The Dynamics of Firm-level Pay: Theory and Evidence from Portugal.*

Rui Castro† Gian Luca Clementi‡

This draft: January 2026

Abstract

We document that Portugal experienced a steep decline in earnings inequality since the mid-1990s, driven primarily by compression in firm pay. To investigate the determinants of this evolution, we develop a tractable model of firm pay heterogeneity that incorporates monopsonistic labor market power into a Hopenhayn-style firm dynamics framework. Our analysis reveals that firms' labor market power has declined over time, with particularly pronounced effects in the market for college-educated workers. We first examine a natural candidate explanation: the substantial increase in the minimum wage since the mid-1990s. Our quantitative analysis suggests, however, that the minimum wage had limited impact, even among non-college workers most directly affected by it. Equilibrium forces, including improved selection on firm productivity and higher wages for non-college workers, reduced both the bite of the minimum wage and its competitive spillovers to higher-wage firms. Our model points instead to a decline in the importance of non-wage job attributes since the mid-1990s as the primary driver of firm pay compression. Consistent with the data, this mechanism generates compression in firm pay particularly at the bottom of the distribution, affecting both non-college and especially college-educated workers.

Keywords: Firm Dynamics, Earnings Inequality, Minimum Wage, Monopsony.
JEL Codes: D33, E24, E25, L11, L25.

*We thank Pedro Bento, Sílvia Gonçalves, Barış Kaymak, Christian Moser, Pedro Raposo, Joana Silva and seminar participants at Banco de Portugal, Cleveland Fed, McGill, NYU, Penn State, UNC at Chapel Hill, Universidade Católica Portuguesa, and University of Winnipeg for helpful comments. All errors are ours. Rui Castro acknowledges financial support from SSHRC and the generous hospitality of Universidade Católica Portuguesa and the Bank of Portugal while working on this project. Both authors thank Portugal's Instituto Nacional de Estatística for the data access.

†Department of Economics and CIREQ, McGill University. Email: rui.castro@mcgill.ca. Web: rui-castro.research.mcgill.ca.

‡Department of Economics, Stern School of Business, New York University and NBER. Email: clem@nyu.edu. Web: http://people.stern.nyu.edu/gclement/.

1 Introduction

Recent empirical work exploiting employer-employee matched datasets has revealed that firms play an important role in shaping earnings inequality dynamics in both developed and developing countries.[1] This paper shows that firms have also played a major role in shaping the evolution of Portugal's earnings inequality, and tries to identify the underlying mechanisms driving this evolution.

Using high-quality administrative data, we document a striking pattern: after rapid growth between the mid-1980s and mid-1990s, earnings inequality has declined steeply. This decline was driven exclusively by compression in firm pay - the component of workers' earnings attributable to the employing firm rather than to individual worker characteristics. In contrast, other components including worker-specific inequality and assortative matching between workers and firms have pushed inequality in the opposite direction.

At the same time, we find no appreciable decline in the heterogeneity of observable firm–level characteristics, such as productivity. Instead, lower firm pay dispersion reflects a decline in the pass-through from firm-level productivity to wages. This pattern points to increased labor market competition among Portuguese firms since the mid-1990s.

We evaluate two primary sources of heightened wage competition: increases in the minimum wage, and a decline in the importance of non-wage job attributes, or job amenities. Our assessment exploits the fact that college and non-college educated workers have been impacted differently by these changes: we provide evidence suggesting that the minimum wage has mostly affected non-college workers, while college workers experienced the largest decline in the valuation of job amenities. These differential effects help us identify and quantify the relative strength of each source of heightened wage competition.

To guide our analysis, we build a tractable theory of firm–level pay heterogeneity by introducing monopsonistic labor market power in the canonical firm dynamics model of Hopenhayn (1992). The model features ex-ante worker heterogeneity in skills (college vs non-college educated workers), and a minimum wage. It delivers skill-specific firm pay schedules that vary with productivity, together with an equilibrium measure of firms over productivity

[1] See for example Abowd et al. (1999); Card et al. (2013); Song et al. (2018); Alvarez et al. (2018).

levels.

A higher minimum wage can compress firm pay through two main channels. The most immediate effect (bite effect) forces low-productivity firms to raise wages, particularly for non-college workers, compressing firm pay at the bottom of the productivity distribution. Additionally, high-productivity firms - especially those paying wages just above the minimum - respond to intensified competitive pressure from the bottom by raising their own wages (spillover effect), as in Engbom and Moser (2022) and Berger et al. (2025), further compressing firm pay.

We find, however, that both effects are quantitatively limited. First, the minimum wage's bite is constrained by equilibrium forces. Non-college wages grew faster in Portugal than college wages (the latter actually declined in real terms), driven by a dramatic increase in the supply of college-educated workers. Higher minimum wages also induce better firm selection on productivity (via both the entry and the exit margins), and reallocate employment toward more productive firms. Together, these forces limit the fraction of constrained firms and workers, reducing the bite of the minimum wage.

Second, constrained firms pose little competitive threat to unconstrained firms. This is due to their relatively low mass in the economy, and also because they cut job vacancies down in response to a higher minimum wage, thus softening the competition. As a result, they elicit minimal wage responses from high-productivity employers. Combined with the limited bite, these findings suggest that the minimum wage made only a modest contribution to the observed compression in firm pay inequality.

Our theory suggests a second, more powerful source of heightened wage competition: changes in the relative importance of non-wage job attributes, such as location or working conditions. We begin by documenting a basic fact: conditional on being unconstrained by the minimum wage, the variance of firm pay in any given period is substantially higher for college-educated workers than for non-college workers. This pattern is consistent with job amenities being more important for college workers - either because they place higher value on amenities, or because of greater variation in the amenities offered to them.

We then find that, while firm pay away from the minimum wage has compressed for both types of worker since the mid-1990s, it did more so for college-educated workers. This

is important because these workers are essentially unaffected by minimum wage increases. We interpret this finding, through the lens of our model, as pointing to a fundamental shift in the nature of wage competition - unrelated to but contemporaneous with the minimum wage increases. Though our theory is agnostic about its root causes, we conjecture that this shift may reflect reduced labor market fragmentation in Portugal, driven by improvements in transportation infrastructure and greater labor force concentration in large coastal urban centers over the past three decades. These changes affected all workers, but particularly the college-educated, who do tend locate in those areas.

Our quantitative analysis shows that a decline in the importance of job amenities intensifies wage competition, raising wage levels across the board (lower markdowns), particularly so at the left tail of the firm pay distribution. This mechanism generates shifts in firm pay schedules consistent with observation, and emerges as the most powerful quantitative force behind firm pay compression in Portugal since the mid-1990s.

Literature Review Our paper contributes to a recent literature whose goal is to quantitatively assess the determinants of inequality dynamics by means of structural equilibrium models. Our conclusion regarding the relevance of the minimum wage in Portugal's experience contrasts with Engbom and Moser (2022), who attribute 45% of the compression in Brazil's labor earnings to the large increase in the real minimum wage experienced by that country over a period of 22 years. They find that a higher minimum wage raises firm pay across all firms, with particularly strong effects on low-productivity firms.

As in our theoretical framework, their model features monopsonistic firms deriving market power from matching frictions in the labor market. They extend Burdett and Mortensen's (1998) wage posting model, and feature on-the-job search and job ladder dynamics, and a fixed distribution of firms. Gouin-Bonenfant (2022) is another related paper built on the same foundation, which abstracts from the minimum wage while adding firm dynamics. The latter focuses on showing how increased firm-level productivity dispersion helps explain the observed decline in the aggregate labor share. Our model differs from these papers in two key respects. First, and most important, we introduce non-wage job amenities as a primary source of market power, on top of matching frictions. Second, we abstract from on-the-job

search: workers need to find a new job every period, so firms do not worry about worker retention and wage posting depends on static considerations. This approach retains the key implications of monopsonistic wage-setting, while achieving parsimony that enables us to incorporate rich firm dynamics. Bilal et al. (2021) develop a rich firm dynamics framework also based on Burdett and Mortensen (1998). As in our paper, and in line with Hopenhayn (1992), firms operate a decreasing returns to scale technology. However they concentrate on the implications for cross-firm labor allocation, and do not have predictions for wages.

Berger et al. (2025) also feature heterogeneous firms in productivity facing a static wage-setting problem. Like in our model, a minimum wage generates heterogeneous effects across the firm distribution - a key departure from the classical representative firm monopsony model due to Robinson (1933). In contrast to our paper, their analysis is normative, seeking to identify the optimal minimum wage. Ours is positive: given Portugal's large minimum wage increase, we seek to understand its role in shaping inequality dynamics. Methodologically, they model a finite number of firms competing à la Cournot, their source of market power, while we consider a continuum of monopsonistically competitive firms. This allows us to provide a richer characterization of firm dynamics, and trace how changes in wage-setting forces propagate through the firm distribution, with implications for wage inequality. Hurst et al. (2022) also examine the minimum wage effects on wage inequality through monopsonistic competition, where firms exercise market power because workers view different firms as offering distinct work environments. Unlike our analysis, they focus on worker heterogeneity and how firms substitute across worker types in response to minimum wages. Since their firms are homogeneous, their model generates no firm pay effects.

The remainder of this paper is structured as follows. In Section 2, we carefully document the evolution of earnings inequality in Portugal and highlight the role played by firms. In Section 3 we introduce a simple static model of the labor market with monopsonistic competition. We show how labor market power shapes the level and slope of firms' wage profiles. In Section 4 we introduce firm dynamics, and characterize the implications of a minimum wage. In Section 5 we present the quantitative model, featuring heterogeneity in worker skills, and describe the main results. Section 6 concludes.

2 Earnings inequality in Portugal: Evidence

2.1 Quadros de Pessoal

The earnings data is from Quadros de Pessoal (QP), an administrative data set covering all private-sector firms employing at least one salaried worker, including those owned by the government. The data is collected annually by the Ministry of Labor to ensure compliance with labor laws. The reporting is for the month of March until 1993 and for October since. The data is unavailable for 1990 and 2001.

QP covers all workers employed in Portugal, except civil servants and independent contractors, as well as temporary rural and domestic workers. It is a matched employer-employee data set, providing comprehensive information on workers and some basic information on firms. We focus on workers aged 18–65, working full-time (at least 140 normal monthly hours), and earning at least the minimum wage as base pay.[2] We exclude firms in farming and fishing, since earnings in these sectors are often not reported on a monthly basis, and focus on for-profit non-financial market-oriented sectors.

Our favored measure of compensation is total monthly earnings, i.e. base salary plus overtime and regular monthly supplements such as productivity and seniority premia, as well as compensating differentials.[3] To avoid outliers, we trim the top 1% earnings observations per year. Further details on the data are in Appendix A.

2.2 The dynamics of earnings inequality

Figure 1 plots the variance of log earnings since 1986, the year Portugal joined the European Union, then known as the European Economic Community.

Inequality has increased dramatically from 1986 until the mid-1990s. The variance grew by about 50% over just a decade. From the mid-1990s onwards, however, inequality has first plateaued and then briskly declined since the financial and sovereign debt crisis of 2008-9. The variance of earnings dropped by more than one fifth between the mid-1990s' peak and

[2]The data on the (general) minimum wage and on the (private consumption, base year 2016) CPI are from https://www.pordata.pt.

[3]When running our analysis using hourly wages instead, the magnitudes change slightly, but the message does not.

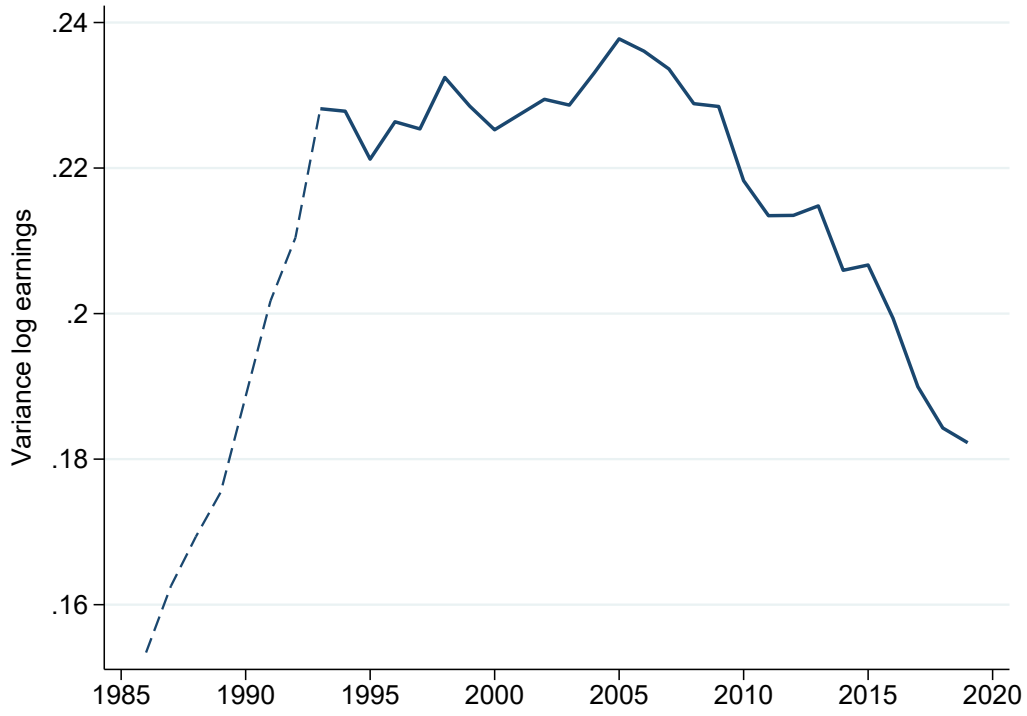


Figure 1: Quadros de Pessoal, 1986-2019

2019. An yearly rate of decline of about 1% is economically significant and comparable to the rate of increase in the same measure of inequality reported by Song et al. (2018) for the United States between 1985 and 2009.

As far as we know, this non-monotonic dynamics is unique in the developed world. Oliveira (2023) has documented the same facts, using the same data. This paper focuses on inequality past the mid-1990s, a period during which, thanks to the reduction in informality and greater compliance, Quadros de Pessoal has significantly increased worker coverage. Hence the dashed lines in the figures pre mid-1990s.

Inspired by the recent empirical literature highlighting the key role that firms have played in shaping inequality in Germany (Card et al., 2013), the U.S. (Song et al., 2018) and Brazil (Alvarez et al., 2018), we ask whether employers have played a substantial role in Portugal as well.

2.2.1 The role of firms: A first look

Figure 2 displays a decomposition of the variance of log earnings into the variance of average earnings between firms and the employment-weighted average of the within-firm earnings variances.

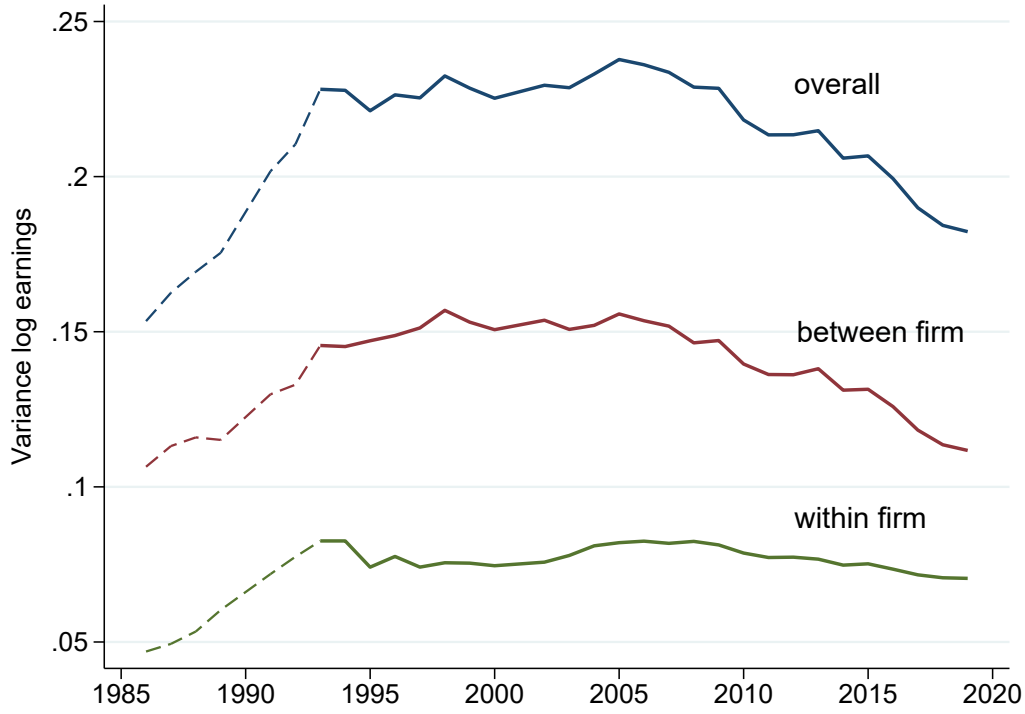


Figure 2: Between-within firm decomposition

The decomposition gives a first indication of the importance of firms in the Portuguese context. The between-firm component accounts for about 60% of the level. This is higher than what Song et al. (2018) report for the U.S. (30% to 40%), but comparable to what Alvarez et al. (2018) report for Brazil. Most importantly for our purpose, it is also the main component behind the dynamics, accounting for about 90% of the decline since the mid-1990s.

A natural question is whether the evolution of the between-firm component is driven by changes over time in the composition of the workforce at the firm level or by variation in intrinsic firm–level characteristics.

2.2.2 AKM regression

To isolate the role played by intrinsic firm–level characteristics, we now resort to the regression framework due to Abowd et al. (1999). The AKM approach extends the prototypical Mincer wage regression specification by adding firm fixed effects to control for observed and unobserved systematic firm-level characteristics, in addition to worker fixed effects and a set of time-varying observables. Our regression equation for the two-way fixed effect (AKM) model is

$$y_{ijt} = \text{WFE}_i + \text{FFE}_j + \mathbf{X}'_{ijt}\beta + \varepsilon_{ijt}, \tag{1}$$

where y_{ijt} is log earnings of worker i at time t, employed at firm $j = j(i,t)$. For the time-varying observables $\mathbf{X}_{ijt}$ we follow Card et al. (2013) and include a full set of year dummies and a worker age cubic polynomial, both fully interacted with schooling attendance dummies.[4] As is standard in the literature, we estimate (1) over six non-overlapping subperiods p of four/five years each. This rolling time window approach is a simple way to capture variation over time in the estimated coefficients.[5] We therefore estimate cross-sectional distributions of WFE_i^p and FFE_j^p per subperiod p, plus the coefficients in β^p.

For estimation purposes, we restrict our sample in a number or ways. First, we focus on the largest connected set of all firms (and their workers) for which every worker in the set has ever worked for. Second, we discard singletons (unique firm-worker observations). See Appendix A for the details.

2.2.3 Variance decomposition

Ignoring the time-varying observables, the AKM regression delivers the following variance decomposition by subperiod p:

$$\text{var}\left(y_{ijt}\right) = \text{var}\left(\text{WFE}_i\right) + \text{var}\left(\text{FFE}_j\right) + 2\text{cov}\left(\text{WFE}_i, \text{FFE}_j\right) + \text{var}\left(\varepsilon_{ijt}\right). \tag{2}$$

[4] As in Card et al. (2013), age is normalized (by deviations from the age of 40) so that the polynomial implies flat experience profiles at age 40. The linear effect of age is dropped as it is captured by the worker fixed effects together with time. Likewise, the independent effect of schooling is also captured by the worker fixed effects.

[5] An alternative would be to extend the AKM model to allow for time varying effects. See Lachowska et al. (2023) and Engbom et al. (2023) for examples of this approach, with unrestricted firm-by-year effects.

The first term captures dispersion in the systematic worker component, reflecting both observable characteristics such as schooling and gender and unobservable traits such as motivation and interpersonal skills. The second term captures dispersion in the systematic firm component, reflecting characteristics such as productivity, size, and management quality. This is commonly called the firm pay component, a wage premium enjoyed by any worker employed at a particular firm. The covariance term measures assortative matching, the extent to which high-wage workers are employed at high-wage firms. The final term captures dispersion in the random component, reflecting permanent or transitory earnings changes due to shocks such as health events or training investments.

Worker fixed effect dispersion can vary across subperiods through two channels: persistent changes in the systematic components of continuing workers, or cohort effects from the characteristics of labor force entrants and exiters. The latter may reflect significant increases in college graduation rates and female labor force participation since the mid-1980s. Similarly, firm pay dispersion may change through cross-period shifts in incumbent firms' pay policies or through cohort effects from the characteristics of entering and exiting firms within each subperiod.

Table 1 presents the results of this decomposition exercise, including the time-varying observables (Xb) omitted from (2), though the terms emphasized in (2) remain the primary focus due to their quantitative significance.

Figure 3a shows the period-by-period dynamics of the most significant components from Table 1. The variance of firm pay is a significant driver of the increase in overall earnings inequality until the mid-1990s and it's by far the key force behind the decline since then. In fact, it is the only major force pushing for earnings compression. It accounts for 180% of the overall decline in inequality and outmatches the impact of greater worker heterogeneity, responsible for a negative contribution of 92%. Leitão et al. (2022) reach a similar conclusion when looking at the same data starting from 2005, when dispersion in worker fixed effects began to stabilize and eventually declined.

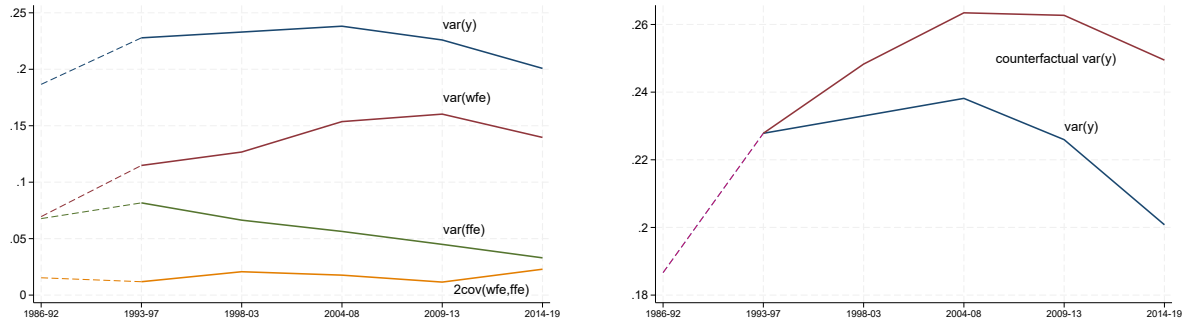
The literature has long emphasized that the cross-sectional variance components in Table 1 might suffer from limited mobility bias, which tends to overstate the importance of firm and worker fixed effects. To address this issue, we have applied the leave-out bias-correction

	Period 2 (1993-1997)		Period 6 (2014-2019)		Change 6–2	
	Value	Share	Value	Share	Value	Share
Var(y)	0.228	100	0.201	100	-0.027	100
Var(WFE)	0.115	50.4	0.14	69.6	0.025	-91.8
Var(FFE)	0.082	35.9	0.033	16.4	-0.049	180
Var(Xb)	0.004	1.9	0.008	3.9	0.003	-12.9
Var(ϵ)	0.014	6.3	0.011	5.3	-0.004	13.4
2Cov(WFE, FFE)	0.012	5.2	0.023	11.4	0.011	-40.9
2Cov(WFE, Xb)	-0.001	-0.5	-0.013	-6.2	-0.011	42.4
2Cov(FFE, Xb)	0.002	0.8	-0.001	-0.4	-0.003	9.7
R^2	0.9370		0.9466			
Nobs	5,189,611		8,890,328			
Nfirms	101,666		152,734			

Notes: The variance components are all weighted by the number of worker-years, and refer to the AKM regression (1). Nobs is the number of worker-years, and Nfirms is the number of firms.

Table 1: Variance Decomposition and Dynamics

procedure proposed by Kline et al. (2020), Appendix C. Although we indeed find evidence of a bias, it remains roughly constant over time and therefore does not significantly impact how the different components vary. We still reach the conclusion that compression firm fixed effects is by far the main driver of the reduction in earnings dispersion.



(a) Main components

(b) Counterfactual earnings variance

Figure 3: Variance Decomposition and Dynamics

An alternative way to appreciate the importance of the compression in firm pay for the overall inequality decline is to evaluate (2) while counterfactually holding the variance of firm fixed effects constant at the mid-1990s level. Figure 3b shows that earnings inequality would have increased rather than declined.

Such outsize role of firm pay as a driver of earnings inequality is in common with the

experiences of Brazil and Ecuador, where inequality has steadily declined, and with those of U.S. and Germany, where inequality has steadily increased.[6]

2.3 Determinants of compression in firm pay since the mid-1990s

What are the economic forces primarily responsible for the greater compression in firm pay? The literature has suggested that firm-level heterogeneity in pay may be intimately tied to the variation in firm-level productivity in a non-competitive labor market. This motivates us to investigate the association between firm pay and firm-level productivity. We postulate for now the following log-linear specification for firm j's pay policy in subperiod p:

$$\text{FFE}_{jp} = \alpha_p + \gamma_p \log z_{jp}, \tag{3}$$

where z_{jp} is the firm's within-period systematic productivity level, and γ_p is the elasticity of firm pay with respect firm–level productivity, commonly referred to as the pass-through from productivity to wages. From the firm pay policy (3) we obtain

$$\text{var}\left(\text{FFE}_{jp}\right) = \gamma_p^2 \text{var}\left(\log z_{jp}\right), \tag{4}$$

which indicates two possible sources for the compression in firm pay: (i) A compression in the distribution of firm–level productivity, and (ii) a decline in the pass-through.

We first ask whether the Portuguese data is consistent with a specification such as (3). We do this by regressing the estimated AKM firm fixed effects on measures of firm–level productivity, following Alvarez et al. (2018) very closely.

2.3.1 Sistema de Contas Integradas das Empresas

We link Quadros de Pessoal with the Sistema de Contas Integradas das Empresas (SCIE), which allows us to compute measures of firm-level productivity. The SCIE provides a broad coverage of private-sector non-financial and for-profit firms operating in Portugal. However, it is only available since 1996.

[6] See Alvarez et al. (2018) for Brazil, Messina and Silva (2019) for Ecuador, Card et al. (2013) for Germany, and Song et al. (2018) for the United States

Between 1996 and 2003, the information is drawn from the Inquérito à Empresa Harmonizado (IEH), a survey compulsory for all large private-sector corporations and a random sample of small corporations (less than 100 employees, or sales less than 5 million euros).

Starting in 2004, it is based upon administrative information, from tax records on a broader set of firms in 2004-5, and since 2006 from Informação Empresarial Simplificada (IES) on the universe of private-sector non-financial and for-profit firms operating in Portugal, drawn in particular from balance-sheet data. This is significantly higher quality information, used by the Public Administration and the Bank of Portugal. The number of available variables was increased with the IES, and SCIE coverage was also extended to include entrepreneurs and independent contractors, although we restrict our sample to corporations. In 2010 there was a major overhaul of the data reporting, for harmonization purposes with the rest of the E.U., which was applied retroactively to some of the variables starting in 2004. Some of the variables are therefore only available since 2004, and some others since 2010. We construct population weights to make the pre and post-IES information comparable. Further details can be found in Appendix B.

2.3.2 Firm–level productivity and firm pay

We compute two alternative measures of firm-level productivity. Our favorite is firm–level TFP, computed for firm j in sector s and year t as

$$z_{jt} = \frac{y_{jt}}{n_{jt}^{\alpha_n^s} m_{jt}^{\alpha_m^s} k_{jt}^{1-\alpha_n^s-\alpha_m^s}},$$

where y is total operating revenue (CPI deflated), n is the number of employees, m is the real value of intermediates, k is the real book value of capital, and the α's are the sector-specific averages of the firm–level revenue shares of total expenditures in labor and materials and services.[7] Unfortunately information on the book value of capital is only available since 2010, and so is firm-level TFP. We therefore rely mostly on real value-added per worker (CPI deflated), which is available since 1996 and hence covers the whole period of firm pay compression. For a given firm, we average each measure of productivity over each subperiod.

[7] We average over observations with strictly positive value added.

To avoid outliers, we trim the top and bottom 1% observations by subperiod.

We first describe how firm–level wages vary with firm–level productivity in line with the existing literature, by positing a uniform schedule across workers. We then consider differential schedules, by worker schooling. Differential schedules contain important information about the sources of firm pay compression.

Uniform schedule We begin by empirically implementing a simple uniform firm pay schedule. Figure 4a plots quantiles of within-period average log real value-added per worker against the corresponding firm fixed effects (normalized by the period 90th percentile value), averaged over each quantile. While high-productivity firms do pay more, the pass-through has clearly weakened over time, especially from the mid-1990s to the mid-2000s.

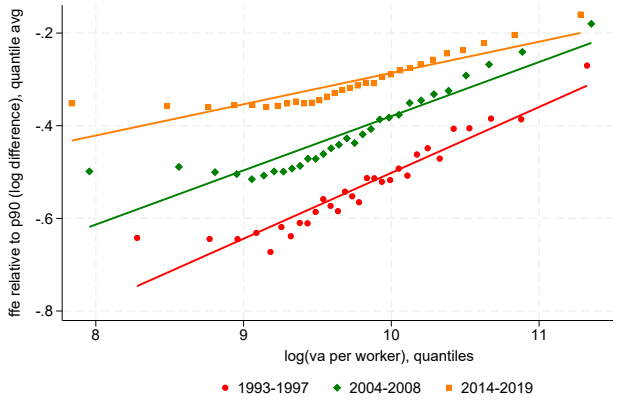
We estimate the firm pay schedule (3) by subperiod, with z_{jp} equal average real value added per worker of firm j in period p. The regression is weighted by worker-years, scaled by the population weights described in Section 2.3.1. Following Alvarez et al. (2018), we also consider specifications with sector and firm size controls likely co-determinants of firm pay.

The results in Table 2 confirm this pattern, revealing a substantial decline in the pass-through. The baseline regression without controls displays a decline from 0.181 to 0.113. These estimates are somewhat lower than but comparable to those estimated by Alvarez et al. (2018) for Brazil – a decline from 0.26 to 0.14 between 1996-2000 and 2008-2012.[8]

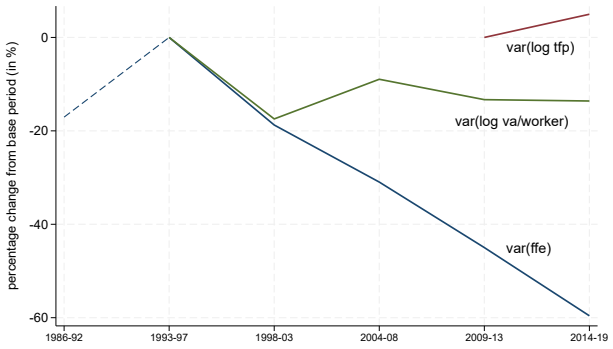
Figure 4b displays the evolution of firm–level productivity dispersion (worker-year and population weighted) alongside the dispersion in firm fixed effects that we ultimately seek to explain. The dispersion in value-added per worker has declined only slightly over time, while the dispersion in firm–level TFP increased very little - our sample is unfortunately quite short, as explained above. The evidence is therefore clear: the compression of firm pay is driven primarily by the weakening of the pass-though.

Although the linear regression facilitates comparison with the existing literature, Figure 4a reveals that the weakening pass-through actually reflects two distinct effects: uniformly higher firm pay at the bottom (relative to the 90th percentile) combined with a lower elastic-

[8]Our figures also align closely with the 0.1-0.15 range reported by Card et al. (2018) for Portugal, though their analysis is restricted to male workers in the late 2000s and employs a different data set for firm–level characteristics.



(a) Firm pay and value added per worker



(b) Variance of firm pay, value-added per worker, and TFP

Figure 4: Determinants of Firm Pay Compression

	1993-97			1998-03			2004-08			2009-13			2014-19		
log va per worker	.181	.15	.118	.183	.149	.13	.153	.131	.122	.132	.121	.111	.113	.103	.093
log size	–	.046	.043	–	.04	.037	–	.033	.031	–	.017	.016	–	.016	.016
sector dummies	no	no	yes	no	no	yes	no	no	yes	no	no	yes	no	no	yes
R^2	.21	.319	.384	.21	.295	.325	.198	.274	.29	.179	.208	.221	.18	.221	.237
Nobs	6,062,306			10,847,210			8,753,664			6,113,613			7,843,322		
Nfirms	21,505			37,357			118,162			94,440			109,068		

Notes: All coefficients are significant at 1% (robust standard errors). Weighted by worker-years and population weights. Nobs is population-adjusted total worker-years and Nfirms is number of firms.

Table 2: Regression of firm pay on firm characteristics

ity at the top. To capture both effects, we estimate the following piecewise linear regression by nonlinear least squares (again weighted by worker-years, and scaled by the population weights):

$$\text{FFE}_{jp} = (\alpha_p + \beta_p \log z_{jp})\, \mathbb{1}_{[\log z_{jp} < \zeta_p]} + (\alpha_p + \beta_p \zeta_p + \gamma_p(\log z_{jp} - \zeta_p))\, \mathbb{1}_{[\log z_{jp} \geq \zeta_p]} + \varepsilon_{jp}. \quad (5)$$

We omit firm size and sector controls here to maintain comparability with the baseline specification in Table 2. This piecewise specification better captures the "hockey stick" pattern evident in the data. Table 3 reports the key estimates, and Figure 5 illustrates the improved fit of the piecewise linear model.[9]

[9] As in Figure 4a, the fitted lines shown are estimated from the plotted quantile means for illustration and do not correspond to the regression coefficients reported in the tables, which are based on the underlying raw data.

The results clearly confirm a flat "blade" portion of the hockey stick. Consistent with our theory, this suggests greater minimum wage incidence among low-productivity, low-pay firms, which manifests as similar firm effects across these firms. An increase in the minimum wage would then produce the pattern displayed in Table 3, with compression in firm pay arising from the bottom due to a more prominent blade. However, this is not the complete explanation for the weakening pass-through. The elasticity of firm pay with respect to productivity Has also declined over time among high-productivity firms (captured by γ_p in (5)) – in fact, by more so than what the simple linear regression conveys. This points to changes in the extent of wage competition, possibly driven by the minimum wage increase but potentially due to other factors as well.

parameters	1993-97	1998-03	2004-08	2009-13	2014-19
β_p	-0.404	0.155	-0.064	-0.035	-0.015
ζ_p	9.085	9.547	9.238	9.196	9.426
γ_p	0.244	0.217	0.180	0.152	0.139
pseudo-R^2	0.294	0.223	0.216	0.192	0.199
Nfirms	21,505	37,357	118,162	94,440	109,068
Nobs	6,062,306	10,847,210	8,753,664	6,113,613	7,843,322

Notes: All coefficients are significant at 1% (robust standard errors). Weighted by worker-years and population weights, Nobs is population-adjusted total worker-years and Nfirms is number of firms. The parameter estimates refer to equation (5).

Table 3: Regression of firm pay on firm characteristics

Differential schedules by skill To obtain pay schedules conditional on worker skills, we first estimate the AKM regression (1) separately for workers with and without a college degree, and then estimate a firm-pay schedule as in (5) by worker type. Figure 6 displays the results with the appropriate piecewise linear regression fit overlaid. Table 4 reports the slope coefficient γ_p for the handle of each hockey stick.

The pay schedules in any given period are clearly heterogeneous across worker types, with two notable features. First, while non-college workers always exhibit a hockey stick blade, college workers don't. Instead, college workers show a region of declining firm pay with value added among low-productivity firms. This suggests that the higher minimum wage incidence among low-productivity firms documented earlier is driven primarily by unskilled

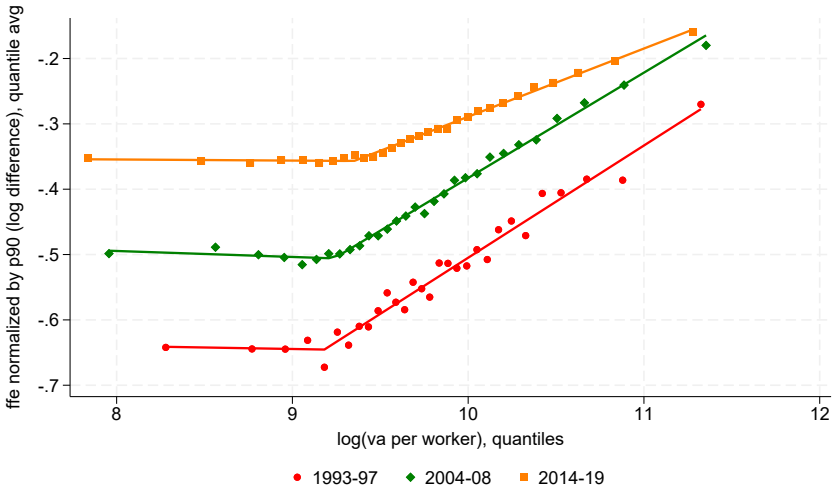


Figure 5: Firm pay and value added per worker, 1993-1997

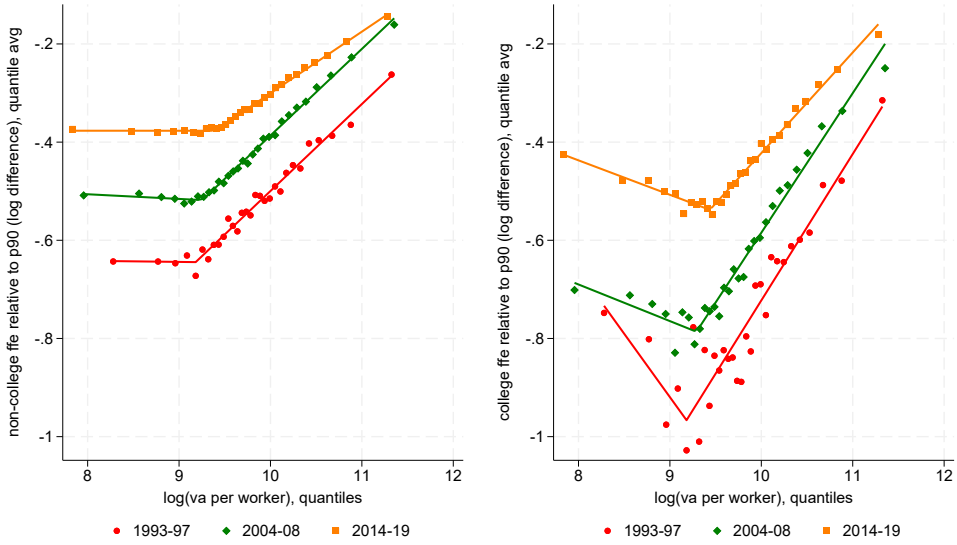


Figure 6: Firm Pay Schedule Conditional on Worker Schooling

workers. In fact, low-productivity firms offer a wage premium to skilled workers, relative to the threshold, with the largest premiums at the least productive firms.[10] Second, above the threshold, the pass-through is systematically lower for non-college workers. This is consistent with a more competitive labor market environment for these workers compared to college-educated workers.

[10]Our model doesn't rationalize this fact. A possible explanation, following Babina et al. (2019), is that low-productivity firms face higher exit risk and must pay higher wages to skilled workers as compensation for job loss risk. We abstract from such considerations in our model.

Pay schedules have also changed significantly over time. Non-college workers experienced a sustained increase in the size of the hockey stick blade, along with a flattening of the handle, especially in the later period. For skilled workers, the association between firm productivity and firm pay has flattened on both sides of the threshold. According to Table 4, the flattening of the pass-through above the productivity threshold has been more pronounced among college-educated workers. This evidence is consistent with a minimum wage bite effect primarily among non-college workers, and points to additional factors - with disproportionate impact on college-educated workers - as sources of compression in firm pay. Our model will exploit these observations.

γ_p	1993-97	1998-03	2004-08	2009-13	2014-19
non-college	0.242	0.221	0.184	0.162	0.147
college	0.280	0.231	0.206	0.182	0.151

Notes: All coefficients are significant at 1%. Weighted by worker-years and population weights. The pass-through estimates refer to equation (5), estimated conditional on schooling.

Table 4: Schooling-specific elasticities of firm pay with respect to value-added per worker, in high value added firms

We now provide a quantitative sense of the potential relevance of these additional sources of firm pay compression. Our aim is to decompose the total variance of firm pay into between and within skill components according to

$$\text{var}(\text{FFE}_{jp}) = n_p^0 n_p^1 (\overline{\text{FFE}}_p^0 - \overline{\text{FFE}}_p^1)^2 + \sum_{s\in\{0,1\}} n_p^s \text{var}(\text{FFE}_p^s), \tag{6}$$

across all firms j in period p, where n_p^s is the share of total worker-years with skill $s \in \{0,1\}$ (denoting non-college and college-educated, respectively) in the population, $\overline{\text{FFE}}_p^s$ is the cross-firm average firm-fixed effect conditional on skill s, weighted by worker-years of skill s, and $\text{var}(\text{FFE}_p^s)$ is the cross-firm variance of firm-fixed effects conditional on skill s, also weighted by worker-years. The first term on the right-hand-side of (6) is the between-group component, and the second term is the within-group component, which is the sum of non-college and college subcomponents.

The challenge in implementing this decomposition lies in obtaining the between compo-

	Period 2 (1993-97)		Period 6 (2014-19)		Change 6–2	
	Values		Values		Values	Share
Total Variance	0.085		0.039		-0.046	100
Between Group Variance	0.002		0.000		-0.002	4.59
Within Group Variance	0.083		0.039		-0.044	95.41
Non-College	0.076		0.029		-0.047	102.74
Population Share		0.957		0.831	-0.126	
Variance		0.080		0.035	-0.045	
College	0.007		0.010		0.003	-7.33
Population Share		0.043		0.169	0.126	
Variance		0.152		0.059	-0.093	

Notes: The components of the variance decomposition are based on (6).

Table 5: Variance Decomposition of Firm Pay into Between and Within Skill Components

nent, which requires renormalizing the skill-specific firm fixed effects to ensure comparability of firm pay levels across skill groups. The AKM estimation produces firm fixed effects normalized with respect to an arbitrary firm, the 90th percentile in Figure 6. To achieve comparability, we must renormalize using firms that offer identical pay across worker types. Importantly, this renormalization only affects the magnitude of the between component; since it produces a (log) additive adjustment, it leaves the within component unchanged. We adapt Card et al.'s (2015) approach for cross-gender comparisons, building on our evidence from Figure 6. This approach amounts to assuming that the lowest type-specific firm fixed effects in Figure 6 are equal to zero, hence comparable and suitable for renormalization. Appendix D provides details. A caveat is that, while the required assumptions are useful for the specific purpose of constructing the between component, they do not emerge from our main analytical framework.

Table 5 displays the outcome of this variance decomposition, between period 2 (1993-97) and period 6 (2014-19). A first observation is that the total variance based on the aggregation of the skill-specific FFE's is remarkably similar to the variance of FFE's from Table 1. Although the heterogeneity in firm pay schedules from Figure 6 suggests the baseline AKM regression model (1) is misspecified, this has a negligible impact on the variance. This is largely because a vast majority of the workforce is non-college educated.

The decomposition reveals that nearly all (95.41%) of the compression in firm pay stems from the within component, with almost all of it (102.74%) attributable to the compression

in firm pay among non-college workers. The within component for college-educated workers actually contributes negatively to overall firm pay compression. The reason is that the fraction of college-educated workers has increased significantly, from 4.3% to 16.9% in terms of worker-years, and these workers exhibit more dispersed firm pay than non-college workers in any given period, particularly in period 6.

Nevertheless, the dispersion in firm pay among college-educated workers has declined dramatically from period 2 to period 6, twice as much as among non-college workers. This suggests an additional important driver of changes in wage competition and firm pay compression. A simple counterfactual exercise allows us to assess its quantitative importance. We computed a counterfactual period 6 firm pay variance in which the variance of firm pay for college-educated workers remained constant at the period 2 level, while their population share still increased substantially as observed in the data. This counterfactual yields a compression in firm pay that is 34% below the actual compression observed (and 36% below the actual compression of the within-skill component - a figure independent of the normalization of firm fixed effects). Our exercise suggests a quantitatively important role for additional factors affecting wage competition, factors that likely impact both non-college and college workers, though to different degrees. These factors appear to be contemporaneous with, yet independent of, the minimum wage increase. Our structural analysis will help isolate their respective contributions.

3 A labor market with monopsonistic competition

We build our quantitative-theoretical framework in three steps. This section develops a simple static model with an exogenous distribution of firms. The next section introduces a minimum wage, firm dynamics, and endogeneizes the firm distribution. The following section presents the full model with heterogeneous worker types.

We introduce a static model of the labor market along the lines of Card et al. (2018) and show how firm-level wage policies vary with market power. We depart from Card et al. (2018) by considering a continuum of firms and solving for equilibrium wages.[11] In Section

[11] In Portugal, collective bargaining agreements at the occupational and industry levels are widespread,

4 we will nest this framework in a firm-dynamics model with endogenous entry and exit, better suited for quantitative analysis.

We assume that a mass $N_f > 0$ of firms produce a homogeneous good by means of the production function $y = zl^\gamma$, where l denotes labor and $0 < \gamma \leq 1$. The variable z is a firm-level (idiosyncratic) productivity index, distributed over a discrete finite set according to the probability density μ. Product demand is infinitely elastic at the unitary price.

A mass $N_w > 0$ of identical individuals sell labor services to firms. The value to a worker of a wage offer w is

$$u(w) = w + \varepsilon,$$

where ε is a match-specific amenity shock distributed according to the Gumbel distribution with location parameter 0 and scale parameter $\sigma > 0$. The shock proxies for non-pecuniary aspects of the job offer that the worker may value, such as proximity, colleagues' characteristics, and effort requirement.

Each worker is matched with two firms, randomly drawn from the population. Since extending an offer is costless, it will be optimal for firms to quote a wage to any of the mass $\bar{m} \equiv 2N_w/N_f$ of applicants.[12]

Confronted with wage offers $\{w_i, w_j\}$, the worker's optimization problem is

$$\max \{w_i + \varepsilon_i, w_j + \varepsilon_j\}.$$

Due to our distributional assumption, the probability that a worker will take a wage offer w_i when the alternative is w_j is

$$\Pr(w_i + \varepsilon_i > w_j + \varepsilon_j) = \frac{1}{1 + e^{\frac{w_j - w_i}{\sigma}}}.$$

For given competitors' wage offers, workers' optimization induces a firm-level labor supply curve $\bar{m}\Phi(w)$, where $\Phi(w)$ is the probability that a wage offer w is accepted.

yet firms retain substantial autonomy in setting wages above the negotiated floors. Firms actively exercise this discretion, resulting in considerable wage flexibility at the firm level (Cardoso and Portugal, 2005; Card and Cardoso, 2022). Accordingly, in our model firms set wages.

[12] This will not be the case in the more general model with minimum wage introduced below.

The firm's optimization problem is

$$\max_{w} z_i[l(w)]^{\gamma} - wl(w) \tag{7}$$
$$\text{s.t. } l(w) = \bar{m}\Phi(w).$$

Necessary condition for optimality is

$$z_i\gamma l^{\gamma-1} - w = \frac{\Phi(w)}{\Phi'(w)}. \tag{8}$$

Since the elasticity of labor supply is $\epsilon_s(w) = \frac{\Phi'(w)w}{\Phi(w)}$, equation (8) can be rewritten as the familiar Lerner's condition relating a monopsonistic firm's wage markdown to the same elasticity:

$$\frac{z_i\gamma l^{\gamma-1} - w}{w} = \frac{1}{\epsilon_s(w)}.$$

Optimality dictates that the wage is pushed up to the point where the marginal gain from adding to labor equals the marginal cost from adding to the wage bill of infra-marginal workers.

Definition. A labor market equilibrium consists of wages $\{w_i\}$ and a map $\Phi(w)$ such that $\forall\ i$, w_i solves problem (7) given $\Phi(w)$, and $\Phi(w)$ is the probability that offering w leads to a hire, i.e.

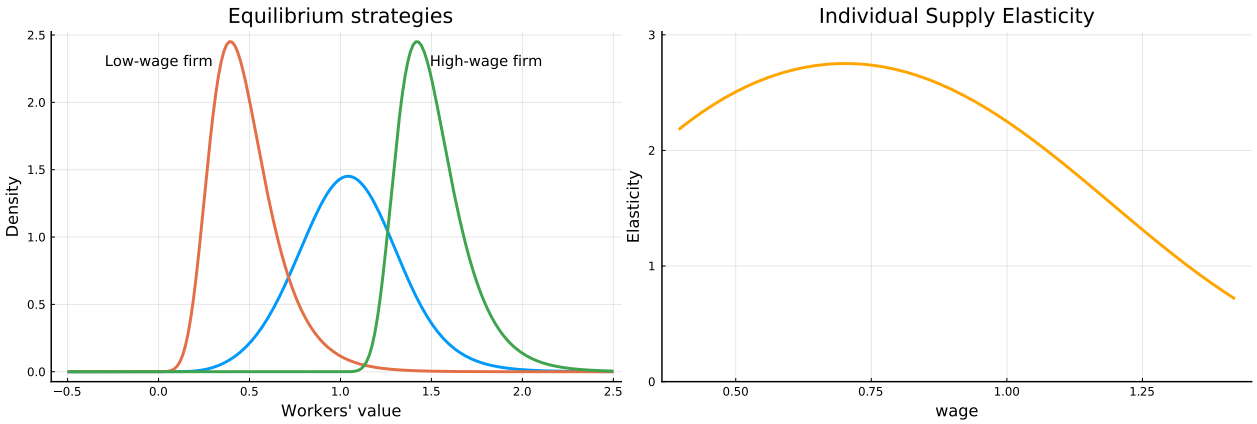
$$\Phi(w) = \sum_{j=1}^{N} \frac{1}{1 + e^{\frac{w_j - w}{\sigma}}} \mu_j.$$

3.1 Equilibrium firm-level pay

We now turn to describing salient features of the equilibrium allocation. When suitable, we illustrate such features by means of a numerical approximation of the equilibrium in a parameterized version of the model.

Consistent with the empirical evidence, our theory predicts that more productive firms are larger and pay higher wages. The reason is straightforward: Higher productivity means higher marginal product of labor and calls for more labor input. Since the mass of contacts (and job offers) is given at $\bar{m}$, hiring more workers requires paying higher wages.

The assumption that all firms have the same mass of contacts with workers is made for tractability. One could think of relaxing it by allowing firms to expend resources in order to contact more job candidates. In such scenario, increasing the contact rate would lead to greater hiring for given wage offer. As long as the cost of increasing the contact rate is increasing and convex, equilibrium would still feature more productive firms paying higher wages. Our assumption calls for that marginal cost to be zero until $\bar{m}$ and infinitely high thereafter.



(a) Equilibrium Strategies

(b) Supply elasticity

Figure 7: Illustration of equilibrium

Recall that the value of a wage offer w to a worker is the random variable $w+\varepsilon$. Figure 7a depicts the distribution of such random variable, in the case of equilibrium wage offers by a low- and a high- productivity firm, respectively. Because of the Gumbel-distributed amenity value, the probability of recruiting a worker is always positive, no matter the wage offered.

When making an offer to a job candidate, recruiters don't know which firm they are competing with. Consistent with rational expectations, we assume that they have perfect foresight of the unconditional equilibrium distribution of workers' values associated with the competing offer, depicted (in blue) at the center of Figure 7a.

The labor supply elasticity faced by firms is also an equilibrium object. Importantly, it varies across the productivity distribution. Two competing forces shape this variation. First, the semi-elasticity $\Phi'(w)/\Phi(w)$ decreases monotonically with w, reflecting decreasing marginal returns to wage increases. The intuition is that high wage firms already "picked all the low-hanging fruit," in the sense of winning most contests against low wage firms and

capturing most workers with relatively favorable amenity draws. Further wage increases attract few additional workers - and wage decreases lose few workers. Second, the elasticity $\epsilon_s(w)$ depends directly on the wage level. Since workers react to absolute wage differences, a 1% wage increase from a higher base wins more contests. Under mild conditions on the productivity density μ - particularly when productivity dispersion is high relative to amenity dispersion σ - the first effect dominates, and high-productivity firms operate on the relatively inelastic portion of the labor supply curve. In this case $\Phi'(w)/\Phi(w)$ decays at a fast enough rate that high wage firms become isolated in the wage space, facing limited local competition (few firms offering similar wages). This isolation allows them to maintain larger wage markdowns.[13] Figure 7b illustrates this pattern. Notably, the elasticity need not decline monotonically with productivity across the entire distribution, particularly at lower productivity levels where competition is more intense.

Both the level and the slope of the map between firm productivity and wages, an element at the center of our analysis, depend upon the parameter σ, which governs the importance of job amenities relative to wages in workers' choices. When $\sigma > 0$ is lower, workers become more sensitive to wage differences, intensifying wage competition among firms. This increases the labor supply elasticity facing each firm, reducing market power and markdowns, thereby pushing wages closer to marginal products and raising labor compensation.

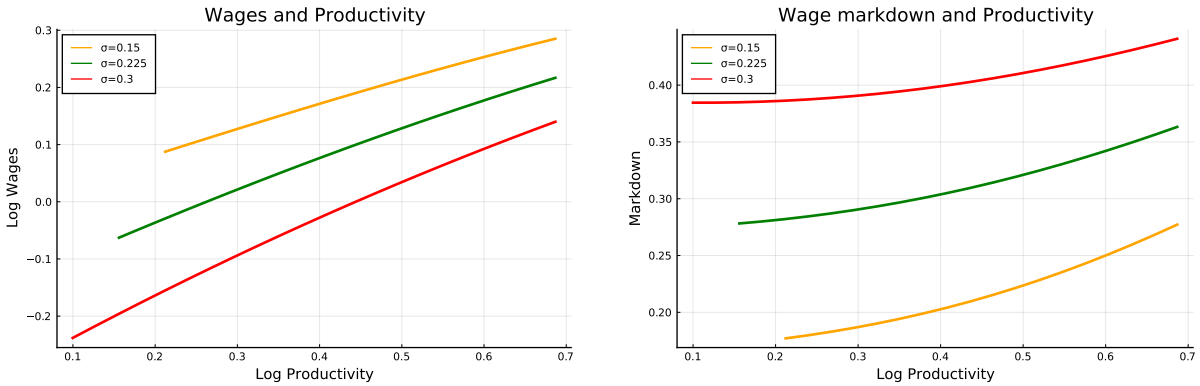
For σ arbitrary large, wages essentially don't matter for workers' choices. In the limit, the probability of the worker accepting a job is 1/2 no matter her wage offers. This corresponds to the least competitive scenario. Markdowns increase with σ. The slope of the wage-productivity map also increases with σ, since the marginal hiring yield of raising the wage is small when amenities are very relevant.

By way of illustration, Figure 8a depicts the equilibrium pay policy for three different levels of σ. A decline in σ is associated with higher wages and a decline in the slope of the map between productivity and wages.

Consistent with that, a decline in σ leads to a rise in firm-level supply elasticity and a decline in markdowns at all productivity levels, affecting low-productivity firms' conduct

[13] Gouin-Bonenfant (2022) identifies a complementary mechanism, where an increase in productivity dispersion shields high-productivity firms from competition, helping explain the decline in the labor share in Canada.

more than their higher-productivity peers'. See Figure 8b.



(a) Wages (b) Markdown

Figure 8: Wages, Markdown, and Productivity

In its simplicity, our model of the labor market is very powerful. Everything else equal, a lower variance of match-specific amenities σ maps into lower firms' market power, a weaker pass-through from productivity to wages, and lower cross-sectional variance of firm pay, with compression occurring primarily from the left tail.

4 Firm dynamics and minimum wage

We embed our model of the labor market in a more general framework that features firm entry and exit, a minimum wage, and unemployment. We will show that, besides the direct compression effect from below, an increase in the minimum wage may affect the variance of firm pay via its indirect effect on the elasticity of labor supply faced by unconstrained firms. In addition, it may also affect the equilibrium measure of firms.

Time is discrete and goes from $t = 0$ to infinity. However, since our analysis will focus on stationary equilibria, we suppress the time notation.

4.1 Workers

We retain a single worker type for moment. Workers, still in exogenous mass $N_w > 0$, have no access to either storage or financial markets and are therefore hand-to-mouth. We assume that their value of non-employment is $b + \varepsilon_b$, where ε_b is an amenity shock also conveniently

distributed according to a Gumbel distribution with location parameter 0 and the same scale parameter $\sigma > 0$. When confronting job offers $\{w_i, w_j\}$, a worker's optimization problem is now

$$\max \{w_i + \varepsilon_i, w_j + \varepsilon_j, b + \varepsilon_b\}.$$

The probability of accepting job i when facing alternative j is $\frac{e^{w_i/\sigma}}{e^{w_i/\sigma}+e^{w_j/\sigma}+e^{b/\sigma}}$.

4.2 Incumbent firms

Incumbent firms are distributed over the ordered discrete set $\{z_i\}_{i=1}^N$ according to the (endogenously determined) density μ. The dynamics of individual productivity are determined by the Markov chain Π, with $\Pr(z_{t+1} = z_j | z_t = z_i) = \pi_{ij} > 0$. Let μ_z denote the mean of the stationary distribution. Firms discount the future at the rate $\beta \in (0, 1)$, incur a random per-period operating cost $\kappa > 0$, expressed in units of the numeraire good, and exit when their expected discounted value is negative. We assume that κ is distributed according to the cdf Ψ, defined over the positive orthant and allowing a density ψ.

The adaptation of our simple labor market protocol to a dynamic framework is seamless. We assume that every period, each worker is matched with two randomly drawn firms, regardless of her prior labor market status. In turn, this means that in every period each firm is matched with a mass $\bar{m} \equiv 2N_w/N_f$ of workers, where N_f is now the endogenously determined mass of operating firms. That is, $N_f = \sum_j \hat{\mu}_j$, where $\hat{\mu}_j$ is the measure of operating firms sharing productivity z_j.

Firms are not compelled to make an offer to each worker they are matched with. We anticipate that, when the minimum wage is high enough with respect to their productivity, they will find it optimal to only extend $m < \bar{m}$ offers.

The optimization problem of an incumbent of productivity z_i can be conveniently stated

in recursive fashion as

$$v(z_i, \Phi, \bar{m}) = \int_0^{+\infty} \max \left\{ 0, \max_{\{m,w\}} \; z_i l(w)^\gamma - w l(w) - \kappa + \beta \sum_{j=1}^{N} \pi_{ij} v(z_j, \Phi, \bar{m}) \right\} d\psi(\kappa) \quad (9)$$

$$\text{s.t. } \; l(w) = m\Phi(w),$$

$$w \geq \underline{w}, \quad (10)$$

$$m \leq \bar{m}. \quad (11)$$

The constraint (10) ensures that firms comply with minimum wage legislation. Condition (11) requires that firms extend offers only to the workers they are matched with.[14]

The recursive formulation emphasizes the time-invariant aggregate state variables: $\Phi(w)$ which remains the probability of hiring a worker when offering a wage w, while $\bar{m}$ is now an endogenous matching rate induced by N_f. We have

$$\Phi(w) = \sum_{j=1}^{N} \left[\left(1 - \frac{m_j}{\bar{m}}\right) \frac{1}{1 + e^{\frac{b-w}{\sigma}}} + \frac{m_j}{\bar{m}} \frac{1}{1 + e^{\frac{b-w}{\sigma}} + e^{\frac{w_j - w}{\sigma}}} \right] \mu_j, \quad \mu_j = \frac{\hat{\mu}_j}{N_f}. \quad (12)$$

Differently from above, a recruiter pursuing a worker that is also matched with a productivity-z_j firm will face a competing offer only with probability $\frac{m_j}{\bar{m}} \leq 1$. It follows that the unconditional probability of not facing a competing wage offer is $\sum_{j=1}^{N} \left(1 - \frac{m_j}{\bar{m}}\right) \geq 0$.

The solution of problem (9) will feature a non-decreasing set of thresholds for the operating cost $\{\kappa_i\}_{i=1}^{N}$ such that firms with productivity z_i will exit if and only if $\kappa > \kappa_i$. The conditional probability of exit for productivity-z_i firms is therefore $1 - \Psi(\kappa_i)$.

4.3 Entry

Entry is modeled along the lines of Clementi and Palazzo (2016). Every period, there is an exogenous mass of potential entrants $M > 0$ that draw signals $\eta > 0$ about their initial productivity from the continuous distribution H.

Conditional on the signal, initial productivity is drawn according to $\Pr(z \leq z_i|\eta) = G_i(\eta)$. For each η, the collection $\{G_i(\eta)\}_{i=1,N}$ is a cumulative distribution function. Crucially, G_i is

[14] Condition (11) resembles Berger et al.'s (2025)'s rationing constraint.

strictly decreasing in η for all i, that is, a higher signal improves the odds of higher initial productivity.

Firms must pay a cost $c_e > 0$, expressed in units of numeraire good, in order to learn their initial productivity. Therefore, a potential entrant will pay c_e and learn its initial productivity if and only if

$$\sum_{i=1}^{N} v(z_i, \Phi, \bar{m}) g_i(\eta) \geq c_e,$$

where g_i is the pdf associated to G_i. It follows that the measure of entrants characterized by initial productivity z_i will be

$$\hat{\mu}_e(z_i) = \Psi(\kappa_i) M \int_{\eta^*} g_i(\eta) dH(\eta).$$

Entry and exit dynamics differ from the original Hopenhayn (1992) model, which features no entry selection, an exogenous entry productivity distribution, and a sharp exit threshold that left-truncates the stationary distribution. Our model generates an endogenous entry distribution determined by the signal threshold η^* and exit rates that decline smoothly with firm-level productivity. This specification eliminates the truncation, better matches observed entry and exit dynamics, and provides greater flexibility in tracing how changes in wage-setting forces propagate through the firm distribution.

4.4 Equilibrium

Definition. A recursive stationary equilibrium consists of

- Policies for wages and job offers $\{w_i, m_i\}_{i=1}^N$,
- Thresholds η^*, $\{\kappa_i\}_{i=1}^N$,
- Firm-level labor supply function $\Phi(w)$,
- Measure of active firms $\hat{\mu}$ and associated mass $N_f \equiv \sum_{i=1}^N \hat{\mu}_i$,

such that

- Wages and job offers solve problem (9) given $\Phi(w)$ and $\bar{m} = 2N_w/N_f$,

- $\hat{\mu}$ is consistent with η^* and $\{\kappa_i\}_{i=1}^N$, i.e. for all i,

$$\hat{\mu}_i = \Psi(\kappa_i)\left[M\int_{\eta^*} g_i(\eta)dH(\eta) + \sum_{j=1}^{N}\pi_{ji}\hat{\mu}_j\right],$$

- $\Phi(w)$ is the probability of hiring a worker, i.e.

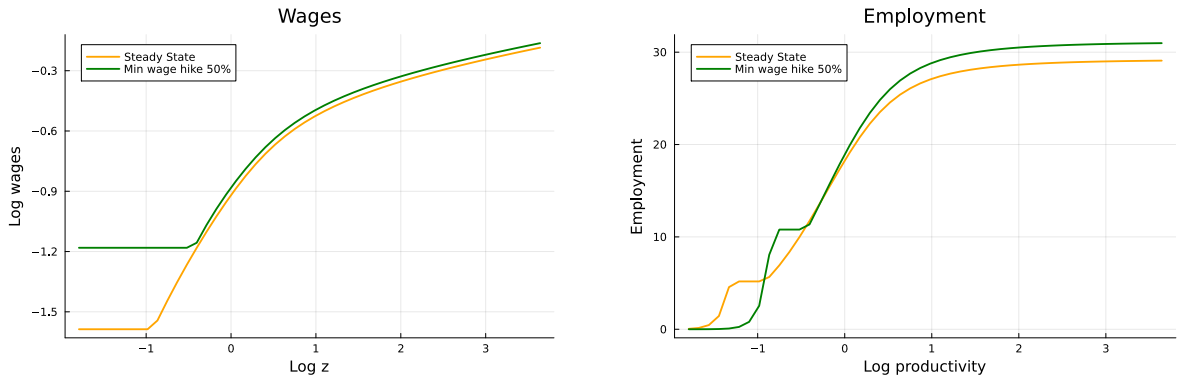
$$\Phi(w) = \sum_{j=1}^{N}\left[\left(1-\frac{m_j}{\bar{m}}\right)\frac{1}{1+e^{\frac{b-w}{\sigma}}} + \frac{m_j}{\bar{m}}\frac{1}{1+e^{\frac{b-w}{\sigma}}+e^{\frac{w_j-w}{\sigma}}}\right]\mu_j, \quad \mu_j = \frac{\hat{\mu}_j}{N_f}.$$

Appendix E provides details on the numerical solution of the equilibrium.

4.5 Raising the minimum wage

We now turn to illustrating the effect of a minimum wage hike, by means of a simple comparative statics exercise. First, we parameterize the model in order to reproduce key features of the Portuguese economy in the period 1993-97, along the lines of the quantitative evaluation carried out in Section 5.

Even though the exercise conducted here is merely illustrative, we make sure that the initial steady-state allocation is consistent with some basic empirical regularities relating to wage dispersion and firm dynamics.



(a) Wages (b) Employment

Figure 9: Comparative statics with respect to the minimum wage

Refer to Figure 9. In equilibrium, firms can be conveniently partitioned in three connected

and ordered subsets of the productivity support. We will refer to them as regions 1, 2, and 3, respectively, following the convention introduced by Berger et al. (2025).

The firms closest to the productivity lower bound (region 3) set $w = \underline{w}$ and $m < \bar{m}$. Their productivity is so low that, if they extended offers to all workers they are matched with, their marginal product of labor would fall short of the minimum wage. In other words, region 3 firms are price takers at the minimum wage and hire only up to the level for which the marginal product of labor equals the minimum wage. Consequently, some workers who meet these firms receive no job offer from them. For a sufficiently high minimum wage, employment at low productivity firms in region 1 falls below the unconstrained monopsonistic optimum.

Moving to the right along the productivity support, we find firms who are still constrained by the minimum wage ($w = \underline{w}$) but whose productivity is large enough to warrant making offers to all workers they are matched with, i.e. $m = \bar{m}$. This is region 2, where firms' marginal product of labor is strictly between the minimum wage and the marginal cost of labor. Their wage markdown is strictly positive, but falls short of the unconstrained optimum, while their employment exceeds the unconstrained optimum.

Further to the right is region 1, where firms' productivity is so high that wage setting results in equating the marginal product of labor to its marginal cost. These are textbook-like unconstrained monopsonists that set $m = \bar{m}$ and $w > \underline{w}$.

Next, we describe and analyze the impact on such initial allocation of a 50% increase in minimum wage – substantially larger than the actual increase in Portugal's minimum wage between 1993 and 2019. This is not an easy feat, as the greater minimum wage will trigger equilibrium changes in the volume and selection of entrants as well as on exit, which in turn will reshape the measure of incumbent firms $\hat{\mu}$ and feed back into firms' decision making. We begin with the incumbents' policies and then move to the distribution of wages.

The conduct of firms initially in region 1 is independent from most of the endogenous elements listed above. When forced to pay a higher wage, they further reduce their offers and contract employment, moving along the demand curve. This captures the standard negative employment effect of a higher minimum wage.

Firms initially in region 2 are also required to pay higher wages. Those among them with relatively low productivity reduce the volume of offers below their upper bound and lower

employment, thereby falling into region 3. It follows that a higher minimum wage widens region 3.

Firms initially in Region 2 but with relatively higher productivity are like the firms that are called to expand employment in the textbook treatment of minimum wage analysis under monopsony. Everything else equal, they will accomplish this by moving upward along the supply curve and reducing their markdown.

Firms in region 1 are not directly affected by the policy change. However, since the equilibrium distribution of firms, their mass, and the wages charged in Regions 3 and 2 are now different, firms in Region 1 face different individual labor supply functions.

Recall that the wage setting equation for such firms is

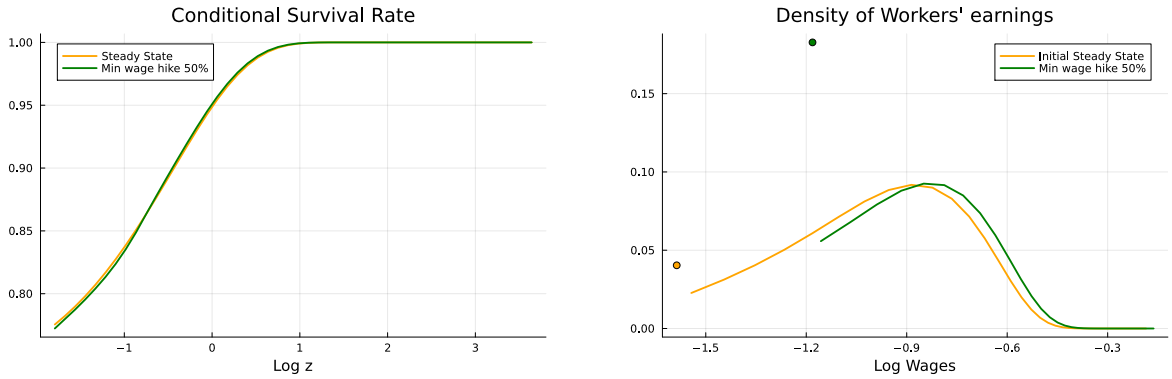
$$\frac{z\gamma\left[\bar{m}\Phi(w)\right]^{\gamma-1}-w}{w}=\frac{\Phi(w)}{w\Phi'(w)},$$

where $\Phi(w)$ is given by (12) and $\bar{m}=2N_w/N_f$. Everything else equal, higher wages by firms in regions 3 and 2 cause the elasticity of labor supply for region 1 firms to rise. This is often referred to in the literature as the spillover effect of a minimum wage hike, that in principle may lead to a propagation of the minimum wage rise to unconstrained firms. However, the region 3 firms that increase their wages are also less likely to extend offers to the workers they match with. This effect tends to reduce the labor supply elasticity for region 1 firms.

We now turn to the effects on wage setting deriving from equilibrium adjustment in the measure of operating firms. An increase in the minimum wage lowers the expected value of entry. The free-entry condition calls for an increase in the entry threshold η^* and a lower mass of entrants. In turn, this results in a smaller equilibrium mass of firms N_f. Every incumbent meets a larger mass of workers $\bar{m}$. Everything else equal, this is an anti-competitive effect that leads to lower wages.

Firm selection improves. Entrants have greater productivity on average, and, as documented in Figure 10a, survival rates decline for low-productivity firms. Better selection on productivity goes in the direction of raising the elasticity of supply faced by Region-1 firms, putting upward pressure on wages.

Figure 10b illustrates the effect of the higher minimum wage on the distribution of work-



(a) Conditional survival probability

(b) Workers' wage density

Figure 10: Equilibrium effects of a minimum wage hike

ers' earnings, our primary object of interest, with dots representing mass points at the minimum wage level. The variance of log wages declines by 3 percentage points.[15]

The wage compression from below is readily visible: the fraction of workers compensated at the minimum wage (the minimum wage bite) rises by more than ten percentage points. The impact on the rest of the wage density is the net result of the equilibrium adjustments in firms' wages policies discussed above, and of the reallocation of workers across firms with different productivity. The firms' employment policies depicted in Figure 9, together with the right-shift of the density μ, suggest that the rise in the minimum wage reallocates workers toward higher-productivity firms.

To quantitatively assess how various mechanisms contribute to the decline in wage dispersion, we perform a decomposition exercise presented in Figure 11. The variance associated with the initial steady state is shown at step 0.

Step 1 isolates the bite effect of the minimum wage increase. We maintain the initial allocation of workers and simply impose that firms offer the maximum of their initial wage and the new minimum wage. Everything else is held constant: the set of operating firms, their employment levels, and all wages above the minimum.

Step 2 adds the pro-competitive spillover effect on region 1 firms. We continue to hold the set of operating firms and their employment levels fixed at initial values, but now allow

[15]This corresponds to a reduction in the variance of firm fixed effects of very similar magnitude The variances of log wages and firm fixed effects differ because the AKM model is misspecified for our economy, since firm–level productivity and therefore wages are time-varying. Some of the dispersion in wages is therefore incorrectly attributed to workers. The bias is not very significant.

firms to set wages optimally.

Step 3 permits all firms to optimize both wages and job offers, while still holding the measure of operating firms constant. The additional effect on the log wage variance at this step stems from reduced job offers by region 3 firms and the resulting reallocation of employment from low to high productivity firms.

Step 4 incorporates the full equilibrium effects, including changes in the measure of operating firms.

This decomposition reveals that the bite effect of higher wages at the bottom of the distribution is the primary driver of wage compression. However, the adjustments arising from firms' optimization and general equilibrium forces partially offset this mechanism. Wage compression is relatively muted, even following a large minimum wage increase.

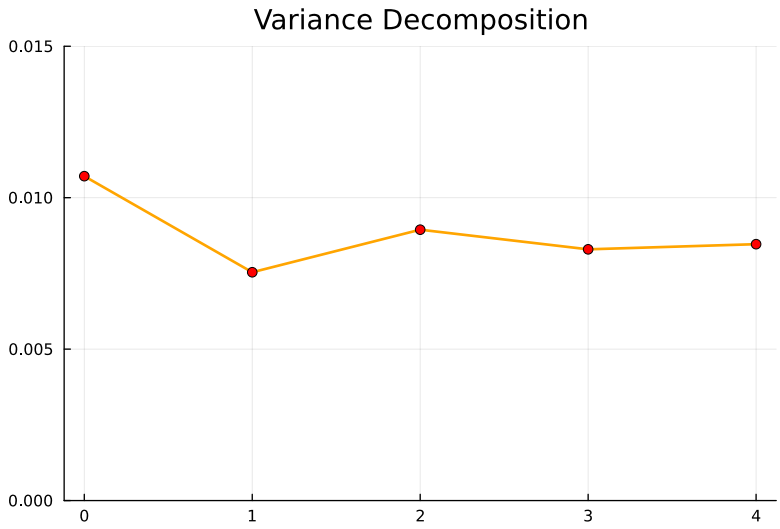


Figure 11: Decomposing the effects of a minimum wage hike on wage dispersion

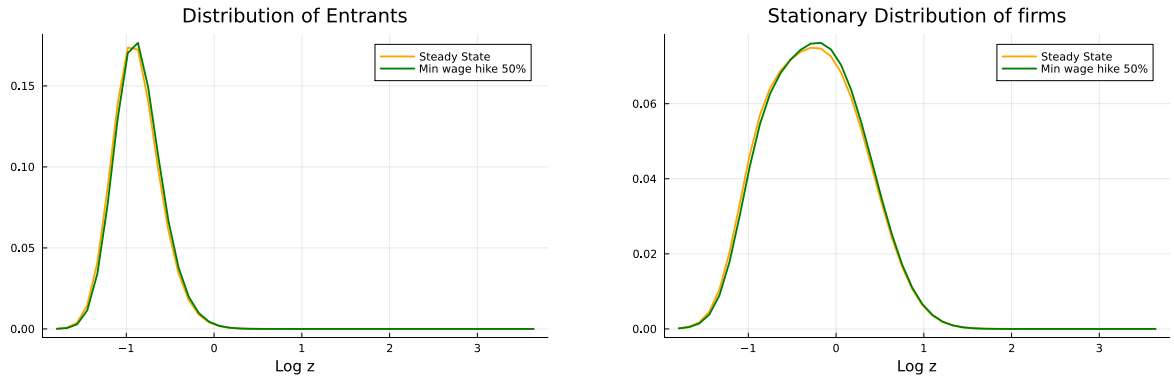
4.5.1 Discussion

While the minimum wage increase does spill over to high-productivity unconstrained firms, it induces a relatively uniform wage increase (Figure 9a). Holding the measure of firms and their employment levels constant, this necessarily increases overall wage inequality, given the minimum wage bite. The spillover effect therefore produces a modest *increase* in the standard deviation of log wages of about 0.25 log points, partially offsetting the initial compression

due to the bite.[16]

The remaining effects do compress wages, though the magnitude is modest. First, constrained firms reduce vacancies, making it easier for unconstrained firms to compete since they face an increased probability of being matched with workers who lack alternative offers. However, because region 1 firms constitute a relatively small fraction of the economy in equilibrium, the resulting compression induced by lower wages is limited.

Second, firm dynamics also contribute negligibly to wage compression. Figure 12 shows the distribution of firms is essentially invariant to the minimum wage increase. While higher minimum wages reduce survival rates among low-productivity firms, the effect is very small. Current profits drop significantly for these firms, but higher expected future profits - driven primarily by reduced competition, due to a lower mass of operating firms - largely offset this decline. Furthermore, since low productivity is estimated to grow relatively quickly, low-productivity firms anticipate escaping quickly the minimum wage constraints. The minimum wage increase therefore has minimal impact on exit or entry decisions.



(a) Distribution of entrants

(b) Distribution of firms over productivity

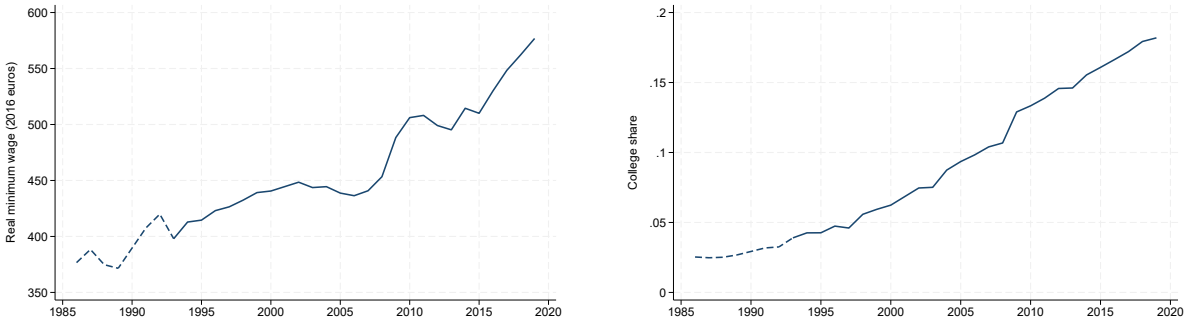
Figure 12: Industry equilibrium

[16]Spillover effects from a large (128%) minimum wage increase have been found to be quite significant in Brazil according to Engbom and Moser (2022), where a one percent increase in the minimum wage is associated with a *decline* in the standard deviation of wages of about 20 log points. They exploit cross-state variation in the bindingness of Brazil's federal minimum wage, and their calculation holds fixed a reference wage percentile sufficiently high to avoid any impact from the minimum wage. Using a distribution regression framework and time-series variation in the minimum wage level, Oliveira (2023) finds significant spillover effects for Portugal from the minimum wage increase since 2006, reaching just above median wages and accounting for nearly 1/2 of the observed wage compression (spillovers reduce dispersion by 3 log points). These estimates are not entirely comparable to our step 2, since they also capture equilibrium effects, and possibly confounding effects.

5 Accounting for the recent decline in Portugal's earnings inequality

We now address the main question of this paper: What accounts for the recent decline in Portugal's firm pay inequality? The candidate explanations are the two mechanisms considered in isolation in the previous sections: a decrease in the importance of amenities, and an increase in the minimum wage. Here we assess their relative importance in the full quantitative model.

Figure 13a illustrates the evolution of Portugal's the general (non-agricultural) monthly real minimum wage since the mid-1980s. The minimum wage is set nationally for annual full-time contracts (14 monthly payments, including two vacation bonuses, and a 40-hour workweek), and covers all workers. By 2019, the minimum wage had risen 40% above its mid-1990s level.



(a) Real minimum wage

(b) Share of college graduates

Figure 13: Minimum wage and share of college graduates among workers

Alongside this minimum wage increase, however, other structural changes occurred that also affect firm pay inequality. Figure 13b shows a dramatic rise in the share of college graduates among all workers, from just under 5% in the mid-1990s to 19% in 2019. This educational expansion matters for three reasons. First, the minimum wage primarily affects non-college workers. Second, the increased relative supply of college-educated workers reduced the college wage premium and raised non-college wages, thereby reducing the bite of the minimum wage. Third, because college-educated workers exhibit greater firm pay variance, this compositional shift mechanically works against overall compression in firm

pay.

To account for these dynamics, we extend the baseline model by distinguishing between two worker types, labeled skilled and unskilled. Let N_i denote the mass of workers of type $i \in \{s, u\}$. We normalize the total mass of workers to 1, so N_i represents the population share of type i workers.

The types differ in both productivity and preferences. Regarding heterogeneity in productivity, we assume that production is a CES aggregate of the services of two types of labor:

$$y = z_i \left[\chi l_s^{\frac{\nu-1}{\nu}} + (1-\chi) l_u^{\frac{\nu-1}{\nu}} \right]^{\frac{\gamma}{\nu-1}},$$

with $\chi \in (0,1)$, $\nu > 0$ and $0 < \gamma \leq \nu$.

Regarding preferences, we allow the variance of amenity shocks to differ across skill levels. Clearly, this assumption is motivated by the evidence presented above suggesting that the variance of firm-level pay vary substantially across levels of educational attainment.

The two labor markets are perfectly segmented. Each firm receives a mass $2N_i/N_f$ of applications in the market for type i. For the sake of brevity, we do not lay out an amended definition of equilibrium. The interested reader is referred to the Appendix.

Our strategy is to calibrate parameters so that the initial steady-state equilibrium of the model matches key empirical moments from the 1993-1997 period. In particular, we target the variance of firm fixed effects by skill, a measure of the minimum wage bite (minimum wage over the median wage), and the college premium.

We then consider a final steady-state in which only parameters related to three sets of forces differ from the initial equilibrium: the two drivers of wage compression ($\underline{w}$, σ_u, σ_s), the race between education and technology (N_s, χ), and aggregate wages and firm turnover (μ_z, c_e). The key element is that the two drivers of wage compression are set to ensure that (i) the minimum wage increases by 30%, as observed in the data, and (ii) the model matches the observed compression in firm pay inequality by skill - i.e., σ_u, σ_s are chosen to account for the residual variance not explained by the minimum wage increase alone. The remaining parameters ensure that the final steady state is consistent with basic observations from the 2014-19 period for aggregate wage growth, the college wage premium, the incidence of the

Parameter	Value	Target	Model	Data
External calibration				
β	0.96	yearly interest rate of 4%		
ν	1.4	Katz and Murphy (1992)		
γ	1.26	returns to scale of 0.9		
ρ_z	0.89	AR(1) for firm–level productivity		
σ_z	0.22	AR(1) for firm–level productivity		
μ_z	0.0	normalization		
N_s	0.0435	share of college educated		
Internal calibration				
σ_u	0.0487	variance of FFE (college)	0.081	0.080
σ_s	0.1776	variance of FFE (non-college)	0.177	0.152
b	0.11	unemployment rate	0.060	0.067
χ	0.226	college premium (log)	0.836	0.837
$\underline{w}$	0.125	ratio of $\underline{w}$ to median wage	0.552	0.567
M	0.0394	avg employees per firm	14.45	14.6
δ	1.716	entrant/incumbent relative size	0.425	0.232
η^*	0.312	entrant/incumbent stdev size	0.656	0.696
κ dist params	several	conditional survival regression		
		intercept	-1.31	-1.503
		slope	-0.34	-0.351

Table 6: Calibration – Initial steady-state targeting 1993-1997 moments

minimum wage, and the approximately constant firm exit rates observed between the initial and the final periods.

Table 6 reports the parameters associated with the initial steady-state. We begin with the ones that remain invariant across steady-states. Some are set externally. According to Acemoglu and Autor (2011), the summary of the evidence for the U.S. suggests an elasticity of substitution between skilled and unskilled labor in the 1.4-2.0 range. Centeno and Novo (2014) provide similar estimates for Portugal, in the 1.5-2.6 range. We set $\nu = 1.4$, which is the estimate in the original Katz and Murphy (1992) study. We also set $\gamma = 1.26$, which implies returns to scale of 0.9. Future profits are discounted at a yearly rate of 4%, which is standard.

The Markov chain Π is set in such a way that the implied stochastic process approximates an AR(1) with autocorrelation $\rho_z = 0.89$ and conditional standard deviation $\sigma_z = 0.22$. We recovered these values when estimating an AR(1) process for total factor productivity in our sample of firms (available from 2012 onwards).

We assume that entering firms draw the productivity signal η from a lognormal distribution with mean $\mu_z - \delta$ and variance $\sigma_z^2/(1-\rho_z^2)$. Initial productivity is then drawn from a discretized lognormal distribution with mean $(1-\rho_z)\mu_z + \rho_z \log(\eta)$ and variance σ_z^2.[17] The shift parameter $\delta > 0$ controls how much less productive entrants are relative to incumbents. We choose it to match the average size of entrants relative to incumbents, and set the mass of potential entrants M to target the average number of employees per firm, across all firms. In the data, entrant sizes are within-sector and refer to the 2004-2019 period, since information on entry is only available from 2004 onwards. We set the entry cost c_e to target the empirical ratio of entrant to incumbent size dispersion.[18]

The distribution of the operating cost is approximated by a piecewise uniform function (see Appendix E.4.1), parameterized so that firm survival rates conditional on size are consistent with the data. We ran a linear regression of log exit rates on (within-sector) employment quintiles both in the data and in the model, and require that the model's regression coefficients match those in the data.

[17] The unconditional distribution of initial log productivity has mean $\mu_z - \rho_z\delta$ and variance $\sigma_z^2/(1-\rho_z^2)$.

[18] Table 6 reports the entry signal threshold η^* rather than the underlying entry cost c_e. The two are related through the free-entry condition. The same applies to Table 7.

Parameter	Value	Target	Model	Data
External calibration				
$\underline{w}$	0.125×1.3	min wage increase, 1993-97 to 2014-19		
N_s	0.1693	share of college educated		
Internal calibration				
η^*	0.366	exit rate	0.067	0.095
μ_z	0.25	ratio of $\underline{w}$ to median wage	0.707	0.64
χ	0.39	college premium (log)	0.535	0.54
σ_u	0.0390	variance of FFE (college)	0.039	0.035
σ_s	0.124	variance of FFE (non-college)	0.070	0.059

Table 7: Calibration, time-varying parameters – Final steady-state, targeting 2014-2019 moments

For the moment we impose the same reservation value for both types of worker, $b_s = b_u = b$, which is consistent with an unemployment rate similar to the data.

We now turn to the parameters that vary across steady-states - Table 7 displays their final steady-state values. The initial level of the minimum wage $\underline{w}$ is calibrated internally, to match the 1993-1997 measure of the minimum wage bite. For the final steady-state, we apply the 30% observed growth rate in the real minimum wage to the initial value.

The variances of amenity shocks σ_s and σ_u are set to generate data-conforming variances of skill-specific firm-fixed effects in each period.

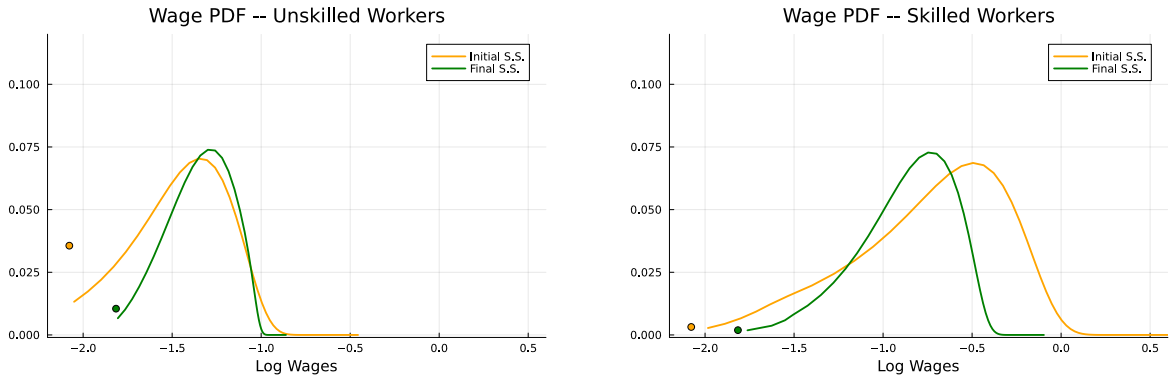
The share of college-educated workers N_s equals its period-specific empirical counterpart (the total mass of workers is normalized to 1 in both steady-states), and the relative productivity of skilled labor χ is set to match the period-specific college premium. Notably, unlike most developed economies, Portugal experienced a decline in the college premium alongside a significant expansion in college-educated labor supply (Centeno and Novo, 2014; Oliveira et al., 2023). Skill-biased technological change was insufficient to offset this supply increase and, in Goldin and Katz's (2008) language, education won the race against technology in Portugal.

Finally, we normalize $\mu_z = 0$ in the initial steady-state, and allow it to increase in the final steady-state to match the 2014-2019 ratio of minimum to median wages. The model is then consistent with the observed growth in aggregate real wages between the two periods (15%).

We note that, combined with the calibrated education-technology dynamics, the model also replicates the observed declined in the real wages of college-educated (-19%). We also allow the entry cost c_e to scale with average productivity, keeping the firm exit rate constant over time as in the data (though the actual level is lower than in the data).

5.1 Results

In terms of bottom-line, the model does fully account for the decline in the variance of firm fixed-effects under the baseline calibration. This is by construction, since we are matching both the skill-specific variances of firm-fixed effects, and the workforce composition. We are mostly interested in evaluating the relative contributions of the two main drivers of firm pay inequality, minimum wage and amenities, and understanding their mechanisms.



(a) Unskilled workers

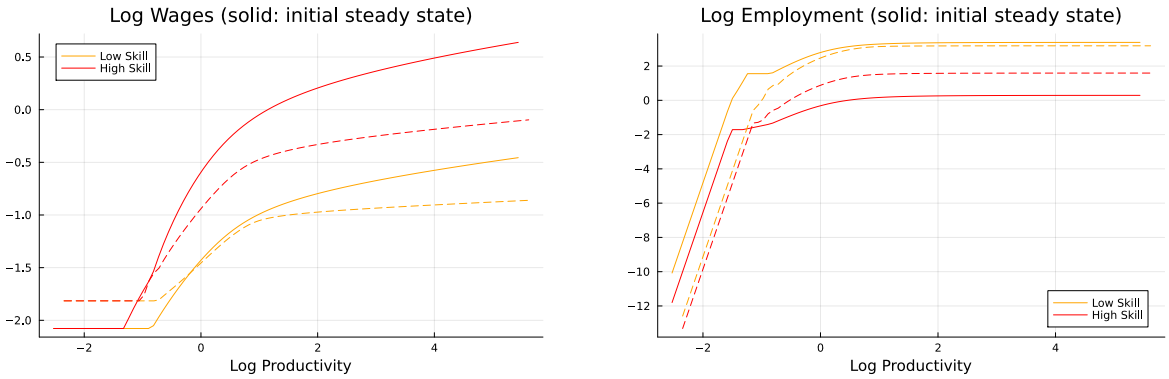
(b) Skilled workers

Figure 14: Wage distributions by skill, main experiment

Figure 14 illustrates the evolution of the wage distributions by worker type, with dots again representing mass points at the minimum wage level.

The distribution dynamics reflect changes in both location and shape. For unskilled workers, a clear rightward shift emerges, driven primarily by aggregate productivity gains between periods (see also Figure 16). The distribution also compresses substantially, predominantly from the left tail. While the minimum wage increase contributes to this compression mostly through the bite effect (Section 4), this mechanism is now more muted due to wage growth at the bottom of the unskilled distribution.

The decline in σ_u (amenities less important) drives this disproportionate increase in un-



(a) Wage Policy

(b) Employment Policy

Figure 15: Wage and Employment Policies (Solid: Initial. Dashed: Final), main experiment

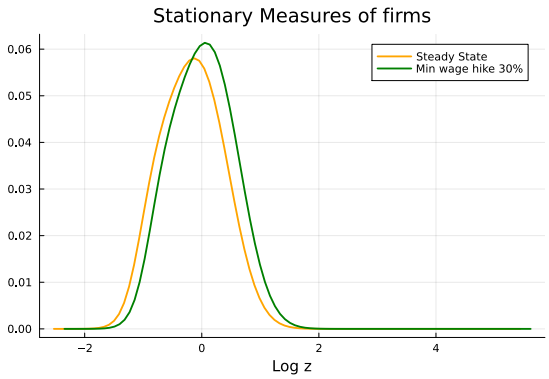


Figure 16: Measure of Firms, main experiment

skilled wages at the bottom. The mechanism operates through labor supply elasticities. As in the static model of Section 3, low-productivity firms face labor supply that is not only more elastic, but also more sensitive to the importance of job amenities. When wage competition intensifies in the final steady state, these firms must drastically reduce markdowns and raise wages. This compression triggers two reinforcing effects. First, marginal low-productivity firms exit the market, producing a more selected firm distribution. Second, employment reallocates from surviving low-productivity firms toward high-productivity firms (Figure 15b). Both effects contribute to reducing the minimum wage bite. Figure 15a confirms this mechanism: optimal wage policies prescribe higher unskilled wages at the bottom for unconstrained low-productivity firms, whereas wages actually decline toward the top - as high-productivity firms also substitute away from unskilled labor (Figure 15b).[19]

[19]Figure 15a is essentially the model counterpart of the empirical Figure 6. The differences are that the empirical figures use normalized firm pay (removing information about wage levels), and plot wages against

For skilled workers, there is a noticeable decrease in average wages, driven primarily by the dramatic expansion in the relative supply of college-educated workers. However the distribution remains well above the minimum wage in both periods, rendering the constraint essentially non-binding for this group. Reinforcing this low bite, firm pay compresses mainly from the bottom, reflecting the decline in σ_s, which our model infers is especially important for skilled workers (see also Figure 15a).

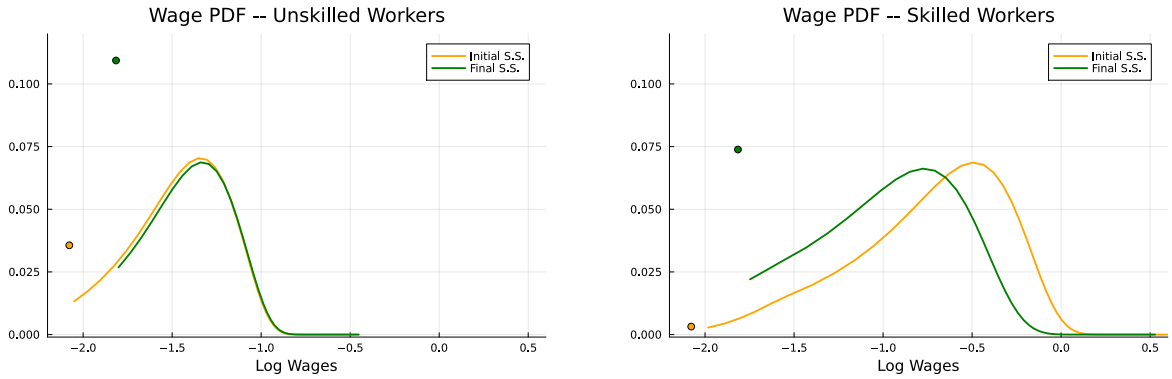
The key implication is that, despite the significant minimum wage increase, equilibrium forces push firms and unskilled workers away from the constraint, while skilled workers remain largely unaffected by it. The primary driver of the firm pay compression appears to be heightened wage competition and reduced firms' labor market power, rather than the minimum wage increase itself.[20]

To isolate these effects, we performed an alternative calibration of the final steady-state, identical to the baseline (Table 7), except that we held σ_u and σ_s constant at their initial steady-state values. In this counterfactual, we didn't target any moment relating to firm pay dispersion - that is, we have ignored the last two rows of the bottom panel in Table 7 - but continued to match the first three moments, while preserving the evolution of the minimum wage and labor force composition, as in the top panel. Our goal is to isolate the relative contribution of the minimum wage increase while preserving the basic dynamics of wage growth and firm exit rates. The key finding is that the variance of firm pay decreases by only 1.9 log points in this experiment, which is 18% below the observed decline of 10.5 log points. This result aligns with our findings from Section 4, which carry through to the full model.

Figure 17 displays the wage distributions by worker type underlying this counterfactual. The bite of the minimum wage is significantly higher, now affecting even skilled workers.

the distribution of jobs across productivity levels (rather than firms over productivity, as in the model figure). The model replicates the key "hockey stick" empirical dynamics, with the exception of firm pay for skilled workers below the productivity threshold, as discussed previously.

[20] A natural question is whether there is evidence of a negligible increase in the minimum wage bite over time, consistent with the equilibrium forces our model highlights. There is no straightforward answer to this question. The fraction of workers with total earnings (base plus supplementary) equal to the minimum wage increases from 2.3% to 3.3%, which appears consistent with our model. However, this likely understates the true bite, since the minimum wage primarily constrains base earnings for a majority of workers (the fraction of workers with base earnings at the minimum wage increases from 3.7% to 17.3%).



(a) Unskilled workers

(b) Skilled workers

Figure 17: Wage distributions by skill (minimum wage counterfactual)

Away from the minimum wage, however, the model fails to replicate the observed compression of firm pay from the bottom of the distribution. Among skilled workers, firm pay variance actually increases slightly. These results indicate that spillover effects from the minimum wage are minimal, for the reasons outlined in Section 4: constrained firms pose little threat to unconstrained firms, and any spillovers are level effects uniform across the productivity distribution. The results also indicate that intensified wage competition is essential to explaining the observed wage compression patterns, not only for college-educated workers, but also for non-college workers.

6 Conclusion

We have documented a substantial decline in earnings inequality in Portugal over the last three decades, driven exclusively by compression in firm pay. This compression reflects primarily a decline in the pass-through from firm–level productivity to wages. Using a tractable model of firm dynamics with monopsonistic labor market competition, we identify the mechanisms underlying this evolution. While minimum wage increases contributed to the decline in firm pay inequality, their impact was quantitatively modest. Equilibrium forces - including faster wage growth for non-college workers, improved firm selection on productivity, and employment reallocation toward high-productivity firms - limited the minimum wage's bite. Moreover, constrained firms posed little competitive threat to unconstrained employers due to their low mass and reduced vacancy posting, generating minimal spillover effects up

the wage distribution.

The primary driver of firm pay compression was instead a fundamental shift in the nature of wage competition: a decline in the importance of non-wage job amenities. This shift intensified competition across the entire productivity distribution, particularly among low-productivity firms. Our quantitative analysis reveals that this mechanism, rather than minimum wage increases, accounts for the bulk of the observed firm pay compression. More broadly, our findings highlight that understanding the evolution of earnings inequality requires careful attention to the structural forces shaping wage competition and firm-level wage-setting behavior, as well as to the equilibrium responses these forces generate.

References

Abowd, J. M., F. Kramarz, and D. N. Margolis (1999): "High Wage Workers and High Wage Firms," *Econometrica*, 67, 251–334.

Acemoglu, D. and D. Autor (2011): "Skills, Tasks and Technologies: Implications for Employment and Earnings," in *Handbook of Labor Economics*, ed. by O. Ashenfelter and D. Card, Elsevier, vol. 4 of *Handbook of Labor Economics*, chap. 12, 1043–1171.

Alvarez, J., F. Benguria, N. Engbom, and C. Moser (2018): "Firms and the Decline in Earnings Inequality in Brazil," *American Economic Journal: Macroeconomics*, 10, 149–189.

Babina, T., W. Ma, C. Moser, P. P. Ouimet, and R. Zarutskie (2019): "Pay, Employment, and Dynamics of Young Firms," Opportunity and Inclusive Growth Institute Working Papers 21, Federal Reserve Bank of Minneapolis.

Berger, D., K. Herkenhoff, and S. Mongey (2025): "Minimum Wages, Efficiency, and Welfare," *Econometrica*, 93, 265–301.

Bilal, A., N. Engbom, S. Mongey, and G. L. Violante (2021): "Labor Market Dynamics When Ideas are Harder to Find," NBER Working Papers 29479, National Bureau of Economic Research, Inc.

Bonhomme, S., K. Holzheu, T. Lamadon, E. Manresa, M. Mogstad, and B. Setzler (2020): "How Much Should we Trust Estimates of Firm Effects and Worker Sorting?" NBER Working Papers 27368, National Bureau of Economic Research, Inc.

Burdett, K. and D. T. Mortensen (1998): "Wage Differentials, Employer Size, and Unemployment," *International Economic Review*, 39, 257–273.

Card, D. and A. R. Cardoso (2022): "Wage Flexibility under Sectoral Bargaining," *Journal of the European Economic Association*, 20, 2013–2061.

Card, D., A. R. Cardoso, J. Heining, and P. Kline (2018): "Firms and Labor Market Inequality: Evidence and Some Theory," *Journal of Labor Economics*, 36, 13–70.

Card, D., A. R. Cardoso, and P. Kline (2015): "Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women *," *The Quarterly Journal of Economics*, 131, 633–686.

Card, D., J. Heining, and P. Kline (2013): "Workplace Heterogeneity and the Rise of West German Wage Inequality," *The Quarterly Journal of Economics*, 128, 967–1015.

Cardoso, A. R. and P. Portugal (2005): "Contractual Wages and the Wage Cushion under Different Bargaining Settings," *Journal of Labor Economics*, 23, 875–902.

Centeno, M. and A. A. Novo (2014): "When Supply Meets Demand: Wage Inequality in Portugal," *IZA Journal of European Labor Studies*, 3, 3–23.

Clementi, G. L. and B. Palazzo (2016): "Entry, Exit, Firm Dynamics, and Aggregate Fluctuations," *American Economic Journal: Macroeconomics*, 8, 1–41.

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

Engbom, N., C. Moser, and J. Sauermann (2023): "Firm pay dynamics," *Journal of Econometrics*, 233, 396–423.

Goldin, C. and L. F. Katz (2008): *The Race Between Education and Technology*, Cambridge, Massachussetts and London, England: Harvard University Press.

Gouin-Bonenfant, É. (2022): "Productivity Dispersion, Between-Firm Competition, and the Labor Share," *Econometrica*, 90, 2755–2793.

Hopenhayn, H. A. (1992): "Entry, Exit and Firm Dynamics in Long Run Equilibrium," *Econometrica*, 60, 1127–50.

Hurst, E., P. J. Kehoe, E. Pastorino, and T. Winberry (2022): "The Distributional Impact of the Minimum Wage in the Short and Long Run," NBER Working Papers 30294, National Bureau of Economic Research, Inc.

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

Kline, P., R. Saggio, and M. Sølvsten (2020): "Leave-Out Estimation of Variance Components," *Econometrica*, 88, 1859–1898.

Lachowska, M., A. Mas, R. Saggio, and S. A. Woodbury (2023): "Do firm effects drift? Evidence from Washington administrative data," *Journal of Econometrics*, 233, 375–395.

Leitão, M., J. Montana, and J. Silva (2022): "The Role of Firms in Wage Inequality Dynamics," CEPR Discussion Papers 17327, C.E.P.R. Discussion Papers.

Messina, J. and J. Silva (2019): "Twenty Years of Wage Inequality in Latin America," *The World Bank Economic Review*, 35, 117–147.

Oliveira, C. (2023): "The minimum wage and the wage distribution in Portugal," *Labour Economics*, 85, 102459.

Oliveira, C., P. Portugal, P. Raposo, and H. Reis (2023): "Inequality in Portugal: 1986–2020," in *Country Studies: Inequalities in Europe and North America*, Institute for Fiscal Studies, IFS Deaton Review, a parallel study to the IFS Deaton Review.

Robinson, J. (1933): *The Economics of Imperfect Competition*, London: Macmillan.

Song, J., D. J. Price, F. Guvenen, N. Bloom, and T. von Wachter (2018): "Firming Up Inequality," *The Quarterly Journal of Economics*, 134, 1–50.

Appendix

A Quadros de Pessoal

We rely upon the following variables:

- Total monthly earnings (*rganho*) is the sum of base earnings, plus overtime pay, plus regular monthly supplements (seniority premiums, performance bonuses, shift differentials, commissions, and meal and transportation allowances).
- Monthly hours (*hnormais*) are normal working-time paid hours.
- Worker age (*idade_Cod*) is the actual worker age at the time of the survey (bottom-coded at 17 and top-coded at 68). To avoid measurement error, age is recomputed as the difference between the current year and the mode of the year of birth implied by reported age.
- Schooling (*habil*) is the highest degree attained. Based upon this variable, we construct three schooling dummies, for the equivalent of 9th grade or less, secondary or post-secondary non-college, and college and over (including Bachelor). We then take the mode of each indicator for each worker.
- Firm–level sectoral codes are at two–digit level, based upon the *C.A.E. - Classificação Portuguesa das Atividades Económicas* classification system. They have been harmonized over time (C.A.E. was revised in 1995, 2003, and 2007). Within an estimation subperiod, the firm–level code is taken to be mode across years.
- Number of employees (*pemp*) equals all persons that have worked for the firm in October, either remunerated by the firm or not (e.g. owner-managers and family members), and including short-term (less than one month) absentees. This is our primary measure for the number of employees, and is the one used to compute firm size.

Our focus on for-profit non-financial market-oriented sectors entails dropping firms in the financial sector, public administration, education, health, social services, families with domestic workers, and international organizations.

For workers reported to be working at different firms in a given year, we kept a single firm observation based upon the following criteria: (i) the firm where the individual works the highest number of normal hours, otherwise (ii) the firm where the individual earns the highest (straight-time) earnings, otherwise (iii) the firm where the individual has the longest tenure. j

B Sistema the Contas Integradas das Empresas (SCIE)

We rely upon the following variables:

- Gross income (*produção*) equals sales revenue plus the change in inventories, plus some additional operating revenue.
- Gross value added (*valor acrescentado bruto a preços de mercado*) equals gross income minus the expenditure on materials and services.
- Materials and services (*consumos intermédios*) equals the cost of materials and services external to the firm.
- Capital (*ativos fixos tangívies + intangíveis + biológicos e propriedades de investimento*) equals the book value of fixed intangible assets, plus intangible assets, plus investment property.
- Number of employees (*pessoal ao serviço*) equals all persons that have worked for the firm during the year, either remunerated by the firm or not (e.g. owner-managers and family members), and including short-term (less than one month) absentees. This variable is used to compute value-added per worker.
- Total labor cost (*gastos com o pessoal* equals total (fixed plus transitory) payments to employees, inclusive of social insurance, pension, and other mandatory contributions such as private insurance, plus hiring and firing costs.

The basic firm identifier we use in our empirical analysis is the one from Quadros de Pessoal. Firms in Quadros de Pessoal come also with the same fictitious fiscal number that

identifies firms in the SCIE. We are therefore able to match each basic firm identifier from Quadros de Pessoal with a SCIE fictitious fiscal number in any given year. We restrict to SCIE firms which are also in Quadros de Pessoal, which for the most part implies dropping independent contractors and firms without paid employees.

SCIE does not provide population weights for the years prior to 2006, i.e. prior to broad coverage associated with IES. As mentioned in the main text, smaller firms are under-represented in the earlier years. We infer population weights by comparing the firm counts in SCIE and in Quadros de Pessoal (which contains the universe of firms for our purpose), according to a stratification of firms based upon the number of employees and sales. All the tables, figures, and regressions in our paper are weighted by these relative counts.

C Leave-Out Bias Correction

Limited mobility bias in the variance components is a potential issue with Portuguese data given the large number of small firms, and the associated low number of worker transitions (Bonhomme et al., 2020; Engbom et al., 2023; Lachowska et al., 2023). We have therefore applied the bias correction proposed by Kline et al. (2020), with the results reported in Table 8 and Figure 18. The solid lines in the figure are the variance components reported in the main text, computed using the standard plug-in approach. The dashed lines are the bias-corrected (KSS) ones. Consistently with the literature, we do find a bias that overstates the importance of firm and worker fixed effects, and understates the importance of sorting. However, since the bias is relatively constant over time, it does not significantly impact the changes in the variance components over subperiods, which is our main object of interest. Our conclusion with the KSS correction remains that the compression in firm fixed effects was the main driver of the wage compression observed in Portugal from 1993-97 to 2014-19, although the correction points to a lower quantitative contribution relative to the plug-in method (135% vs 180%), and finds some positive contribution for assortative matching (44%).

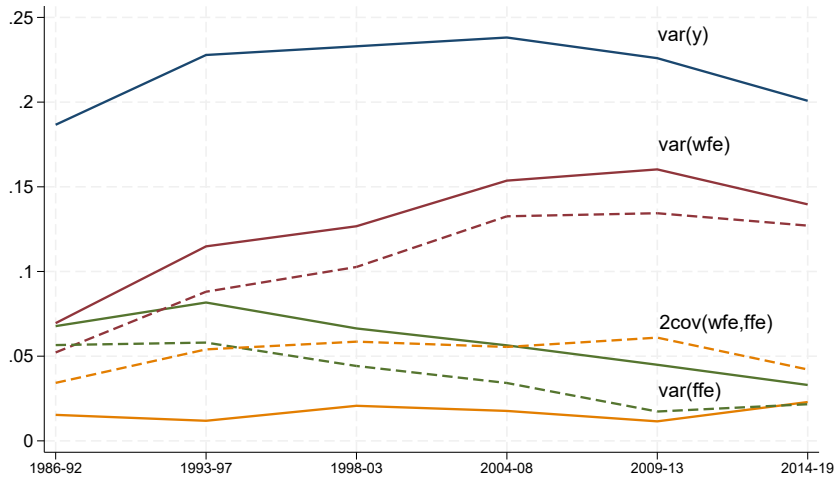


Figure 18: Variance Decomposition and Dynamics, plug-in (solid) vs KSS (dashed)

	Period 2 (1993-1997)		Period 6 (2014-2019)		Change 6–2	
	Value	Share	Value	Share	Value	Share
Var(y)	0.228	100	0.201	100	-0.027	100
Var(WFE)						
plug-in	0.115	50.4	0.14	69.6	0.025	-91.8
KSS	0.088	38.7	0.127	63.3	0.039	-144.3
Var(FFE)						
plug-in	0.082	35.9	0.033	16.4	-0.049	180
KSS	0.058	25.5	0.022	10.8	-0.036	134.6
2Cov(WFE, FFE)						
plug-in	0.012	5.2	0.023	11.4	0.011	-40.9
KSS	0.054	23.7	0.042	21	-0.012	43.8

Notes: the variance components refer to the AKM regression (1).

Table 8: Variance Decomposition, plug-in and bias-corrected (KSS)

D Renormalization of the Skill-Specific Firm Fixed Effects

We adapt Card et al.'s (2015) approach to ensure that the cross-type (in their case, cross-gender) firm fixed effects estimated in Section 2.3.2 are comparable. In a nutshell, their framework views firm profits as increasing in firm-level productivity, and firm pay determined by profit sharing with workers, with sharing coefficients that vary by worker type (e.g. due to differences in bargaining power). Their key insight is that low productivity firms, who might not share any profits with their workers, provide a natural benchmark. These are firms with zero true firm fixed effects, where wages equal worker fixed effects. Renormalizing relative to such firms eliminates the arbitrary normalization and recovers the true firm fixed effect.

Our adaption accommodates the specific patterns that emerge from our evidence in Figure 6. As in Card et al. (2015), the approach relies on correct specification of the AKM model. Assume that a firm's total surplus $S(z) = E(z) + R(z)$ can be decomposed into a non-risky component $E(z)$ from established operations (e.g. due to product market power, customer loyalty) and a volatile component $R(z)$ from risky operations (e.g. due to product and market exploration, innovation). Assume that $E(z)$ and $R(z)$ both increase with value-added per worker z, however $E(z)$ increases faster so that $r(z) \equiv R(z)/S(z)$ decreases with z - more productive firms generate higher surplus, with a larger share coming from established operations.

Now posit that firm j's pay to an unskilled worker ($s = 0$) in a given period is

$$FFE_j^0 = \gamma^0 \max\{0, S(z^j) - \tau^0\},$$

where τ^0 is a threshold above which the firm starts sharing surplus from established operations. The firm pay to a skilled worker ($s = 1$) is instead

$$FFE_j^1 = \gamma^1 \max\{0, S(z^j) - \tau^1\} + \beta^1 \max\{0, r(z^j) - \tau_r^1\},$$

meaning that firms offer skilled workers an additional compensation related to the share of risky operations, according to the sharing coefficient β^1 and the threshold τ_r^1. We further

assume that there is a non-empty set of z such that $FFE_j^1 = 0$. A sufficient condition is that $z_r < z_S$, where $S(z_S) = \tau^1$ and $r(z_r) = \tau_r^1$, so that there is a middle-range region of productivity levels where neither component of firm pay meets the sharing threshold. This specification implies that the lowest type-specific firm fixed effects in Figure 6 are equal to zero, and therefore suitable for renormalization. Note that the logic would go through if these firms offered the same positive level of firm pay to any worker.

These assumptions allow us to proceed as follows. For unskilled workers, and for each period p, we renormalize their firm fixed effects using the average across all firms below the estimated productivity threshold ζ_p^0. This part is entirely analogous to Card et al. (2015). For skilled workers, we use instead the average within a narrow band (5% in practice) around the estimated threshold ζ_p^1. The results are not sensitive to bandwidth choice. Note that, for this identification to work, the minimum wage must not bind for different workers within a given firm, particularly among firms not sharing surplus - otherwise these workers' wages would not equal their respective portable wage components (the AKM model would be misspecified).

While we maintain this interpretation of Figure 6, for the specific purpose of constructing the between component, our main analysis will provide a different interpretation. In particular, it features a perfectly competitive product market, monopsony power which varies by worker type and governs the pass-through from productivity to firm pay, and a minimum wage that binds at low-productivity firms. We also abstract from risk considerations and the associated skilled worker compensation (footnote 10). We acknowledge this caveat when interpreting the between component of our decomposition.

E Numerical approximation of the stationary equilibrium

In stationary equilibrium, the measure of firms over productivity is an invariant vector $\hat{\mu}$ of size N, where N is the number of points in the grid for productivity. Let $N_f = \sum_{i=1}^{N} \hat{\mu}_i$ the mass of incumbent firms and $\mu_i \equiv \hat{\mu}_i / N_f$.

E.1 Solving for the optimal firm policies, given μ and $\bar{m}$

Given μ and $\bar{m}$, the firm policies $\{m_i, w_i\}_{i=1}^N$ are such that for all i,

$$\begin{aligned} m_i, w_i = \arg\max_{\{m,w\}} \; & z_i l(w)^\gamma - w l(w) \\ \text{s.t.} \;\; & l(w) = m\Phi(w) \\ & w \geq \underline{w} \\ & m \leq \bar{m}, \end{aligned}$$

where $\bar{m} = 2N_w/N_f$ and

$$\Phi(w) = \sum_{j=1}^{N} \left[\left(1 - \frac{m_j}{\bar{m}}\right) \frac{1}{1 + e^{\frac{b-w}{\sigma}}} + \frac{m_j}{\bar{m}} \frac{1}{1 + e^{\frac{b-w}{\sigma}} + e^{\frac{w_j - w}{\sigma}}} \right] \mu_j.$$

The support of productivity levels can be partitioned into three connected subsets. In region 1, $w_i = \underline{w}$ and $m_i < \bar{m}$. The mass of offers m_i satisfies

$$z_i \gamma [m_i \Phi(\underline{w})]^{\gamma - 1} = \underline{w}.$$

For firms in region 2, $w_i = \underline{w}$ and $m_i = \bar{m}$. In such region,

$$\underline{w} < z_i \gamma [\bar{m} \Phi(\underline{w})]^{\gamma - 1} < \underline{w} + \frac{\Phi(\underline{w})}{\Phi'(\underline{w})}.$$

For firms in region 3, $m_i = \bar{m}$ and $w_i > \underline{w}$. The wage w_i is pinned down by the condition

$$z_i \gamma [\bar{m} \Phi(w_i)]^{\gamma - 1} = w_i + \frac{\Phi(w_i)}{\Phi'(w_i)}.$$

E.2 A fixed point for μ

For every i, denote by $R(z_i, \mu, \bar{m})$ the vector of operating profits conditional on productivity z_i, the distribution μ, and the matching rate $\bar{m}$ induced by N_f. This is the value of the optimization problem described in the previous section. The incumbent's Bellman equation

is

$$v(z_i, \mu, \bar{m}) = \int \max \left\{ 0, R(z_i, \mu, \bar{m}) - \kappa + \beta \sum_j \pi_{ij} v(z_j, \mu, \bar{m}) \right\} \psi(\kappa) d\kappa.$$

As long as the function $R(z_i, \mu, \bar{m}) + \beta \sum_j \pi_{ij} v(z_j, \mu, \bar{m})$ is strictly monotone, there exists $\hat{\kappa}(z_i)$ such that

$$\hat{\kappa}(z_i) = R(z_i, \mu, \bar{m}) + \beta \sum_j \pi_{ij} v(z_j, \mu, \bar{m}). \tag{13}$$

It follows that one can rewrite

$$v(z_i, \mu, \bar{m}) = \int^{\hat{\kappa}(z_i)} \left\{ R(z_i) - \kappa + \beta \sum_j \pi_{ij} v(z_j, \mu, \bar{m}) \right\} \psi(\kappa) d\kappa.$$

In turn, using (13), one can write

$$\begin{aligned} v(z_i, \mu, \bar{m}) &= \int^{\hat{\kappa}(z_i)} \{\hat{\kappa}(z_i) - \kappa\} \psi(\kappa) d\kappa \\ &= \Psi[\hat{\kappa}(z_i)] \hat{\kappa}(z_i) - \int^{\hat{\kappa}(z_i)} \kappa \psi(\kappa) d\kappa \\ &= \Psi[\hat{\kappa}(z_i)] \left[\hat{\kappa}(z_i) - E(\kappa | \kappa \leq \hat{\kappa}(z_i)] \right. . \end{aligned}$$

Substituting back in (13), one obtains the condition

$$\hat{\kappa}(z_i) - \beta \sum_j \pi_{ij} \Psi[\hat{\kappa}(z_j)] \left[\hat{\kappa}(z_j) - E(\kappa | \kappa \leq \hat{\kappa}(z_j)] \right. = R(z_i, \mu, \bar{m}). \tag{14}$$

For $i = 1, ..., N$, the system of equations (14) identifies the vector of cost thresholds $\hat{\kappa}(z_i)$ consistent with the distribution μ.

Recall that $\Psi[\hat{\kappa}(z_i)]$ is the probability of survival, conditional on z_i. Invoking the law of large number, $\Psi[\hat{\kappa}(z_i)]$ is also the fraction of firms that survive, conditional on a productivity draw z_i.

Let $\mathcal{E}(z_i)$ denote the mass of entrants that obtained a draw z_i. Conditional a signal threshold η^*, we have

$$\mathcal{E}(z_i) = M \int_{\eta^*} g_i(\eta) h(\eta) d\eta.$$

In turn, the threshold is pinned down by the free-entry condition

$$\sum_i g_i(\eta^*)\Psi[\hat{\kappa}(z_i)]\left[\hat{\kappa}(z_i) - E(\kappa|\kappa \leq \hat{\kappa}(z_i)\right] = c_e.$$

Stationarity requires that

$$\mu(z_i) = \Psi[\hat{\kappa}(z_i)]\left(\sum_j \pi_{ji}\mu(z_j) + \mathcal{E}(z_i)\right).$$

Let $\hat{\mathbf{\Psi}}$ denote a $N \times N$ diagonal matrix such that $\hat{\mathbf{\Psi}}_{ii} = \Psi[\hat{\kappa}(z_i)]$. We can rewrite the above condition as

$$\mu = (\Pi\hat{\mathbf{\Psi}})^T\mu + \hat{\mathbf{\Psi}}\mathcal{E}$$

It follows that the stationary measure must satisfy

$$\mu = \left[I - (\Pi\hat{\mathbf{\Psi}})^T\right]^{-1}\hat{\mathbf{\Psi}}\mathcal{E},$$

where $\hat{\mathbf{\Psi}}$ depends on μ by (14) and $\mathcal{E}$ depends itself on $\hat{\mathbf{\Psi}}$.

E.3 Non-parametric specification of the operating cost distribution

Instead of assuming a parametric form for the distribution Ψ, we retrieve it non-parametrically by imposing that in the initial steady state, the model be consistent with the vector of conditional exit rates X_i. The conditional exit rate (i.e. the probability of exit conditional on a given productivity level) is simply

$$X_i = \sum_i \pi_{ij}\left\{1 - \Psi[\hat{\kappa}(z_j)]\right\}.$$

The unconditional exit rate is

$$\sum_i \mu_i \sum_j \pi_{ij}\left\{1 - \Psi[\hat{\kappa}(z_j)]\right\}$$

We select values for Ψ_j such that

$$\sum_i \pi_{ij} \{1 - \Psi_j\} = \mathrm{X}_i$$

for all i.

The vector of probabilities Ψ_j can then be used to back out μ and solve for the equilibrium $R(z)$. Finally, the thresholds for the fixed cost $\hat{\kappa}_i$ can be backed out from the system of equations

$$\hat{\kappa}_i - \beta \sum_j \pi_{ij} \Psi_j \left[\hat{\kappa}_j - E(\kappa | \kappa \leq \hat{\kappa}_j\right] = R(z_i).$$

Notice that the above conditions yield no restrictions regarding the distribution $\Psi(\kappa)$ for $\kappa < \hat{\kappa}_1$. We will assume that $\kappa \in \Re^+$ and that the distribution features an atom for $\kappa = 0$. The execution will reveal that none of our conclusions depend upon this assumption.

E.4 Algorithm – The initial steady state

1. Choose a vector of conditional exit rates $X_{(N\times 1)}$
2. Given X, compute the vector $\hat{\Psi}_{(N\times 1)}$ using $\hat{\Psi} = \mathbf{1} - \Pi^{-1} X$. For every j, the element $\hat{\Psi}_j$ is the cdf of Ψ evaluated at the cost threshold $\hat{\kappa}_j$
3. Choose the threshold for the signal η^*
4. The choice of η^* pins down the vector of entrants $\mathcal{E}$
5. Denote as $\hat{\mathbf{\Psi}}_{(N\times N)}$ the diagonal matrix with the elements of $\hat{\Psi}_{N\times 1}$ on the diagonal. Then we can determine μ by computing

$$\mu = \left[I - (\Pi \hat{\mathbf{\Psi}})^T\right]^{-1} \hat{\mathbf{\Psi}}^T \mathcal{E},$$

6. Once we have μ, we can determine the vector of revenues R
7. Once we have the vector R, we can compute the thresholds for exit on the cost space using the system of linear equations

$$\hat{\kappa}_i - \beta \sum_j \pi_{ij} \hat{\Psi}_j \left[\hat{\kappa}_j - E(\kappa | \kappa \leq \hat{\kappa}_j) \right] = R_i.$$

Denote as ψ the pdf associated to the cdf Ψ. Rewrite the latter, which is equation (14) recovered above, as

$$\hat{\kappa}_i - \beta \sum_j \pi_{ij} \hat{\Psi}_j \hat{\kappa}_j + \beta \sum_{j=1} \pi_{ij} \int_{\hat{\kappa}_0}^{\hat{\kappa}_j} \kappa \psi(\kappa) d\kappa = R_i, \tag{15}$$

where $\hat{\kappa}_0$ is arbitrary.

For the purpose of computing the integral in (15), we need to make assumptions about the density ψ. We will assume that it is piece-wise uniform.

E.4.1 Piece-wise uniform representation of the density ψ

We posit that

$$\psi(\kappa) = \psi_s \quad \forall\, \kappa \in [\hat{\kappa}_{s-1}, \hat{\kappa}_s], \quad s = 1, ..., N.$$

We also assume that the lower bound of the support is $\kappa_0 = 0$, and we allow for an atom at $\kappa_0 = 0$, i.e. $\Psi_0 \geq 0$. Then, the sequence of constants ψ_s are pinned down by the condition

$$\int_{\hat{\kappa}_{s-1}}^{\hat{\kappa}_s} \psi_s d\kappa = \Psi_s - \Psi_{s-1}, \quad \forall s > 0 \tag{16}$$

It follows that

$$\psi_s = \frac{\Psi_s - \Psi_{s-1}}{\hat{\kappa}_s - \hat{\kappa}_{s-1}}$$

and that

$$\int_{\hat{\kappa}_{s-1}}^{\hat{\kappa}_s} \psi_s \kappa d\kappa = \frac{1}{2} \left(\Psi_s - \Psi_{s-1} \right) \left(\hat{\kappa}_s + \hat{\kappa}_{s-1} \right)$$

$$\hat{\kappa}_i - \beta \sum_j \pi_{ij} \hat{\Psi}_j \hat{\kappa}_j + \frac{\beta}{2} \sum_{j=1} \pi_{ij} \sum_{s=1}^{j} \left(\Psi_s - \Psi_{s-1} \right) \left(\hat{\kappa}_s + \hat{\kappa}_{s-1} \right) = R_i,$$

Let $\mathcal{A}$ denote the diagonal matrix $(N \times N)$ with generic element $\mathcal{A}_{ss} = \Psi_s - \Psi_{s-1}$. Then, $\mathcal{A}\kappa$ is a vector $(N \times 1)$ with generic element s equal to $(\Psi_s - \Psi_{s-1})\hat{\kappa}_s$.

$$\mathcal{A} = \begin{pmatrix} \Psi_1 - \Psi_0 & 0 & \dots & 0 \\ 0 & \Psi_2 - \Psi_1 & \dots & 0 \\ \vdots & & & \vdots \\ 0 & \dots & \dots & \Psi_N - \Psi_{N-1} \end{pmatrix}.$$

Let $\mathcal{B}$ denote the matrix $(N \times N)$ with element $\mathcal{B}_{s+1,s} = \Psi_s - \Psi_{s-1}$ and all other elements equal to 0. Then, $\mathcal{B}\kappa$ is a vector $(N \times 1)$ with generic element s equal to $(\Psi_s - \Psi_{s-1})\hat{\kappa}_{s-1}$.

$$\mathcal{B} = \begin{pmatrix} 0 & 0 & \dots & \dots & 0 \\ \Psi_2 - \Psi_1 & 0 & \dots & \dots & 0 \\ 0 & \Psi_3 - \Psi_2 & & & \vdots \\ \vdots & & & & \vdots \\ 0 & \dots & \dots & \Psi_N - \Psi_{N-1} & 0 \end{pmatrix}.$$

Let $\mathcal{T}$ denote a lower-triangular matrix were all the non-zero elements are 1. That is, for every s, $\mathcal{T}_{sr} = 1$ if $r \leq s$ and $\mathcal{T}_{sr} = 0$ otherwise.

$$\mathcal{T} = \begin{pmatrix} 1 & 0 & \dots & \dots & 0 \\ 1 & 1 & 0 & \dots & 0 \\ \vdots & 1 & 1 & & \vdots \\ \vdots & 1 & & 1 & 0 \\ 1 & 1 & \dots & 1 & 1 \end{pmatrix}.$$

Then, the system of equations above is linear and can be written in matrix form as

$$\kappa - \beta\Pi\hat{\mathbf{\Psi}}\kappa + \frac{\beta}{2}\Pi\mathcal{T}(\mathcal{A} + \mathcal{B})\kappa = R$$

It follows that

$$\kappa = \left[I - \beta\Pi\hat{\mathbf{\Psi}} + \frac{\beta}{2}\Pi\mathcal{T}(\mathcal{A} + \mathcal{B})\right]^{-1} R.$$

The cost of entry c_e is given by free-entry condition:

$$\sum_i g_i(\eta^*)\Psi_i \left[\hat{\kappa}_i - E(\kappa|\kappa \leq \hat{\kappa}_i)\right] = c_e.$$

The latter can be rewritten as

$$\sum_i g_i(\eta^*)\Psi_i\hat{\kappa}_i - \sum_i g_i(\eta^*) \int_{\hat{\kappa}_0}^{\hat{\kappa}_i} \kappa\psi(\kappa)d\kappa = c_e.$$

The algorithm also yields the discrete approximation for the cdf $\Psi(\kappa)$ and the associated density $\psi(\kappa)$, defined over the support $[\hat{\kappa}_0, +\infty)$, which can be used for comparative statics.

E.5 Comparative Statics

1. Guess η^*.

   (a) Make an initial guess for the vector of cost thresholds $\tilde{\kappa}_0$. Associated with it, there will be a cdf $\widetilde{\Psi}_0$.

   (b) Figure out the vector of entrants $\mathcal{E}_0$ and the stationary distribution of firms is

   $$\mu_0 = \left[I - (\Pi\widetilde{\mathbf{\Psi}}_0)^T\right]^{-1} \widetilde{\mathbf{\Psi}}_0^T\mathcal{E}_0,$$

   (c) Compute the revised guess of the vector of cost thresholds $\tilde{\kappa}_1$ by solving

   $$\tilde{\kappa}_{1i} = R_i + \beta\sum_j \pi_{ij}\Psi(\tilde{\kappa}_{0j}) \left[\tilde{\kappa}_j - E(\kappa|\kappa \leq \tilde{\kappa}_j)\right],$$

   (d) Repeat the process until convergence, i.e. $||\tilde{\kappa}_T - \tilde{\kappa}_{T-1}||$ is small enough

2. Check that the free-entry condition holds

   $$\sum_i g_i(\eta^*)\Psi(\tilde{\kappa}_{Ti}) \left[\tilde{\kappa}_{Ti} - E(\kappa|\kappa \leq \tilde{\kappa}_{Ti})\right] = c_e$$

3. If not, then revise the guess for η^* and repeat