on its lagged value, $\omega_{i,t-1}$:

$$\omega_{i,t} = \rho(\omega_{i,t-1})\,\omega_{i,t-1} + \nu_{i,t}, \qquad \rho(\omega) = \frac{\bar{\rho}}{1+\exp(-\eta\omega+\phi)}. \tag{7}$$

[39] With $\eta > 0$ in (7), persistence rises with the lagged persistent component: workers with a low $\omega_{i,t-1}$ mean-revert quickly, while workers with a higher $\omega_{i,t-1}$ hold their positions longer. This state dependence helps the model match the counterfactual wage-growth profile, which exhibits strong mean reversion at the bottom of the initial wage distribution and much flatter growth further up. Following Karahan and Ozkan (2013), I allow the standard deviations of both shocks to vary with age:

$$\nu_{i,t} \sim \mathcal{N}\big(0, \sigma_\nu^2(\tilde{a}_{i,t})\big), \quad \sigma_\nu(\tilde{a}) = \sigma_\nu e^{k_\nu \tilde{a}}, \qquad \xi_{i,t} \sim \mathcal{N}\Big(0, \sigma_\xi^2(\tilde{a}_{i,t})\Big), \quad \sigma_\xi(\tilde{a}) = \sigma_\xi e^{k_\xi \tilde{a}}. \tag{8}$$

Initial conditions. At $t = 0$, the log latent wage is given by:

$$z_{i,0} \;=\; \alpha_i \;+\; g_i(\tilde{a}_{i,0}) \;+\; \mu_0 \;+\; \omega_{i,0} \;+\; \xi_{i,0}. \tag{9}$$

The time effect is normalized to $\mu_0 = 0$, so that time effects are measured relative to the 2006 level. The remaining components are initialized as follows. I draw α_i from the permanent-heterogeneity distribution described above. I draw initial normalized age $\tilde{a}_{i,0}$ from the empirical 2006 age distribution. I set $\omega_{i,0} = 0$ for workers aged 18; for older workers, $\omega_{i,0}$ reflects persistent innovations accumulated since age 18.[40] The transitory component is drawn from the age-dependent distribution in (8).

5.1.2 The Wage Floor Function

The wage floor function converts latent wages into actual wages by applying the statutory minimum M_t and allowing for spillovers above it. The schedule has three regions: a binding region in which workers are paid the statutory minimum, a spillover region in which actual wages exceed latent wages, and an

[39] I set $\bar{\rho}$ to a value close to, but strictly below, one, ruling out unit-root behavior and keeping the persistent component mean reverting. The parameter η governs how persistence changes with the lagged component, while ϕ shifts the level of the persistence function. When $\eta = 0$, the process nests a standard AR(1) with constant persistence.

[40] For older workers, I construct $\omega_{i,0}$ by iterating an annual approximation to (7) from age 18 to their age in 2006, using persistence $\sqrt{\rho(\omega)}$ and innovation standard deviation $\sigma_\nu(\tilde{a})/\sqrt{2}$.

unaffected region in which actual wages equal latent wages:

$$W(Z,M)=\begin{cases} M, & Z\le Z_1,\\ M+(Z_2-M)H(u), & Z_1<Z<Z_2,\\ Z, & Z\ge Z_2,\end{cases} \tag{10}$$

$$H(u)=(1-q)u+q\big[(m_1-2)u^3+(3-m_1)u^2\big],$$
$$Z_1=\chi_1 M,\quad Z_2=\chi_2 M,\quad u=\frac{Z-Z_1}{Z_2-Z_1},\quad m_1=\frac{Z_2-Z_1}{Z_2-M}.$$

The spillover region flexibly blends a linear segment with a cubic Hermite segment, nesting the linear and nonlinear cases, with q governing the blend.[41] The middle region allows the model to match the spillovers documented in Section 4.1.2, whatever their source: firms preserving internal pay differentials, upward pressure on reservation wages, or the renegotiation of collectively bargained wage floors (Section 2.3). Figure 9 illustrates the schedule for a given statutory minimum.

Actual wages are observed with classical measurement error:

$$\ln W_{i,t}^{\text{obs}} \;=\; \ln W(Z_{i,t},M_t)+\varepsilon_{i,t},\qquad \varepsilon_{i,t}\sim\mathcal{N}(0,\sigma_\varepsilon^2). \tag{11}$$

Measurement error absorbs reporting error and smooths the model's mass at the minimum wage, accommodating the small share of recorded wages below the statutory floor.

5.2 Estimation

For each candidate parameter vector Φ, I simulate a panel of workers from $t=0$ (2006) onwards to mimic the empirical panel.[42] A simulated worker is represented by

$$\mathscr{S}_i=\Big(\alpha_i,\omega_{i0},\{\widetilde{a}_{it},\nu_{it},\xi_{it},\varepsilon_{it}\}_{t=0}^{T_i},\{d_{it}\}_{t=1}^{T_i}\Big)\longrightarrow\{z_{it}\}_{t=0}^{T_i}\xrightarrow{\{M_t^{(r)}\}}\{w_{it}^{(r)}\}_{t=0}^{T_i}.$$

[41] When $q=0$, the transition is linear. When $q=1$, the cubic joins the flat minimum-wage region with slope zero and the unaffected region with slope one; $m_1=(Z_2-Z_1)/(Z_2-M)$ ensures the latter condition. Intermediate values of q blend the two cases and generally retain kinks.

[42] The simulated panel abstracts from attrition. A natural next step is to model survivorship explicitly, allowing the probability of remaining observed in the panel to vary with age and position in the wage distribution.

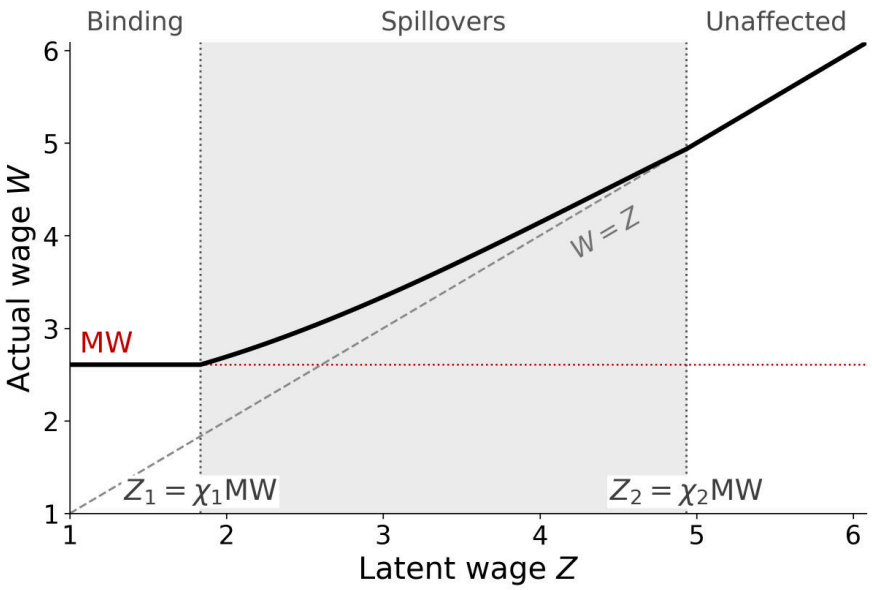


Figure 9: The Wage Floor Function $W(Z, M)$

Notes: The estimated wage-floor mapping at $M_{2006} = €2.6065$. Actual wages bind at $W = M$ for latent wages $Z \leq Z_1 = \chi_1 M$, rise through the spillover region between Z_1 and $Z_2 = \chi_2 M$, and equal latent wages above Z_2. The mapping uses the fitted nonlinear curvature in the spillover region. Both thresholds are marked by dotted vertical lines; the shaded band marks spillovers and the dashed diagonal is $W = Z$. This schedule excludes recording error.

Here, $\{M_t^{(r)}\}$ denotes the minimum-wage path under regime r. I map the same latent wage paths, z_{it}, into observed wages under two minimum-wage regimes: one follows the actual 2006–2014 path, while the other keeps the real minimum wage fixed at its 2006 level. These paths provide the model counterparts to the empirical treatment effects,[43] counterfactual wage growth, and other cross-sectional and individual dynamic moments, which I compare with their counterparts in the data.

The data moments include treatment effects and counterfactual wage-growth profiles (Section 4.1), the 2006 and 2010 wage distributions to capture the reform's cross-sectional compression (Section 2.3), and rates of growing out of the minimum wage (Section 2.4). I also match minimum-wage shares by age to capture age-specific incidence and the mean log-wage profile by age to capture lifecycle wage growth, an important force behind growing out of the minimum wage. The appendix reports the loss function and further estimation details, the complete set of parameter estimates, and the fit to matched and validation moments (Appendix C.4; Figure A22).

[43]Because the draws are held fixed across the two regimes, the only thing that differs is the policy, so the difference between them is the model's treatment effect, and the frozen-minimum-wage regime gives the model counterpart to the adjusted counterfactual growth in the data.

5.3 Policy Counterfactuals

I conduct three related policy simulations, all comparing a floor permanently fixed at its 2010 level with one permanently fixed at its 2006 level. I first compare the compression of current versus lifetime earnings. Mirroring the empirical exercise, I then examine the long-run gains of workers initially at the minimum wage and how those gains vary by starting age. Within this second exercise, I decompose the roles of lifecycle wage growth, persistent shocks, and transitory shocks in moving workers away from the minimum wage. In all three exercises, I set the time effects and measurement error to zero.[44]

5.3.1 Current versus Lifetime-Income Compression

This section presents the paper's main finding and answers the question *"Does the Minimum Wage Reduce Income Inequality?"* from both a current and a lifetime perspective.

A permanent minimum-wage increase delivers a permanent wage boost to workers who remain at the bottom of the wage distribution throughout their working lives. Concretely, a worker who spends their entire working life at the minimum wage receives not only a 16.2% increase in current earnings but also a 16.2% increase in lifetime earnings, absent employment effects. But most low-wage workers move through the distribution over time. The policy may therefore compress current earnings much more strongly than lifetime earnings.

Two main factors determine lifetime exposure to a minimum-wage increase: the fraction of the working life spent at the minimum wage, which depends strongly on income mobility; and the extent of spillovers above the floor, which determines whether workers remain exposed after moving beyond the minimum wage.[45]

I begin by comparing the effects of a permanently higher minimum wage on the distributions of current and lifetime earnings. For current earnings, I draw the model's time-zero distribution and apply the 2006 and 2010 minimum wages to that distribution. For lifetime earnings, I simulate entire individual earnings paths from ages 18 to 64 and sum wages without discounting.[46] The results barely

[44] Allowing aggregate wages to grow while holding the real floor fixed would cause workers gradually to outgrow the floor and would reduce their policy exposure. Classical measurement error is excluded because the exercises concern true wages, the welfare-relevant outcome, rather than recorded wages; the estimated measurement error is small in any case.

[45] The historical shock decomposition in Section 5.3.2 shows how different sources of upward income mobility matter.

[46] Following the convention adopted by the GRID project, I also measure lifetime earnings over the prime-age window of ages 25 to 55, for which continuous employment is a more plausible approximation because the wage process does not model labor-market entry or exit. The results barely change.

change over an empirically relevant range of discount rates.[47] Figure 10 plots the distribution of permanent earnings, the annualized version of lifetime earnings, to make its values directly comparable with current earnings. Workers who earn the minimum wage throughout their working lives under the low 2006 real minimum-wage regime therefore have permanent earnings equal to the minimum wage of approximately €2.60 per hour.

The effects of the policy differ sharply across the two distributions. Workers in the bottom 2.5 percent of the current-earnings distribution—approximately corresponding to the mass of workers at the minimum wage in the model's time-zero 2006 distribution—receive the full 16.2% contemporaneous increase. By contrast, among workers in the bottom 2.5 percent of the baseline lifetime-earnings distribution, mean lifetime earnings rise by 7.3%. Note that for the poorest 1 percent of workers in terms of lifetime earnings, wages are at the minimum wage for around 30% of their working lives under the lower €2.60 minimum-wage regime (the 2006 minimum wage is around 55% of the median). This share rises to almost 50% under the higher regime (around 61% of the median).

I then shut down the spillover effects of the minimum-wage increase and find that the same workers' mean lifetime earnings rise by just 3.8%, roughly half the gain with spillovers and about a quarter of the current gain of 16.2%. This suggests that in Anglo-Saxon labor markets, such as the United States and the United Kingdom, where smaller spillovers have been documented (Cengiz et al., 2019; Giupponi et al., 2024), the minimum wage may compress lifetime inequality less than its effect on the current distribution suggests. In Continental European labor markets such as Portugal, where stronger spillovers have been documented (Oliveira, 2023), or in labor markets with substantial informality such as Brazil, where spillovers are even stronger (Engbom and Moser, 2022), lifetime compression may be greater. Likewise, countries with greater upward income mobility than Portugal would tend to see smaller lifetime effects, since workers move out of the bottom more quickly.

5.3.2 Long-Run Wage Trajectories by Starting Age

I next use the fitted model to ask how the long-run effects vary with the age at which a worker is initially at the floor. Whereas the empirical analysis pools workers of different ages and therefore estimates an average effect across the age distribution, these simulations isolate the role of starting age. For each

[47] In the limit where future earnings receive no weight, the lifetime ranking becomes the current-earnings ranking and the effects at the bottom coincide.

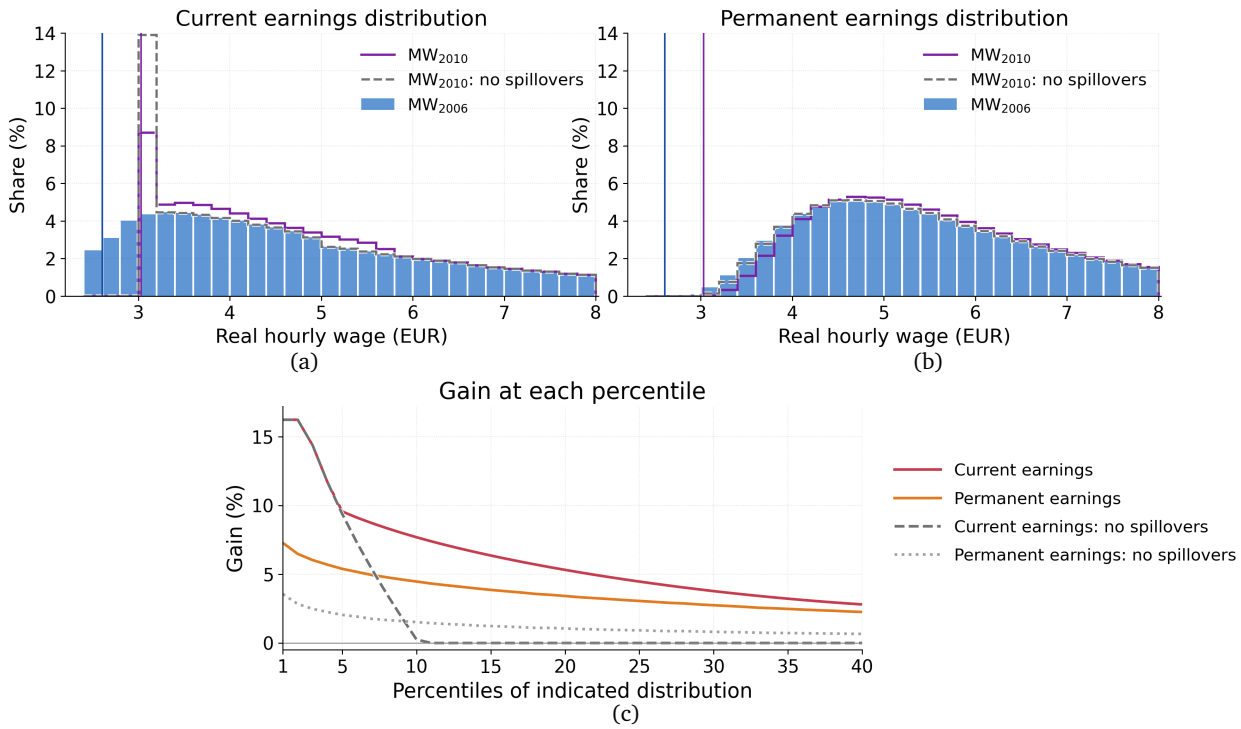


Figure 10: Minimum-Wage Compression of Current and Lifetime Income

Notes: Current hourly wages (panel a) and average lifetime wages over ages 18–64 (panel b) under the 2006 floor (blue), the 2010 floor with spillovers (purple), and the 2010 floor without spillovers (grey dashed). Current ages follow the empirical pre-reform age distribution; lifetime averages use 24 equally weighted biennial dates. Both experiments hold latent wage draws and parameters fixed across policy regimes and exclude recording error and time effects. Histogram shares use the full population denominator, including observations outside the displayed range. Panel c reports changes at percentiles 1–40; these compare distribution quantiles, whereas the text reports mean gains for a fixed group selected by baseline lifetime rank.

starting age $a_0 \in \{18, 25, 50\}$, I draw workers from the model at exactly that age, retain those initially earning at or very close to the minimum wage under the low-floor regime, and follow the same selected workers, with the same shocks, under the permanent low- and high-floor regimes defined above.

The panels follow workers at two-year intervals over three age ranges: 18–64, the conventional prime-age window of 25–55, and 50–64, which isolates the remaining gains for older workers. Figure 11 plots mean log wages in the two regimes. For workers starting at age 18, the gap begins at full pass-through but closes almost completely by mid-career as most workers move above the floor. For workers starting at age 50, who have little lifecycle growth left, the gap remains much wider.

Historical Shock Decomposition. What drives workers' growing out of the minimum wage? Workers may initially earn the minimum wage because they are early in their careers, or because they have experienced an adverse wage shock. I use a historical shock decomposition to quantify how lifecycle wage growth and persistent and transitory wage movements contribute to the fading gains.

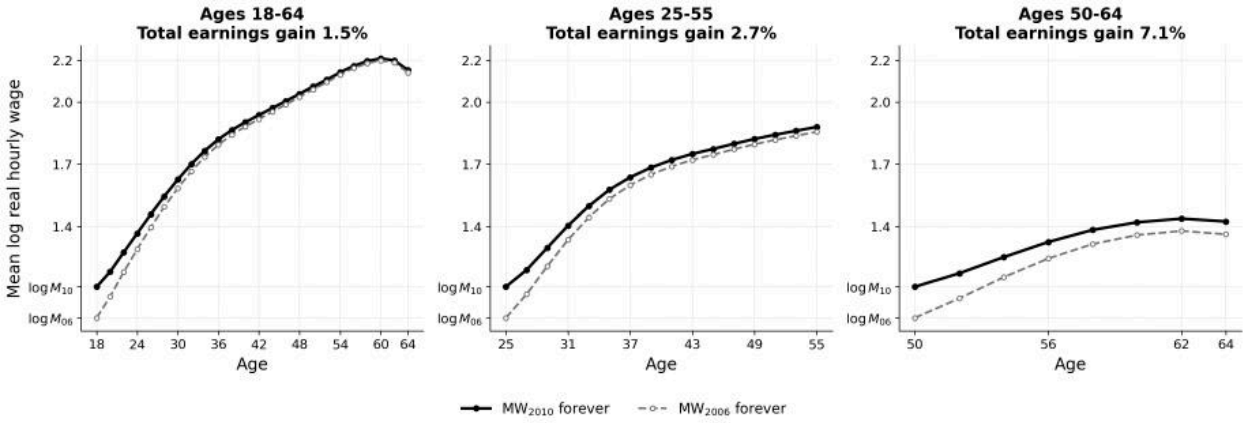


Figure 11: Long-Run Wage Paths of Workers Initially Earning the Minimum Wage

Notes: Workers initially earning at or very close to the low minimum wage are followed at two-year intervals over ages 18–64, 25–55, and 50–64. The same workers and shocks are compared under permanently high (M_{2010}, black) and low (M_{2006}, grey dashed) floors. Paths show mean log wages; titles report gains in total earnings calculated from wage levels. Selection is by initial wages, rather than lifetime-earnings rank. Recording error and time effects are excluded.

In the benchmark, I hold the lifecycle, persistent, and transitory components fixed at their initial values, $g_i(\widetilde{a}_{i,0})$, $\omega_{i,0}$, and $\xi_{i,0}$, respectively. For workers initially at the minimum wage at age 18, as well as for the other starting ages, I apply the wage-floor function to each latent wage path under both minimum-wage regimes and compare mean lifetime earnings. Workers initially at the minimum wage therefore remain there throughout the simulated horizon and receive the full 16.2% lifetime earnings gain. I then allow each component to evolve along its simulated path and measure its contribution to the fading lifetime gain.

Figure 12 illustrates the buildup of wage paths under the low 2006 minimum-wage regime. Starting from each worker's initial latent wage, held constant over time, I first allow lifecycle growth, then persistent wage movements, and finally transitory wage movements.

For workers initially at the minimum wage at age 18, allowing lifecycle wage growth reduces the lifetime gain from 16.2% to 3.2%, accounting for about 88% of the total fading. Adding persistent wage movements reduces the gain to 2.0%, and adding transitory wage movements reduces it to 1.5%, accounting for about 8% and 4% of the total fading, respectively.

However, this is only one ordering. Because the effects of the components need not be additive, their contributions can depend on the order in which they are introduced. I use a Shapley decomposition to average these contributions across all orderings. Appendix C.5 describes the calculation and reports the detailed results. Lifecycle wage growth accounts for around two-thirds of the fading gains for workers

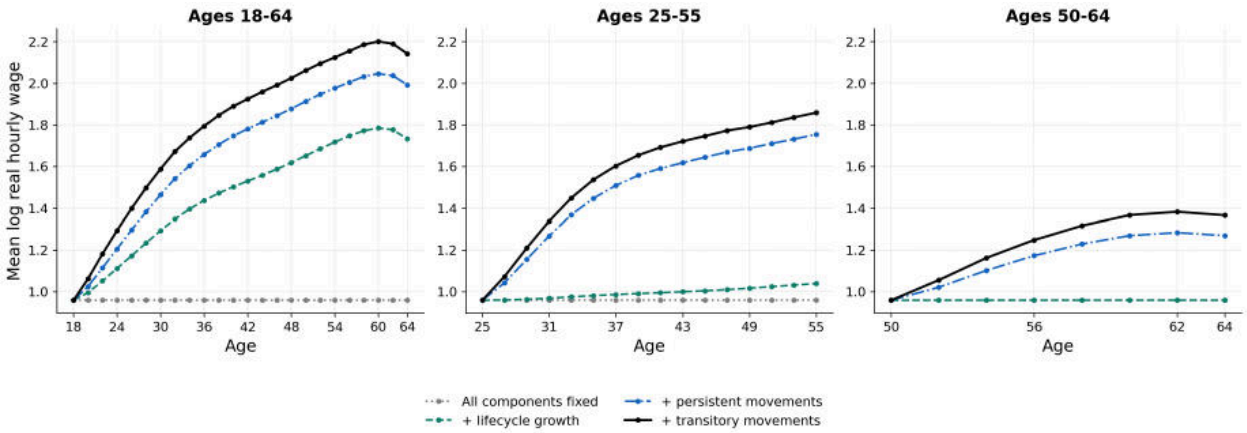


Figure 12: Historical Shock Decomposition: Building Up the Mean Wage Path

Notes: Mean log-wage trajectories under the 2006 minimum wage for workers initially earning at or very close to that minimum. The four curves show the frozen benchmark, lifecycle growth, lifecycle growth plus persistent movements, and the full wage process including transitory movements. Lifecycle growth combines the common and individual-specific profiles. Persistent evolution includes initial-state reversion and innovations. Parameters and wage-update draws are held fixed, with recording error and time effects excluded. The curves describe wage paths; Appendix Table 8 allocates the fading of gains in total earnings.

initially at the minimum wage at age 18, with a smaller contribution at older starting ages. For the age-18 cohort, persistent and transitory wage movements account for 21.5% and 11.9% of the fading gains, respectively.

6 The Optimal Minimum Wage Revisited

This section asks whether a considerably weaker lifetime equity motive necessarily implies a lower optimal minimum wage. In particular, Berger et al. (2025) show that the equity motive is central to the case for a high minimum wage. In their model, a utilitarian planner chooses a minimum wage of $11, approximately 54 percent of the median wage, whereas maximizing efficiency alone yields $7.35, approximately 36 percent of the median. Up to this level, there is no efficiency–equity trade-off, as raising the minimum wage improves both. Above this $7.35 level, the planner faces an efficiency–equity trade-off, and the optimal minimum wage is reached when the marginal equity gain exactly offsets the marginal efficiency loss. Other things equal, the larger the equity gains, the higher the optimal minimum wage. I introduce household idiosyncratic type risk into the model of Berger et al. (2025), generating income mobility that allows current gains to initially low-wage workers to fade on average.

6.1 The Permanent-Type Benchmark

In this subsection, I briefly explain the oligopsonistic benchmark model of Berger et al. (2025). Households view working for different firms as imperfect substitutes, modeled through a nested CES structure with greater substitutability between firms within local labor markets than across them. Consequently, a firm offering a higher wage does not attract all workers, giving firms some market power over the workers they employ. This market power allows firms to restrict employment and pay wages below workers' marginal products. A moderate minimum wage can therefore improve efficiency by narrowing these wage markdowns and increasing employment. Firms differ in productivity, which remains fixed over time. The extensions that follow modify only the household side, leaving the firm side unchanged; I refer the reader to Berger et al. (2025) for the full firm-side specification.

There is a unit mass of households, indexed by $i \in [0,1]$, and twelve household types, $z \in \{1,\dots,12\}$, with population shares π_z. Let z_i denote household i's permanent type. For each household i, type is fixed over time, so $z_{i,t+1} = z_{it}$; this assumption is relaxed in the subsections that follow. Types differ in productivity, preferences for leisure, and their shares of aggregate profits and capital ownership.

With permanent types, each household has no income fluctuations over time in steady state. And so its consumption, equal to labor earnings plus net capital income plus profits, is constant over time. There is therefore no need to borrow or save in bonds to smooth consumption, something households would like to do when their income fluctuates over time. To see this more clearly, I use an equivalent steady-state representation with households' physical capital held by a fund. Its risk-neutral manager discounts future payouts using the fixed discount factor β, and households' stakes in the fund are nontradable, both as in McKay et al. (2016). The fund rents capital to firms, collects capital income and profits, and makes investment decisions. In steady state, investment replaces depreciated capital, and the fund distributes net capital income plus profits to households according to the benchmark's type-specific ownership shares.

Household i chooses consumption c_{it}, bond holdings $b_{i,t+1}$, and a labor index N_{it}, taking firm-specific wages, hours constraints, and fund payouts as given. Period utility is separable in consumption and labor:

$$U_z(c,N) = u(c) - v_z(N). \tag{12}$$

Conditional on its permanent type, the household faces no idiosyncratic or aggregate shocks, so its problem is deterministic. For household i of permanent type z_i, the optimization problem is

$$\max_{\{c_{it},b_{i,t+1},N_{it}\}_{t\geq 0}} \sum_{t=0}^{\infty} \beta^t U_{z_i}(c_{it}, N_{it}), \tag{13}$$

subject to its borrowing limit, a no-Ponzi condition, and the budget constraint

$$c_{it} + b_{i,t+1} = (1+r)b_{it} + \underbrace{E_{z_i}(N_{it}; M)}_{\text{labor earnings}} + \underbrace{d_{z_i}(M)}_{\text{fund payout}}. \tag{14}$$

Here, M denotes the minimum wage. $E_{z_i}(N_{it}; M)$ denotes maximum labor earnings attainable at the chosen labor index.[48]

This representation exactly reproduces the steady-state equilibrium of the permanent-type benchmark, as shown in Appendix D.1. At the matched permanent steady state, each household i starting from $b_{i0} = 0$ optimally maintains zero bond holdings: $b_{i,t+1} = b'_{z_i}(0; M) = 0$.

6.2 Type Mobility

I now allow household i's type, z_{it}, to change over time while keeping the cross-sectional distribution of types fixed. Households thus churn within this distribution, generating income mobility. The index i continues to identify the same household as its type changes. Let P denote the transition matrix over the twelve types. I leave P general for now, requiring it to preserve those shares:

$$\pi^\top P = \pi^\top. \tag{15}$$

I analyze two extreme but useful cases that differ in households' ability to smooth consumption: hand-to-mouth households and complete consumption insurance.

[48] In a second stage, households allocate hours across firms within and across local labor markets to maximize earnings for the chosen labor index, subject to firm-specific hours constraints. Appendix D.1 defines this earnings function, retaining rationing under a binding minimum wage. See Berger et al. (2022) for a microfoundation of the nested-CES structure.

6.2.1 Hand-to-Mouth Households

Households are not allowed to borrow or save in bonds, so they must consume their income each period and are hand-to-mouth. Thus, $b_{it} = b_{i,t+1} = 0$, and their problem is effectively static, conditional on their current type z_{it}. In this experiment, productivity, preferences for leisure, and fund-payout shares follow current type.[49] Their budget becomes

$$c_{it} = E_{z_{it}}(N_{it}; M) + d_{z_{it}}(M). \tag{16}$$

At zero bond holdings, households' consumption and labor choices (policy functions) coincide with those in the permanent benchmark. With both type-specific choices and the distribution of types unchanged, the cross-sectional distributions of consumption, labor, and period utility, as well as aggregate equilibrium outcomes, remain the same as in the permanent benchmark.

In the permanent benchmark, workers' types are fixed, so current gains from a higher minimum wage are permanent and never fade. With mobility, a worker who is a low type today need not remain a low type tomorrow, so those initial gains need not last forever.

An Invariance Result Allowing households' types to change over time leaves the optimal minimum wage unchanged under a *utilitarian planner*, regardless of the degree of mobility. A utilitarian planner sums expected lifetime utilities across all households. Equivalently, it sums period utilities across households at each date, then adds these totals over time using the common discount factor. Mobility changes which households occupy different positions in the distribution, but leaves each date's total unchanged because the cross-sectional distribution of period utility across households remains the same. The sum of expected lifetime utilities therefore remains unchanged at each minimum wage.

Income mobility changes the dispersion of expected lifetime utilities across individuals, as the currently rich and poor need not remain in their positions, even when the sum of expected lifetime utilities is unchanged. A planner who values not only the sum of expected lifetime utilities but also their distri-

[49] The fund manager's fixed discount factor β and nontradable ownership claims, both as in McKay et al. (2016), together with payouts that follow current type, are strong simplifying assumptions. With idiosyncratic risk, households' stochastic discount factors need not equal β, as the ratio of future to current marginal utility of consumption need not be one. If fund claims were tradable, households might adjust their holdings to smooth consumption. As shown below, with fixed type shares and hand-to-mouth behavior, these assumptions preserve the cross-sectional allocation of outcomes and provide a useful illustration where a weaker lifetime equity motive need not lower the optimal minimum wage.

bution across individuals may therefore choose a different minimum wage.[50]

To show this more formally, let $V_i(M)$ denote individual i's expected discounted lifetime utility, conditional on their initial type:

$$\begin{aligned} V_i(M) &= \mathbb{E}\left[\sum_{t=0}^{\infty} \beta^t U_{z_{it}}\big(c_{it}(M), N_{it}(M)\big) \,\middle|\, z_{i0}\right] \\ &= \sum_{t=0}^{\infty} \beta^t \sum_{z=1}^{12} \pi^{(t)}_{z|z_{i0}} U_z\big(c_z(M), N_z(M)\big). \end{aligned} \tag{18}$$

Here, $\pi^{(t)}_{z|z_0} = (P^t)_{z_0,z}$ is the share of households initially in type z_0 that are in type z at date t. A utilitarian planner chooses the minimum wage to maximize the sum of expected discounted lifetime utilities:

$$W_U(M) = \int_0^1 V_i(M)\, di, \tag{19}$$

Substituting Equation (18) and grouping households by their initial type gives

$$\begin{aligned} W_U(M) &= \int_0^1 V_i(M)\, di \\ &= \int_0^1 \left[\sum_{t=0}^{\infty} \beta^t \sum_{z=1}^{12} \pi^{(t)}_{z|z_{i0}} U_z\big(c_z(M), N_z(M)\big)\right] di \\ &= \sum_{t=0}^{\infty} \beta^t \sum_{z=1}^{12} \left[\int_0^1 \pi^{(t)}_{z|z_{i0}}\, di\right] U_z\big(c_z(M), N_z(M)\big) \\ &= \sum_{t=0}^{\infty} \beta^t \sum_{z=1}^{12} \underbrace{\left[\sum_{z_0=1}^{12} \pi_{z_0}\, \pi^{(t)}_{z|z_0}\right]}_{(\pi^\top P^t)_z = \pi_z \text{ by stationarity}} U_z\big(c_z(M), N_z(M)\big) \\ &= \sum_{t=0}^{\infty} \beta^t \sum_{z=1}^{12} \pi_z\, U_z\big(c_z(M), N_z(M)\big) \\ &= \frac{1}{1-\beta} \sum_{z=1}^{12} \pi_z\, U_z\big(c_z(M), N_z(M)\big). \end{aligned} \tag{20}$$

Note that risk aversion is already incorporated into U_z. A utilitarian planner is sensitive to consumption

[50] Assuming positive expected lifetime utilities, an example is a Sen-type planner with the social welfare function

$$W_S(M) = \bar{V}(M)\big[1 - G_V(M)\big], \tag{17}$$

where $\bar{V}(M)$ is the sum of expected lifetime utilities and $G_V(M)$ is their Gini index.

dispersion across individuals because marginal utility of consumption is diminishing (and to dispersion in hours because labor disutility is convex). It is indifferent to the dispersion of expected lifetime utilities across individuals, provided their sum remains unchanged.

The Insurance Value of the Minimum Wage With uninsured idiosyncratic risk, the minimum wage's role partly shifts from redistribution to insurance. The IID case makes the distinction particularly clear. If $P_{zz'} = \pi_{z'}$, every individual faces the same distribution of future types, regardless of their initial type. In the long-run average without discounting, individuals spend the same fractions of time in each state and have the same average outcomes. There are then no permanently rich and permanently poor groups between which the policy redistributes. Its redistributive component instead provides insurance: it shifts resources from high-income to low-income states within individuals' lives. Initial types still affect discounted lifetime prospects because early periods receive greater weight.

The insurance value of the minimum wage has been studied by Adamopoulou et al. (2027), who show that wage floors limit how far wages can fall after negative firm productivity shocks. More generally, the redistribution and insurance roles of taxes and transfers have been studied extensively in the New Dynamic Public Finance literature (e.g., Farhi and Werning, 2013; Golosov et al., 2016) and in empirical work (Stepner, 2023). Here, I study these two roles in determining the optimal minimum wage.

6.2.2 Complete Consumption Insurance

I next consider the opposite polar case, in which households can fully insure consumption against idiosyncratic type risk. When households can smooth consumption, consumption is more closely tied to permanent income than to current income alone, making the policy's effects on permanent income differences particularly relevant. Insurance is purchased before the initial type draw, so it covers both initial differences and subsequent changes in type. With equal initial claims on aggregate resources and no aggregate uncertainty, consumption is the same across individuals and states, while labor supply remains type specific. Contingent transfers insure both labor earnings and fund payouts, and clear in the aggregate.[51]

[51]The comparison is between stationary equilibria. Insurance contracts condition on exogenous type, not chosen hours, so households retain their labor-supply incentives. Purchasing insurance only after the initial type draw would leave initial differences in lifetime resources uninsurable.

The optimal minimum wage is \$9.77, or 47.6% of the baseline median wage, compared with about \$11 (53.6%) in the hand-to-mouth benchmark.[52] Full insurance therefore lowers the optimum, but a substantial minimum wage remains desirable. The intuition is that the minimum wage no longer insures consumption against low-income states: private insurance already performs that role. It still changes employment, labor allocation, and firms' monopsony distortions. These effects determine the optimum when the consumption-insurance motive is absent. Moreover, equalizing consumption changes marginal utility and hence labor supply and equilibrium wages. The insured optimum therefore need not coincide with the \$7.35 efficiency optimum evaluated along the original, uninsured allocation path.

With stationary type shares, the fully insured allocation is independent of the persistence of individual types. Mobility changes who receives each period's income, while contingent claims absorb those changes in consumption. The contrast with the hand-to-mouth result shows why weaker lifetime inequality alone does not determine the optimal policy: households' ability to insure consumption matters as well.

6.2.3 Discussion

Income mobility makes permanent (or lifetime) income differences between individuals smaller than their current income differences. A utilitarian planner cares about consumption dispersion between individuals because of diminishing marginal utility. Whether consumption is more closely tied to permanent or current income depends on households' ability to smooth income fluctuations.

When individuals cannot smooth consumption, consumption tracks current income. The more income mobility there is, the smaller permanent (or lifetime) income differences between individuals become, reducing the lifetime equity motive. But, at the same time, the more their income paths fluctuate, increasing the insurance value of the policy. In the hand-to-mouth case, with cross-sectional allocations fixed, the decrease in redistributive value and the increase in insurance value offset each other, leaving the optimal minimum wage unchanged.

In the other extreme, when there is complete consumption insurance, consumption tracks ex-ante permanent income. The policy therefore has a weaker lifetime equity motive and provides no additional consumption insurance, since individuals are already privately insured. Nevertheless, even when

[52] Both ratios use the no-minimum-wage benchmark median of \$20.52, keeping the denominator fixed across policy regimes.

initial differences are insured, labor supply remains type dependent. Households have an incentive to work more in high-wage periods and less in low-wage ones, other things equal, reflecting an intertemporal labor-supply mechanism (MaCurdy, 1981). Equal consumption therefore does not recover the homogeneous-worker case.

These two extremes, with either hand-to-mouth households or complete consumption insurance, help us understand the different roles of minimum wage policy and their implications for the optimal minimum wage level. Nevertheless, an empirically relevant intermediate case is one of incomplete markets, where households can save but face limits on borrowing. In this case, households may optimally choose $b' \neq 0$, unlike in the permanent-type benchmark. Solving for equilibrium therefore requires characterizing the wealth distribution across households; this remains ongoing work.

7 Conclusion

This paper shows that the lifetime distributional effects of the minimum wage are considerably smaller than cross-sectional evidence suggests. Using a high-quality administrative panel covering private-sector employment in Portugal, I study a permanent 16.2% real minimum-wage increase and follow its effects on individual wage and earnings trajectories for eight years. Current gains are large but dissipate relatively quickly, even though the floor remains permanently higher. The evidence points to upward income mobility as the principal mechanism: workers grow out of minimum-wage exposure and generally spend only a small fraction of their observed working lives at the minimum wage. I find little evidence that adverse effects on the determinants of future earnings, such as employment, hours, or job promotion, explain this dissipation.

To quantify lifetime incidence beyond the empirical horizon, I estimate a statistical wage process with an explicit minimum-wage floor and spillovers on wages above it, matching empirical treatment effects and other moments characterizing individual wage dynamics. A permanent 16.2% increase raises earnings at the bottom of the current earnings distribution by the full 16.2%, whereas mean lifetime earnings for workers in the bottom 2.5 percent of the baseline lifetime earnings distribution rise by 7.3%. Income mobility limits lifetime incidence, but these gains remain sizable because the lowest lifetime earners still spend a meaningful share of their working lives near the floor. Spillovers further sustain the gains by raising wages above the minimum, allowing workers to benefit even after they grow

out of direct minimum-wage exposure. Without these spillovers, the lifetime gain falls to 3.8%.

Finally, I investigate whether a weaker lifetime equity motive implies a lower optimal minimum wage. I introduce household idiosyncratic risk into the model of Berger et al. (2025), generating income mobility. Income mobility reduces permanent income differences and thus weakens the policy's lifetime equity motive, but it also creates an insurance role when individuals cannot insure against income fluctuations. With hand-to-mouth households, the weaker redistributive role is offset by the insurance role, leaving the utilitarian optimal minimum wage unchanged. However, when individuals can fully insure consumption, the optimal minimum wage is lower. Smaller lifetime distributional effects do not, by themselves, establish that the minimum wage should be lower. Households' ability to smooth consumption is therefore central to this assessment.

References

Acemoglu, Daron and Jörn-Steffen Pischke, "Minimum Wages and On-the-Job Training," in "Worker Well-Being and Public Policy (Research in Labor Economics, Vol. 22)," Emerald, 2003, pp. 159–202.

Adamopoulou, Effrosyni, Francesco Manaresi, Omar Rachedi, and Emircan Yurdagul, "Minimum Wages and Insurance within the Firm," *Journal of Labor Economics*, 2027. Forthcoming.

Ahlfeldt, Gabriel M., Duncan Roth, and Tobias Seidel, "Optimal Minimum Wages," Discussion Paper 16913, CEPR 2022.

Albagli, Pinjas, Rui Costa, and Stephen Machin, "Minimum Wage Effects: Adjustment through Labour Market Dynamics and Alternative Work Arrangements," CEP Special Paper 49, Centre for Economic Performance, London School of Economics 2025.

Alexandre, Fernando, Pedro Bação, João Cerejeira, Hélder Costa, and Miguel Portela, "Minimum Wage and Financially Distressed Firms: Another One Bites the Dust," *Labour Economics*, 2022, *74*, 102088.

Arellano, Manuel, Richard Blundell, and Stéphane Bonhomme, "Earnings and Consumption Dynamics: A Nonlinear Panel Data Framework," *Econometrica*, 2017, *85* (3), 693–734.

Arulampalam, Wiji, Alison L. Booth, and Mark L. Bryan, "Training and the New Minimum Wage," *The Economic Journal*, 2004, *114* (494), C87–C94.

Ashenfelter, Orley, "Estimating the Effect of Training Programs on Earnings," *Review of Economics and Statistics*, 1978, *60* (1), 47–57.

_ and Štěpán Jurajda, "Minimum Wages, Wages, and Price Pass-Through: The Case of McDonald's Restaurants," *Journal of Labor Economics*, 2022, *40* (S1), S179–S201.

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

Barrela, Rodrigo, Eduardo Costa, and Pedro Portugal, "On the Asymmetrical Sensitivity of the Distribution of Real Wages to Business Cycle Fluctuations," IZA Discussion Paper 16911, IZA Institute of Labor Economics 2024.

Berger, David W., Kyle F. Herkenhoff, and Simon Mongey, "Labor Market Power," *American Economic Review*, 2022, *112* (4), 1147–1193.

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

Björklund, Anders, "A Comparison between Actual Distributions of Annual and Lifetime Income: Sweden 1951–89," *Review of Income and Wealth*, 1993, *39* (4), 377–386.

Bönke, Timm, Giacomo Corneo, and Holger Lüthen, "Lifetime Earnings Inequality in Germany," *Journal of Labor Economics*, 2015, *33* (1), 171–208.

Bossler, Mario and Thorsten Schank, "Wage Inequality in Germany after the Minimum Wage Introduction," *Journal of Labor Economics*, 2023, *41* (3), 813–857.

Bowlus, Audra J. and Jean-Marc Robin, "Twenty Years of Rising Inequality in U.S. Lifetime Labour Income Values," *Review of Economic Studies*, 2004, *71* (3), 709–742.

_ and _ , "An International Comparison of Lifetime Inequality: How Continental Europe Resembles North America," *Journal of the European Economic Association*, 2012, *10* (6), 1236–1262.

Butcher, Tim, Richard Dickens, and Alan Manning, "Minimum Wages and Wage Inequality: Some Theory and an Application to the UK," 2012. CEP Discussion Paper No. 1177, London School of Economics.

Caires, Filipe B., "Internal Organization of Firms and Minimum Wage Spillovers," 2025. Working paper, European University Institute.

Caliendo, Marco and Linda Wittbrodt, "Did the Minimum Wage Reduce the Gender Wage Gap in Germany?," *Labour Economics*, 2022, *78*, 102228.

Card, David, "Who Set Your Wage?," *American Economic Review*, 2022, *112* (4), 1075–1090.

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

_ **and Ana Rute Cardoso**, "Wage Flexibility Under Sectoral Bargaining," *Journal of the European Economic Association*, 2022, *20* (5), 2013–2061.

Cardoso, Ana Rute, "Long-Term Impact of Minimum Wages on Workers' Careers: Evidence from Two Decades of Longitudinal Linked Employer-Employee Data," *Scandinavian Journal of Economics*, 2019, *121* (4), 1337–1380.

Carneiro, Anabela, Paulo Guimarães, and Pedro Portugal, "Real Wages and the Business Cycle: Accounting for Worker, Firm, and Job Title Heterogeneity," *American Economic Journal: Macroeconomics*, 2012, *4* (2), 133–152.

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.

Centeno, Mário, Cláudia Duarte, and Álvaro A. Novo, "The Impact of the Minimum Wage on Low-Wage Earners," *Banco de Portugal Economic Bulletin*, 2011, *17* (3), 107–121. Autumn.

Clemens, Jeffrey and Michael Wither, "The Minimum Wage and the Great Recession: Evidence of Effects on the Employment and Income Trajectories of Low-Skilled Workers," *Journal of Public Economics*, 2019, *170*, 53–67.

Coviello, Decio, Erika Deserranno, and Nicola Persico, "Minimum Wage and Individual Worker Productivity: Evidence from a Large US Retailer," *Journal of Political Economy*, 2022, *130* (9), 2315–2360.

de Almeida Vilares, Hugo and Hugo Reis, "Who's Got the Power? Wage Determination and Its Resilience in the Great Recession," CEP Discussion Paper 1885, Centre for Economic Performance, London School of Economics 2022.

De Nardi, Mariacristina, Giulio Fella, and Gonzalo Paz-Pardo, "Nonlinear Household Earnings Dynamics, Self-Insurance, and Welfare," *Journal of the European Economic Association*, 2020, *18* (2), 890–926.

Derenoncourt, Ellora and Claire Montialoux, "Minimum Wages and Racial Inequality," *Quarterly Journal of Economics*, 2021, *136* (1), 169–228.

Dias, Mónica Costa and José Varejão, "Estudo de Avaliação das Políticas Ativas de Emprego," Relatório Final, Faculdade de Economia da Universidade do Porto / CEF.UP 2012.

Dickens, Richard and Alan Manning, "Spikes and Spill-overs: The Impact of the National Minimum Wage on the Wage Distribution in a Low-Wage Sector," *Economic Journal*, 2004, *114* (494), C95–C101.

Dube, Arindrajit and Attila Lindner, "Minimum Wages in the 21st Century," in Christian Dustmann and Thomas Lemieux, eds., *Handbook of Labor Economics*, Vol. 5, Elsevier, 2024, pp. 261–383.

_ , Laura Giuliano, and Jonathan Leonard, "Fairness and Frictions: The Impact of Unequal Raises on Quit Behavior," *American Economic Review*, 2019, *109* (2), 620–663.

Dustmann, Christian, Attila Lindner, Uta Schönberg, Matthias Umkehrer, and Philipp vom Berge, "Reallocation Effects of the Minimum Wage," *Quarterly Journal of Economics*, 2022, *137* (1), 267–328.

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

Farhi, Emmanuel and Iván Werning, "Insurance and Taxation over the Life Cycle," *The Review of Economic Studies*, 2013, *80* (2), 596–635.

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

Golosov, Mikhail, Maxim Troshkin, and Aleh Tsyvinski, "Redistribution and Social Insurance," *American Economic Review*, 2016, *106* (2), 359–386.

Gopalan, Radhakrishnan, Barton Hamilton, Ankit Kalda, and David Sovich, "State Minimum Wages, Employment, and Wage Spillovers: Evidence from Administrative Payroll Data," *Journal of Labor Economics*, 2021, *39* (3), 673–707.

Gottschalk, Peter and Robert Moffitt, "The Growth of Earnings Instability in the U.S. Labor Market," *Brookings Papers on Economic Activity*, 1994, *1994* (2), 217–272.

Gregory, Terry and Ulrich Zierahn, "When the Minimum Wage Really Bites Hard: The Negative Spillover Effect on High-Skilled Workers," *Journal of Public Economics*, 2022, *206*, 104582.

Grigsby, John, Erik Hurst, and Ahu Yildirmaz, "Aggregate Nominal Wage Adjustments: New Evidence from Administrative Payroll Data," *American Economic Review*, 2021, *111* (2), 428–471.

Guvenen, Fatih, "An Empirical Investigation of Labor Income Processes," *Review of Economic Dynamics*, 2009, *12* (1), 58–79.

_ , Fatih Karahan, Serdar Ozkan, and Jae Song, "What Do Data on Millions of U.S. Workers Reveal about Lifecycle Earnings Dynamics?," *Econometrica*, 2021, *89* (5), 2303–2339.

_ , Greg Kaplan, Jae Song, and Justin Weidner, "Lifetime Earnings in the United States over Six Decades," *American Economic Journal: Applied Economics*, 2022, *14* (4), 446–479.

_ , Luigi Pistaferri, and Giovanni L. Violante, "Global Trends in Income Inequality and Income Dynamics: New Insights from GRID," *Quantitative Economics*, 2022, *13* (4), 1321–1360.

Haefke, Christian, Marcus Sonntag, and Thijs van Rens, "Wage Rigidity and Job Creation," *Journal of Monetary Economics*, 2013, *60* (8), 887–899.

Hansen, Nikolaus, "The CMA Evolution Strategy: A Tutorial," arXiv:1604.00772 2016.

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

Hazell, Jonathon and Bledi Taska, "Downward Rigidity in the Wage for New Hires," *American Economic Review*, 2025, *115* (12), 4183–4217.

Heckman, James J., Robert J. LaLonde, and Jeffrey A. Smith, "The Economics and Econometrics of Active Labor Market Programs," in "Handbook of Labor Economics," Vol. 3, Elsevier, 1999, pp. 1865–2097.

Hirsch, Barry T., Bruce E. Kaufman, and Tetyana Zelenska, "Minimum Wage Channels of Adjustment," *Industrial Relations*, 2015, *54* (2), 199–239.

Hurst, Erik, Patrick J. Kehoe, Elena Pastorino, and Thomas Winberry, "The Macroeconomic Dynamics of Labor Market Policies," *Journal of Political Economy*, 2026. Forthcoming.

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

Karahan, Fatih and Serdar Ozkan, "On the Persistence of Income Shocks over the Life Cycle: Evidence, Theory, and Implications," *Review of Economic Dynamics*, 2013, *16* (3), 452–476.

Kerr, Sari Pekkala, William R. Kerr, and Louis J. Maiden, "Careers of Minimum Wage Workers," Working Paper 34405, National Bureau of Economic Research 2025.

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

Luca, Dara Lee and Michael Luca, "Survival of the Fittest: The Impact of the Minimum Wage on Firm Exit," Working Paper 25806, National Bureau of Economic Research 2019.

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

MaCurdy, Thomas E., "An Empirical Model of Labor Supply in a Life-Cycle Setting," *Journal of Political Economy*, 1981, *89* (6), 1059–1085.

Manning, Alan, "The Elusive Employment Effect of the Minimum Wage," *Journal of Economic Perspectives*, 2021, *35* (1), 3–26.

Martins, Pedro S., "30,000 Minimum Wages: The Economic Effects of Collective Bargaining Extensions," *British Journal of Industrial Relations*, 2021, *59* (2), 335–369.

Mayneris, Florian, Sandra Poncet, and Tao Zhang, "Improving or Disappearing: Firm-Level Adjustments to Minimum Wages in China," *Journal of Development Economics*, 2018, *135*, 20–42.

McKay, Alisdair, Emi Nakamura, and Jón Steinsson, "The Power of Forward Guidance Revisited," *American Economic Review*, 2016, *106* (10), 3133–3158.

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

Meghir, Costas and Luigi Pistaferri, "Earnings, Consumption and Life Cycle Choices," in Orley Ashenfelter and David Card, eds., *Handbook of Labor Economics*, Vol. 4B, Elsevier, 2011, chapter 9, pp. 773–854.

Neumark, David and William Wascher, "Minimum Wage Effects on Employment and School Enrollment," *Journal of Business & Economic Statistics*, 1995, *13* (2), 199–206.

_ and _ , "Minimum Wages and Training Revisited," *Journal of Labor Economics*, 2001, *19* (3), 563–595.

Oliveira, Carlos, "The Minimum Wage and the Wage Distribution in Portugal," *Labour Economics*, 2023, *85*, 102459.

Phelan, Brian J., "Hedonic-Based Labor Supply Substitution and the Ripple Effect of Minimum Wages," *Journal of Labor Economics*, 2019, *37* (3), 905–947.

Portugal, Pedro and Ana Rute Cardoso, "Disentangling the Minimum Wage Puzzle: An Analysis of Worker Accessions and Separations," *Journal of the European Economic Association*, 2006, *4* (5), 988–1013.

Renkin, Tobias, Claire Montialoux, and Michael Siegenthaler, "The Pass-Through of Minimum Wages into U.S. Retail Prices: Evidence from Supermarket Scanner Data," *Review of Economics and Statistics*, 2022, *104* (5), 890–908.

Rinz, Kevin and John Voorheis, "The Distributional Effects of Minimum Wages: Evidence from Linked Survey and Administrative Data," CARRA Working Paper 2018-02, U.S. Census Bureau 2018.

Schanzenbach, Diane Whitmore, Julia A. Turner, and Sarah Turner, "Raising State Minimum Wages, Lowering Community College Enrollment," *Review of Economics and Statistics*, 2026, *108* (5), 1448–1457.

Smith, Ralph E. and Bruce Vavrichek, "The Wage Mobility of Minimum Wage Workers," *Industrial and Labor Relations Review*, 1992, *46* (1), 82–88.

Solon, Gary, Robert Barsky, and Jonathan A. Parker, "Measuring the Cyclicality of Real Wages: How Important Is Composition Bias?," *Quarterly Journal of Economics*, 1994, *109* (1), 1–25.

Stepner, Michael, "The Insurance Value of Redistributive Taxes and Transfers," 2023. Working paper, University of Toronto, June 6, 2023.

Stewart, Mark B., "Wage Inequality, Minimum Wage Effects, and Spillovers," *Oxford Economic Papers*, 2012, *64* (4), 616–634.

A Institutions and Data

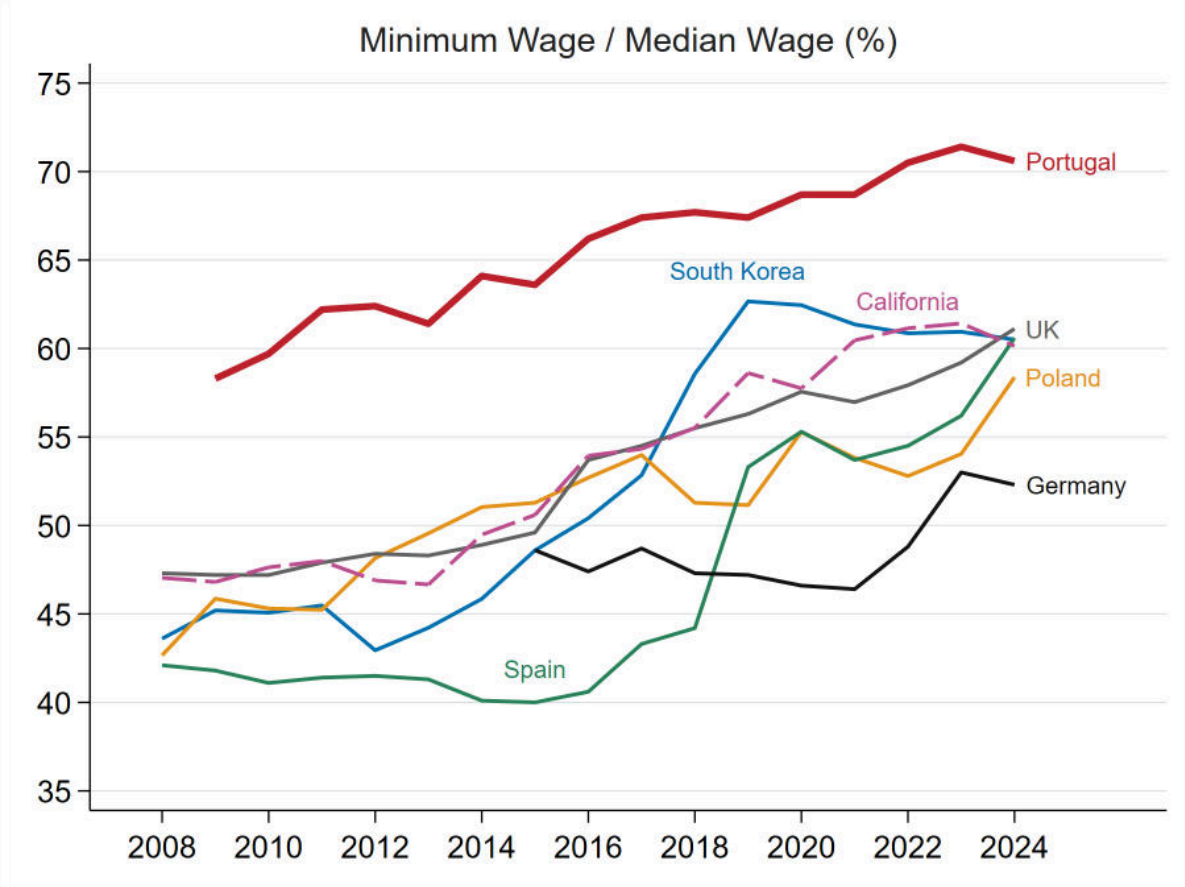


Figure A1: Minimum Wages Relative to Median Wages: International Evidence, 2008–2024

Notes: The Kaitz ratio is the minimum wage divided by the median wage, expressed as a percentage. Sources: Eurostat (Portugal, Spain, Germany, and the United Kingdom through 2019); OECD, "Minimum relative to average wages of full-time workers," median (South Korea, Poland, and the United Kingdom from 2020), accessed September 2026; California Department of Industrial Relations, "History of California Minimum Wage"; Economic Policy Institute, State of Working America Data Library.

A.1 Macroeconomic Trends over the Analysis Window

Figure A2 plots real GDP per capita, the median real hourly wage, and the Kaitz index over 1998–2015, all indexed to 2006, together with the recession chronology.

A.2 Percentiles of the 2006 Cross-Sectional Wage Distributions

Table 4 reports the percentiles of the 2006 real hourly wage distribution, as a reference for translating the euro thresholds used throughout the paper into positions in the distribution.

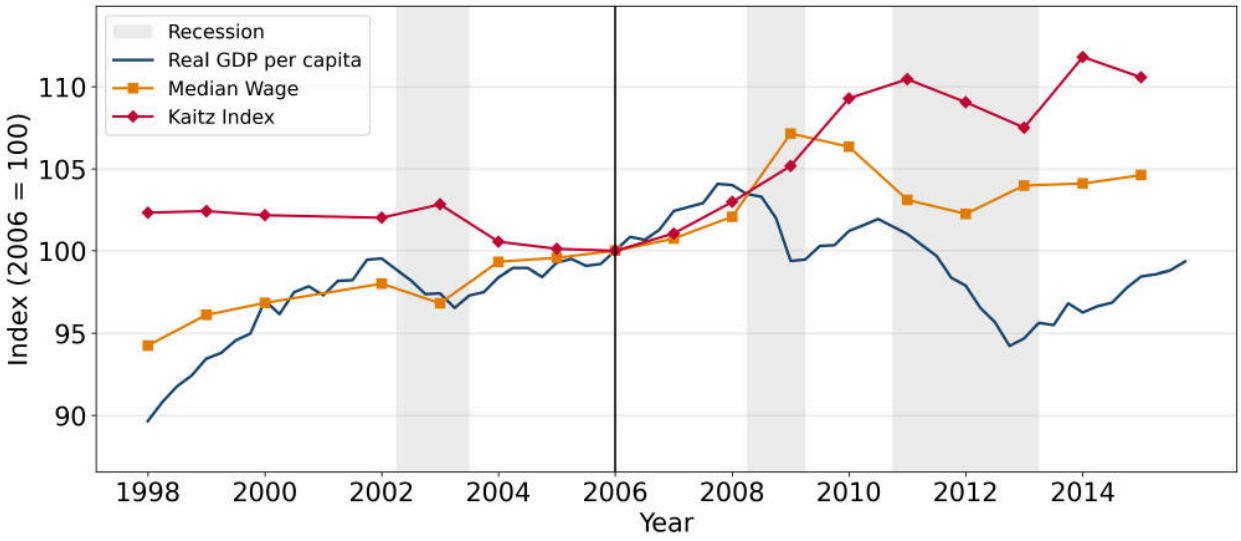


Figure A2: Macroeconomic Trends in Portugal, 1998–2015

Notes: All series are indexed to 2006 = 100; for GDP the base is the first quarter of 2006. Real GDP per capita: quarterly real GDP (chain-linked volume, seasonally adjusted; Banco de Portugal) divided by the resident population on January 1 (Eurostat), interpolated across quarters. Median wage: median real hourly total wage in the QP cross-section (October reference month), annual and dated at the beginning of the year, retaining the analysis sample restriction to real hourly wages above €2.45; the wage series therefore start in 2002. Kaitz index: ratio of the real hourly minimum wage to the median real hourly wage. Shaded bands mark recession quarters (Banco de Portugal chronology); the vertical line marks 2006, the start of the reform.

Table 4: Percentiles of the 2006 Hourly Wage Distribution

Percentile	5	10	15	20	25	30	35	40	45
Wage (€/h)	2.70	2.99	3.19	3.37	3.54	3.76	3.98	4.19	4.46
Percentile	50	55	60	65	70	75	80	85	90
Wage (€/h)	4.76	5.09	5.48	5.97	6.62	7.44	8.59	10.28	12.84

Notes: Real hourly total wage (2019 euros) in the cleaned 2006 QP cross-section: one observation per worker aged 18–64 with a valid positive wage (N = 2,054,749). No additional lower wage cutoff is applied. The 95th percentile is €17.79/h. The 2006 real minimum wage is €2.61/h; the €5/h control threshold is around the 54th percentile; the €8/h top bin edge is around the 78th percentile.

B Empirical Analysis

B.1 Wage Bin Construction

Wage bins are defined using the 2006 wage distribution, expressed in 2019 euros. The cutoff points are 2.4, 2.61, and then every €0.20 from 2.8 up to 8.0, yielding 28 bins. The real minimum wage in 2006 was approximately €2.605 per hour (in 2019 terms), so the first bin (2.4–2.61) straddles it; given very low non-compliance, it effectively captures all minimum wage workers. Results are robust to alternative

bin definitions, including extending the top of the grid from €8.0 (around the 78th percentile of the 2006 wage distribution) to €9.4 (the 83rd percentile). In particular, bin assignment is robust to using the base wage (base pay plus regular benefits and seniority payments, the components exposed to the minimum wage floor) rather than the total hourly wage, while the outcome variable remains the total wage, including bonuses.

In panel (a) of Figure 8, the minimum-wage group is the initial [€2.40,€2.61) group. In panel (b), the teal, orange, and purple series use the separate [€3.20,€3.40), [€4.00,€4.20), and [€4.80,€5.00) mean log-wage paths. These lie near the 20th, 40th, and 50th percentiles, respectively (Table 4). Initial log-wage levels are anchored at the log 2006 statutory floor for the minimum-wage group and the log group midpoint for the higher-wage groups.

B.2 General Specification and Regression Periods

The specification in Equation (1) of the main text restricts the analysis to the case where the POST period starts exactly at the reform date, equivalently the event-time distance from the reform to the POST period endpoint equals the horizon h. The general specification introduces an event-time parameter τ that allows the POST period to be shifted relative to the reform:

$$\begin{aligned} \mathscr{Y}_{i,t,t+h} = & \sum_{k=1}^{K} \gamma_k^{(h)} \mathbf{1}\{b_{k-1} < w_{i,t} \leq b_k\} + \lambda^{(h)} X_{i,t} \\ & + \text{Post}_t \cdot \left[\sum_{k=1}^{K} \delta_k^{(h,\tau)} \mathbf{1}\{b_{k-1} < w_{i,t} \leq b_k\} + \tilde{\lambda}^{(h,\tau)} X_{i,t} \right] + \varepsilon_{i,t+h}, \end{aligned} \tag{21}$$

where $\mathscr{Y}_{i,t,t+h}$ is a generic outcome over the h-year period. For example:

- Wage growth: $\mathscr{Y}_{i,t,t+h} = \log w_{i,t+h} - \log w_{i,t}$.
- Earnings growth: $\mathscr{Y}_{i,t,t+h} = \log \text{earnings}_{i,t+h} - \log \text{earnings}_{i,t}$.
- Hours growth: $\mathscr{Y}_{i,t,t+h} = \log \text{hours}_{i,t+h} - \log \text{hours}_{i,t}$.
- Employment status (conditional on i employed at t): $\mathscr{Y}_{i,t,t+h} = \mathbf{1}\{i \text{ employed at } t+h\}$.
- Firm retention (conditional on i employed at t): $\mathscr{Y}_{i,t,t+h} = \mathbf{1}\{i \text{ at same firm at } t+h\}$.

The PRE period runs from $2006-h$ to 2006; the POST period runs from $2006+\tau-h$ to $2006+\tau$. The main analysis fixes $\tau = h$ (POST period begins at the reform), and I write $\delta_k^{(h)} \equiv \delta_k^{(h,h)}$ and $\beta_k^{(h)} \equiv \beta_k^{(h,h)}$.

Placebo regressions set $\tau < 0$, shifting the POST period into the pre-reform period. Bin assignment is always by the initial hourly wage $w_{i,t}$, regardless of the outcome variable $\mathscr{Y}_{i,t,t+h}$, since policy exposure depends on where the worker sits in the wage distribution at the start of the period.

Table 5: Regression Windows by Horizon and Event Time

	h	τ	PRE period	POST period	Purpose
Main $h=2$	2	2	2004–2006	2006–2008	$\beta_k^{(2)}$
Main $h=4$	4	4	2002–2006	2006–2010	$\beta_k^{(4)}$
Main $h=6$	6	6	Synthetic	2006–2012	$\beta_k^{(6)}$
Main $h=8$	8	8	Synthetic	2006–2014	$\beta_k^{(8)}$
Placebo $h=2$	2	-1	2004–2006	2003–2005	$\beta_k^{(2,-1)} \approx 0$
Placebo $h=2$	2	-2	2004–2006	2002–2004	$\beta_k^{(2,-2)} \approx 0$
Placebo $h=1$	1	-1	2005–2006	2004–2005	$\beta_k^{(1,-1)} \approx 0$
Placebo $h=1$	1	-2	2005–2006	2003–2004	$\beta_k^{(1,-2)} \approx 0$
Placebo $h=1$	1	-3	2005–2006	2002–2003	$\beta_k^{(1,-3)} \approx 0$
Placebo $h=3$	3	-1	2003–2006	2002–2005	$\beta_k^{(3,-1)} \approx 0$

Notes: Each row corresponds to a separate regression stacking one PRE and one POST observation per worker. The PRE period always ends at 2006; the POST period ends at $2006+\tau$. Main regressions have $\tau = h$; placebos have $\tau < 0$. The $h=6$ and $h=8$ PRE growth rates are synthetic (Appendix B.3).

B.3 Construction of the Synthetic PRE Growth

The analysis panel begins in 2002, so the longest raw PRE window ending in 2006 covers four years. The six- and eight-year comparisons therefore use a *synthetic* PRE growth that chains each worker's actual four-year pre-reform growth with a bin-level pre-reform continuation of length $h-4$. For each worker i observed in both 2002 and 2006, the synthetic h-year growth is

$$\widetilde{g}_i^{(h)} \;=\; g_i^{2002\to2006} \;+\; \widehat{\gamma}^{(h-4)}_{b(i,2006)} \;+\; \widehat{\lambda}^{(h-4)} X_{i,2006}, \qquad h \in \{6,8\},$$

where $g_i^{2002\to2006} = \log w_{i,2006} - \log w_{i,2002}$ is the worker's actual first-leg growth, $b(i,2006)$ is the wage bin containing the worker's 2006 wage (the end of the first leg), and $\widehat{\gamma}_b^{(h-4)} + \widehat{\lambda}^{(h-4)} X_{i,2006}$ is predicted *pre-reform* growth over the next $h-4$ years, evaluated at the worker's 2006 wage bin and characteristics. The two-year continuation for $h=6$ is estimated from 2004–2006 growth, and the four-year continuation for $h=8$ from 2002–2006 growth. Workers whose 2006 wage lies above the top

estimated bin receive the mean growth of the control region.[53]

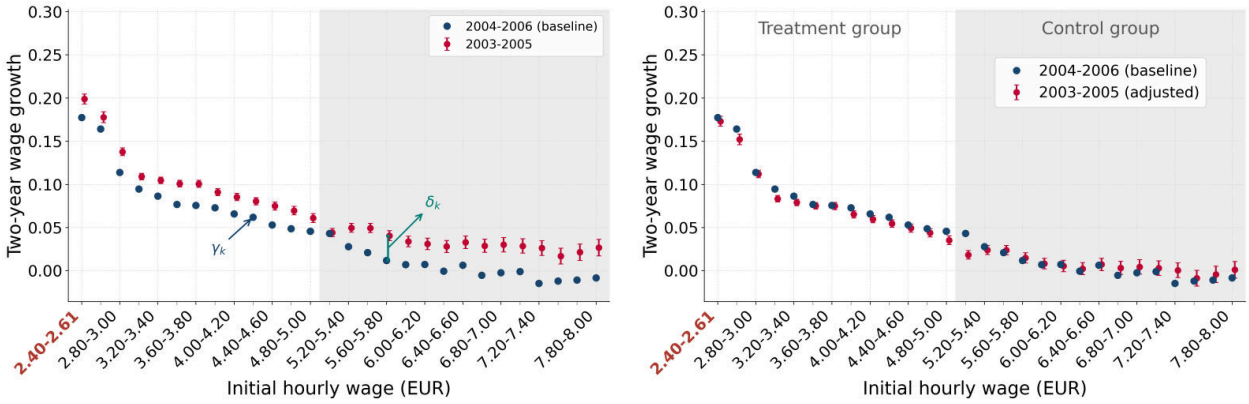
Both legs of the chain are measured under the constant pre-reform real minimum wage, so $\widetilde{g}_i^{(h)}$ is an h-year growth rate free of reform-period wage dynamics. The synthetic PRE observations are stacked with the actual POST observations (2006 → 2006+h growth) and run through Equation (1), with bin assignment by each observation's window-start wage: the 2002 wage for synthetic PRE rows and the 2006 wage for POST rows. The treatment effect $\beta_k^{(h)}$ therefore compares the h-year growth of bin-k workers with and without the reform, purged of the common control-region shift exactly as at the shorter horizons.

B.4 Illustration of the Control-Group Adjustment

The pre-reform parallel-shifts evidence in Section 4.1.1 reports growth rates after a control-group adjustment that subtracts a count-weighted mean across high-wage bins. Figure A3 illustrates the adjustment in two panels. The unadjusted panel (left) plots bin-level two-year individual wage growth for the analysis baseline 2004–2006 (blue dots, the $\gamma_k^{(2)}$ from Equation (1)) and for the placebo period 2003–2005 (orange squares with 99% confidence intervals). Above the control threshold (€5/h, vertical dashed line), the orange placebo sits roughly two percentage points above the blue baseline, the cycle-related component of growth, which differs across the two periods. The adjusted panel (right) keeps the blue baseline at its raw values but subtracts from the orange series the count-weighted high-bin mean of the placebo–baseline difference (the $\delta_C^{(h)}$ object from Equation (2)). The cycle-related offset disappears in the high bins; the two series now sit essentially on top of each other across the entire distribution, including the treatment region, which is what makes the parallel-shifts assumption credible at the placebo periods.

Figure A4 collects the adjusted two-year pre-reform stability profiles and their corresponding placebo treatment effects. The effects retain their estimated values and 99% confidence intervals around the zero-effect reference.

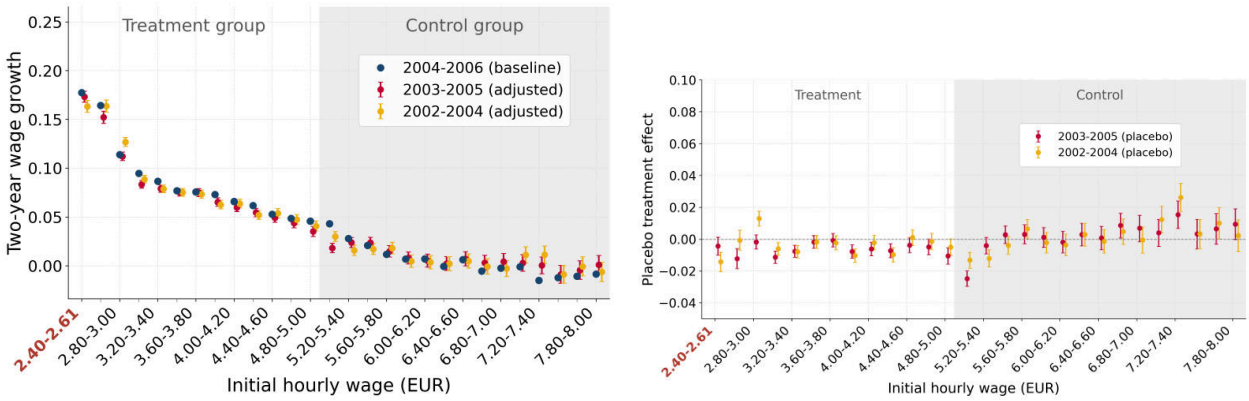
[53] The construction assumes that, conditional on the worker's 2006 wage bin and characteristics, their 2002 wage bin does not further predict the continuation. Persistent worker heterogeneity can violate this Markov restriction: upward movers from low initial bins may subsequently grow faster than other workers who reached the same 2006 bin. The resulting continuation would be understated for initially low-wage workers, biasing the six- and eight-year treatment effects upward. The documented fading is therefore conservative with respect to this approximation.



(a) Unadjusted (b) Adjusted

Figure A3: Illustration of the Control-Group Adjustment ($h = 2$)

Notes: Baseline bin-level two-year wage growth in 2004–2006 (navy), compared with the 2003–2005 placebo (crimson). Left: raw comparison; γ_k identifies a baseline coefficient and δ_k the between-window shift at one illustrative bin. Right: the baseline is kept raw and the placebo is net of its count-weighted mean difference in the control bins, as in Equation (2). The shaded region identifies initial wages above €5/h. The 99% intervals quantify uncertainty in the between-window contrast, rather than the raw growth levels.



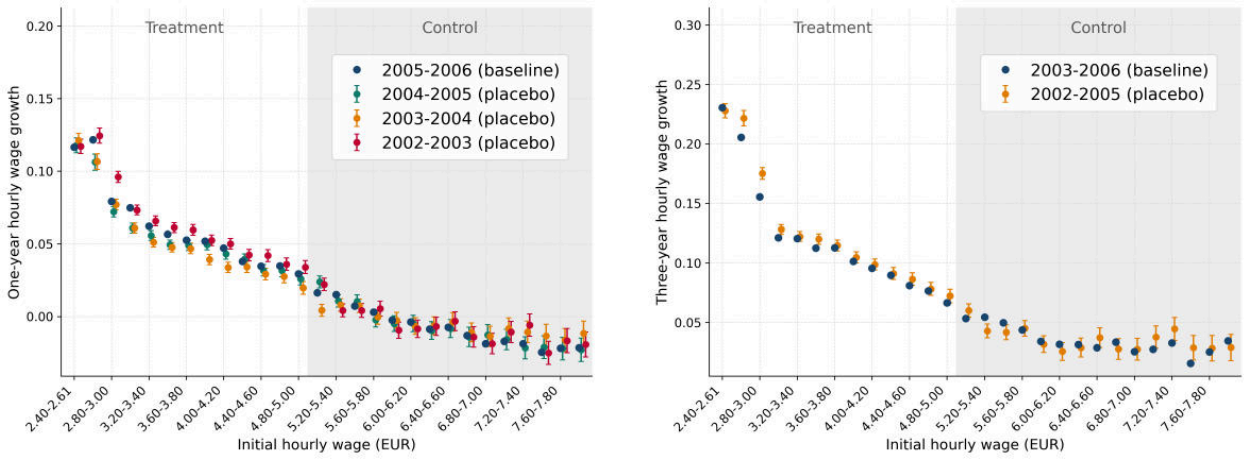
(a) Two-Year Growth Profiles (b) Placebo Treatment Effects

Figure A4: Individual Wage Growth Before the Reform: Two-Year Placebo Comparisons

Notes: Panel (a): the 2004–2006 baseline two-year wage-growth profile and the 2003–2005 and 2002–2004 placebo profiles, adjusted for their count-weighted mean difference from the baseline in the control bins. Panel (b): the corresponding placebo treatment effects, with 99% normal confidence intervals using bin-level standard errors. Initial total hourly wages above €5/h define the shaded control region.

B.5 Pre-Reform Wage Growth at One- and Three-Year Windows

For completeness I repeat the pre-reform diagnostics of Section 4.1.1 at one-year and three-year horizons. Not only is the shape of individual growth rates approximately stable across two-year windows, but also across one-year and three-year windows when the real minimum wage is unchanged, as seen in Figure A5: after the control adjustment, the placebo and baseline profiles lie on top of each other across the wage distribution. The left panel compares the one-year baseline (2005–2006) with the three earlier one-year windows in 2002–2005; the right panel compares the three-year baseline (2003–2006) with the 2002–2005 placebo.



(a) One-Year (b) Three-Year

Figure A5: Adjusted Wage Growth at One- and Three-Year Windows Across the Pre-Reform Period (2002–2006)

B.6 Growth Profiles in Other Recessions

This appendix shows the growth-rate profiles behind the other-recessions argument of Section 3.2: the adjacent one-year windows around COVID (2019–2020 vs 2020–2021) and around the 1993 recession (1992–1993 vs 1993–1994), and the adjacent two-year windows around the early-2000s slowdown (2000–2002 vs 2002–2004), for total and base hourly wages. In each figure the left panel plots the raw profiles and the right panel applies the control-region adjustment of Section 3; both outcomes are binned on the initial total hourly wage, so the total- and base-wage figures are comparable bin by bin.

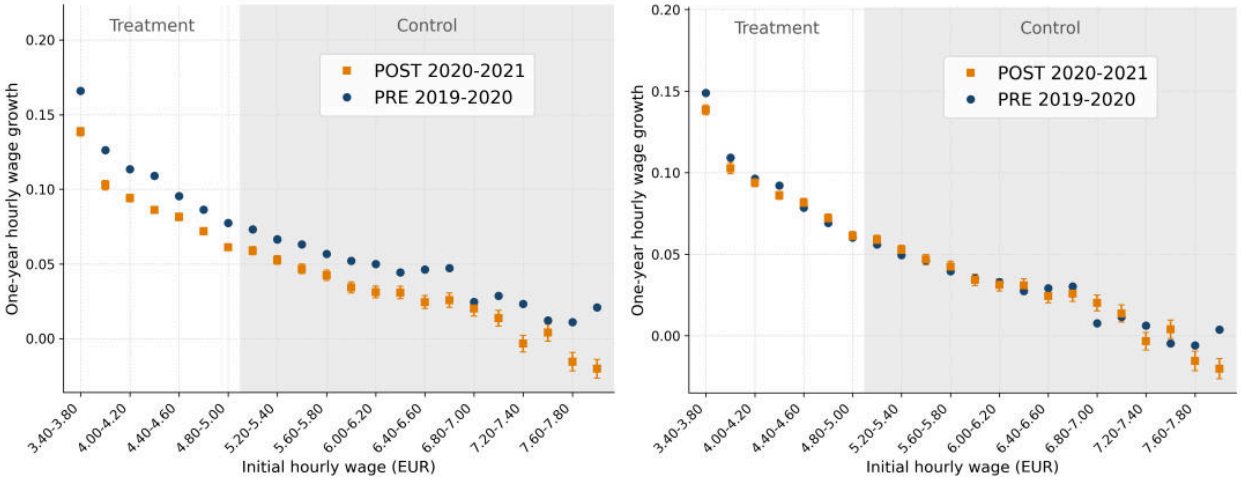


Figure A6: One-Year Total-Wage Growth Around COVID: 2019–2020 vs 2020–2021

Notes: One-year growth of the real total hourly wage by initial total-wage bin, 2019–2020 (navy) vs 2020–2021 (orange). Left: raw profiles. Right: after the control-region adjustment. Bands are 99% confidence intervals.

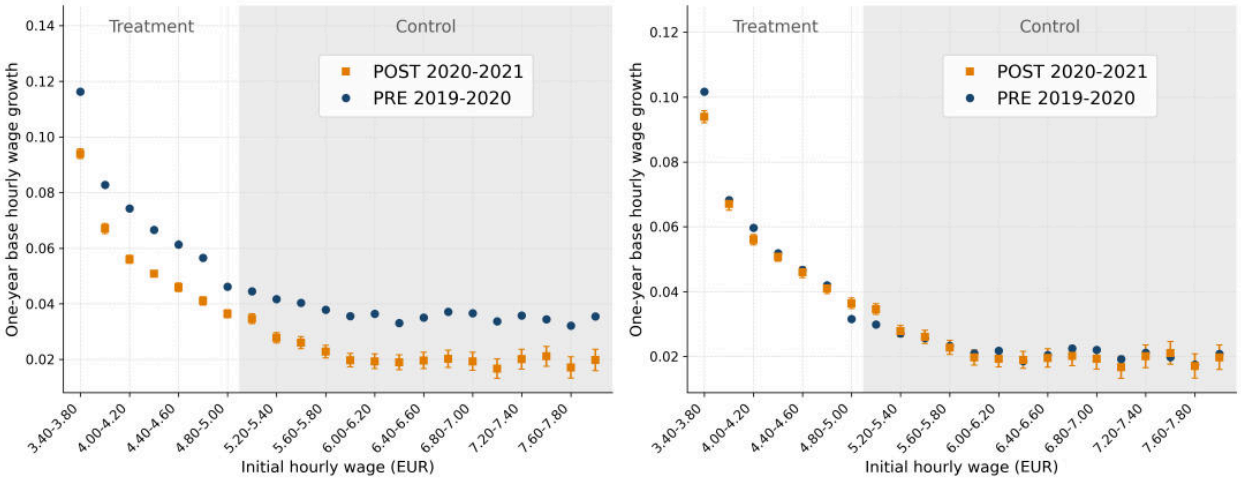


Figure A7: One-Year Base-Wage Growth Around COVID: 2019–2020 vs 2020–2021

Notes: One-year growth of the real base hourly wage by initial total-wage bin, 2019–2020 (navy) vs 2020–2021 (orange). Left: raw profiles. Right: after the control-region adjustment, using workers above €5/h as controls. Bands are 99% confidence intervals.

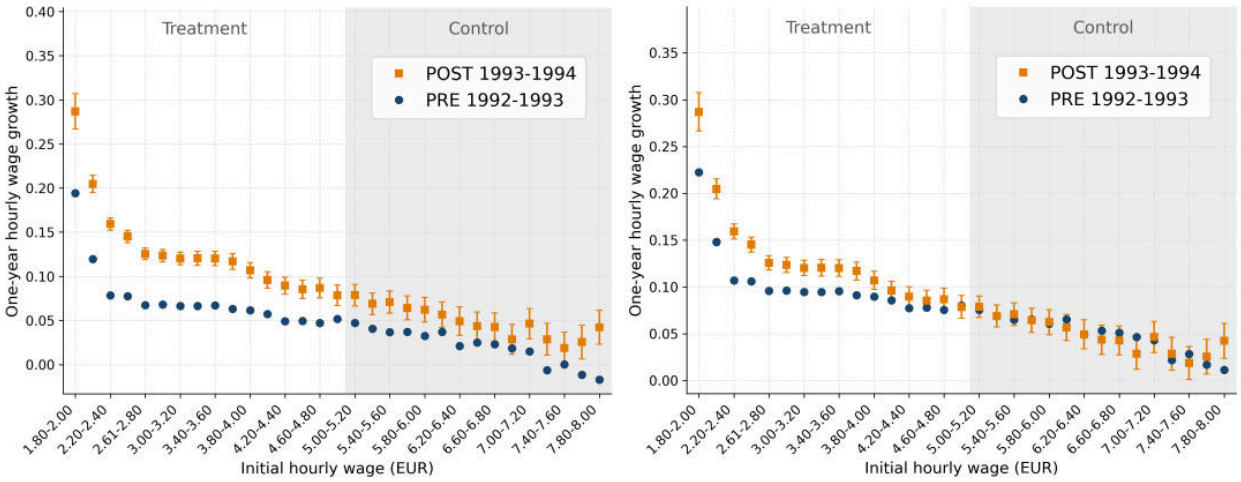


Figure A8: One-Year Total-Wage Growth Around the 1993 Recession: 1992–1993 vs 1993–1994

Notes: One-year growth of the real total hourly wage by initial total-wage bin, 1992–1993 (navy) vs 1993–1994 (orange). Left: raw profiles. Right: after the control-region adjustment. Bands are 99% confidence intervals.

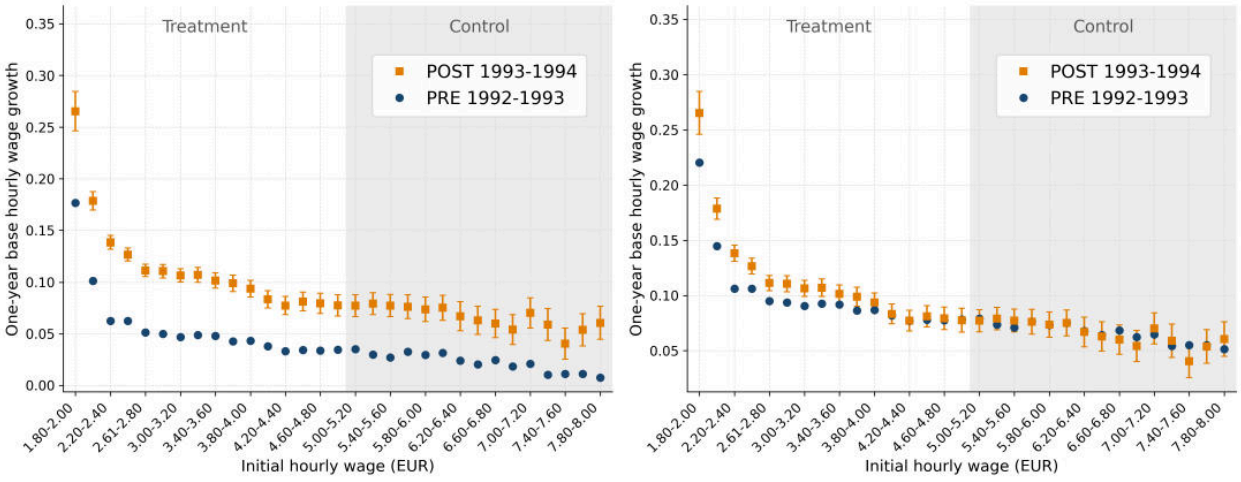


Figure A9: One-Year Base-Wage Growth Around the 1993 Recession: 1992–1993 vs 1993–1994

Notes: One-year growth of the real base hourly wage by initial total-wage bin, 1992–1993 (navy) vs 1993–1994 (orange). Left: raw profiles. Right: after the control-region adjustment. Bands are 99% confidence intervals. The real minimum wage fell by 6.1% across 1992–1993 and rose by 5.8% across 1993–1994, so the adjusted 1993–1994 profile is expected to display, and does display, a positive tilt at the bottom.

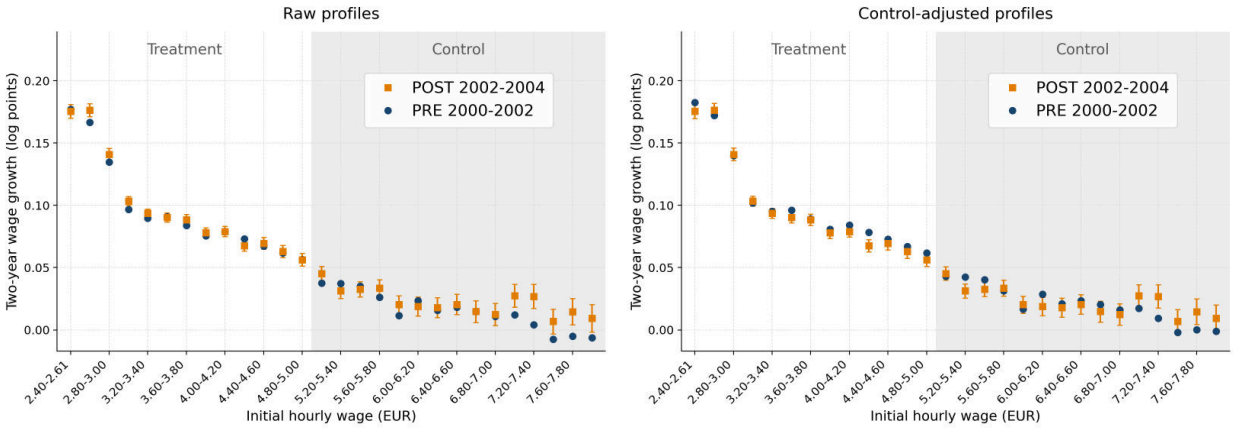


Figure A10: Two-Year Total-Wage Growth Around the Early-2000s Slowdown: 2000–2002 vs 2002–2004

Notes: Two-year log growth of the real total hourly wage by initial total-wage bin, 2000–2002 (navy) vs 2002–2004 (orange), without annualization. Left: raw profiles. Right: after the control-region adjustment, using workers with initial total hourly wages between €5 and €8 as controls. Bands are 99% confidence intervals for the between-window gap, with standard errors clustered by worker. The real minimum wage rose by 1.1% across 2000–2002 and fell by 0.1% across 2002–2004.

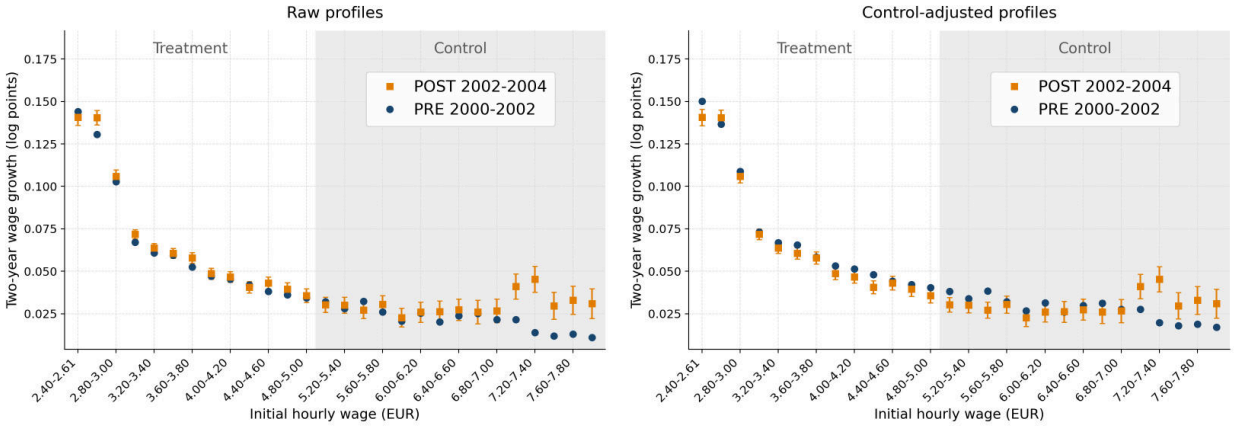


Figure A11: Two-Year Base-Wage Growth Around the Early-2000s Slowdown: 2000–2002 vs 2002–2004

Notes: Two-year log growth of the real base hourly wage by initial total-wage bin, 2000–2002 (navy) vs 2002–2004 (orange), without annualization. Left: raw profiles. Right: after the control-region adjustment, using workers with initial total hourly wages between €5 and €8 as controls. Bands are 99% confidence intervals for the between-window gap, with standard errors clustered by worker. Bins and sample conventions match Figure A10.

B.7 Hours over the Life Cycle

The longitudinal profile of mean monthly hours, estimated from within-worker changes controlling for calendar-year effects,[54] ranges from 171.8 to 175.8 hours over ages 18–64, a difference of about 2.4%. The range is even narrower over ages 25–55: 174.4 to 175.8 hours, or about 0.8%. Together with the small estimated hours response to the reform (Table 3), this supports normalizing hours to one in the lifetime calculations.

B.8 Other Outcomes

This section reports the bin-by-bin versions of the outcomes summarised in Table 3: earnings, hours, employment, firm retention, reallocation quality, occupational and industry retention, educational upgrading, and promotion. Each figure shows the short two-year horizon in the top row and the longer eight-year horizon in the bottom row. Within each row, the left panel plots the adjusted pre-period profile against the raw post-period profile with 99% confidence intervals, while the right panel plots the bin-level treatment effect with 99% confidence intervals. Earnings gains closely track hourly-wage gains. Most other effects are small, with the exceptions and qualifications discussed in Section 4.2.

Earnings. Figure A12 replaces the hourly-wage outcome with log earnings (hourly wage × hours). The earnings treatment-effect profile closely tracks the hourly-wage profile (Figure 6). Gains are strongest for workers initially earning the minimum wage and weaken moving up the initial wage distribution. These findings are consistent with a small response through hours worked.

Hours. Figure A13 tests the hours margin directly, replacing the outcome with the growth of monthly hours. The estimated effects are small and generally statistically indistinguishable from zero. The close correspondence between wage and earnings gains is therefore consistent with a limited response through hours worked.

Employment. Figure A14 shows the probability of remaining in observed private-sector employment at short and longer horizons. The short-term effect for initial minimum-wage workers is statistically indistinguishable from zero, while the longer-term estimate is negative. These results do not rule out

[54] I estimate

$$H_{i,t+1} - H_{i,t} = \delta_{a_{i,t}} + \tau_{t+1} + \varepsilon_{i,t},$$

using consecutive-year observations from 2002–2014, where $H_{i,t}$ denotes monthly hours and $a_{i,t}$ denotes age. Age and calendar-year effects are estimated jointly, with the year effects normalized to have a worker-pair-weighted mean of zero. The longitudinal profile cumulates the estimated age effects from age 18, anchored at mean monthly hours among 18-year-olds in the pooled 2002–2006 data.

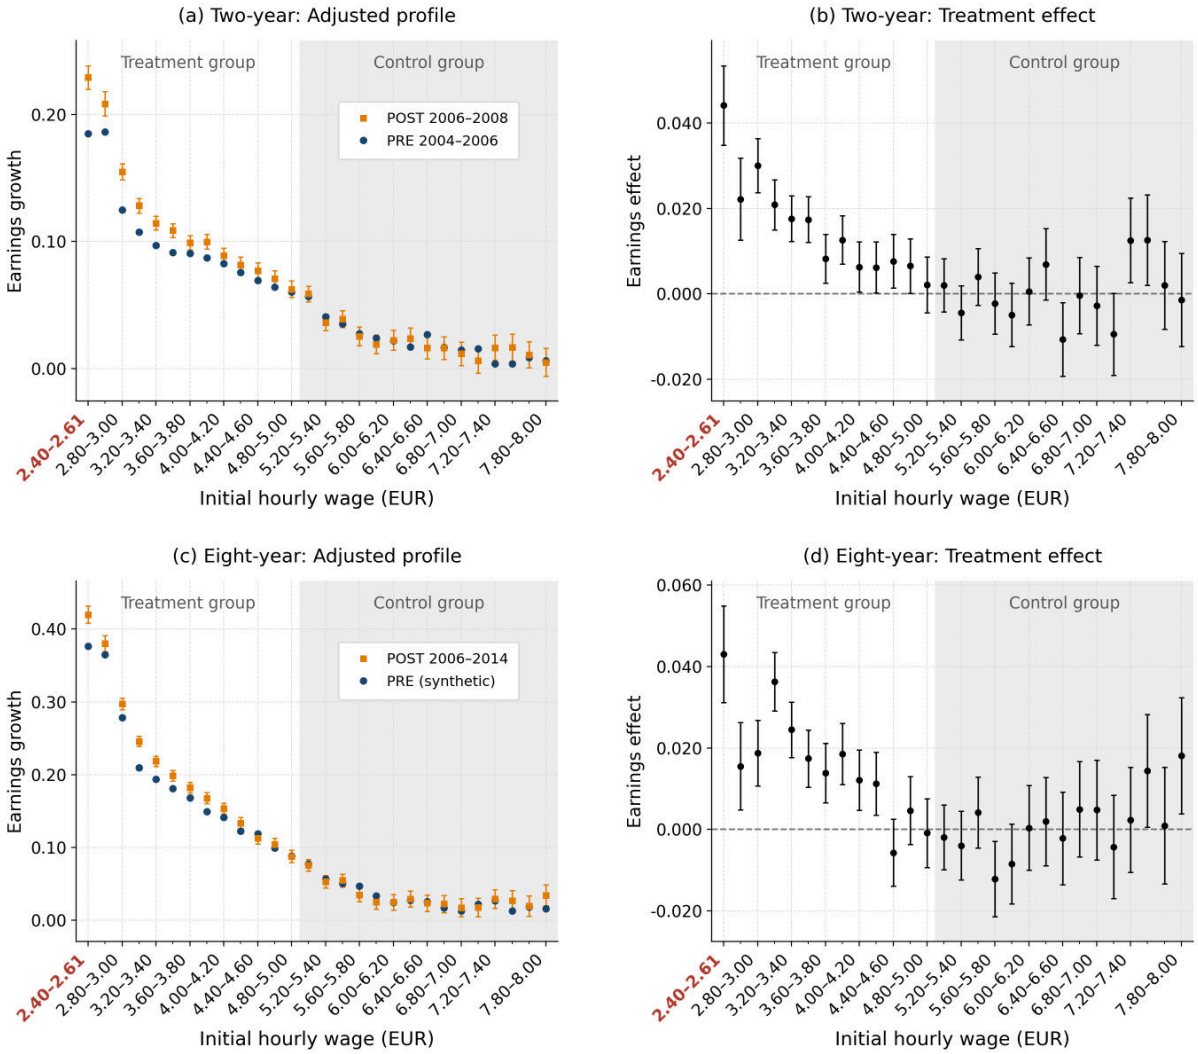


Figure A12: Earnings Growth (Wage × Hours): Short- and Longer-Term Effects

Notes: Outcome: log earnings growth (hourly wage times hours). Top: 2006–2008; bottom: 2006–2014 with a synthetic counterfactual. Left panels compare observed and adjusted counterfactual profiles; right panels show their difference. Initial 2006 total hourly wages define bins, with wages above €5/h forming the control group. Intervals are 99% normal confidence intervals for the estimated contrast, displayed around the observed profile and the treatment effect. Vertical scales vary across panels.

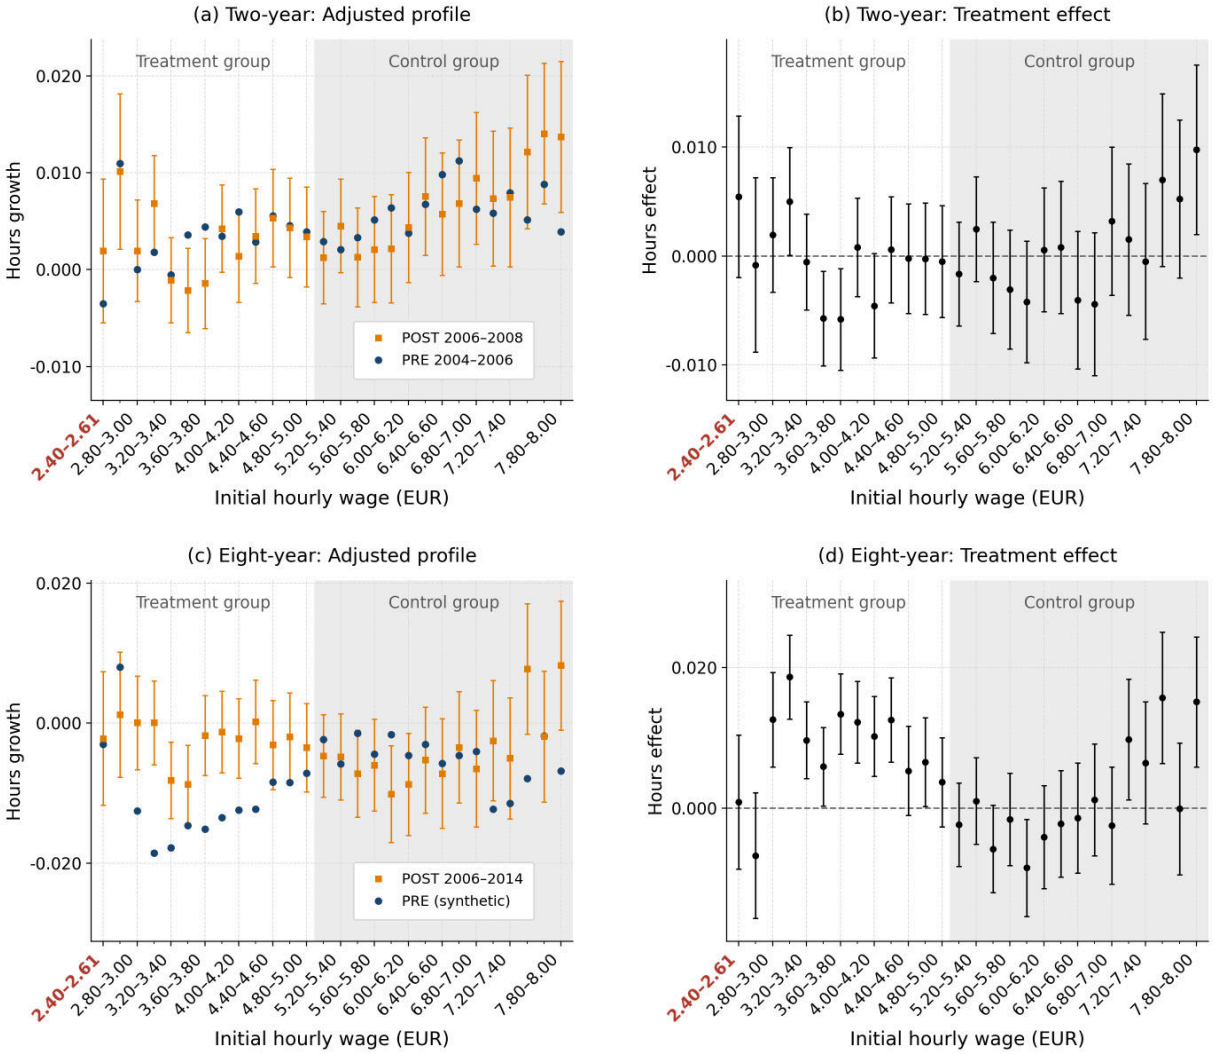


Figure A13: Hours Growth: Short- and Longer-Term Effects

Notes: Outcome: log growth of monthly hours. Top: 2006–2008; bottom: 2006–2014 with a synthetic counterfactual. Left panels compare observed and adjusted counterfactual profiles; right panels show their difference. Initial 2006 total hourly wages define bins, with wages above €5/h forming the control group. Intervals are 99% normal confidence intervals for the estimated contrast, displayed around the observed profile and the treatment effect. Vertical scales vary across panels.

selection among workers whose wages remain observed.

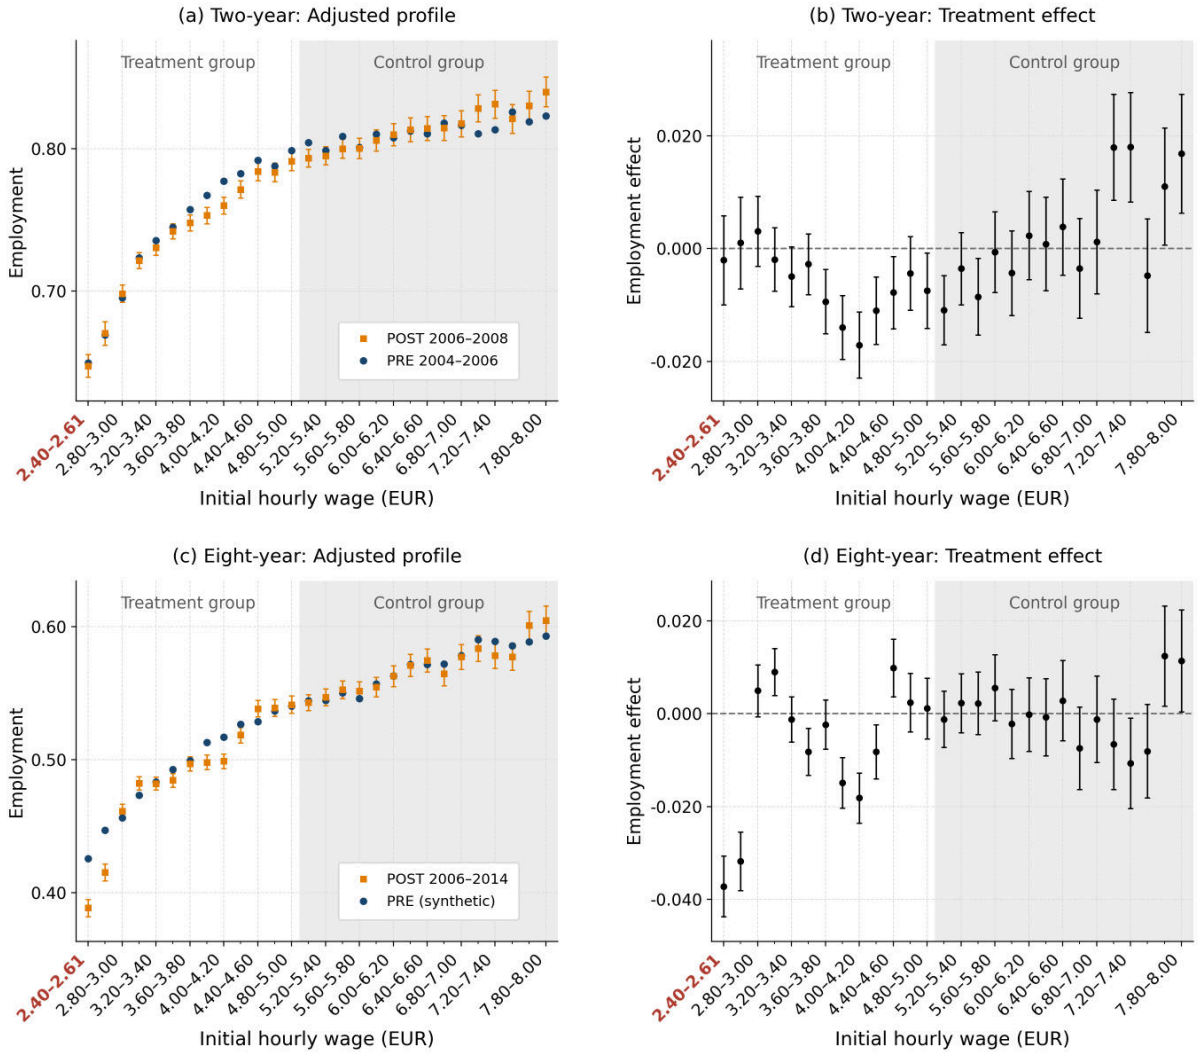


Figure A14: Employment (Extensive Margin): Short- and Longer-Term Effects

Notes: Outcome: remaining in observed private-sector employment at the endpoint, conditional on employment in 2006. Top: 2006–2008; bottom: 2006–2014 with a synthetic counterfactual. Left panels compare observed and adjusted counterfactual profiles; right panels show their difference. Initial 2006 total hourly wages define bins, with wages above €5/h forming the control group. Intervals are 99% normal confidence intervals for the estimated contrast, displayed around the observed profile and the treatment effect. Vertical scales vary across panels.

Firm retention. Figure A15 reports the probability of remaining at the same firm at short and longer horizons. The short-term effect for initial minimum-wage workers is small and statistically indistinguishable from zero, while the longer-term estimate is negative.

Reallocation quality. Figure A16 examines whether workers who change firms move to more productive ones. The estimated effects are small and statistically indistinguishable from zero, providing little

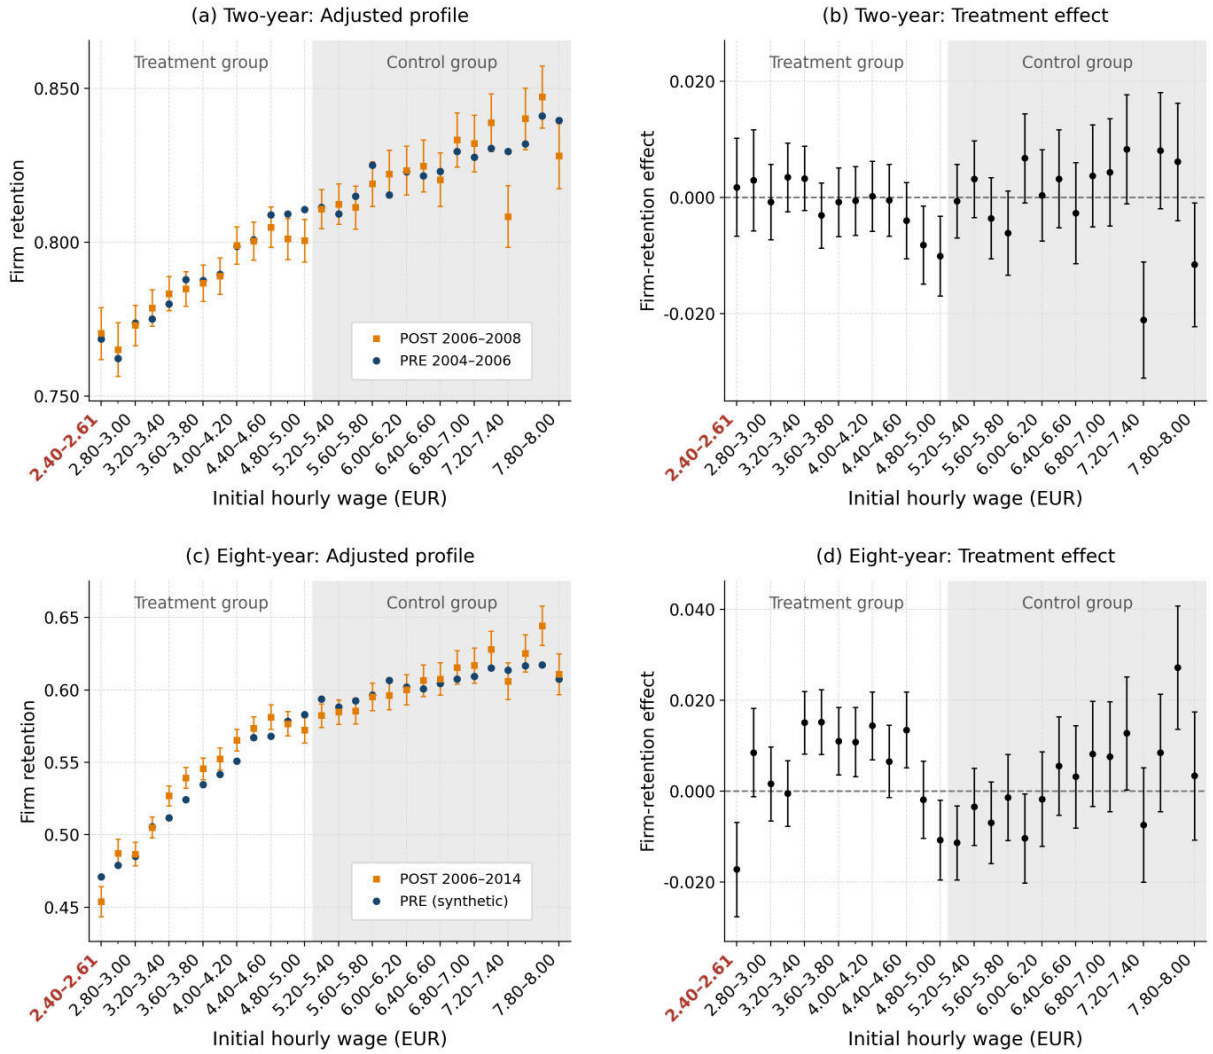


Figure A15: Firm Retention: Short- and Longer-Term Effects

Notes: Outcome: remaining with the initial employer at the endpoint. Top: 2006–2008; bottom: 2006–2014 with a synthetic counterfactual. Left panels compare observed and adjusted counterfactual profiles; right panels show their difference. Initial 2006 total hourly wages define bins, with wages above €5/h forming the control group. Intervals are 99% normal confidence intervals for the estimated contrast, displayed around the observed profile and the treatment effect. Vertical scales vary across panels.

evidence that the reform shifted reallocation toward more or less productive firms.

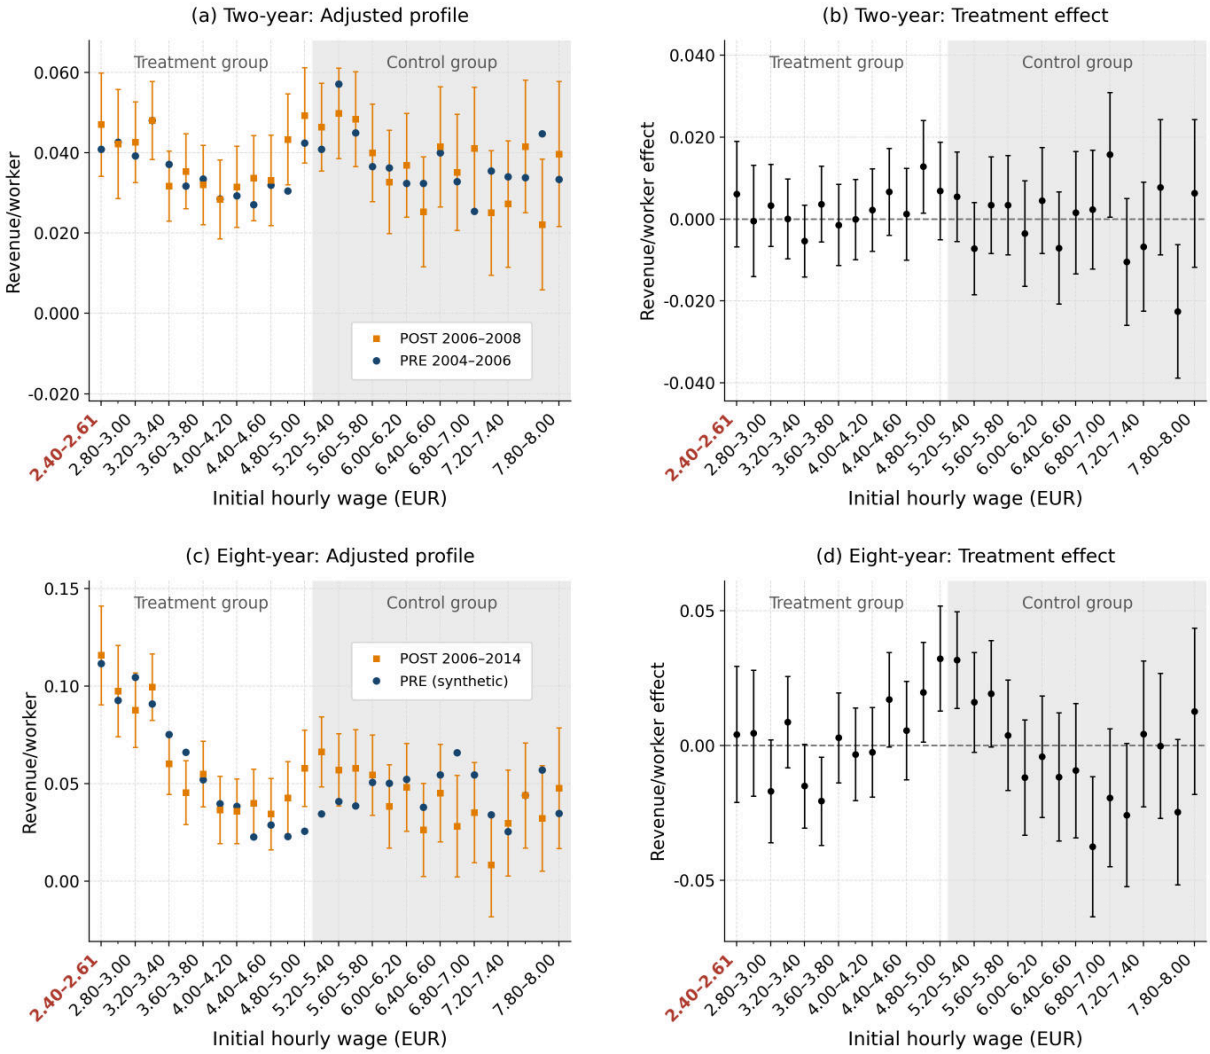


Figure A16: Reallocation Quality (Firm Revenue per Worker): Short- and Longer-Term Effects

Notes: Outcome: log change in employer revenue per worker among workers changing firms. Top: 2006–2008; bottom: 2006–2014 with a synthetic counterfactual. Left panels compare observed and adjusted counterfactual profiles; right panels show their difference. Initial 2006 total hourly wages define bins, with wages above €5/h forming the control group. Intervals are 99% normal confidence intervals for the estimated contrast, displayed around the observed profile and the treatment effect. Vertical scales vary across panels.

Occupational and industry retention. Figures A17 and A18 show the probability of staying in the same four-digit occupation and industry. The short-term effects for initial minimum-wage workers are small, positive, and statistically significant, suggesting that these workers became slightly more likely to remain in their occupation and industry. At longer horizons, both effects are close to zero.

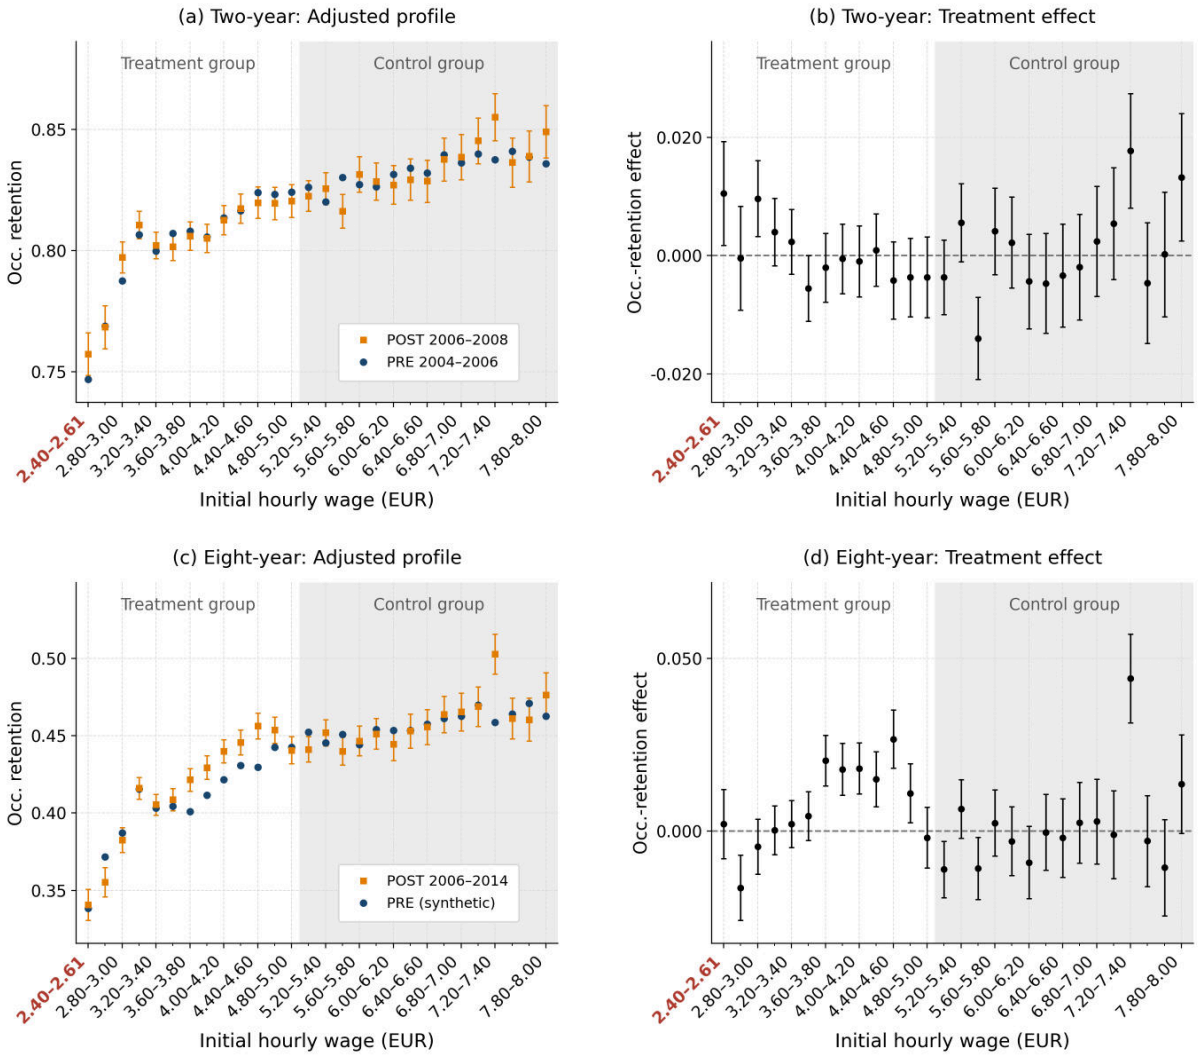


Figure A17: Occupational Retention (4-digit): Short- and Longer-Term Effects

Notes: Outcome: retaining the initial four-digit occupation. Top: 2006–2008; bottom: 2006–2014 with a synthetic counterfactual. Left panels compare observed and adjusted counterfactual profiles; right panels show their difference. Initial 2006 total hourly wages define bins, with wages above €5/h forming the control group. Intervals are 99% normal confidence intervals for the estimated contrast, displayed around the observed profile and the treatment effect. Vertical scales vary across panels.

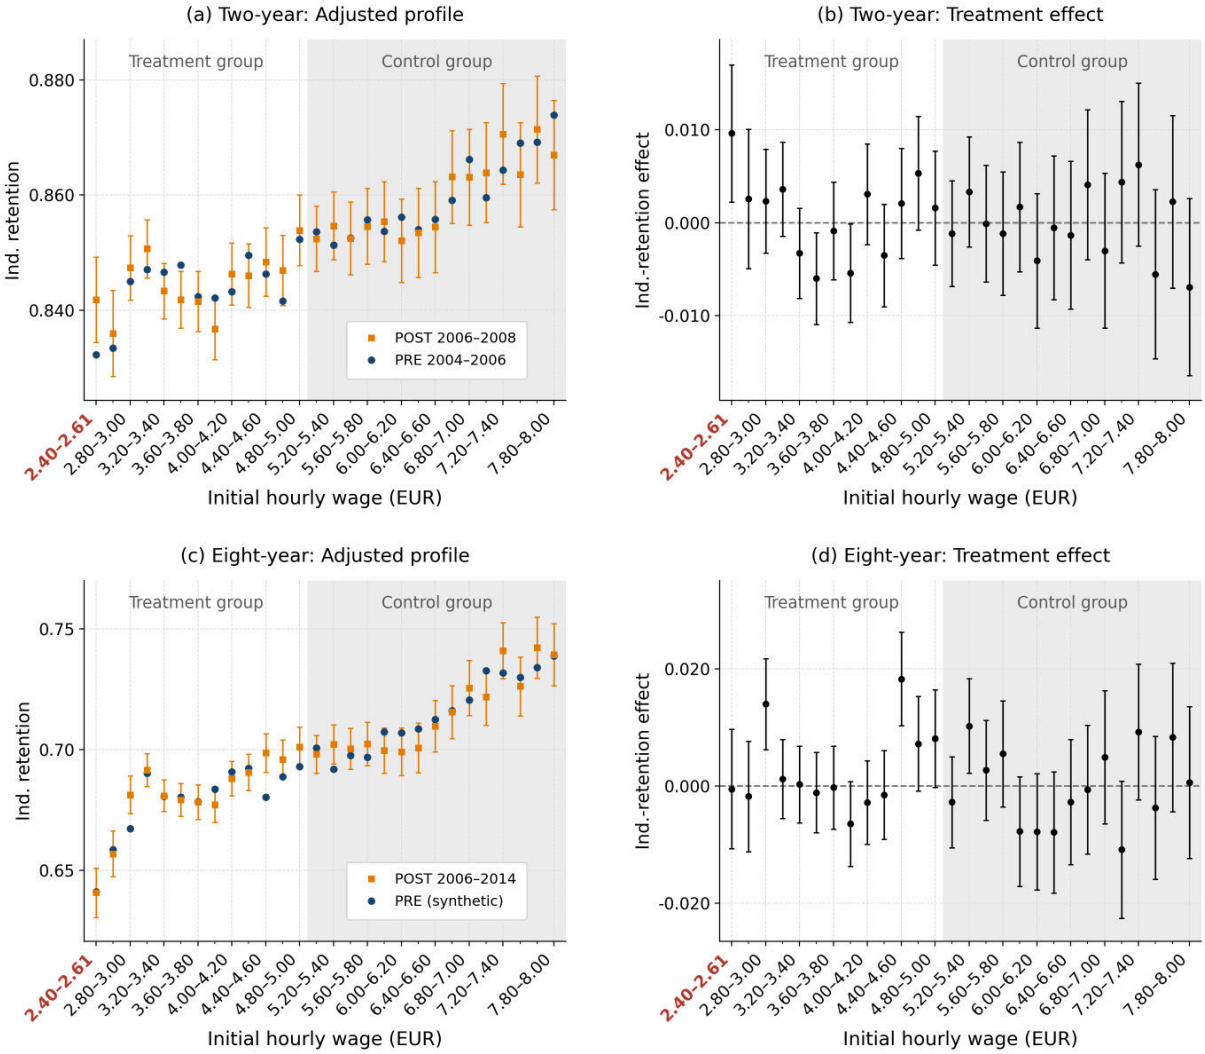


Figure A18: Industry Retention (4-digit): Short- and Longer-Term Effects

Notes: Outcome: retaining the initial four-digit industry. Top: 2006–2008; bottom: 2006–2014 with a synthetic counterfactual. Left panels compare observed and adjusted counterfactual profiles; right panels show their difference. Initial 2006 total hourly wages define bins, with wages above €5/h forming the control group. Intervals are 99% normal confidence intervals for the estimated contrast, displayed around the observed profile and the treatment effect. Vertical scales vary across panels.

Educational upgrading and promotion. Figures A19 and A20 report effects on educational upgrading and promotion. The short- and longer-term profiles complete the set of career-progression outcomes summarised in Table 3.

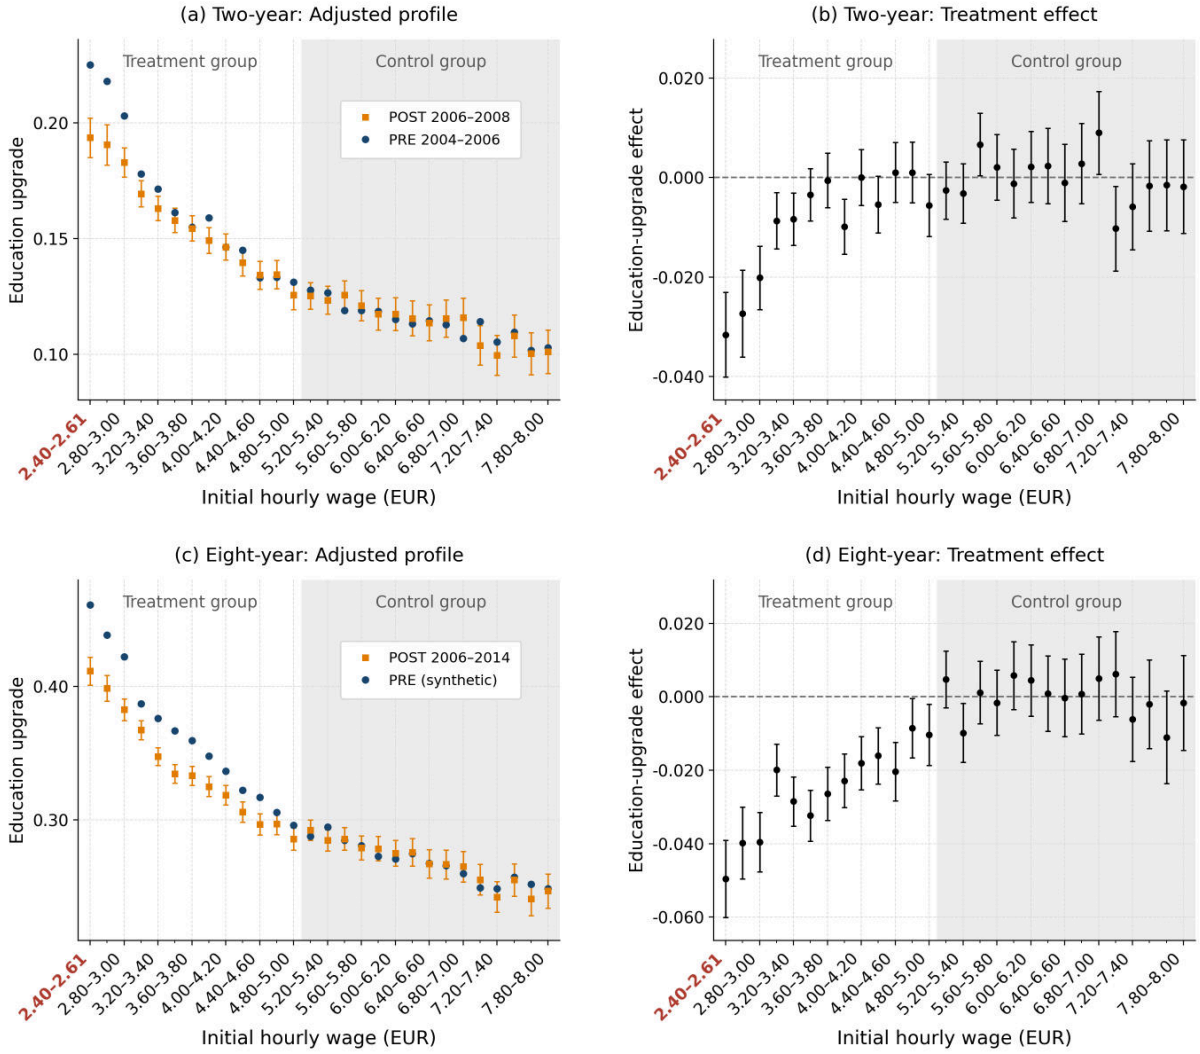


Figure A19: Educational Upgrading: Short- and Longer-Term Effects

Notes: Outcome: an increase in recorded educational-attainment category; reporting changes complicate interpretation. Top: 2006–2008; bottom: 2006–2014 with a synthetic counterfactual. Left panels compare observed and adjusted counterfactual profiles; right panels show their difference. Initial 2006 total hourly wages define bins, with wages above €5/h forming the control group. Intervals are 99% normal confidence intervals for the estimated contrast, displayed around the observed profile and the treatment effect. Vertical scales vary across panels.

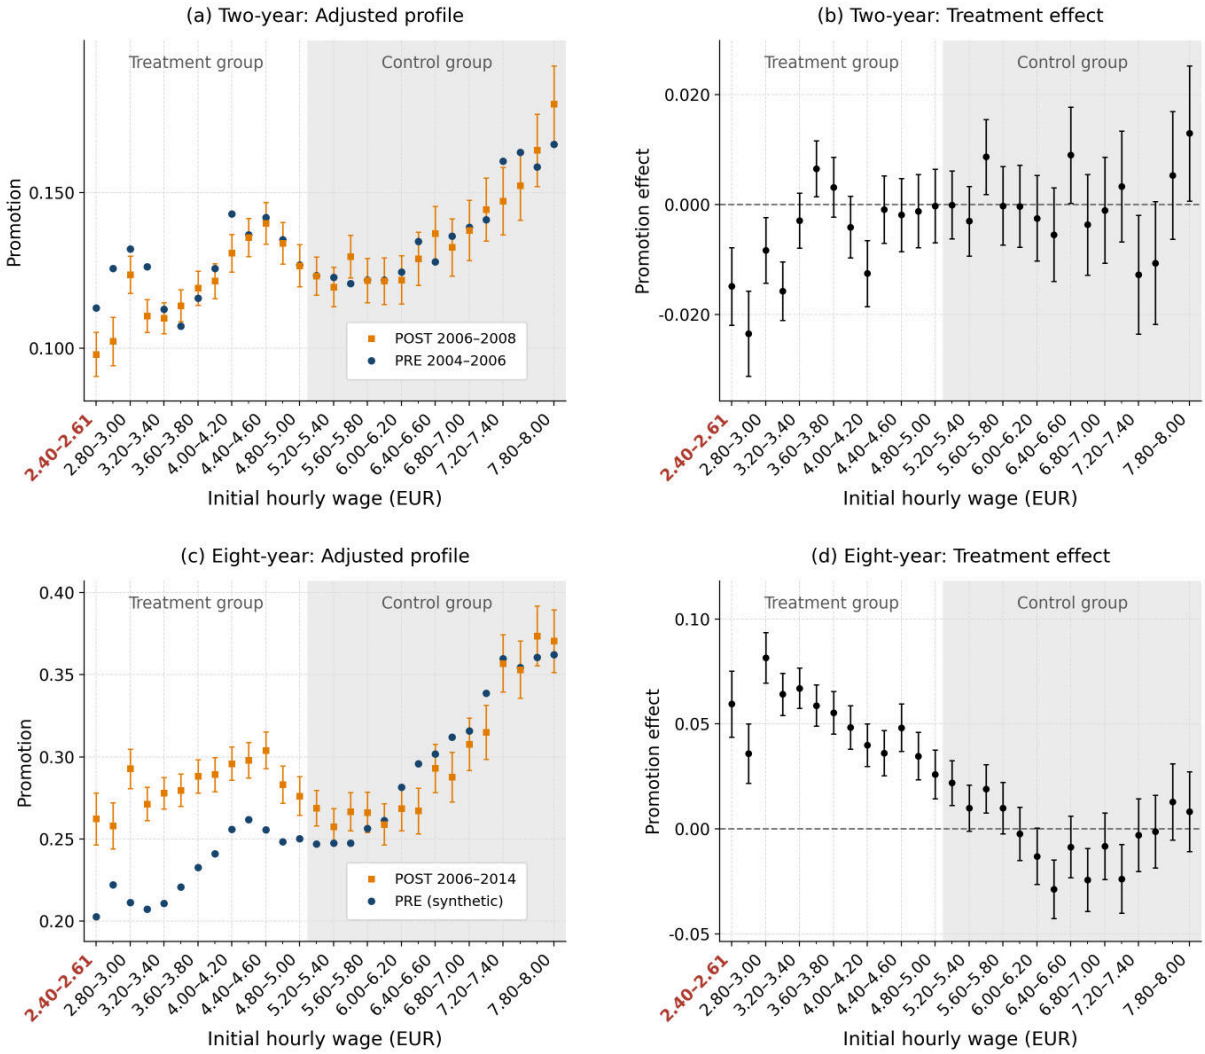


Figure A20: Job Promotion: Short- and Longer-Term Effects

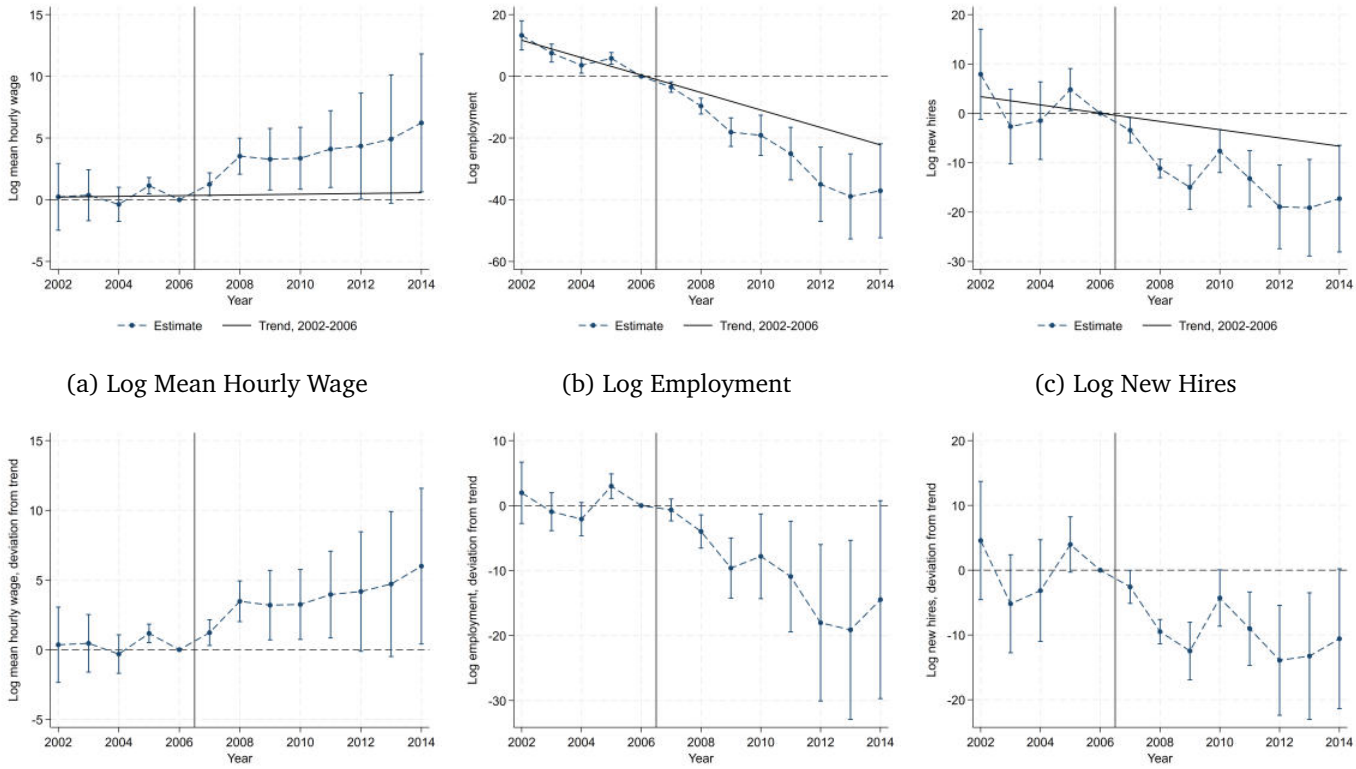
Notes: Outcome: a recorded promotion within the horizon among workers remaining with their employer. Top: 2006–2008; bottom: 2006–2014 with a synthetic counterfactual. Left panels compare observed and adjusted counterfactual profiles; right panels show their difference. Initial 2006 total hourly wages define bins, with wages above €5/h forming the control group. Intervals are 99% normal confidence intervals for the estimated contrast, displayed around the observed profile and the treatment effect. Vertical scales vary across panels.

B.9 Employment Effects Including Entrants: A Regional Exposure Design

The individual design conditions on employment in 2006, so it captures only the separation margin. To assess employment effects that include the entry margin, I follow a regional-age differential exposure approach, as in Albagli et al. (2025), combined with the minimum wage exposure measure of Dustmann et al. (2022).[55] I interact the mainland NUTS II regions with nine five-year age bands, yielding 45 region-age cells. Exposure is the cell's pre-policy wage-bill shortfall relative to the 2011 minimum wage, the endpoint of the announced path, held fixed at its 2002–2006 mean so that it never responds to the reform. I estimate event-study regressions of cell-level outcomes on exposure interacted with year dummies, weighting cells by employment and absorbing cell and year fixed effects, with standard errors clustered by cell. Because exposure correlates with pre-existing cell trends, Figure A21 reports deviations from a linear trend fitted on the pre-reform years.

Wages in more exposed cells rise steadily after 2006, confirming the design's first stage. Employment and new hires decline in exposed cells after 2008; the declines coincide with the financial and sovereign debt crises, and they shrink or disappear once cell-specific trends are allowed, so they are best read as upper bounds on the disemployment effect of the reform (Section 3.2).

[55] Dustmann et al. (2022) exploit purely regional variation across 401 German districts; with fewer than 10 NUTS II regions, regional variation alone cannot support estimation, so I interact region with age to recover enough treatment units.



(a) Log Mean Hourly Wage (b) Log Employment (c) Log New Hires

(d) Log Mean Hourly Wage, Detrended (e) Log Employment, Detrended (f) Log New Hires, Detrended

Figure A21: Regional Exposure Design: Event Studies for Wages, Employment, and New Hires

Notes: Event-study coefficients on exposure interacted with year (relative-gap exposure measure). Top row: raw coefficients, with the linear trend fitted on the pre-reform years projected forward. Bottom row: deviations from that trend, re-centered at 2006. New hires are workers with tenure below one year at their employer. Bands are 95% confidence intervals clustered by cell; on the detrended series they reuse the raw-coefficient standard errors and are indicative.

C Statistical Wage Process Details

C.1 Model Components and Interpretation

This section provides further detail on aspects of the wage process in Section 5.1 that are only briefly discussed in the main text.

Wage stickiness helps capture the concentration of wage changes around zero, motivated by administrative evidence on wage adjustment in Grigsby et al. (2021) and the leptokurtic distribution of earnings changes documented by Guvenen et al. (2021).[56] Near the bottom of the distribution, when the real minimum wage is unchanged, latent-wage movements that leave a worker in the binding region generate no change in actual pay, adding a further source of rigidity. Figure A26 shows the sharper concentration around zero among workers initially in the minimum-wage bin. Recording error smooths

[56] Grigsby, Hurst, and Yildirmaz study nominal base wages; the evidence here and the model concern real hourly wages.

these mass points in observed wages.

In the spirit of Guvenen (2009), I allow systematic differences in wage growth across workers, correlated with permanent pay differences. This heterogeneity allows workers with similar current wages to follow different expected future wage paths. I draw individual slopes according to

$$s_i = \sigma_s \left[\rho_{\alpha s} \alpha_i^{\text{std}} + \sqrt{1-\rho_{\alpha s}^2}\, \varepsilon_i^s \right], \qquad \varepsilon_i^s \sim \mathcal{N}(0,1), \tag{22}$$

where α_i^{std} denotes standardized permanent heterogeneity and ε_i^s is independent of α_i. A negative $\rho_{\alpha s}$ associates lower permanent pay with faster early-career growth, a source of catch-up for low-paid workers.

The worker-specific component of the lifecycle profile is $s_i h_\tau(\widetilde{a})$, where

$$h_\tau(\widetilde{a}) = \tau \left[1 - e^{-\widetilde{a}/\tau} \right], \qquad \widetilde{a} = a - 18.$$

As $\tau \to 0^+$, this component vanishes and all workers share the common lifecycle profile. As $\tau \to \infty$, $h_\tau(\widetilde{a}) \to \widetilde{a}$, recovering linear divergence according to individual slopes. For finite τ, $h'_\tau(\widetilde{a}) = e^{-\widetilde{a}/\tau}$, so differences in growth rates decline with age. Their accumulated contribution to wage-level differences remains: the heterogeneous-profile contribution to the log-wage gap between two workers approaches $(s_i - s_j)\tau$.

Recall from (7) that the persistence coefficient of ω_{it} is

$$\rho(\omega) = \frac{\bar{\rho}}{1 + \exp(-\eta\omega + \phi)}.$$

I set $\bar{\rho}$ to 0.98, close to but strictly below one, which rules out unit-root behavior and keeps the persistent component mean reverting. The parameter ϕ shifts the level of persistence, while η governs how steeply it rises with the lagged persistent component. At the estimated η and ϕ, the two-year persistence coefficient is 0.47 at $\omega = -1$, 0.71 at $\omega = 0$, and 0.87 at $\omega = 1$. When $\eta = 0$, the process nests a standard AR(1) with constant persistence $\bar{\rho}/(1 + e^{\phi})$.

This state dependence is motivated by the nonlinear earnings dynamics documented by Arellano et al. (2017). Mean reversion in the data remains strong conditional on age. In the model, persistence depends on the lagged persistent component, while the standard deviations of persistent innovations

and transitory shocks vary with age, in the spirit of Karahan and Ozkan (2013).

For workers older than 18 in 2006, I construct ω_{i0} using an annual approximation to the two-year process. Starting from zero at age 18, I roll the persistent component forward using annual persistence $\sqrt{\rho(\omega)}$ and innovation standard deviation $\sigma_\nu(\widetilde{a})/\sqrt{2}$. Subsequent transitions follow the two-year specification.

On the wage-floor function, q governs the degree of nonlinearity in spillovers. When $q = 0$, the spillover segment is linear, with slope $(Z_2 - M)/(Z_2 - Z_1)$, and the wage-floor mapping generally has kinks at both thresholds. When $q = 1$, the segment is cubic and joins the binding and unaffected regions smoothly: its slope is zero at Z_1, while setting $m_1 = (Z_2 - Z_1)/(Z_2 - M)$ ensures a unit slope at Z_2. The estimate $q = 0.3485$ lies between these cases, so the mapping retains kinks, which contribute to the kink in current-wage gains across percentiles in Figure 10(c). In earlier specification comparisons, estimating Z_1 and Z_2 directly as fixed real-wage thresholds common across minimum-wage regimes yielded a worse fit than scaling them with the minimum wage.

The different components can partly substitute for one another in fitting observed wage dynamics. For example, permanent heterogeneity and highly persistent shocks can generate similar wage paths over the observed horizon. Nevertheless, to the extent that alternative decompositions imply similar lifetime wage paths, estimated lifetime gains should be largely unaffected, even if the relative importance of the components in explaining growing out of the minimum wage changes.

C.2 Estimation and Weighting

I estimate the wage process by matching simulated wage distributions and wage-growth profiles to their empirical counterparts. The targets combine information about the level and dispersion of wages, lifecycle wage growth, movement out of the minimum wage, and the response to the reform. Table 6 lists the estimation and validation moments in the order shown in the figures in Appendix C.4.

For the reform-period estimation moments, I simulate a panel observed in 2006, 2008, 2010, 2012, and 2014. Separate cohort simulations cover the longer age windows used for the lifetime-inequality target and policy counterfactuals. For each horizon, I retain workers whose age at the endpoint is no greater than 64, matching the empirical age restriction. The simulation does not otherwise model attrition. Within each initial wage bin, growth under the frozen minimum wage provides the model

counterpart to counterfactual wage growth. The treatment effect is the actual-minus-frozen growth difference, net of the corresponding difference for the high-wage control group (initial hourly wages of at least €5). These calculations use the same initial bins and two-, four-, six-, and eight-year horizons as the empirical analysis.

The cross-sectional targets are the 2006 and 2010 wage distributions,[57] the mean log-wage profile by age, age-specific wage dispersion and minimum-wage incidence, and the aggregate inequality measures in the table. The mean age profile is smoothed with a quartic polynomial before estimation. To isolate the cross-sectional effect of the reform, the model's 2010 distribution holds population composition fixed while changing the minimum wage and aggregate time effect.[58]

An external benchmark also disciplines how much annual inequality persists over a career. Representative estimates across studies put lifetime income inequality at roughly 60–80 percent of annual income inequality (Björklund, 1993; Bowlus and Robin, 2004, 2012; Bönke et al., 2015). The comparisons differ in income concepts, inequality measures, populations and age coverage, and some estimates lie outside this range. I target the ratio of the Gini of cumulative wages to the average age-specific Gini at two-thirds, within this indicative range. This is a calibration assumption, rather than a moment estimated from the Portuguese panel. It is evaluated for a cohort followed from age 18 to 60, with the same upper age limit as in Bönke et al. (2015).[59] The reform's effect on lifetime earnings remains an output of the estimated process.

[57] For estimation, the "2006 distribution" denotes the pre-reform distribution pooled over 2002–2006. Pooling averages small annual fluctuations across otherwise very similar wage distributions, during a period when the real minimum wage was approximately constant.

[58] I apply the 2010 time effect and floor to the initial latent wages, retaining the initial composition and recording-error draws. Using the simulated panel's 2010 observations would instead age the original workers and exclude those who cross the age limit; it would not represent a new cross section. Histogram shares use the full population denominator, including workers outside the displayed wage range.

[59] Its age window is distinct from the lifetime-policy exercises, which cover ages 18–64 and, as a prime-age comparison, ages 25–55.

Table 6: Estimation and Validation Moments

Moment	Number
Estimation moments	
Dynamic moments (Figure A22)	
Treatment effects by initial wage: four horizons	112
Counterfactual growth by initial wage: four horizons	112
Minimum-wage persistence: four horizons	4
Wage distributions and age profiles (Figure A23)	
2006 wage-bin shares	28
2010 wage-bin shares	28
Mean log wage by age	47
Standard deviation of log wages by age	47
Minimum-wage incidence and inequality (Figure A24)	
Aggregate initial minimum-wage share	1
Minimum-wage share by age	47
2006 Gini	1
2010-to-2006 Gini ratio	1
2010-to-2006 p_{50}/p_{10} ratio	1
2010-to-2006 p_{90}/p_{10} ratio	1
Lifetime-to-average-annual Gini ratio	1
Total estimation	**431**
Validation moments	
75th, 90th and 95th wage percentiles (Figure A25)	3
Counterfactual wage changes: all workers (Figure A26)	140
Counterfactual wage changes: initial minimum-wage bin (Figure A26)	140
Total validation	**283**

Notes: Rows follow the order of the estimation and validation figures. Only the estimation moments enter the objective; validation moments are excluded. Age profiles contain one value for each of the 47 ages from 18 to 64. Wage profiles use 28 bins, from [€2.40,€2.61) to [€7.80,€8.00); the first bin contains the 2006 minimum wage. Each treatment-effect and counterfactual-growth profile contributes 28 values at each of four horizons, giving 112 moments. For validation, each wage-change distribution has 70 bins at each of the two- and four-year horizons, giving 140 values for all workers and 140 for workers initially in the minimum-wage bin.

For each parameter vector, I calculate the discrepancy between every simulated moment and its target. I divide these discrepancies by positive scales, square them, and average within each block of related moments. The loss is the sum of these block averages:

$$Q(\Phi) = \sum_{k} \frac{1}{n_k} \sum_{j=1}^{n_k} \left[\frac{m_{kj}^{\text{model}}(\Phi) - m_{kj}^{\text{target}}}{s_{kj}} \right]^2 .$$

Here, n_k is the number of moments in block k, and s_{kj} is the scale applied to its jth discrepancy. Scaling makes discrepancies in different units comparable. Averaging within a block prevents its contribution from rising mechanically with the number of moments it contains. I avoid percentage errors because

some targets are zero or close to zero. The treatment-effect profiles use a common scale across horizons; the counterfactual-growth profiles do likewise. The scales incorporate the block factors and overall normalization used in estimation.

I minimize this loss over bounded parameter values using the Covariance Matrix Adaptation Evolution Strategy (CMA-ES), a stochastic, population-based method for nonconvex optimization (Hansen, 2016). At each iteration, it draws candidate parameter vectors from a multivariate normal distribution and adapts its mean, covariance matrix, and overall step size using the better-performing candidates. In contrast, standard differential evolution generates new candidates from scaled differences between members of an existing population and recombines their coordinates, without explicitly estimating a covariance matrix. In my optimization experiments, CMA-ES reached comparably low loss values much faster than differential evolution. I follow the initial search with a smaller-step CMA-ES refinement.[60] Using the same random draws across parameter vectors reduces simulation noise. Search evaluations use 50,000 workers; candidate solutions are compared using 300,000 workers and the chosen solution is checked with an independent random seed. The reported fit uses 1.2 million workers.

In earlier experiments, simulated annealing was also slower to reach comparably low loss values. Both alternative algorithms sometimes failed to attain these values within the computational budget. The loss is relatively flat along some parameter combinations, consistent with partial substitutability between ingredients of the wage process, as discussed in the final paragraph of Appendix C.1.

C.3 Estimated Parameters

Table 7 reports the 31 estimated parameters. The common age profile rises sharply early in the career, is relatively flat through middle age, and declines late in the working life. Individual growth rates have standard deviation $\sigma_s = 0.0161$ and are negatively correlated with permanent pay differences ($\rho_{\alpha s} = -0.935$). Thus workers with lower permanent pay tend to have faster early-career growth. Heterogeneous profiles and a negative intercept–slope correlation have precedents in Guvenen (2009), although the correlation here is substantially stronger than his reported estimates.

The persistent component mean-reverts more quickly at low values. At a lagged component of zero, its two-year autoregressive coefficient is 0.712; it rises with the lagged state. This state dependence is

[60] The search adds a shape penalty to the fitting loss, so the full criterion is $\mathscr{L}(\Phi) = Q(\Phi) + \mathscr{P}(\Phi)$. The penalty discourages a weak early-career rise or an implausible late-career rise in mean wages, excessive dispersion from heterogeneous lifecycle slopes, and a nonmonotone wage-floor mapping. It is zero at the reported estimates.

consistent with the broader evidence on nonlinear earnings persistence in Arellano et al. (2017), rather than a direct replication of their specification. At age 18, the persistent-innovation and transitory-shock standard deviations are 0.411 and 0.156 log points. The former declines modestly with age and the latter rises modestly. Allowing shock dispersion to vary over the lifecycle follows Karahan and Ozkan (2013), but their estimated age patterns are more flexible. Differences in wage versus earnings measures, samples, and model frequency preclude treating those studies' coefficients as numerical targets for this model.

The probability that the latent wage remains unchanged over a two-year period is 0.574. Recording error has standard deviation 0.0387 log points, approximately 3.9 percent. These estimates refer to latent and recorded real hourly wages, so the inaction probability is not directly comparable to the frequency of nominal pay changes. With the minimum wage held fixed, recording error leaves only 14.5 percent of workers whose latent wages remain unchanged with recorded log-wage growth between -0.01 and 0.01; the remaining 85.5 percent record changes outside this band.

The wage-change distributions in Figure A26 provide an additional validation check and do not enter estimation. Within the plotted band between -0.025 and 0.025 log points, the model closely reproduces the overall share of very small wage changes at the two- and four-year horizons, but underestimates this share among workers initially in the minimum-wage bin. This shortfall may reflect overstated recording error among low-paid workers.

Table 7: Estimated Parameters

Parameter	Estimate
Permanent heterogeneity	
α_{loc}	1.1273
α_{scale}	0.1791
α_{skew}	-0.6433
α_{tail}	1.1132
Common Hermite age profile	
$g(4)$	0.3742
$g(7)$	0.4528
$g(12)$	0.5408
$g(17)$	0.5680
$g(22)$	0.5476
$g(32)$	0.5521
$g(37)$	0.5685
$g(42)$	0.5341
$g(46)$	0.3467
Heterogeneous lifecycle profiles	
σ_s	0.0161
$\rho_{\alpha s}$	-0.9351
τ	47.4975

Parameter	Estimate
Persistent shock	
η	1.0474
ϕ	-0.9777
$\sigma_{\nu,18}$	0.4112
κ_ν	-0.0332
Transitory shock	
$\sigma_{\xi,18}$	0.1560
κ_ξ	0.0354
Wage-floor function	
χ_1	0.7034
χ_2	1.8933
q	0.3485
Stickiness and measurement	
θ	0.5740
σ_ε	0.0387
Time effects	
μ_{2008}	0.0776
μ_{2010}	0.0096
μ_{2012}	-0.0007
μ_{2014}	0.0750

Notes: The table reports the 31 estimated parameters of the statistical wage process. Age is normalized as $\widetilde{a} = a - 18$, and the common profile satisfies $g(0) = 0$. The nine reported values are its free coordinates at the listed normalized ages. Shock standard deviations are reported at age 18; κ_ν and κ_ξ describe their change per decade of age. All transition parameters are at the model's two-year frequency.

C.4 Model Fit and Validation

Fit to estimation targets. Figure A22 compares treatment effects, counterfactual wage growth, and minimum-wage persistence with the targets used in estimation. Figures A23 and A24 show the remaining targets: the wage distributions and age profiles, minimum-wage incidence, aggregate inequality, and the external lifetime-to-annual Gini benchmark. Together, these figures cover all moments used in estimation.

Validation outside the objective. The wage-change distributions and upper-tail wage percentiles below are excluded from the loss. The wage-change distributions assess whether matching mean growth and minimum-wage persistence also reproduces the concentration of small changes and the tails of individual wage changes, for all workers and for those initially in the minimum-wage bin. The 75th, 90th, and 95th percentiles provide an additional check on the upper tail of the wage distribution. These are checks on untargeted features of the same data, rather than validation on a separate sample.

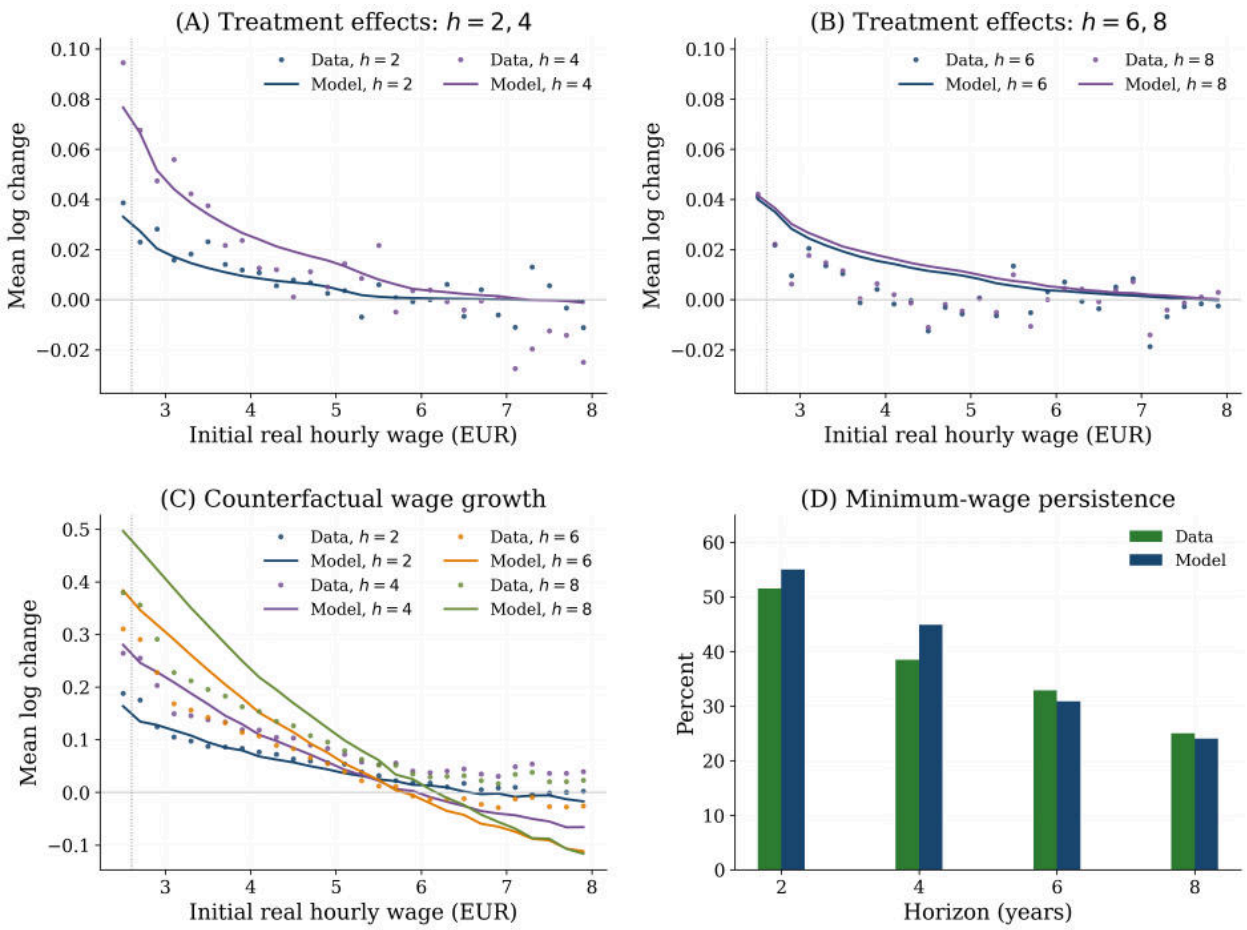


Figure A22: Fit to Dynamic Moments

Notes: Panels A and B report treatment-effect profiles at the two- and four-year horizons and the six- and eight-year horizons, respectively. Panel C reports counterfactual wage-growth profiles at all four horizons. Points denote the data and solid lines the model. Panel D reports the share of workers at the minimum wage in 2006 who are also at the minimum wage after two, four, six, and eight years, with bars for the data and model. Here, being at the minimum wage means having a recorded real hourly wage no more than 5% above the statutory floor in the relevant year ($w_{it} \leq 1.05 M_t$), including wages recorded below the floor. All displayed moments enter estimation.

The comparisons use recorded wages on both sides. Recording error affects initial-bin membership and subsequent wage changes, and can generate apparent exits from the minimum-wage bin without an equivalent underlying wage increase. Its inclusion is therefore relevant to interpreting both persistence and the wage-change distributions. Recording error is excluded from the lifetime-policy and historical-decomposition exercises.

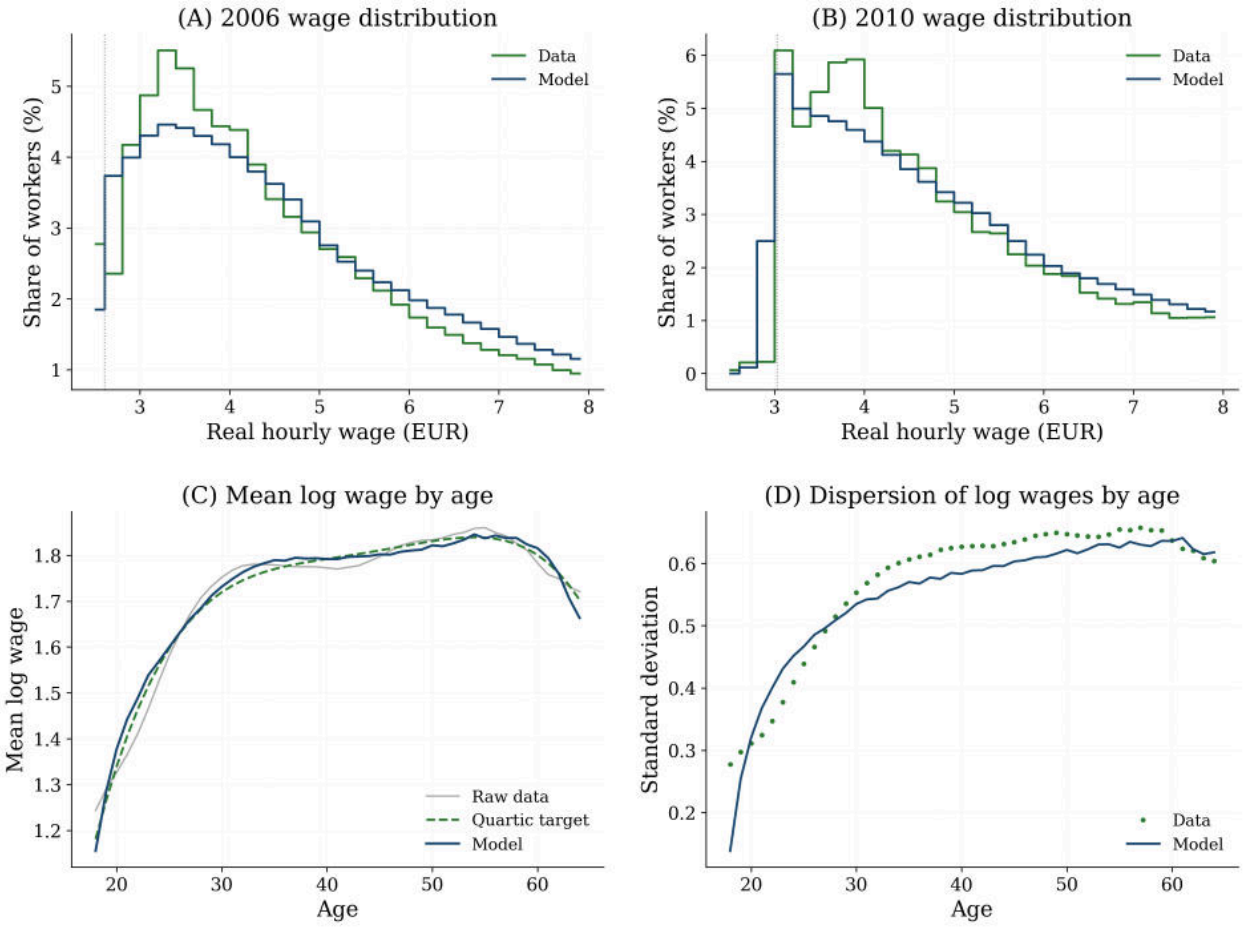


Figure A23: Fit to Wage Distributions and Age Profiles

Notes: Panels A and B compare the model with the 2006 and 2010 wage distributions. Panel C compares the model's mean log-wage profile by age with the raw data and the quartic target used in estimation. Panel D compares the standard deviation of log wages at each age. All displayed targets enter estimation; the raw mean age profile is shown for context.

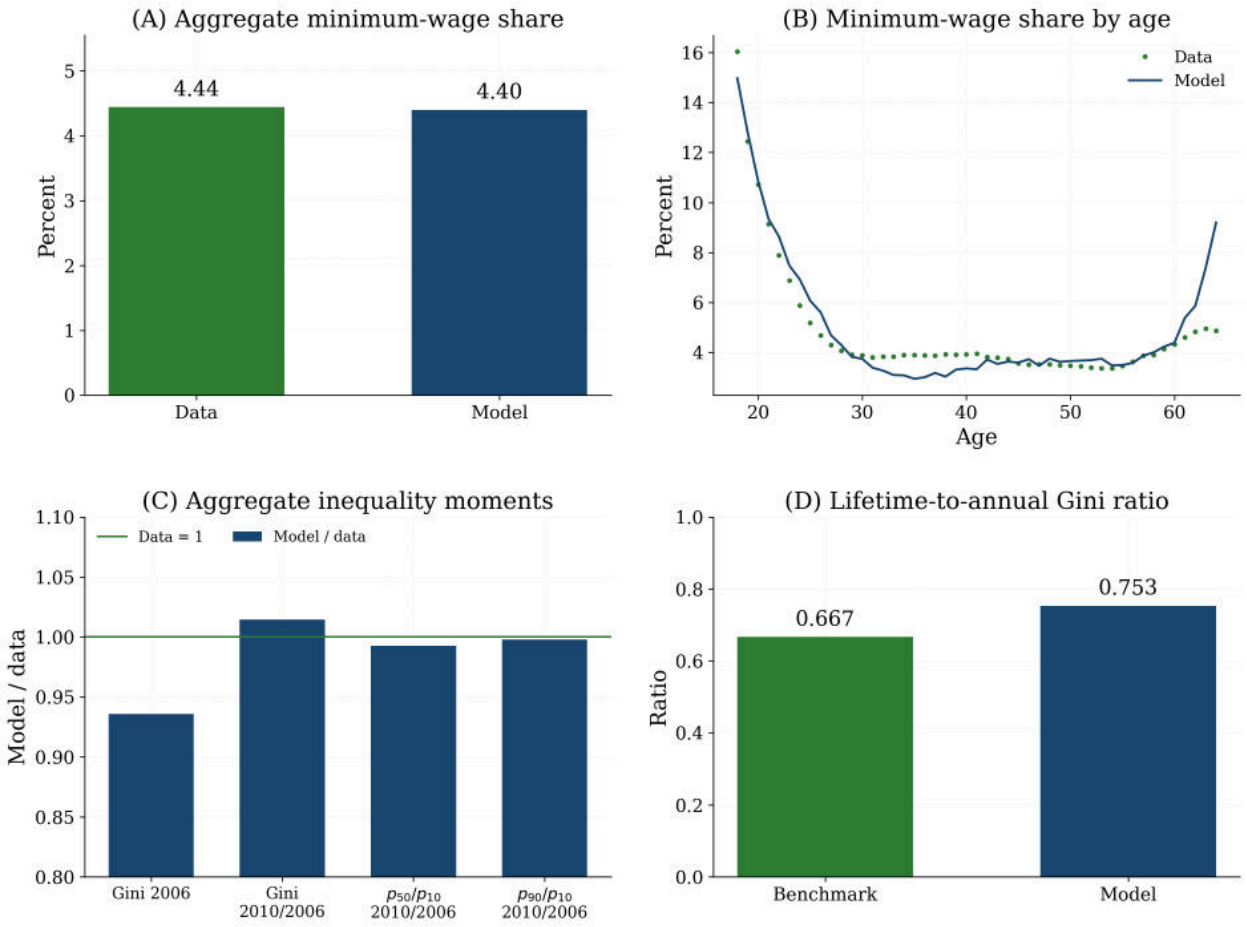


Figure A24: Fit to Minimum-Wage Incidence and Inequality Moments

Notes: Panels A and B compare the aggregate minimum-wage share and the share at each age. For these shares, being at the minimum wage means having a recorded real hourly wage no more than 5% above the 2006 statutory floor ($w \leq 1.05M_{2006}$), including wages recorded below the floor. Panel C reports model-to-data ratios for the 2006 Gini and the 2010-to-2006 ratios of the Gini, p_{50}/p_{10}, and p_{90}/p_{10}; the data value is normalized to one. Panel D compares the model's lifetime-to-average-annual Gini ratio for a cohort followed from age 18 to 60 with the external two-thirds benchmark, motivated by the comparisons in Björklund (1993); Bowlus and Robin (2004, 2012); Bönke et al. (2015). I use an external benchmark because I do not have social security data, and absence from QP need not imply zero earnings. A lifetime measure constructed from QP is very sensitive to whether unobserved years are assigned zero earnings, following the convention of retaining zero-earnings periods in lifetime earnings measures (e.g., Bönke et al., 2015), or excluded when averaging across observed years. All displayed moments enter estimation.

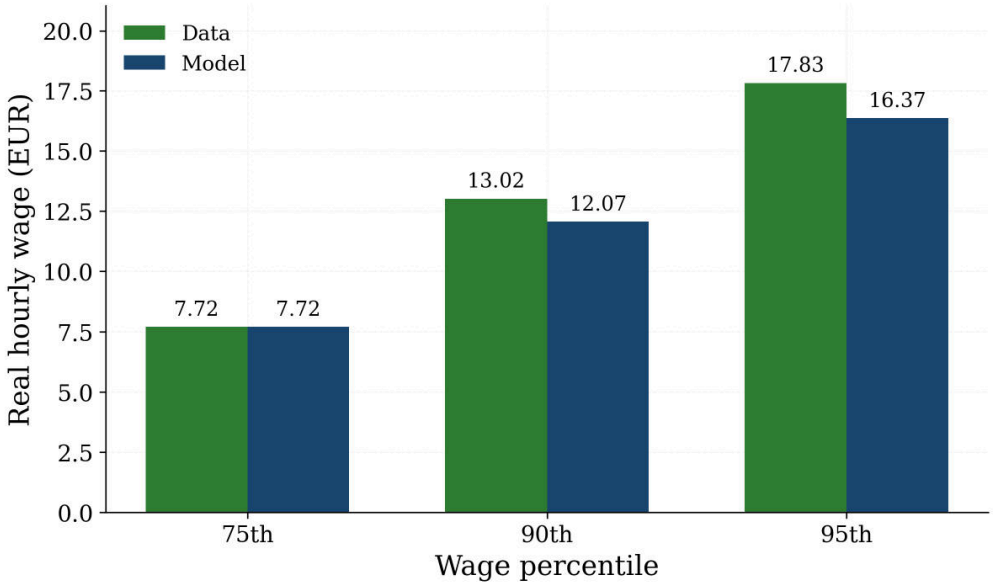


Figure A25: Validation: Upper-Tail Wage Percentiles

Notes: The 75th, 90th, and 95th percentiles of recorded real hourly wages in the data and model. The data use the pooled pre-reform cross section labelled 2006. These three percentiles do not enter estimation.

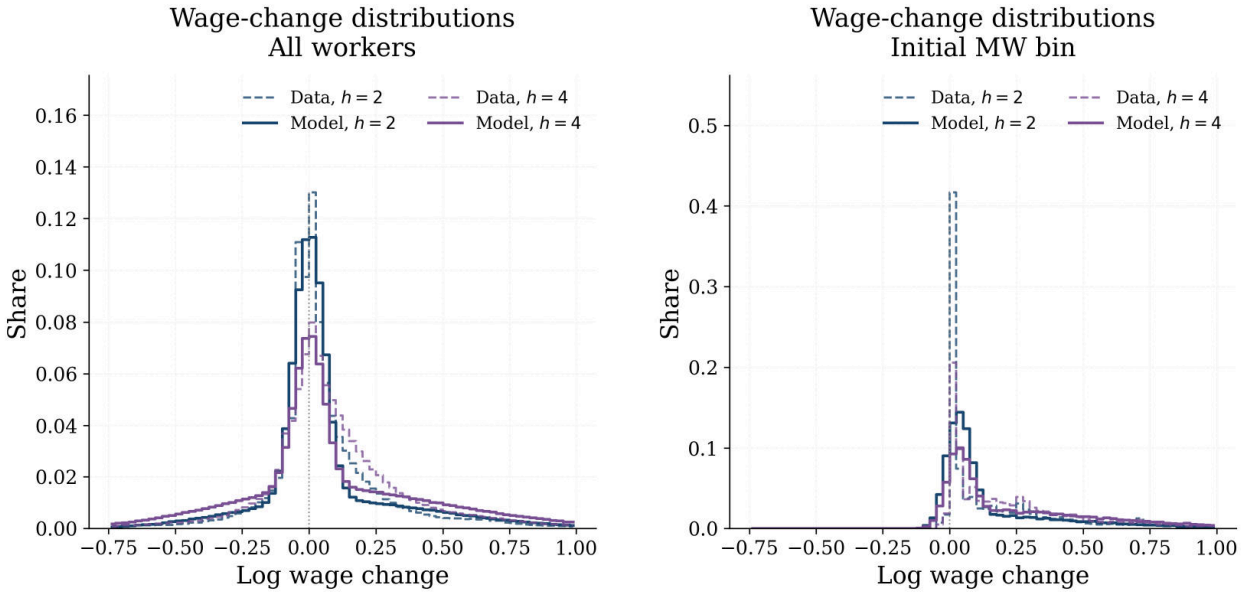


Figure A26: Validation: Distributions of Counterfactual Wage Changes

Notes: Two- and four-year counterfactual wage-change distributions for all workers (left) and workers initially in the minimum-wage bin (right). Dashed lines show the data and solid lines the model. These distributions do not enter estimation.

C.5 Historical Shock Decomposition: Procedure and Contributions

The decomposition follows workers initially in the minimum-wage bin at ages 18, 25, and 50 over the respective age windows 18–64, 25–55, and 50–64. The wage-trajectory exercise uses an initial simulated population of $N = 1.2$ million workers at each starting age; the decomposition uses 300,000 workers at each starting age. Each comparison uses the same workers under permanently low and high minimum wages, holding estimated parameters, simulated shocks, and wage-update draws fixed.

Let $G(S)$ denote the percentage increase in cumulative earnings from the higher minimum wage when only the components in S are allowed to evolve. Components outside S remain at their initial values. When all components are frozen, workers remain exposed to the floor throughout the simulated horizon, giving $G(\varnothing) = 16.2\%$. Allowing components to evolve changes their wage paths and therefore their exposure to the minimum wage.

Lifecycle evolution follows each worker's estimated profile $g_i(\widetilde{a}_{i,t})$. Persistent evolution includes both reversion from the initial state $\omega_{i,0}$ and subsequent innovations. Transitory evolution replaces the initial draw $\xi_{i,0}$ with the simulated subsequent draws. The permanent fixed effect α_i remains unchanged. Time effects and recording error are excluded throughout.

For each ordering, a component's contribution is the reduction in $G(S)$ when that component is introduced. The Shapley contribution averages this reduction across all orders. The calculation treats the common and individual-specific lifecycle terms separately, averaging over the 24 orderings of the four underlying components; their contributions are then added to report lifecycle growth as a single category.

Table 8 reports each category's share of the total reduction from the frozen benchmark to the full wage process.

Table 8: Sources of the Fading Minimum-Wage Gain: Shapley Decomposition

	Starting age		
Contribution to fading gains	18	25	50
Lifecycle growth	66.6%	12.2%	-3.7%
Persistent component	21.5%	76.7%	84.3%
Transitory component	11.9%	11.1%	19.4%

Notes: Contributions are shares of the fading cumulative-earnings gain for workers initially in the low-floor wage bin [€2.40,€2.61), followed over ages 18–64, 25–55, and 50–64. Gains compare permanently high and low minimum wages (M_{2010} and M_{2006}). Lifecycle growth combines the common and individual-specific contributions from the four-component Shapley calculation. Negative shares indicate that a component increases the policy gain on average. Parameters, simulated shocks, and wage-update draws are held fixed; recording error and time effects are excluded.

A negative contribution means that allowing the component to evolve increases the cumulative policy gain. For older workers, lifecycle changes modestly strengthen exposure to the floor, while persistent and transitory wage movements account for the fading gains.

D Optimal Minimum Wage Details

D.1 The Permanent Benchmark and the Fund Representation

This appendix derives the stationary mapping used in Section 6. Its notation s, m corresponds to z, M in the main text. Population mass is one, so π_s denotes both a type's share and its mass. The comparison concerns a matched stationary allocation.

Capital ownership in the permanent benchmark. Let π_s and κ_s denote type s's population and aggregate-profit shares, both summing to one. Berger et al. (2025) assign the same shares to initial capital ownership. For capital per person a_{st} and labor earnings $\ell_s(\mathbf{n}_t; m)$, the original budget is

$$c_{st} + a_{s,t+1} = \ell_s(\mathbf{n}_t; m) + (1 + R_t - \delta)a_{st} + \frac{\kappa_s}{\pi_s}\Pi_t, \qquad a_{s0} = \frac{\kappa_s}{\pi_s}K_0. \tag{23}$$

Here R_t is the capital rental rate. In an interior stationary equilibrium, the household Euler equation and capital-market clearing give

$$1 = \beta(1 + R - \delta), \qquad r \equiv R - \delta = \beta^{-1} - 1, \qquad K = \frac{\alpha(1-\gamma)Y}{R}. \tag{24}$$

The Euler equation determines the required return; firms' capital demand and market clearing determine K, rather than its ownership distribution.

Fund accounting and investment. Operating firms retain separate wage and employment decisions. Fund investment I_t and distributable income D_t satisfy

$$K_{t+1} = (1-\delta)K_t + I_t, \qquad D_t = R_t K_t + \Pi_t - I_t. \tag{25}$$

A price-taking manager maximizes $\sum_{t\geq 0} \beta^t D_t$, taking rental prices and operating profits as given, subject to feasibility and a transversality condition. Its interior investment condition is $1 = \beta(1+R_{t+1}-\delta)$. The stationary required return therefore agrees with (24), and capital-market clearing recovers the benchmark capital stock. Requiring only $I = \delta K$ would not determine K.

In the permanent steady state, set the nontradable payout entitlement to $d_s = \kappa_s D/\pi_s$. Then

$$D = (R-\delta)K + \Pi, \qquad c_s = \ell_s(\mathbf{n}_s; m) + \frac{\kappa_s}{\pi_s} D, \tag{26}$$

which reproduces the original stationary household budget.

Bond holdings and stationary equivalence. The household chooses consumption, bonds, and firm-specific hours as in (13), subject to the labor rationing constraints, a borrowing bound, a no-Ponzi condition, and budget (14). Bonds are claims on other households, separate from the capital held by the fund. With no government debt or outside lending, bond-market clearing is

$$B_t \equiv \int b \, d\mu_t = 0. \tag{27}$$

Here μ_t is the population distribution over individual states. For common bond holdings b_{st} within each permanent type, this is $\sum_s \pi_s b_{st} = 0$.

Starting from $b_0 = 0$ in the matched stationary allocation, the constant path $b_t = 0$ satisfies the budget, labor optimality conditions, and transversality condition. The interior bond Euler equation is

$$u'(c_t) = \beta(1+r)u'(c_{t+1}), \tag{28}$$

With constant consumption and $\beta(1+r) = 1$, this condition holds. Concave preferences and the bench-

mark convex choice set establish optimality of zero additional borrowing or saving at this allocation. To preserve the original feasible saving choices locally, write $a_s = \bar{a}_s + b$, with $\bar{a}_s = \kappa_s K/\pi_s$, and translate the capital-holding bound accordingly. For example, $a' \geq 0$ becomes $b' \geq -\bar{a}_s$, rather than $b' \geq 0$. This matches the selected stationary allocation; it does not establish equality of saving policies at other asset states or along transition paths.

The reduced labor-earnings function. For feasible labor-index values, the main text's earnings function is

$$E_s(N; m) = \max_{\mathbf{n} \in \mathcal{N}_s(m):\, N_s(\mathbf{n}) = N} \ell_s(\mathbf{n}; m),$$

The set $\mathcal{N}_s(m)$ includes the firm-specific hours ceilings. At a fixed labor index, labor disutility is fixed and utility is increasing in consumption, so maximizing earnings gives the optimal allocation across firms. This retains rationing and does not treat the CES index as raw total hours. At zero bonds, the remaining labor choice solves

$$\max_{\mathbf{n} \in \mathcal{N}_s(m)} \left\{ u\left(\ell_s(\mathbf{n}; m) + \frac{\kappa_s}{\pi_s} D \right) - v_s\big(N_s(\mathbf{n})\big) \right\}. \tag{29}$$