

Minimum Wages and the Distribution of Firm Wage Premia*

Marcelo Bergolo Rodrigo Ceni Mathias Fondo Damián Vergara

Abstract

This paper leverages a large minimum wage reform in Uruguay to study the effects of minimum wages on the distribution of firm wage premia. We document a substantial decrease in wage inequality after the minimum wage reform, almost completely driven by a reduction of between-firm inequality. AKM variance decompositions show that the relative variance of firm fixed effects substantially decreased after the reform. A time-varying AKM model reveals this pattern was driven by a compression in the distribution of firm fixed effects, with low-paying firms increasing their fixed effect after the minimum wage increase. Both firm-level and worker-level difference-in-differences models show that the minimum wage reform had a causal effect on the compression in the distribution of firm wage premia. The results suggest that minimum wages can increase the supply of “good jobs” not only through reallocation effects but also by “making bad jobs better”.

*First version: February 2025. This version: February 2025. Bergolo: Department of Economics, Universidad de la Republica, marcelo.bergolo@fcea.edu.uy. Ceni: Department of Economics, Universidad de la Republica, rodrigo.ceni@fcea.edu.uy. Fondo: Department of Economics, Universita di Bologna, mathias.fondo@unibo.it. Vergara: Department of Economics, University of Michigan and IZA, damiv@umich.edu. We thank David Card, Patrick Kline, Raffaele Saggio, and seminar participants at the Global Minimum Wages Conference, the 3rd M-NEW Annual Workshop, Princeton University, Universidad de Chile, and Universita di Bologna for very helpful comments and suggestions. Usual disclaimers apply.

1 Introduction

Firms affect the levels and trends in wage inequality (Kline, 2024). This conclusion usually follows from the observation that, after fitting the canonical two-way fixed-effects model of Abowd et al. (1999) (AKM), a non-trivial portion of the variance of log wages can be explained by the variance of the estimated firm fixed effects (F-FEs). Based on this empirical regularity, recent work investigates the drivers of the cross-sectional variance in firm wage premia, concluding that some combination of firm-level heterogeneity in productivity, rent-sharing elasticities, and compensating differentials may rationalize equilibrium dispersion in firm pay when labor markets have competitive frictions (Card et al., 2018; Sorkin, 2018; Lamadon et al., 2022; Haanwinckel, 2024; Morchio and Moser, 2024).

Most of the discussion about the determinants of F-FEs, however, focuses on idiosyncratic firms' characteristics, thus abstracting from the effect of labor market policies on the ability of firms to set wages. Including labor market policies in the analysis can broaden the set of candidates for rationalizing equilibrium dispersion in firm pay for similar workers and inform about its policy implications. Moreover, the role of labor market policies may offer a rationale for the differential role of F-FEs in explaining wage inequality across countries and within countries across time. For example, Card et al. (2013) finds that the increased importance of F-FEs in recent decades in Germany coincides with a decrease in the prevalence of collective bargaining. Conversely, Song et al. (2019) conjectures that the stable role of F-FEs in the US is related to the small role of labor market institutions.

This paper takes a step in that direction by quantifying the effects of minimum wages on the distribution of firm wage premia. There are two channels through which minimum wages can affect the effect of firms on aggregate wage inequality. First, there is a "composition" channel. Minimum wages can induce employment reallocation across employers and between employment and non-employment (Dustmann et al., 2022; Engbom and Moser, 2022; Berger et al., 2025). Then, for a fixed set of F-FEs, reallocation effects impact how the cross-sectional dispersion in firm wage premia maps to the overall variance of wages. Second, there is "causal" channel. Recent evidence shows that minimum wages can affect how rents are split between workers and employers (through, for example, profit effects) and the levels of the rents to be split (through, for example, price and productivity effects) (Dube and Lindner, 2024). Through the lens of the AKM framework, this evidence suggests that minimum wages can causally affect the level of F-FEs, thus affecting their distribution even in the absence of employment reallocation.

The main contribution of this paper is to show that the quantitative relevance of the causal channel can be substantial. We test this hypothesis using matched employer-employee data from Uruguay built from administrative social security records for the period 1997-2013. The raw data is monthly and contains information on earnings and hours, allowing us to precisely measure job transitions and hourly wages. Important for our research question, Uruguay prominently increased the national minimum wage

in 2005 after several years of being non-binding. In 2005, the nominal minimum wage increased by almost 80%, while annual inflation rates were around 5%. The government kept implementing nominal increases above inflation after 2005: in 2013, the real minimum wage was more than three times larger than the one prevailing in 2004. These increases had implications for the “bite” of the minimum wage. Between 1997 and 2004, the Kaitz index (minimum wage as a share of the median wage) was, on average, 9.5%. The index rose sharply to 24% in 2005 and kept increasing after the reform, reaching 43% in 2013.

We use these data and reform to study the effect of the minimum wage on wage inequality, focusing on the role of firms in mediating its effect. The paper proceeds in two steps. We first present a set of descriptive analyses of the evolution of wage inequality around the minimum wage reform, which points to an important role in the distribution of F-FEs for explaining the observed trends in wage inequality. We then perform empirical analyses at the firm- and worker-level exploiting the variation induced by the reform to show that the minimum wage had a causal effect on the documented inequality dynamics.

We start with a descriptive analysis of the evolution of wage inequality around the minimum wage reform. The reform coincided with a sharp and substantial decrease in wage inequality. The standard deviation of log hourly wages grew from 77 log points in 1997 to 82 log points in 2004. In 2005, it decreased to 76 log points – reverting, in one year, more than seven years of increased inequality – and it kept falling throughout the post-reform period, reaching 62 log points in 2013. The evolution of different wage percentiles reveals that the decrease in wage inequality after the minimum wage reform was driven by a higher wage growth of low percentiles, implying that the overall decrease in wage inequality was materialized through wage compression at the bottom of the distribution. Interestingly, almost all of the reduction in wage inequality observed after 2005 was explained by a decrease in between-firm (rather than within-firm) inequality, and the trend was observed both in populations of stayers and new hires.

These stylized facts are complemented with AKM estimates using overlapping 5-year periods around the minimum wage reform (1997-2001, 2001-2005, 2005-2009, and 2009-2013). Bias-corrected variance decompositions a la [Kline et al. \(2020\)](#) show that the levels of all components of wage inequality fell after 2005, but the decrease was particularly large for the firm component. As a result, the relative importance of firms to explain wage inequality decreased after the minimum wage reform. Relative to the pre-reform period, the share explained by worker fixed effects increased in the post-2005 period, and the share explained by the covariance term remained constant. By contrast, the share explained by F-FEs decreased from almost 30% in the pre-2005 period to 22% in 2005-2009 and 18% in 2009-2013.

A final descriptive analysis follows from applying recent econometric techniques known as the Time-Varying AKM ([Engbom et al., 2023](#); [Lachowska et al., 2023](#), TV-AKM) that generalize the AKM specification to allow for time-varying F-FEs. This procedure endows us with estimates of F-FEs that vary at the firm-by-year level. We use tax data on corporate tax records available for 2009-2013 to compute value added per worker and show, consistent with previous literature, that F-FEs are positively correlated

with firm productivity (Kline, 2024). We show that this positive relationship is attenuated but remains large and significant after controlling by industry and, more strikingly, by firm indicators, suggesting that within-firm across-time variation in value added per worker predicts drifts in the F-FEs. To proxy for heterogeneous rent-sharing behavior, we also compute measures of firm-level labor shares (measured as total gross wages over value added) and find that, conditional on value added per worker, F-FEs are positively correlated with firms’ labor shares, with a steep slope that is also robust to the inclusion of industry and firm indicators. This result – which, to the best of our knowledge, has not been documented in the literature – suggests that variation in rent-sharing elasticities plays an important role in rationalizing the dispersion in F-FEs on top of the known role of heterogeneous productivities.

Furthermore, the time-varying F-FEs indicate that, before the 2005 reform, the short-run within-firm persistence in pay premia was substantial: binned scatter plots of F-FEs in year t against F-FEs in year $t - 1$ lie almost exactly on the 45-degree line, with some generalized cyclical component in recession years, echoing the findings documented in Engbom et al. (2023) and Lachowska et al. (2023). However, we document a trend break for low-paying firms after the reform. Starting in 2005, the same binned scatter plots start displaying a “hockey stick” pattern: the F-FEs of firms at the bottom of the distribution display large annual growth rates while firms up in the distribution remain close to the 45-degree line. Through time, this differential growth pattern generates a sizable compression in the support of the F-FEs distribution, as firms with low initial pay premia catch up with firms up in the distribution. These patterns are not explained by composition effects as the same behavior is observed in longer horizons when restricting to the balanced panel of firms that are active throughout our complete analysis period.

These analyses show that, after the minimum wage increase, the wage distribution was compressed because of a corresponding compression in the distribution of F-FEs. While the timing and mechanisms in play suggest that the minimum wage reform played a role in these dynamics, the descriptive facts are not sufficient to conclude a causal relationship. Importantly, the reform was enacted after a recession, so business cycle dynamics and mean reversion could also explain the observed outcomes. Given these confounders, and motivated by a simple extension of the model proposed in Card et al. (2018), the final sections empirically test whether the minimum wage causally affected the distribution of F-FEs.

We proceed with two designs. First, a firm-level design exploits variation in exposure to the NMW reform using the GAP design proposed in Card and Krueger (1994) and recently used by Draca et al. (2011), Dustmann et al. (2022), and Derenoncourt and Weil (Forthcoming) to estimate a difference-in-differences (DID) model. The (continuous) treatment variable is the increase in labor costs that firms would have to incur to comply with the new NMW based on pre-reform hourly wages, so firms with a higher gap between observed and counterfactual labor costs are said to be more exposed to the reform. The results show that more exposed firms had similar trends in real hourly wages relative to less exposed firms before 2005, but that real hourly wages increased sharply and permanently in the treated firms

starting in 2005. Moreover, we find that almost all of the increase in the firm-level average log real hourly wage is explained by an increase in the time-varying F-FE. We also show that, in the medium run, firms accommodate the NMW increase by changing the composition of workers, as the average worker fixed effect also increases in relatively more exposed firms after the NMW reform (Butschek, 2021; Clemens et al., 2021). However, that effect is an order of magnitude smaller than the effect on the F-FE.

Firm-level designs can identify firm-level responses to minimum wage reforms but are insufficient to assess the aggregate consequences of the policy change as they have trouble dealing with between- and within-firm spillovers and are not robust to the presence of extensive margin responses given the restrictions imposed on firm survival. Then, we complement our firm-level design with a worker-level design that closely follows Dustmann et al. (2022). We estimate the effect of the NMW reform on 2-year growth rates of different outcomes of interest using a DID design that sorts workers into wage bins based on their baseline real hourly wages. This design mimics a standard DID model where the unit of analysis is the initial wage bin, so treatment effects are recovered by comparing outcomes before and after the NMW reform between low real hourly wage workers (directly exposed to the NMW reform) and workers higher in the wage distribution. The focus on 2-year growth rates allows us to deal with the well-known fact that mean reversion implies that wage growth is faster at the bottom of the wage distribution despite having weaker employment prospects than the rest of the wage distribution (Ashenfelter and Card, 1982).

The results show that the real hourly wage growth by wage bin was constant before 2005 in all wage bins but sharply increased in 2005 and 2006 at the bottom of the wage distribution, especially for workers earning below the new NMW before the reform. We find evidence of wage spillovers to workers earning more than the new NMW that decrease in the distance to the NMW and dissipate when wages reach about 1.5 times the NMW level. We estimate no effect on the probability of remaining employed by wage bin, suggesting that the increase in wages was not accompanied by disemployment effects. We find evidence of reallocation effects by replicating Dustmann et al. (2022) exercise and showing that job switchers at the bottom of the distribution experience an increase in their F-FEs after the NMW reform, with the effects concentrated at the very bottom of the distribution. We also find evidence of productive reallocation in the pre-reform recession, consistent with the “cleansing” hypothesis based on Schumpeter (1939) work. However, the main result in this section is to show that changes in F-FEs are much larger for job-stayers and display even larger wage spillovers up above the wage distribution. A simple decomposition exercise shows that this “causal” effect on the F-FEs of stayers accounts for more than 75% of the total wage effect across all exposed wage bins, while the relative importance of reallocation effects is concentrated at the bottom bin and accounts for around 10-20% of its wage effect.

In sum, this paper shows that public policy, in particular the minimum wage, can affect the wage-setting power of firms, having sizable implications for the aggregate effects of policies on wage inequality. Our results warn that the within-firm persistence in wage setting power documented in Engbom et al.

(2023) and Lachowska et al. (2023) can be broken by changes in labor market institutions. Importantly, a long tradition in economics documents wage compression effects of the minimum wage (e.g., Lee, 1999; Autor et al., 2016; Dube, 2019; Fortin et al., 2021; Giupponi et al., 2024). Our results suggest that changes in the cross-sectional profile of F-FEs are an important driver of those patterns. Under the interpretation that F-FEs are a valid measure of job quality, our results suggest that minimum wages can increase the supply of “good jobs” (Acemoglu, 2001; Rodrik and Stantcheva, 2021), not only through reallocation dynamics as in Dustmann et al. (2022) but also by “making bad jobs better”.

2 Institutional Setting

Uruguay is an upper-middle-income country in South America with a current population of almost 3.5 million. Uruguay performs well above the regional standards in terms of income, corruption, human development, and employment informality.¹ The country’s employment structure is centered around the services sector, with an approximate employment share of 70%. Agriculture and manufacturing account for about 10% and 20% of total employment, respectively. The public sector concentrates around 15% of employment, with special participation in education, healthcare, and public administration.

2.1 The National Minimum Wage and the 2005 reform

The National Minimum Wage (NMW) in Uruguay has a long history. However, for several decades, it was a non-binding policy. Established by Decree 1,534/969 on November 28, 1969, the NMW in Uruguay sets a nationwide wage floor for all private sector employees aged 18 and older. During the 1970s and 1980s, the real NMW substantially decreased since large annual inflation rates completely undermined its nominal raises. In the 1990s and early 2000s, inflation rates were controlled, but nominal NMW increases were also muted. The lack of relevant hikes in the NMW mainly relates to two reasons. First, the NMW was used as a reference point for calculating several social benefits provided by the central government. Second, personal income tax brackets were also indexed to the value of the NMW. Both factors led governments to avoid increasing the NMW. By 2003, the real NMW had fallen to just a quarter of its 1969 value, rendering it largely ineffective as a labor market institution.

A recession in the early 2000s generated substantial real wage decreases, especially at the bottom of the wage distribution (see Section 4). Motivated by this fact, the new government elected in 2004

¹All these statistics were accessed on February, 2025. According to the World Bank, Uruguay’s GDP per capita (PPP) in 2021 US dollars was \$31,019 in 2023 (<https://data.worldbank.org/indicator/NY.GDP.PCAP.PP.KD>). According to the ILO, 95% of the population is covered by at least one social protection benefit, and the share of informal employment is around 32%. On average, these numbers in South America amount to 61% and 52%, respectively (<https://rshiny.ilo.org/dataexplorer15/>). Transparency International ranks Uruguay 13rd out of 180 countries in terms of corruption (<https://www.transparency.org/en/cpi/2019/results>), while the United Nations Development Programme ranks Uruguay 52nd among 193 countries in terms of human development (<https://hdr.undp.org/data-center/country-insights#/ranks>).

reinstated the NMW as a key policy instrument. Importantly, the government detached the NMW from the determination of social benefits and income tax brackets, freeing it from the constraints imposed by other fiscal policies. This shift triggered a period of robust real growth for the NMW, beginning in 2005.

Figure 1 shows the evolution of the NMW between 1993 and 2022. Panel (a) shows log differences relative to 1997 for the nominal NMW and the Consumption Price Index (CPI), while Panel (b) shows the annual growth rates of the same variables. Between 1993 and 2001, the NMW exactly traced the evolution of the CPI, meaning that the NMW stayed constant in real terms. Between 2002 and 2004, the real NMW experienced a decrease given the increased inflation rates. In 2005, however, the nominal NMW increased by almost 80%, with an annual inflation rate on the order of 5%. Large nominal increases in the NMW above inflation rates persisted in the decade after the reform, stabilizing to approximately track inflation only around 2013. Panel (c) shows that this reform substantially increased the real NMW. The real increase in 2005 was around 70%, and in 2013, the real NMW was more than three times larger than its value in 2004. This real increase was accompanied by an increase in the “bite” of the NMW. Between 1997 and 2004, the Kaitz index (minimum wage as a share of the median wage) was, on average, 9.5%. The index rose to 24% in 2005 and kept increasing after the reform, reaching 43% in 2013.²

The main goal of this paper is to study the effects of this large NMW reform on wage inequality, focusing on the role of firms as mediators of the policy. The 2005 reform, however, introduced additional policies and institutions with potential impacts on the distribution of earnings. The two most important changes were the restoration of a collective bargaining system, and a progressive income tax reform. In what follows, we briefly discuss these institutions, their interaction with the NMW reform, and the concerns they may introduce to our causal analyses of the NMW reform.

Collective bargaining. In 1943, Law No. 10,449 established a contractual pay scheme that brings together representatives from workers, employers, and the government to negotiate industry-specific minimum wages above the NMW that are allowed to vary with occupation. These collective bargaining agreements (CBAs) experienced two long periods of inactivity, first, during the military dictatorship (1973–1984) and second, during liberal democratic governments between 1992 and 2005. Negotiations were gradually restored starting in 2005, with a proper functioning after Law No. 18,566 was sanctioned in 2009. This law expanded the sectoral coverage and established concrete rubrics for the bargaining rounds. These agreements apply universally to all workers and firms in an industry, regardless of union membership or participation in negotiations. Collective bargaining directly interacts with the NMW, as the NMW works as the lower bound for the industry-level negotiations.

The presence of industry-specific wage floors above the NMW may affect our analysis below because the NMW may not be the relevant restriction for firms and workers participating in industries that set wage floors above the NMW. A few considerations follow. First, while the process of restoring sectoral

²We compute the median wage using administrative records only available between 1997 and 2013 (see Section 3).

negotiations started in 2005, it was not until 2009 that the system was operating widely. This timing contrasts with the NMW, whose biggest change happened in 2005. Our results below show instant effects on our outcomes of interest around 2005, suggesting that the policy that predominantly drives the action is the NWM. Second, the NMW moves in tandem with the collectively bargained wage floors as the NMW sets a salient reference point for sectoral negotiations, especially for industries that negotiate wage floors close to the NMW. Then, increases in the NMW should co-move with changes in industry-specific wage floors. Third, if anything, this measurement problem should work against finding causal effects of the NMW because “control” firms and workers (i.e., not exposed to changes in the NMW because they have larger baseline wages) may experience wage increases dictated by their respective contracts. Fourth, our firm-level regressions control by firm fixed effects, and our individual-level regressions control by industry indicators. These fixed effects attenuate concerns related to confounders related to industry attachment.

Income tax reform. The government elected in 2004 also introduced a big income tax reform in 2007 through Law No. 18,083, which replaced a linear labor income tax (*Impuesto a las Retribuciones Personales*) with a progressive tax schedule (*Impuesto a la Renta de las Personas Físicas*). The post-2007 income tax applies to all forms of individual labor income, including wage earners and the self-employed, using a progressive tax scheme with six income brackets, annually adjusted by CPI, with marginal tax rates ranging from 0% to 25%. It is well known that changes in income taxes can lead to changes in the distribution of pre-tax earnings because of different types of labor supply responses (Saez et al., 2012). However, a few comments follow. First, as is the case with the collective bargaining reform, the timing of the tax reform prevents it from rationalizing the immediate effects documented around 2005. Second, below we document that the decrease in wage inequality exhibited after 2005 is mainly explained by faster wage growth at the bottom and an increased pay premia of initially low-paying firms, a pattern that is unlikely to be tied to changes in marginal tax rates at the top. Third, we mostly focus on the inequality of hourly wages, thus abstracting from labor supply responses that relate to the hours’ margin (and, reassuringly, we find negligible effects on hours’ inequality after the reform). Finally, many of our analyses focus on low-wage firms and workers, which are unlikely to be representative of agents exposed to the most salient changes in personal income taxes at the top. Then, we believe it is unlikely that the tax reform is a relevant confounder of our NMW analysis.

3 Data

This section describes the data sources, variable definitions, and cleaning procedures. Additional details can be found in Appendix A.

3.1 Data sources

Our main data source is the individual-level administrative records from the Uruguayan Social Security Administration (*Banco de Previsión Social*, SSA). These records provide matched employer-employee labor histories constructed from the payroll tax forms that firms are required to file monthly to report social security contributions to the SSA. The dataset covers all formal workers who reported any positive earnings to the SSA for at least one month between 1997 and 2013 and includes monthly information on workers' gross earnings, hours worked, days worked within the month, and sociodemographic characteristics, including gender and birth date. The dataset also contains firm-level information, most importantly firm identifiers and industry indicators (5-digit ISIC codes).

We also access Corporate Income Tax (CIT) data from the Uruguayan Tax Authority (*Dirección General Impositiva*), which covers the universe of medium-sized and large private firms in Uruguay between 2009 and 2017. All firms in Uruguay, except for self-employed businesses and other small-sized firms subject to special regimes, must file an annual tax return to pay the CIT. These data provide information on balance sheets and income statements, allowing us to measure value added at the firm level. We merge the employer-employee labor history data with the CIT data using an auxiliary dataset from the SSA that merges worker identifiers from the income tax data with those from the SSA. This key allows us to use the income tax records to obtain the tax identifier of the firm in which the individual is employed and thus merge the firm-level CIT records to the SSA labor histories.

The tax data has two important limitations. First, the auxiliary dataset used to recover tax identifiers comprises a non-random subset of workers from the SSA records, resulting in our inability to match the CIT records for the universe of firms present in the SSA data. Using this procedure, we can match the tax records of 72% of the firm-year observations that appear in our final estimation sample between 2009 and 2013, which account for 78% of the total employment in our sample (after the sample restrictions).³ Second, we only access the corporate tax records for 2009 onwards, which implies that we cannot use this data in our main causal estimates of the effects of the 2005 reform. Therefore, balance sheet variables are only used in descriptive exercises to validate our procedures.

3.2 Main variables

The main outcome used in our empirical analysis is the real hourly wage, which is built by dividing annual earnings by annual hours worked in each employer-employee match, properly adjusted by inflation. Monthly variables are annualized by adding up all monthly entries within a worker-firm-year combination. Our final dataset is a worker-level unbalanced panel that records the annual outcomes for the primary

³An individual appears in this auxiliary SSA dataset if they have enabled a dependent to access certain social security benefits. The most common scenario is when formal workers provide health insurance to other household members, typically their children. Other programs, such as conditional cash transfers, are also included.

employer of each worker, defined as the employer that delivered the highest earnings each year.

Annual earnings comprise all forms of labor compensation, including regular salary (which generally accounts for more than 90% of total earnings), the 13th salary (paid half in July and half in December), and any additional payments made to the worker after the termination of the match, before taxes and social security contributions. Our measure of annual hours is the total hours worked in a year, including legal vacation, holidays, and lunch breaks. When employers report both weekly hours worked and total days worked within a month, monthly hours can be computed without any imputation. This case represents roughly 75% of our baseline yearly panel observations. When employers report weekly hours but days worked are missing, the worker is assumed to have worked the complete month in the firm. This case represents roughly 5% of our baseline yearly panel observations. Finally, when employers report days worked but weekly hours worked are missing, the worker is assumed to work the standard labor contract of 8 hours per day. This case represents roughly 20% of our baseline yearly panel observations.

A typical concern about data on days and hours worked is the quality of the reported information. In Uruguay, employers are required to report hours worked in the SSA records, and the reported information is subject to periodic controls by the SSA. Firms that are found to be providing erroneous information are subject to penalties and fines. Then, firms have an incentive to report hours worked accurately. Despite this institutional feature, we assess the hour data quality following [Lachowska et al. \(2022\)](#), who propose a sequence of validation tests of administrative data on hours. The results from these evaluations suggest that the quality of the Uruguayan records is generally high (see [Appendix A.1](#) for details).

Finally, we closely follow [Harasztosi and Lindner \(2019\)](#) to process the CIT records and build, at the firm-by-year level, measures of value added per worker and within-firm labor shares. Value added is broadly defined as the value of production (revenue plus operational grants) minus production costs (all operating expenses excluding labor costs). The firm-level labor share is the ratio between total gross wages and value added, where total gross wages are computed by aggregating all employees' labor expenses using the SSA records before imposing our sample restrictions. We also provide results using an alternative measure of labor share equal to total gross wages over total gross wages plus pre-tax profits. More details on the computation of value added can be found in [Appendix A.2](#).

3.3 Annual panel and estimation sample

We follow standard practices (e.g., [Sorkin, 2018](#); [Lachowska et al., 2020](#)) in the treatment of matched employer-employee data to process the monthly records to build an annual panel of hourly wages at the worker level with firm identifiers, where the relevant firm for each worker-year pair is the primary employer of the worker based on annual earnings. We provide details on the steps taken to go from the raw monthly records to the annual worker-level panel in [Appendix A.3](#). One advantage of our data is that the raw entries are recorded at the monthly level so we can observe exactly when job matches are

created and destroyed. This implies we don't need to implement adjustments (needed in quarterly or annual data) to annualize outcomes when the start and end dates of job matches are not observed.

Descriptive statistics for the yearly worker-level panel are displayed in Panel (a) of Table 1. Pooling all years, we have an unbalanced panel of 10,901,288 worker-year observations, with 1,720,159 different workers that work in 155,329 different firms. The average log hourly wage decreased in the 2001-2005 period (relative to the 1997-2001 period) but then increased in the post-2005 periods. The standard deviation of log wages decreased after 2005. The average worker is between 36 and 37 years old, and between 58% and 61% of the observations pertain to male workers.⁴

Since our analysis below mostly uses the data to estimate AKM and TV-AKM models, we implement further (standard) restrictions to the sample to make it suitable for the empirical analysis. First, to avoid labor market transitions given by education or retirement decisions, we only keep job matches where the worker is always between 20 and 60 years old. Second, we focus on workers with strong labor market attachment by dropping worker-year observations where the worker worked less than 400 hours per year in the primary employer. Third, we drop outliers in hours worked (workers with average weekly hours greater than 60, <0.1%) and hourly wages (smaller than 5 Uruguayan pesos, <0.5%). Fourth, we only keep firms with more than 4 valid observations (i.e., after the age and labor market attachment restrictions are imposed). This restriction ensures a minimum firm size to estimate within- and between-firm variances and the corresponding F-FEs. Panel (b) of Table 1 shows that when pooling all the years, the number of observations is reduced by 47% after imposing all these restrictions. Not surprisingly, the most important restriction is the one pertaining to firm size since in Uruguay (as in many developing economies) the firm-size distribution is skewed to smaller firms. Indeed, we see that, after imposing the restrictions, the number of surviving workers is in the order of 60% relative to the baseline panel, while the number of surviving firms is below a fifth. While the age and gender composition of this sample is comparable to the baseline panel, we see that the restricted sample is slightly higher-wage and displays a slightly smaller wage variance.

Finally, it is well known that AKM models are identified by job switchers in the largest connected set of firms (Card et al., 2013; Kline, 2024). Panel (c) of Table 1 shows descriptive statistics for the largest connected set after imposing the abovementioned restrictions. 99.9% of the observations of the yearly panel with restrictions are part of the largest connected set when pooling all years, which is not surprising given the firm-size restriction (mechanically, larger firms are more likely to be connected). We use this worker-level panel as the input for all our exercises below.⁵

⁴The sample described in Panel (a) of Table 1 is the one used to compute median wages in Panel (c) of Figure 1.

⁵When looking at specific 5-year windows, the connected set restriction has a larger impact. This happens because firms can be connected through moves that happen outside the interval considered. Except for the baseline AKM decompositions by 5-year period, all exercises below are based on the dataset that pools all years.

4 Stylized Facts

We present four stylized facts using our processed annual worker-level panel to motivate our analysis.

Panel (a) of Figure 2 plots the annual standard deviation of log hourly wages and the real hourly NMW. The evolution of wage inequality closely mirrors the evolution of the real NMW. The correlation between the two series is -0.98. The figure shows that, between 1997 and 2002, the real NMW remained constant while the standard deviation of log hourly wages increased by 5%, from 77 log points to 81 log points. In the years 2003 and 2004, the real NMW decreased while wage inequality slightly increased. However, in 2005, the real NMW increased by 70%, and the standard deviation of log hourly wages decreased by 8% (from 82 log points to 76 log points), more than reverting the increase in inequality experienced in the previous 8 years. This pattern persisted during the post-reform period. In 2013, the real hourly NMW was more than three times larger than the one prevailing in 2004, and the standard deviation of log hourly wages reached 62 log points, 76% of the level prevailing in 2004. In Figure B.1 of Appendix B, we extend Panel (a) of Figure 2 by including the standard deviation of log monthly earnings and log monthly hours. The figure reveals that baseline inequality in monthly hours is an order of magnitude smaller (around 0.3 log points) and does not exhibit a change after 2005. Consequently, the standard deviation of log monthly earnings closely follows the standard deviation of log hourly wages.

Panel (b) of Figure 2 plots the evolution of different wage percentiles relative to their values in 1997. Real wages decreased before 2005 across the whole wage distribution, with lower percentiles (p5, p10, and p25) being especially hurt by the recession experienced in the first half of the 2000s. Starting in 2005, however, hourly wages at the bottom grew much faster than the rest of the wage distribution. By 2006, the p5 and p10 of the wage distribution had already recovered their pre-crisis real values despite the stronger dip before the reform, and the faster growth rates persisted until 2013. On the contrary, higher percentiles exhibited much slower growth rates, achieving full recovery to pre-crisis levels only around 2009-2010. These patterns suggest that the overall decrease in inequality documented in Panel (a) was materialized by wage compression at the bottom of the distribution. Indeed, Figure B.2 of Appendix B shows that lower-tail percentile ratios (p50/p5, p50/p10, and p50/p25) experienced mild trends before 2005 but exhibited dramatic decreases after the 2005 reform. This pattern contrasts with upper-tail percentile ratios (p95/p50, p90/p50, and p75/p50), which exhibit much smoother and smaller decreases after the 2005 reform, suggesting that dynamics at the top are not driving the aggregate patterns.

Panel (c) of Figure 2 provides a standard employment-weighted bias-corrected variance decomposition of the aggregate trend depicted in Panel (a), where the total variance in log wages is split into a between-firm component and a within-firm component.⁶ To ease the interpretation of the levels, the components

⁶The bias correction follows standard practices (e.g., Krueger and Summers, 1988). The between-firm component is adjusted by the square of the standard error to avoid attributing sampling error to the true between-firm variance component.

are expressed in standard deviations. Two aspects of the figure are worth discussing. First, before the reform, the between-firm component was substantially larger than the within-firm one. On average, between 1997 and 2004, the within-firm component was equal to 0.48 log points, while the between-firm component amounted to 0.63 log points. Second, almost all of the decrease in inequality observed starting in 2005 can be attributed to a decrease in between-firm inequality. Indeed, the decrease in the between-firm component is large enough to completely close the gap between components in 2012 and even revert it in 2013. It is well known that changes in between-firm inequality may arise either from changes in firms' wage premia or by changes in sorting patterns of workers to firms (Song et al., 2019). The next sections explore which of these stories is more likely to rationalize this aggregate pattern.

Finally, Panel (d) of Figure 2 shows the evolution of the standard deviation of log hourly wages separately for new hires and incumbent workers. In this simple exercise, a new hire is a worker who either appears for the first time in the sample (after 1997), the primary employer is a different firm relative to the previous year, or appears in the sample after disappearing for at least one year. Two aspects of the figure are worth discussing. First, the cross-sectional inequality of incumbent workers is larger than the one exhibited for new hires every year, as illustrated by the different levels of the axes. Second, the slight increase in inequality before 2005 and the sharp decrease in inequality after 2005 are experienced by both series. While several composition stories can rationalize these patterns, the fact that the decrease in inequality is experienced by incumbent workers suggests that the firms in which they work may be changing due to the NMW reform.

5 AKM Analysis

The descriptive facts presented in the previous section suggest that firms were important mediators of the change in inequality after the NMW reform of 2005. In the rest of the paper, we formally test this claim. This section and the following use the AKM (Abowd et al., 1999) and TV-AKM frameworks to descriptively address the role of firms. Sections 7, 8, and 9 take the analysis one step further to assess whether the NMW had a causal impact on the role of firms in mediating aggregate inequality.

We first estimate standard AKM regressions using the estimation sample described in Section 3:

$$\log w_{it} = \alpha_i + \psi_{j(i,t)} + X'_{it}\beta + \varepsilon_{it}, \quad (1)$$

where w_{it} is the hourly wage of worker i in year t , α_i is a worker fixed-effect, $j(i,t)$ denotes the firm where worker i is employed in year t , $\psi_{j(i,t)}$ is the firm fixed effect (F-FE), X_{it} is a vector of time-varying covariates that include a (properly normalized) polynomial in age and year fixed-effects, and ε_{it} is the error term.⁷ Following previous work (e.g., Card et al., 2013), we estimate equation (1) separately for

⁷As discussed in Card et al. (2018), controlling by age is key but can lead to standard collinearity problems between age,

overlapping 5-year periods: 1997-2001, 2001-2005, 2005-2009, and 2009-2013.

Identification restrictions for having unbiased estimates of the fixed effects are well understood in the literature (see [Kline, 2024](#) for an extensive discussion). Assuming the model is correctly specified (i.e., fixed effects are time-invariant and worker and firm components are additively separable), the identification argument is summarized by two conditions. First, the model is identified by job switchers within a connected set of firms since, in the absence of these switches, worker and firm fixed effects are not separately identified. As a consequence, F-FEs are only identified up to a constant since switches only identify differences in pay premia. Second, the error needs to satisfy strict exogeneity:

$$\mathbb{E}[\varepsilon_{it} | i, j(i, s) = \iota, X_{is} = x] = 0, \quad (2)$$

for all $s \in \{1, \dots, T\}$. In terms of the mobility network, strict exogeneity implies that firm switches can be selected on time-varying covariates and F-FEs, but they cannot be selected on time-varying shocks ε_{it} .

While these restrictions are inherently untestable, a typical practice to assess their plausibility is to report “event studies” of log hourly wages of switchers around job transitions ([Card et al., 2013, 2018](#)). We define a switcher within a 5-year period as a worker that stays for (at least) 2 consecutive years in a firm and switches to a different firm and stays for (at least) 2 consecutive years. Then, we compute employment-weighted quartiles of firms based on the hourly wages of switchers’ coworkers and plot the residualized (against year fixed effects and age controls) log hourly wage of switchers that experience transitions between firms in different quartiles. Under the assumptions of the model, there should be no trends before the switch, and the switch should generate a discrete and symmetric change in wages depending on the origin and destination firms (based on the coworkers’ wages quartiles). [Figure B.4](#) of [Appendix B](#) shows these event studies separately for each of the four 5-year periods. Pre-switch trends in hourly wages are essentially flat, and wage changes around moves are remarkably symmetric across quartiles, suggesting that the AKM assumptions may provide a reasonable approximation in our setting.

Another diagnostic to assess the restrictiveness of the assumption of additive separability between worker and firm components consists of comparing the adjusted R-squared of the AKM regression with the adjusted R-squared of a regression that replaces the worker and firm fixed effects with “match” fixed effects, i.e., a set of indicators $\varphi_{ij(i,t)}$ unique to each job match between a worker and a firm. Averaging across the four 5-year periods, the adjusted R-square of the AKM regression is 0.915, while the adjusted R-square of the “match effect” model is 0.936. The mild increase in regression fit suggests there is little loss in assuming the additive representation. In regressions using the full period 1997-2013, the adjusted R-squares are 0.851 and 0.895, respectively. The overall decrease in fit relative to the average across

year, and cohort effects, biasing the estimates of the worker fixed effects. The recommendation is to control by a polynomial centered around the age where the life-cycle pattern of hourly wages flattens. In our data, this happens approximately at age 52 (see [Figure B.3](#) of [Appendix B](#)), so we include a third-degree polynomial in $\widehat{\text{Age}}_{it} = (\text{Age}_{it} - 52)$.

individual 5-year periods suggests that time-varying changes in F-FEs may be important. Indeed, the adjusted R-square of the TV-AKM model discussed in the next section that allows for time-varying F-FEs is 0.887, a fit that almost reaches the one of the general match effects model.

Using the AKM model estimates, it is standard to report a variance decomposition of log wages into worker and firm components. Abstracting from the role of time-varying covariates, we can write:

$$\mathbb{V}(\log w_{it}) = \mathbb{V}(\alpha_i) + \mathbb{V}(\psi_{j(i,t)}) + 2 \cdot \mathbb{C}(\alpha_i, \psi_{j(i,t)}) + \mathbb{V}(\varepsilon_{it}), \quad (3)$$

where $\mathbb{V}(\cdot)$ denotes the variance operator and $\mathbb{C}(\cdot)$ denotes the covariance operator.⁸ Each variance component describes the share of wage inequality explained by heterogeneity in worker and firm fixed effects. The covariance component is usually interpreted as a measure of worker sorting: the share of wage inequality explained by high- α_i workers working in high- $\psi_{j(i,t)}$ firms. When the AKM identification assumptions hold, the estimated fixed effects are unbiased but the variance decomposition is not, a problem known as “limited mobility bias” (Andrews et al., 2008; Bonhomme et al., 2023). In simple words, fixed effects are unbiased but noisy estimates, so variance components square the standard error, spuriously attributing noise to the relative importance of firms. Since the identification of F-FEs is based on job switchers, this estimation error is inversely proportional to the number of movers per firm (Bonhomme et al., 2023; Lachowska et al., 2023). To deal with this problem, we report bias-corrected variance decompositions based on the leave-out method proposed in Kline et al. (2020).⁹

5.1 Results

Table 2 shows the results of the variance decomposition using the log hourly wage as the dependent variable. Panel (a) shows the bias-corrected results. Panel (b) shows the (biased) naive plug-in decompositions for reference. As expected, the bias correction decreases the relative importance of F-FEs and increases the relative importance of the covariance term. The bias, however, does not seem substantial enough to produce negative covariance terms. This feature reflects the non-trivial number of movers per firm present in our dataset, possibly facilitated by our restriction based on firm size. Table B.2 of Appendix B shows similar results using monthly earnings as the dependent variable. For the variance components, the tables also report standard deviations so values can have a log point interpretation.

Consistent with Figure 2, Table 2 shows that the variance of wages increased between 1997-2001 and 2001-2005 to then exhibit a substantial decrease in the 2005-2009 and 2009-2013 periods. The AKM

⁸In principle, this decomposition should also include variance and covariance terms for the portion explained by $X'_{it}\beta$. Our results below, however, residualize against time-varying covariates to focus on the worker, firm, and sorting components.

⁹The bias correction of Kline et al. (2020) further restricts the sample of firms to the leave-one-out largest connected set (LOO-LCS), a sample of firms that remain connected after randomly extracting one worker in each firm. Table B.1 of Appendix B compares the panel restricted to the LCS to the panel restricted to the LOO-LCS. Around 10% of the worker-year observations are lost in the LOO-LCS, with no detectable differences in average wages, age, or gender composition.

decomposition shows that firms played an important role in the decrease in inequality observed after 2005. While all components faced a decrease in levels, the decrease was particularly large for the firm component. This implies that while the relative contribution of worker fixed effects increased after 2005 and the relative contribution of the covariance term remained constant, the relative contribution of firms decreased from around 30% in the pre-2005 period to 22% in the 2005-2009 period and to 18% in the 2009-2013 period.¹⁰ These results suggest that the descriptive patterns discussed in Section 4 are related to a change in the distribution of F-FEs. We also observe that the total variance explained by the worker and firm fixed effects decreased throughout the period, suggesting that the NMW increase may have given low-wage workers a premium above the one that firms pay to all workers.

Table B.3 of Appendix B presents an analog variance decomposition using monthly hours as the dependent variable. The levels of dispersion are an order of magnitude smaller than the ones observed for hourly wages (around 0.3 log points), with no difference between the years before and after the reform. Compared to the results reported in Lachowska et al. (2024), the level of hours inequality is slightly smaller (around 0.3 log points versus 0.35 log points), but the share explained by our two-way fixed-effects model is much larger (around 70% versus 35%). The relative importance of workers and F-FEs is symmetric and around five times larger than the role of the covariance terms.

6 TV-AKM Analysis

The previous analysis shows that the role of F-FEs in explaining wage inequality decreased after the minimum wage reform. It is difficult, however, to dissect precisely what is driving this evolution. The set of firms considered in each 5-year period and their employment shares may vary, implying that this evolution may reflect composition effects. The overlap of firms between periods, however, is substantial, which also allows for the possibility of within-firm changes between periods. This latter possibility is difficult to assess within the AKM framework because F-FEs are identified up to a constant, so the levels of F-FEs are not comparable between periods. Additional restrictions must be imposed for the required normalization within each period to be comparable across regressions.

To sidestep this normalization problem, we estimate the Time-Varying AKM model (TV-AKM) recently proposed by Engbom et al. (2023) and Lachowska et al. (2023). The TV-AKM model relaxes the assumption of constant F-FEs by allowing for within-firm drifts in wage premia. Then, the evolution of the role of firms can be studied using the complete sample (1997-2013) without needing arbitrary normalizations on the estimated F-FEs. As an additional benefit, using the 1997-2013 period instead of 5-year periods increases the potential number of movers per firm and expands the set of firms considered

¹⁰Kline (2024) compares different bias-corrected variance decompositions and finds that developing economies generally exhibit a higher standard deviation of F-FEs. Our pre-reform estimates align with this pattern. However, the reduction experienced after the 2005 reform puts Uruguay closer to the levels observed in developed economies.

in the largest connected set, thus attenuating limited mobility concerns ubiquitous in AKM regressions.

Formally, the estimating equation is given by:

$$\log w_{it} = \alpha_i + \psi_{j(i,t)t} + X'_{it}\beta + \varepsilon_{it}, \quad (4)$$

where, now, the F-FEs are indexed by t . All other objects are defined as in equation (1) with the detail that, naturally, year fixed effects are excluded from the vector X_{it} since F-FEs now absorb the aggregate time variation. As thoroughly discussed in Engbom et al. (2023) and Lachowska et al. (2023), the identification argument is analogous to the one pertaining to the standard AKM model by reinterpreting firm-year combinations as different firms, meaning that stayers play the role of switchers between firm (j, t) and $(j, t + 1)$. While firm stayers play no role in identifying F-FEs in the AKM model, they play a key role in the TV-AKM by pinning down the within-firm change in F-FEs. The strict exogeneity assumption is also required, but now conditional on the set of time-varying firm wage premia.

How much fit is gained by allowing the F-FEs to be time-variant? Figure B.5 of Appendix B shows projections between the AKM F-FEs effects estimated using the complete sample (1997-2013) and the TV-AKM F-FEs. Panel (a) shows a binned scatter plot representation of a projection of the TV-AKM F-FEs into the AKM F-FEs. Consistent with Engbom et al. (2023) and Lachowska et al. (2023), we find a projection slope of 0.99, with all bins aligned with the 45-degree line. That is, omitting time variation in F-FEs does not importantly affect estimates of the average pay premia along the period. Panel (b), however, shows a binned scatter plot representation of a projection of the AKM F-FEs into the TV-AKM F-FEs. Similar to what is found by Engbom et al. (2023), we estimate a much smaller projection slope of 0.62, suggesting substantial within-firm variation across time in pay premia.

Correlating F-FEs with firm characteristics. We use the CIT data described in Section 3 and Appendix A.2 to compute, for 2009-2013, two variables at the firm-by-year level and explore whether they correlate with the time-varying F-FEs. First, we compute value added per worker to test whether F-FEs are correlated with firm productivity. Second, we compute firm-level labor shares that equal total gross wages over value added to test if, conditional on productivity, firms that share more rents display larger F-FEs. This variable aims to proxy a firm-level rent-sharing elasticity that may be heterogeneous across firms and time. Similar results are found when measuring the firm-level labor share as total gross wages over total gross wages plus pre-tax profits (see Figure B.6 of Appendix B).

Figure 3 presents the results. Panel (a) shows a binned scatter plot that relates F-FEs with value added per worker. Panel (b) shows a binned scatter plot that relates F-FEs with the firm-level labor share (conditional on value added per worker). In each figure, the blue curve shows the raw correlations, the red curve controls by industry fixed effects (5-digit codes), and the green curve controls by (time-invariant) firm indicators. In the latter case, slopes are built from within-firm comparisons across time.

Panel (a) shows that F-FEs are positively correlated with firm productivity. Consistent with what has been found in other contexts (see [Kline, 2024](#) for a survey), the positive correlation displays a “hockey stick” pattern where the relationship is flat at the bottom to then give form to a positive log-linear relationship. When including industry fixed effects, the projection slope is attenuated but remains large and significant, suggesting that differences in productivity across industries explain a minor part of the baseline correlation. A novel insight of our analysis is that, after including firm indicators, the slope decreases to around 44% of the baseline estimate but remains significant. This implies that within-firm changes in rents across time have predictive power for changes in firm wage premia.

Panel (b) shows that, conditional on value added per worker, firms that share more rents with workers display larger F-FEs. This correlation is steeper than the ones displayed in Panel (a). This result – which, to the best of our knowledge, has not been documented in the literature – suggests that variation in rent-sharing elasticities plays an important role in explaining the dispersion in F-FEs. Interestingly, the inclusion of industry and firm indicators does not attenuate the estimated slope (it is even steeper with firm indicators), suggesting that within-firm changes in rent-sharing across time seem important to predict changes in F-FEs. Related to our research question, one can think of labor market policies such as the minimum wage to have an effect on the rent-sharing policy in the firm. This descriptive fact helps motivate our research question that asks whether minimum wages can causally affect the F-FEs.

6.1 Results

The TV-AKM model allows us to work directly with the estimated F-FEs for each firm and year to dynamically analyze the role of firms on wage inequality. We therefore develop several descriptive exercises using the complete set of estimated F-FEs as the unit of analysis.

Revisiting dispersion in F-FEs. A first descriptive exercise consists of extending the AKM decomposition reported in [Table 2](#) by manually aggregating the estimated F-FEs year by year instead of by 5-year periods. This strategy allows the assessment of different aggregations of the estimates, informing about the importance of the composition effects. [Figure B.7](#) in [Appendix B](#) shows naive standard deviations (i.e., non-biased-corrected) of F-FEs for each year in our sample. Assuming that the limited mobility bias is relatively constant across time ([Card et al., 2013](#); [Bonhomme et al., 2023](#)), the figure is valid to get insights on the drivers of the trends of F-FEs dispersion.

The figure plots four series: the employment-weighted standard deviation of F-FEs using all firms in the panel (unbalanced panel), the employment-weighted standard deviation of F-FEs using firms that appear in each of the 17 years of our panel (balanced panel), and corresponding versions that do not weight by employment. The unbalanced-weighted series incorporates all mechanisms through which firms can mediate impacts in aggregate inequality (changes in composition and within-firm F-FEs levels), thus being closer in spirit to the results of [Table 2](#). The balanced-weighted series omits the effects of firm

entry and exit, leaving the composition of firms fixed. The unbalanced-unweighted series abstracts from changes in the employment composition across firms. Finally, the balanced-unweighted series omits both composition effects in terms of active firms and employment. All four series exhibit a substantial decrease after the 2005 reform that mimics the aggregate evolution of wage inequality documented in Figure 2, a pattern that is especially interesting for the balanced-unweighted series: holding constant the set of firms and the employment composition across them, dispersion in F-FEs decreased after the 2005 reform.

Persistence in F-FEs. The across-time comparability of time-varying F-FEs has been used by previous literature to assess how persistent F-FEs are (Engbom et al., 2023; Lachowska et al., 2023). Two main findings emerge from these papers. First, F-FEs are persistent within firms, especially in the short run. Second, F-FEs are cyclical, with firms paying smaller premia in recessions and vice versa. We use our time-varying F-FEs estimates to revisit these conclusions in our data. Figure B.8 of Appendix B shows, for each pair of consecutive years in our sample, the within-firm short-run persistence of the estimated F-FEs. Each plot shows a raw binned scatter plot of the F-FE in year t against the F-FE in year $t - 1$. Figure 4 shows the plots for four selected pair of years: the endpoints of our sample (1997-1998 and 2012-2013) and the pre- and post-NMW reform years (2003-2004 and 2004-2005).

In the pre-reform period, F-FEs are very persistent and slightly cyclical, mimicking the findings of Engbom et al. (2023) and Lachowska et al. (2023). In 1997-1998, 1998-1999, 1999-2000, and 2000-2001, the scatterplots lie almost exactly on the 45-degree line. As a consequence, the support of the F-FEs distributions is very stable. In 2001-2002 and 2002-2003, there is a decrease in the F-FEs levels across the whole distribution (a “year fixed effect”), possibly reflecting the effects of the recession. In 2003-2004, the cyclical shock stabilizes, and F-FEs are located again on top of the 45-degree line.

Starting in 2005, however, the perfect persistence breaks and the binned scatter plots start displaying a “hockey stick” pattern. In 2004-2005, firms at the bottom of the distribution of the F-FEs located far above the 45-degree line, meaning they experienced faster growth in their F-FEs relative to firms up in the distribution that remained located on top of the 45-degree line. In the years that follow, the “hockey stick” pattern persists, although it gets attenuated through time. Close to the end of our analysis period, the short-run behavior returns to its baseline persistence, with binned scatter plots located again close to the 45-degree line along the whole distribution. Importantly, because of several years of “hockey stick” behavior, the support of the F-FEs was much more compressed in 2012 relative to before the 2005 reform.

To rule out that these patterns are an artificial consequence of transitory shocks and mean-reversion, Figure 5 shows long-run changes in the distribution of F-FEs using the balanced sample of firms for which we have F-FEs estimates for each of the 17 years of our sample. The figure plots the average F-FE between 2005 and 2013 against the average F-FE between 1997 and 2004. The hockey-stick pattern becomes clearer, with firms in the bottom half of the distribution showing a sizable departure from the 45-degree line, while firms up in the distribution showing a much smaller and evenly distributed increase in

pay premia over their baseline level. As a further exploration of the long-run changes in the distribution of F-FEs, we define unweighted deciles of firms based on their F-FE level in 2004 and analyze their evolution in time, holding the deciles fixed. Figure 6 reports the results for both an unbalanced panel of firms (firms that were active in 2004) and a balanced panel of firms (firms that were active during the 17 years of our sample). Both figures show that, before the reform, trends looked parallel across deciles and exhibited cyclical behavior but that, starting in 2005, low-paying firms exhibited a much faster growth in their average F-FE, catching up with the average values of upper deciles. At the end of the period, the cross-sectional dispersion in F-FEs was much smaller because of the compression of the distribution.

These findings suggest that the substantial persistence in F-FEs documented in previous work may not extend to all settings and policy contexts. If the NMW reform is causing these trend breaks around 2005, the results suggest that F-FEs are, indeed, idiosyncratically persistent, but can be affected by policy reforms. The rest of the paper addresses this causal question using two different research designs.

7 The Causal Effect of the Minimum wage

Before presenting our empirical strategies for testing the causal effect of the NMW reform on the F-FEs, we provide a brief conceptual discussion to develop intuition. We do this by presenting a very stylized model that builds on the discrete choice model presented in [Card et al. \(2018\)](#) (CCHK).

7.1 Model

Workers are heterogeneous in skill and firms are heterogeneous in productivity. To simplify exposition, assume there are low- and high-skill workers indexed by $s \in \{l, h\}$, and low- and high-productivity firms indexed by $j \in \{1, 2\}$. Workers have heterogeneous preferences for firms (conditional on skill) and firms operate with linear technologies in labor, so production of firm j takes the form $Y_j = R_j (\theta_l L_j + \theta_h H_j)$, where R_j is the firm-specific productivity shifter, θ_s is the relative productivity of workers of skill s , and (L_j, H_j) is the number of low- and high-skill workers hired by firm j , respectively.

CCHK shows that, if there are no strategic interactions between firms, a set of parametric assumptions yields an equilibrium equation for log wages that has an AKM representation:

$$\log w_{sj} = \alpha_s + \psi_j \equiv \log \theta_s + \log(1 + R_j), \quad (5)$$

where, for simplicity, we have normalized to 1 constants that do not affect our argument below. For completeness, Appendix C derives equation (5) from first principles to explicitly state details on the baseline structure, the derivations, and the additional restrictions.

Wage setting norms. To analyze the effects of introducing a binding minimum wage on the struc-

tural equation (5), we extend the CCHK model to incorporate firm-specific wage setting norms that potentially mediate within-firm wage spillovers (e.g., Dube et al., 2019; Giupponi and Machin, 2024).

To fix ideas, assume that, in the absence of a minimum wage, $w_{l1} < w_{l2} < w_{h1} < w_{h2}$. From equation (5), it follows that $w_{hj}/w_{lj} = \theta_h/\theta_l$ for every j . Workers and firms internalize this equilibrium condition, which gives rise to a wage-setting norm of relative wages within the firm. We assume that, in addition to being heterogeneous in productivity, firms are heterogeneous in the extent to which they respect the wage-setting norm, which we parametrize by the parameter β_j . β_j affects wage-setting behavior as follows. Suppose that, because of some regulation (e.g., the NMW), w_{lj} exogenously increases to $w_{lj}(1 + a_j)$, for some a_j . The firm-specific wage-setting norm restricts the behavior of the high-skill wage, which is forced to increase to $w_{hj}(1 + \beta_j a_j)$. If $\beta_j = 0$, there is no within-firm spillover and we are back to the baseline CCHK model. If $\beta_j = 1$, there is full passthrough to high-skill wages so the new wages respect the benchmark θ_h/θ_l . In a case with $\beta_j \in (0, 1)$, there will be a partial spillover to the high-skill wage.

Introducing a minimum wage. We denote by $x^{\bar{w}}$ the value of a variable x when the minimum wage equals \bar{w} . Values of x when there is no binding minimum wage are denoted by x^0 .

Assume that the government introduces a minimum wage $\bar{w} \in (w_{l1}^0, w_{l2}^0)$, so \bar{w} only binds low-skill workers in the low-paying firm. Since there are no strategic interactions between firms, wages in firm $j = 2$ are not affected by the minimum wage, so $w_{s2}^{\bar{w}} = w_{s2}^0$, for $s \in \{l, h\}$. Likewise, the firm-specific component is fixed in firm $j = 2$, so $\psi_2^{\bar{w}} = \psi_2^0$. It follows from equation (5) that $\alpha_s^{\bar{w}} = \alpha_s^0$, for $s \in \{l, h\}$.

On the contrary, wages in firm $j = 1$ are affected by the minimum wage. The new equation for low-skill wages in firm $j = 1$ is given by $\log w_{l1}^{\bar{w}} = \alpha_l^{\bar{w}} + \psi_1^{\bar{w}} + e_{l1} \equiv \log \bar{w}$, where e_{l1} is a potential residual generated by the introduction of the minimum wage. Using the restriction $\alpha_l^{\bar{w}} = \alpha_l^0$, we can write $\log \bar{w} - \log w_{l1}^0 = \psi_1^{\bar{w}} + e_{l1} - \psi_1^0$. While the behavior of firm $j = 2$ restricts this wage change to not be caused by a change in the worker fixed effect, it remains an open question whether it is driven by a change in the F-FE or the residual term.

Both terms can be separately identified using the wage equation of high-skill workers in firm $j = 1$. We know that $\log w_{h1}^{\bar{w}} - \log w_{h1}^0 = \psi_1^{\bar{w}} + e_{h1} - \psi_1^0$, where, by definition, $L_1 e_{l1} + H_1 e_{h1} = 0$. We note that $\log \bar{w} - \log w_{l1} = \log(1 + a_1(\bar{w}))$: the wage increase factor $a_1(\bar{w}) > 0$ after the minimum wage introduction is firm-specific and depends on the minimum wage. Putting all terms together, we have that:

$$\psi_1^{\bar{w}} + e_{l1} - \psi_1^0 = \log(1 + a_1(\bar{w})), \quad (6)$$

$$\psi_1^{\bar{w}} + e_{h1} - \psi_1^0 = \log(1 + \beta_1 a_1(\bar{w})). \quad (7)$$

Using $L_j e_{lj} + H_j e_{hj} = 0$ and solving for $\psi_1^{\bar{w}}$ yields:

$$\psi_1^{\bar{w}} = \psi_1^0 + \left(\frac{L_1}{H_1 + L_1} \log(1 + a_1(\bar{w})) + \frac{H_1}{H_1 + L_1} \log(1 + \beta_1 a_1(\bar{w})) \right), \quad (8)$$

which implies that:

$$e_{s1} = \log(1 + a_1(\bar{w})(1\{s = l\} + 1\{s = h\}\beta_1)) - \left(\frac{L_1}{H_1 + L_1} \log(1 + a_1(\bar{w})) + \frac{H_1}{H_1 + L_1} \log(1 + \beta_1 a_1(\bar{w})) \right). \quad (9)$$

The relative size of the change in the firm fixed effect and the residual is governed by β_1 and the within-firm employment shares. When $\beta_1 = 1$, all of the minimum wage increase is loaded into the firm fixed effect. Conditional on (L_1, H_1) , the relative importance of the residual decreases with β_1 .

7.2 Discussion

Several aspects of the model are worth discussing.

First, for simplicity, the model assumes that there are only two types of firms and workers and that the minimum wage only binds in firm $j = 1$. However, what is necessary for the argument to go through is the existence of firms where workers earning the minimum wage can counterfactually earn higher wages. The intuition comes directly from the switchers' logic in AKM. Assume that a low-skill worker moves from a firm that is constrained by the minimum wage to a firm that is not. Since the unconstrained firm is not changing its wage policy, the wage gain obtained in the constrained firm is not portable to the unconstrained firm and, therefore, cannot be loaded in the worker fixed effect. It is easy to see that if the minimum wage binds for low-skill jobs in all firms in the economy, then what is loaded in each component of the wage equation is no longer identified because the counterfactual wages of a worker earning the minimum wage would increase in tandem across all firms. This intuition implies that, conceptually, the minimum wage is a job (rather than a worker) attribute: the economy is populated by minimum wage jobs (taken more often by certain groups of workers, in our model, the low-skilled) rather than by an intrinsic population of minimum wage workers.¹¹

Second, the conclusion that worker fixed effects are unaffected by the minimum wage relies on the lack of responses in firm $j = 2$, which, in turn, relies on the absence of strategic interactions between firms. In a model with strategic interactions, the different components of the wage equation of unconstrained firms could react to the minimum wage, for example, because of general equilibrium effects generated by competition forces. A model with such characteristics can be found in [Morchio and Moser \(2024\)](#), who develop an equilibrium model based on [Burdett and Mortensen \(1998\)](#) that yields an equilibrium wage equation that has an AKM representation. A careful inspection of their identification result reveals that strategic interactions only affect the microfoundation of the F-FEs but do not affect the worker fixed

¹¹In our estimation sample, among all workers that earn around the minimum wage at least one year, less than 25% earn around the minimum wage every year. The remaining group oscillates between jobs that pay the minimum wage and jobs that pay above. This pattern is not only explained by pure life-cycle wage dynamics (maybe, within-firm) as we observe that a third of workers that earn around the minimum wage at least one year changes firm, and, among all those firm switches, almost 70% represent discrete changes from minimum wage jobs to non-minimum wage jobs and vice-versa.

effects which are only proportional to individual productivity (see their Proposition 1). This suggests that the logic developed in our model can go through in settings with strategic interactions.

Third, the model predicts that wages correlate with firm productivity and that there is wage dispersion conditional on worker skill and firm productivity driven by heterogeneity in β_j . These properties are consistent with the empirical patterns showed in Section 6, in particular, with the fact that F-FEs correlate value added per worker and with within-firm labor shares conditional on value added per worker.

Fourth, conditional on β_j , the employment composition of the firm also affects the passthrough of the minimum wage introduction to the F-FE. When $L_j/(L_j + H_j) \rightarrow 1$, the minimum wage shock is completely transmitted to the F-FE even if $\beta_j \rightarrow 0$. Then, employment segregation mediates how the shock maps into the distribution of F-FEs. Empirically, wage spillovers to jobs earning more than the minimum wage can be used to disentangle the role of wage-setting norms versus employment segregation.

Fifth, the conceptual exercise applies to other policies that affect firms' wage-setting behavior. Moreover, one can think of β_j as a function of other policies and labor market institutions, giving even more role to policy to affect firm wage premia. Examples of these policies and institutions include collective bargaining, unions, and pay transparency standards.

8 Causal Analysis: Firm-Level Design

Sections 4, 5, and 6 document a large decrease in wage inequality after 2005 mediated by changes in the F-FEs. In the remainder of the paper, we provide evidence that the NMW had a causal role in this evolution. This section presents a firm-level analysis. Section 9 presents a worker-level analysis.

8.1 The GAP design

Our firm-level exercise follows the GAP design first proposed by [Card and Krueger \(1994\)](#) and recently used by [Draca et al. \(2011\)](#), [Dustmann et al. \(2022\)](#), and [Derenoncourt and Weil \(Forthcoming\)](#). The idea of the GAP design is to proxy the firm-level exposure to the NMW reform by estimating how much firms should increase the labor costs to comply with the policy, using data from the pre-reform period.

Concretely, for each firm j , we compute the following statistic:

$$\overline{GAP}_j = \frac{1}{8} \sum_{t=1997}^{2004} GAP_{jt}, \quad (10)$$

where:

$$GAP_{jt} = \frac{\sum_{i \in j} h_{ij} \max\{0, 1.25 \cdot NMW_{2005} - w_{it}\}}{\sum_{i \in j} h_{it} w_{it}}, \quad (11)$$

where NMW_{2005} is the real hourly national minimum wage in 2005. For each firm j in a given year t , $\sum_{i \in j} h_{it} w_{it}$ are the total labor costs of the firm, with h_{ij} the hours worked by worker i in firm j and w_{ij} the corresponding hourly wage. Then, GAP_{jt} measures how much firm j labor costs should increase in year t as a share of the total labor costs to ensure that everyone earns at least 125% of the counterfactual minimum wage. The firm-level time-invariant exposure measure (“treatment”) is the average of GAP_{jt} across all pre-reform years (1997 to 2004). The rationale for this design is that firms that need to engage in higher payroll adjustments to comply with the minimum wage are the ones that will be more affected by the reform in 2005, so we can exploit variation in the level of exposure to recover the causal effect of the minimum wage reform using a difference-in-differences (DID) design.

Two comments follow. First, we choose a wage threshold equal to 1.25 times the NMW in 2005 because the NMW experienced non-trivial increases in the years after 2005 (see Figure 1), so firms that, before the reform, paid slightly above the 2005 value were possibly also exposed. Second, we follow the implementation in [Dustmann et al. \(2022\)](#) and [Derenoncourt and Weil \(Forthcoming\)](#) and use a time-invariant GAP measure that averages over several pre-reform years. The reason is twofold. First, having a time-invariant measure based on pre-reform variables is cleaner for the purposes of a DID design because post-reform GAPs are likely contaminated by the treatment effects of the 2005 reform. Second, the pre-reform GAP measure averages over several years to better deal with mean reversion because considering a single year may be picking a transitory shock in exposure, eventually leading to a spurious treatment indicator. This concern is particularly relevant in our context, where a recession was experienced in the pre-reform years, possibly leading to a higher short-run volatility.

Estimating equation. Using the GAP exposure measure described in equations (10) and (11), we estimate the following DID specification:

$$y_{jt} = \gamma_j + \lambda_t + \sum_{t=1997}^{2013} \beta_t \cdot \overline{GAP}_j + \varepsilon_{jt}, \quad (12)$$

where y_{jt} is the outcome of interest in firm j in year t , γ_j and λ_t are firm and year fixed effects, and ε_{jt} is the error term. The coefficients β_t recover the causal treatment effect of the NMW reform on y_{jt} under the parallel trends assumption: firms with different GAP levels would have trended in parallel, absent the NMW increase in 2005. While this assumption is not testable, we use estimates of β_t for $t \leq 2004$ to assess its plausibility. Since \overline{GAP}_j is a continuous treatment variable, β_t is interpreted as the effect of going from a 0% GAP to a 100% GAP (i.e., a situation in which the firm needs to double the labor costs to pay at least 1.25 times the 2005 NMW to all workers). We cluster the standard errors at the 2-digit industry level. We follow the standard practice of normalizing β_{2004} to 0. To avoid confounding treatment effects with composition effects, we estimate equation (12) using a balanced panel of firms that were active between 1997 and 2013 (i.e., during 8 years before and 8 years after the reform). This restriction leaves us

with 4,158 firms. We also report results focusing on 7, 6, and 5-year windows around the NMW reform, which allows us to increase the number of included firms by 7%, 15%, and 24%, respectively.

To summarize the results, we also estimate the following pooled regression using the balanced panel:

$$y_{jt} = \gamma_j + \lambda_t + \beta \cdot \overline{GAP}_j \cdot \text{Post}_t + \varepsilon_{jt}, \quad (13)$$

where $\text{Post}_t = 1\{t \geq 2005\}$ and all other variables are defined as in equation (12). In this regression, β measures the average treatment effect of the policy relative to the pre-reform window.

We estimate regressions (12) and (13) for three different outcomes. First, we use the hours-weighted average log hourly wage at the firm level, built from the worker-level dataset used in Sections 5 and 6. Second, we use the time-varying F-FE estimated using the TV-AKM framework, ψ_{jt} . This is our main outcome of interest because we aim to test whether the NMW reform of 2005 had a causal impact on the F-FEs. Finally, we use the hours-weighted firm-level average worker fixed effect estimated using the TV-AKM framework, $\bar{\alpha}_{jt} \equiv \left(\sum_{i \in j} h_{it}\right)^{-1} \sum_{i \in j} h_{it} \alpha_i$. With this outcome, we can test whether firms adapt to the NMW reform by changing the α -composition of workers.

8.2 Results

Figure 7 shows the results. We plot the estimated $\{\beta_t\}_{t=1997}^{2013}$ coefficients with their corresponding 95% confidence intervals. The three series represent different regressions that use different dependent variables: the average log hourly wage of the firm, the time-varying F-FE, and the average worker fixed effect, the latter two estimated with the TV-AKM model described in Section 6.

The figure shows that more exposed firms significantly increased the log hourly wage after the 2005 reform relative to less exposed firms. While there is a mild differential trend between more and less exposed firms several years before the reform (high-exposure firms had steeper downward trends in the recession years, consistent with Panel (b) of Figure 2), the difference in trends disappears in pre-reform years closer to 2005 and the change exhibited around 2005 is sharp, large, and immediate, providing suggestive evidence in favor of the parallel trends assumption. Reassuringly, Figure B.9 of Appendix B shows very similar results (and with better-behaved pre-trends) when considering the estimation samples with more firms but shorter windows around the 2005 reform. This sort of “first stage” confirms that the NMW reform had an economic impact on wages of exposed firms.

Figure 7 shows that more exposed firms also experienced faster growth in their TV-AKM F-FEs. Indeed, almost all of the increase in the average log hourly wage due to the NMW reform can be attributed to the increase in the F-FE. This result suggests that the descriptive patterns exhibited in previous sections can be causally linked to the NMW reform. The figure also shows that, as time passes, the average worker fixed effect increases in more exposed firms, suggesting that more exposed firms accommodate part of

the labor cost shock by changing the composition of workers within the firm, as in [Butschek \(2021\)](#) and [Clemens et al. \(2021\)](#). The results, however, suggest that the effect on the composition of workers is quantitatively an order of magnitude below the effect on the firm wage premium.

To put magnitudes in context, [Table 3](#) shows estimated β coefficients from equation (13). As a reference, 1.25 times the log real hourly NMW in 2005 was equal to 3.51. The baseline estimate of the effect on the average log hourly wage (Column (1)) is 3.14 log points. This coefficient is large and significant despite being smaller than the benchmark wage given that real wages also increased in control firms due to the business cycle (see [Figure 2](#)). The baseline estimates of the effect of the NMW reform on the F-FE and the average worker fixed effect (Columns (2) and (3)) are 2.69 and 0.39, respectively, confirming that the majority of the wage increase can be attributed to the increase in the F-FE.

As an additional robustness check, [Table 3](#) reports estimates of β that control for stricter time fixed effects. Columns (4) to (6) include year fixed effects that are allowed to vary by firm size. We do this by computing quintiles based on the average annual employment in the pre-reform period. Columns (7) to (9) include year fixed effects that are allowed to vary by firm age, where, again, quintiles are used to classify firms. Finally, Columns (10) to (12) include year fixed effects that vary by both the size and the age quintiles. Reassuringly, results are remarkably stable across specifications.

9 Causal Analysis: Worker-Level Design

The firm-level design presents three main limitations. First, it is by design constrained to the sample of firms observed in the balanced panel. As such, it is silent about the effects of the NMW on the extensive margin of firms. Second, the firm-level design is not well-suited to estimate impacts on employment – because potential reallocation effects a la [Dustmann et al. \(2022\)](#) violate the SUTVA assumption – and on structural residuals (that the model presented in [Section 7](#) suggests can be important) – because they mechanically aggregate to zero within each firm. Third, since the firm-level regression considers average wages, it is silent about within-firm wage spillovers. As such, the firm-level design can inform the causal effects on firm-specific behavior but can be misleading to understand the aggregate consequences of the reform on different equilibrium outcomes of interest. The worker-level design overcomes these limitations.

9.1 Empirical design

We implement [Dustmann et al. \(2022\)](#) design as follows. We define bins of real hourly wages and allocate worker-year combinations to these bins. We consider 11 bins indexed by k , with real hourly wage thresholds b_k so that a worker-year observation belongs to bin k if $b_{k-1} < w_{it} \leq b_k$, with w_{it} the real hourly wage. We set b_{11} equal to 48 Uruguayan pesos. As a reference, the real NMW in 2005 and 2006 was around 16.5 and 20 Uruguayan pesos, respectively, so the upper bound is almost 3 times the real

value in 2005.¹² The first two bins are below the NMW in 2005. The third bin is between the NMW in 2005 and the NMW in 2006. Bins 4 to 11 are all above the NMW in 2006.

Then, for $t \in \{2001, \dots, 2006\}$, we estimate the following DID regression:

$$y_{it} - y_{it-2} = \sum_k 1[b_{k-1} < w_{it-2} \leq b_k] \gamma_k + \lambda_t + \delta' X_{it-2} + \sum_k 1[b_{k-1} < w_{it-2} \leq b_k] 1[t \in \{2002, 2006\}] \beta_{kt} + \varepsilon_{it}, \quad (14)$$

where $y_{it} - y_{it-2}$ is the two-year growth of an outcome of interest, γ_k are bin fixed effects, λ_t are year fixed-effects, X_{it-2} is a vector of worker-level controls in $t-2$ that includes gender, age, tenure, firm size, and firm industry (5-digit codes), β_{kt} are the analogs of event-study coefficients that vary by bin, and ε_{it} is the error term. The standard errors are clustered at the (lagged) industry level.

We restrict the sample to workers that were employed in $t-2$. We consider four outcomes. First, we use log real hourly wages, so the dependent variable is the real wage growth. Second, we use an indicator of being employed, so the dependent variable measures whether the worker remains employed. Third, we use the time-varying F-FE estimated using the TV-AKM framework ψ_{jt} , so the dependent variable measures the change in the employer-specific wage premia. When using this outcome, we further split the sample between stayers and switchers to decompose how much of the estimated effect on the F-FE is due to reallocation effects a la [Dustmann et al. \(2022\)](#) and how much is due to a causal effect on F-FEs that benefit incumbents. Finally, we use the worker-specific residual of the TV-AKM model, which the model presented in Section 7 suggests can be affected by the reform when within-firm spillovers are limited. The employment status regression does not condition on the employment status at t . The other three outcomes restrict the estimation to workers that remain employed at t .

This design essentially mimics the intuition of a DID model where the unit of observation is the wage bin but allows the event indicators to vary by wage bin, a feature possible given the use of worker-level data. The cost of this decision is the preclusion from using individual fixed effects, which implies that the model's fixed effects do not account for within-bin heterogeneity across workers, an issue we address with the inclusion of a rich vector of controls that vary at the individual level. We focus on growth rates rather than levels because of the well-known differences in wage and employment dynamics at different parts of the wage distribution ([Ashenfelter and Card, 1982](#)). Specifically, mean reversion dynamics typically imply higher wage growth and weaker employment attachment at the bottom of the wage distribution regardless of the minimum wage policy, so not properly accounting for these differences may lead to spurious conclusions with respect to the effects of the NMW reform. Focusing on outcome changes attenuates these concerns as the treatment effect is estimated on changes in the growth rates.

We estimate equation (14) for $t \in \{2001, 2006\}$. Assessing differences in the outcomes' changes

¹²As a reference, [Dustmann et al. \(2022\)](#) consider bins up to 2.3 times the minimum wage value they consider.

between 2001-1999, 2002-2000, 2003-2001, and 2004-2002 works as a standard “pre-trends” diagnostic in event studies: the parallel trend assumption implies that, absent the NMW reform, there should be no differential growth rates in outcomes across bins relative to a baseline year. Then, the DID model tests for changes in growth rates by wage bin in 2005-2003 and 2006-2004 relative to a baseline year. We don’t use years beyond 2006 because the two-year lagged wages starting in 2007 are affected by the reform and, therefore, are no longer comparable to the baseline period.

To properly interpret the coefficients, note that our implementation of equation (14) makes two normalizations. First, as it is implicitly stated in the equation, we normalize $\beta_{k2001} = 0$ for all k , meaning that the effects on outcome changes are estimated relative to the bin-specific two-year growth rates in 2001 (up to a constant common to all bins captured by the year fixed effect). Second, we normalize $\beta_{11t} = 0$ for all t , so within a year, event study coefficients represent differences with respect to the highest bin. The intuition is that wages far up in the distribution are not “treated” by the reform, an assumption supported by related literature where wage spillovers of minimum wage reforms are found to be decreasing in the distance to the minimum wage (e.g., [Cengiz et al., 2019](#); [Engbom and Moser, 2022](#)).

9.2 Results

Figure 8 shows the results. We plot the estimated $\{\beta_{tk}\}_{t=2002,k=1}^{2006,11}$ coefficients with their corresponding 95% confidence intervals. Each plot represents a different outcome variable. Within each plot, the x-axis represents different wage bins, and the 5 sets of event study coefficients represent different years.

Panel (a) shows estimates using the two-year real hourly wage growth as the dependent variable. Relative to 2001, the relative bin-specific two-year real wage growth rates did not change in 2002, 2003, and 2004, providing support to the parallel trends assumption. In 2005 and 2006, however, the estimates show a significant change in the real hourly wage growth, especially at the bottom of the distribution. Three aspects of this result are worth discussing. First, the effect is sharp starting in 2005 and is larger for wage bins at the bottom, especially the ones below the NMW levels of 2005 and 2006, suggesting that the NMW reform played a causal role in this increase. Second, we estimate significant wage spillovers above the NMW levels, however, the spillovers are decreasing in the distance to the NMW and converge in bins with wage levels around 1.5 times the NMW, validating our assumption that bins high up the distribution are valid controls for bottom wage bins. Third, the effects are larger in 2006, consistent with the fact that the NMW exhibited two consecutive non-trivial increases.

Panel (b) shows that the wage effects were not accompanied by an impact on the relative probabilities of remaining employed across wage bins. Relative to 2001, relative bin-specific probabilities of remaining employed did not change in the years 2002, 2003, 2004, 2005, and 2006. This result suggests that, at the aggregate level, the NMW increase did not destroy jobs on top of the natural turnover rates specific to each wage level. This figure also shows that the possible effect of the pre-reform recession on employment

probabilities was evenly distributed across wage bins, thus being absorbed by the year fixed effect.

Panel (c) shows the change in the F-FE for job switchers. Consistent with [Dustmann et al. \(2022\)](#), we find evidence of reallocation effects: relative to 2001, workers at the bottom of the wage distribution who switch jobs are more likely to land a job in a firm with a higher wage premium after the NMW reform of 2005. This reallocation effect is concentrated at the very bottom of the distribution. This figure also shows that the pre-reform recession induced productive reallocation in the Uruguayan economy, consistent with “cleansing” narratives built on [Schumpeter \(1939\)](#). While reallocation effects were absent in 2002, they were positive in 2003 and 2004 (and of a similar order of magnitude of the post-reform ones), suggesting that the recession damaged jobs in low-paying firms that were absorbed by better-paying firms.

Panel (d) shows the change in the F-FE for job stayers. We find that the NMW reform of 2005 had a sizable effect on F-FEs that benefited wage growth for incumbent workers at the bottom of the distribution, but also with spillovers above the NMW level that are more pronounced than the ones exhibited in Panel (a). This result supports the conceptual framework developed in previous sections and is consistent with the main argument developed throughout this paper: the NMW reform affected inequality mainly through a causal effect on the F-FEs that benefits incumbent workers.

Finally, Panels (e) and (f) show the change in the TV-AKM worker-level residual for the sample of switchers and stayers. Low-wage workers who switch jobs, if anything, seem to experience a decrease in their wage residuals, suggesting they held rents in their former jobs. The pattern for job stayers is different and more precisely estimated. The model presented in Section 7 shows that if wage spillovers are limited, increases in wages after the NMW reform can be rationalized by a change in residuals with no change in the F-FEs. We find evidence of an increase in the residuals of stayers around the NMW levels, but the effect is modest and an order of magnitude smaller than the change in the F-FE.

To assess the quantitative implications of our results, we provide the following decomposition of the wage effects between reallocation, causal, and residual effects. For the bottom four wage bins in the years 2005 and 2006, the bin-specific estimated change in wage growth is decomposed as follows:

$$\Delta(\log w_{it} - \log w_{it-2}) = \underbrace{\omega^{sw} \Delta^{sw}(\psi_{jt} - \psi_{jt-2})}_{\text{Reallocation}} + \underbrace{\omega^{st} \Delta^{st}(\psi_{jt} - \psi_{jt-2})}_{\text{Causal}} + \underbrace{\Delta(R_{it} - R_{it-2})}_{\text{Residual}}, \quad (15)$$

where ω^{sw} is the share of switchers within the bin, $\omega^{st} = 1 - \omega^{sw}$ is the corresponding share of stayers, $\Delta^{sw}(\psi_{jt} - \psi_{jt-2})$ is the change in the F-FE growth for switchers (Panel (c)), $\Delta^{st}(\psi_{jt} - \psi_{jt-2})$ is the change in the F-FE growth for stayers (Panel (d)), and $\Delta(R_{it} - R_{it-2})$ is the overall change in residuals (weighted average between Panels (e) and (f)). Based on this decomposition, we compute the share of the change in wage growth $\Delta(\log w_{it} - \log w_{it-2})$ that can be attributed to each of the three components: reallocation, causal, and residual.¹³

¹³We note this coefficient identity does not hold exactly because coefficients on controls and fixed effects also vary by

Table 4 shows the results. For each bin, the causal effect of the NMW reform on the F-FEs of stayers is the most important driver of the increase in wage growth. This result is mainly explained by job switching being much less common relative to staying at the firm. At the bottom bin, the reallocation component accounts for 20% in 2005 and 10% in 2006, while the causal change for stayers is 84% and 90%. Changes in residuals are negligible in this bin. In bins 2, 3, and 4, reallocation effects become negligible while the residual component becomes relatively more important. The causal effect remains the most important component, amounting always for at least 75% of the change in wage growth.

10 Conclusion

This paper shows that minimum wage reforms can importantly affect the distribution of F-FEs, not only because of the reallocation effects on firms and employment but mainly because they causally affect the wage premia of low-paying firms. Under the interpretation that F-FEs are good proxies of job quality, the results are consistent with the minimum wage increasing the supply of good jobs by making “bad jobs better”. We show this empirically using administrative employer-employee data from Uruguay using a battery of descriptive analyses that build on the AKM and TV-AKM frameworks, and two causal designs using both firms and workers as units of observation.

On a conceptual level, our results suggest that the long tradition of wage compression estimates after minimum wage reforms can be rationalized by the ability of the minimum wage to limit the role of firms in affecting wage inequality. As a consequence, our results conclude that F-FEs can be affected by policy, even if they are rooted in firms’ fundamentals that may be persistent across time. These conclusions provide support to the conjecture that differences in labor market institutions rationalize the heterogeneous impact of firms on inequality, both across countries and within-countries across time.

Two fruitful avenues of future research naturally follow from our empirical analysis. First, our results have implications for the development of labor market models that rationalize wage dispersion across firms for similar workers since they indicate that the minimum wage plays an important role in the determination of firm wage premia. Second, other labor institutions such as collective bargaining, unions, or pay transparency could also have effects on the distribution of F-FEs. It would be interesting to test this hypothesis using suitable reforms on the aforementioned policies.

Bibliography

Abowd, John M, Francis Kramarz, and David N Margolis, “High wage workers and high wage firms,” *Econometrica*, 1999, 67 (2), 251–333.

regression. We, therefore, use the change in wage growth implied by the reallocation, causal, and residual estimates as the share denominator, which closely aligns with the estimates reported in Panel (a).


- Acemoglu, Daron**, “Good jobs versus bad jobs,” *Journal of Labor Economics*, 2001, 19 (1), 1–21.
- Andrews, Martyn J, Len Gill, Thorsten Schank, and Richard Upward**, “High wage workers and low wage firms: Negative assortative matching or limited mobility bias?,” *Journal of the Royal Statistical Society Series A: Statistics in Society*, 2008, 171 (3), 673–697.
- Ashenfelter, Orley and David Card**, “Time series representations of economic variables and alternative models of the labour market,” *The Review of Economic Studies*, 1982, 49 (5), 761–782.
- Autor, David, Alan Manning, and Christopher L Smith**, “The contribution of the minimum wage to US wage inequality over three decades: A reassessment,” *American Economic Journal: Applied Economics*, 2016, 8 (1), 58–99.
- Berger, David, Kyle Herkenhoff, and Simon Mongey**, “Minimum wages, efficiency and welfare,” *Econometrica*, 2025, 93 (1), 265–301.
- Bonhomme, Stéphane, Kerstin Holzheu, Thibaut Lamadon, Elena Manresa, Magne Mogstad, and Bradley Setzler**, “How much should we trust estimates of firm effects and worker sorting?,” *Journal of Labor Economics*, 2023, 41 (2), 291–322.
- Burdett, Kenneth and Dale T Mortensen**, “Wage differentials, employer size, and unemployment,” *International Economic Review*, 1998, pp. 257–273.
- Butschek, Sebastian**, “Raising the bar: Minimum wages and employers’ hiring standards,” *American Economic Journal: Economic Policy*, 2021.
- Card, David, Ana Rute Cardoso, Joerg Heining, and Patrick Kline**, “Firms and labor market inequality: Evidence and some theory,” *Journal of Labor Economics*, 2018, 36 (S1), S13–S70.
- **and Alan B Krueger**, “Minimum wages and employment: A case study of the fast-food industry in New Jersey and Pennsylvania,” *American Economic Review*, 1994, 84 (4), 772–793.
- **, Jörg Heining, and Patrick Kline**, “Workplace heterogeneity and the rise of West German wage inequality,” *The Quarterly Journal of Economics*, 2013, 128 (3), 967–1015.
- Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer**, “The effect of minimum wages on low-wage jobs,” *The Quarterly Journal of Economics*, 2019, 134 (3), 1405–1454.
- Clemens, Jeffrey, Lisa B Kahn, and Jonathan Meer**, “Dropouts need not apply? The minimum wage and skill upgrading,” *Journal of Labor Economics*, 2021, 39 (S1), S107–S149.
- Derenoncourt, Ellora and David Weil**, “Voluntary Minimum Wages: The Local Labor Market Effects of National Retailer Policies,” *Quarterly Journal of Economics*, Forthcoming.

- Draca, Mirko, Stephen Machin, and John Van Reenen**, “Minimum wages and firm profitability,” *American Economic Journal: Applied Economics*, 2011, 3 (1), 129–51.
- Dube, Arindrajit**, “Minimum wages and the distribution of family incomes,” *American Economic Journal: Applied Economics*, 2019, 11 (4), 268–304.
- **and Attila S Lindner**, “Minimum wages in the 21st century,” *Handbook of Labor Economics Vol. 5*, 2024.
- , **Laura Giuliano, and Jonathan Leonard**, “Fairness and frictions: The impact of unequal raises on quit behavior,” *American Economic Review*, 2019, 109 (2), 620–663.
- Dustmann, Christian, Attila Lindner, Uta Schönberg, Matthias Umkehrer, and Philipp Vom Berge**, “Reallocation effects of the minimum wage: Evidence from Germany,” *Quarterly Journal of Economics*, 2022, 137 (1), 267—328.
- Engbom, Niklas and Christian Moser**, “Earnings inequality and the minimum wage: Evidence from Brazil,” *American Economic Review*, 2022, 112 (12), 3803–3847.
- , – , **and Jan Sauermann**, “Firm pay dynamics,” *Journal of Econometrics*, 2023, 233 (2), 396–423.
- Fortin, Nicole M, Thomas Lemieux, and Neil Lloyd**, “Labor market institutions and the distribution of wages: The role of spillover effects,” *Journal of Labor Economics*, 2021, 39 (S2), S369—S412.
- Giupponi, Giulia and Stephen Machin**, “Company wage policy in a low-wage labor market,” *Working Paper*, 2024.
- , **Robert Joyce, Attila Lindner, Tom Waters, Thomas Wernham, and Xiaowei Xu**, “The Employment and Distributional Impacts of Nationwide Minimum Wage Changes,” *Journal of Labor Economics*, 2024, 42 (S1), S293–S333.
- Haanwinckel, Daniel**, “Supply, Demand, Institutions, and Firms: A Theory of Labor Market Sorting and the Wage Distribution,” *Working Paper*, 2024.
- Harasztosi, Péter and Attila Lindner**, “Who Pays for the minimum Wage?,” *The American Economic Review*, 2019, 109 (8), 2693–2727.
- Kline, Patrick**, “Firm Wage Effects,” *Handbook of Labor Economics*, 2024.
- , **Raffaele Saggio, and Mikkel Sølvsten**, “Leave-out estimation of variance components,” *Econometrica*, 2020, 88 (5), 1859–1898.
- Krueger, Alan and Lawrence Summers**, “Efficiency wages and the inter-industry wage structure,” *Econometrica*, 1988, pp. 259–293.

- Lachowska, Marta, Alexandre Mas, and Stephen A Woodbury**, “Sources of displaced workers’ long-term earnings losses,” *American Economic Review*, 2020, *110* (10), 3231–3266.
- , – , and **Stephen A. Woodbury**, “How reliable are administrative reports of paid work hours?,” *Labour Economics*, 2022, *75*, 102131.
- , – , **Raffaele Saggio, and Stephen A Woodbury**, “Do firm effects drift? Evidence from Washington administrative data,” *Journal of Econometrics*, 2023, *233* (2), 375–395.
- , – , – , and – , “Work hours mismatch,” *Working Paper*, 2024.
- Lamadon, Thibaut, Magne Mogstad, and Bradley Setzler**, “Imperfect competition, compensating differentials and rent sharing in the US labor market,” *American Economic Review*, 2022, *112* (1), 169–212.
- Lee, David**, “Wage inequality in the United States during the 1980s: Rising dispersion or falling minimum wage?,” *The Quarterly Journal of Economics*, 1999, *114* (3), 977–1023.
- Morchio, Iacopo and Christian Moser**, “The gender pay gap: Micro sources and macro consequences,” *Working Paper*, 2024.
- Rodrik, Dani and Stefanie Stantcheva**, “Fixing capitalism’s good jobs problem,” *Oxford Review of Economic Policy*, 2021, *37* (4), 824–837.
- Saez, Emmanuel, Joel Slemrod, and Seth H Giertz**, “The elasticity of taxable income with respect to marginal tax rates: A critical review,” *Journal of Economic Literature*, 2012, *50* (1), 3–50.
- Schumpeter, Joseph A**, “Business Cycles,” *McGraw-Hill*, 1939.
- Song, Jae, David J Price, Fatih Guvenen, Nicholas Bloom, and Till Von Wachter**, “Firming up inequality,” *The Quarterly Journal of Economics*, 2019, *134* (1), 1–50.
- Sorkin, Isaac**, “Ranking firms using revealed preference,” *The Quarterly Journal of Economics*, 2018, *133* (3), 1331–1393.


Figures and Tables

Figure 1: Evolution of the National Minimum Wage



(a) Log difference relative to 1997


(b) Annual growth rates




(c) Real minimum wage and Kaitz index

Notes: These figures plot the evolution of the national minimum wage between 1993 and 2022. In each figure, the black vertical line in 2005 accounts for the year of the minimum wage reform, and the red vertical lines in 1997 and 2013 account for the period for which we have access to the administrative records. CPI series comes from official reports of the Uruguayan National Institute of Statistics. In Panel (a), each dot represents the log difference relative to 1997. In Panel (b), each dot represents nominal annual growth rates. In Panel (c), the real minimum wage is normalized to be 1 in 1997, and the Kaitz index is computed as the hourly minimum wage divided by the median hourly wage computed using the administrative records, which is why the Kaitz index series is limited to the period 1997-2013. We compute the median hourly wage using our annual panel of formal workers before imposing restrictions specific to the AKM model (i.e., before dropping observations based on age, firm size, labor market attachment, and connected set).


Figure 2: Stylized Facts on Wage Inequality Around the Minimum Wage Reform



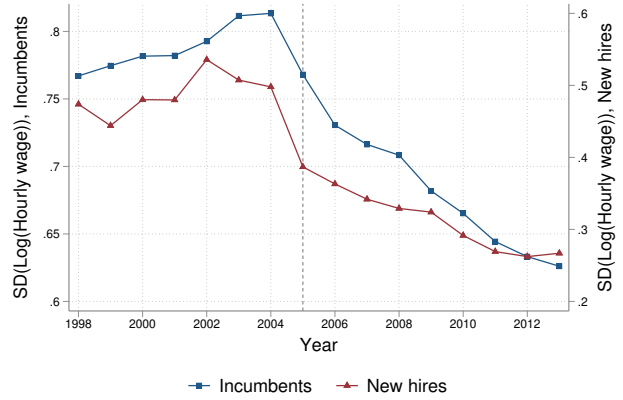
(a) Wage inequality and the real hourly minimum wage



(b) Wage growth by percentile



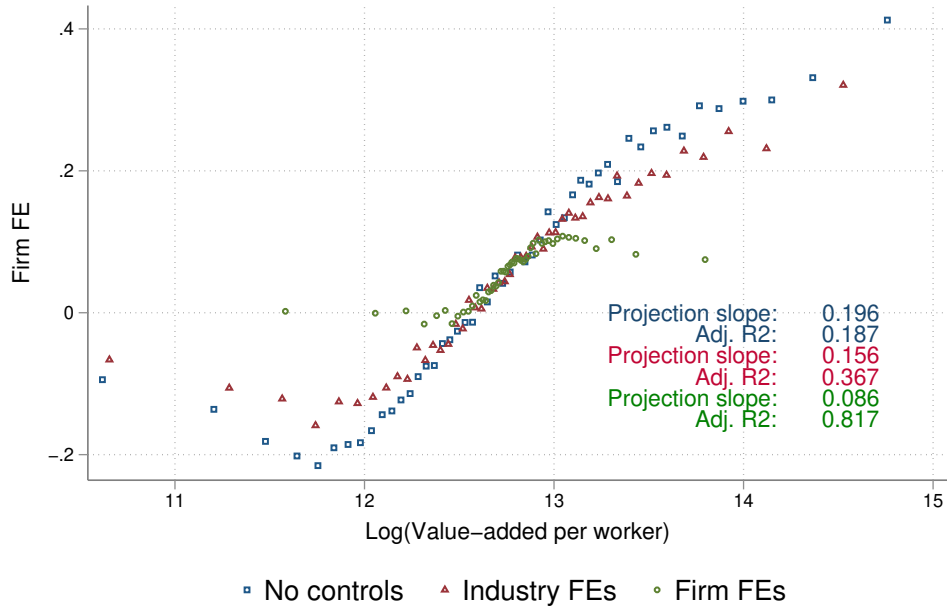
(c) Between- and within-firm inequality



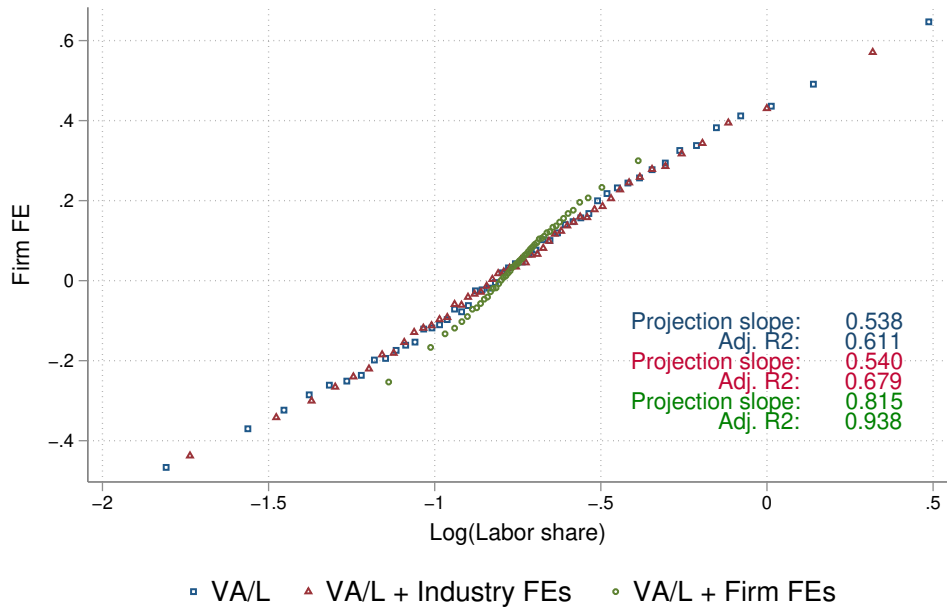
(d) Incumbents versus new hires

Notes: This figure presents stylized facts using the processed yearly panel described in Section 3 (Panel (c) of Table 1). In all figures, SD accounts for Standard Deviation. Panel (a) plots the standard deviation of log hourly wages and the real hourly national minimum wage. Panel (b) plots the evolution of different percentiles of the wage distribution relative to their levels attained in 1997. Panel (c) plots an employment-weighted bias-corrected variance decomposition of the total variance of log hourly wages into a between-firm component and a within-firm component. Panel (d) plots the standard deviation of log hourly wages separately for new hires and incumbent workers, where a new hire is defined as a worker who either appears for the first time in the sample (after 1997), the primary employer is a different firm relative to the previous year, or appears in the sample after disappearing for at least one year. By construction, the year 1997 is omitted in Panel (d) since the new hire definition relies on the information from the previous year.

Figure 3: Correlation Between TV-AKM F-FEs and Firm Attributes (2009-2013)




(a) Value added per worker (VA/L)




(b) Labor share (conditional on VA/L)

Notes: This figure plots binned scatter plots of the firm fixed effects estimated with the TV-AKM model against log value added per worker (Panel (a)) and log labor share conditional on value added per worker (Panel (b)). Value added is computed using corporate income tax records and is normalized in Panel (a) by the full-time equivalent workers in each firm, which is computed using the SSA records. Firm-specific labor shares are computed by dividing total gross wages (computed using the SSA data) by total value added. Details on the computation of these variables can be found in Appendix A.2. Given the limited availability of the tax data, these plots only consider the years 2009-2013 and focus on a balanced panel of firms observed in the 5 years. In Panel (a), the blue series includes no controls, the red series controls by industry indicators, and the green series controls by firm indicators. In Panel (b), the blue series controls by value added per worker, the red series controls by value added per worker and industry indicators, and the green series controls by value added per worker and firm indicators. Projection slopes and adjusted R-squares are recovered from OLS regressions. Panels winsorize at the percentiles 1 and 99 of the value added per worker and labor share distributions, respectively.


Figure 4: Within-Firm Short-Run Persistence in F-FEs: Selected Years



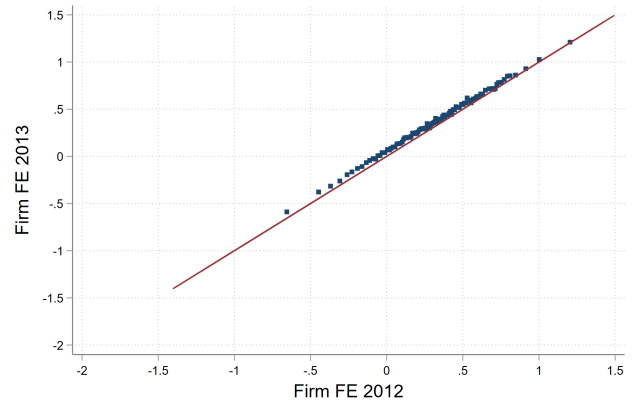
(a) 1997-1998



(b) 2003-2004



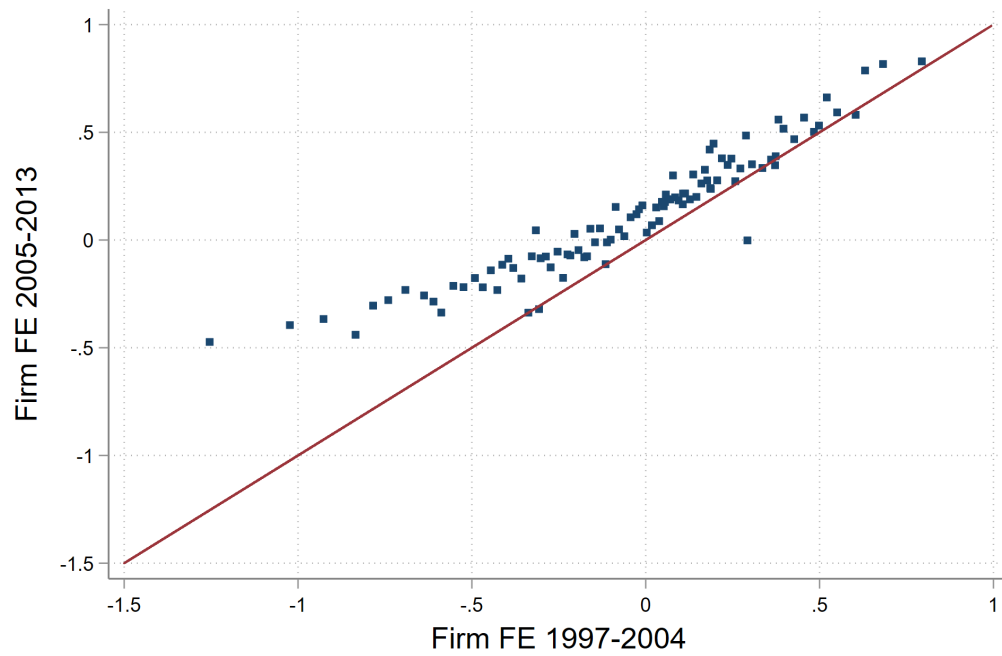
(c) 2004-2005



(d) 2012-2013


Notes: This figure presents binned scatter plots of the firm fixed effects (F-FEs) estimated with the TV-AKM model in year t against the F-FEs estimates in year $t - 1$. This figure presents results for four selected pairs of years: 1997-1998, 2003-2004, 2004-2005, and 2012-2013. Figure B.8 in Appendix B shows the results for all pairs of years. Observations are weighted by employment in year $t = 1$.

Figure 5: Long-Run Persistence in F-FEs: 1997-2004 vs 2005-2013




Notes: This figure presents a binned scatter plot of the average firm fixed effect (F-FEs) estimated with the TV-AKM model in years 2005-2013 against the average F-FEs estimates in years 1997-2004. This figure only considers firms in the balanced panel (i.e., observed in our data in all 17 years between 1997 and 2013). Observations are weighted by the average employment in the years 1997-2004.

Figure 6: Evolution of Average F-FE by Decile of F-FE in 2004



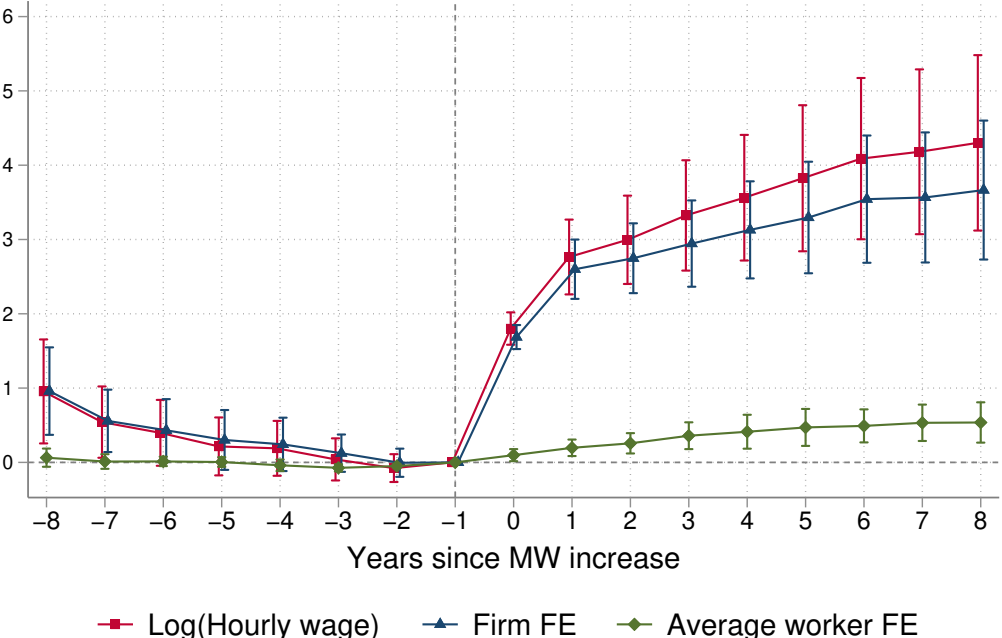
(a) Unbalanced panel



(b) Balanced panel


Notes: This figure plots the evolution of the firm fixed effect (F-FEs) estimated with the TV-AKM model by decile of F-FE in 2004. Deciles are determined based on the F-FE level in 2004 and are fixed for the rest of the periods. Panel (a) considers an unbalanced panel, meaning that firms are required to exist in 2004 but not in the other years. Panel (b) considers a balanced panel, meaning that firms are observed in our data in all 17 years between 1997 and 2013.

Figure 7: Firm-Level Design: Hourly Wages, Firm FE, and Average Worker FE




Notes: This figure plots the estimated event study coefficients β_t of the firm-level equation (12) with their corresponding 95% confidence intervals. The regression considers a balanced panel of firms observed in our data in all 17 years between 1997 and 2013. Each series corresponds to a different regression with a different dependent variable. Panel (a) uses the hours-weighted average log real hourly wage at the firm level, computed using the processed yearly panel described in Section 3 (Panel (c) of Table 1). The blue series uses the firm fixed effect (F-FEs) estimated with the TV-AKM model. The green series uses the hours-weighted firm-level average worker fixed effect estimated with the TV-AKM model. Standard errors are clustered at the 2-digit industry level.


Figure 8: Worker-Level Design: Hourly Wages, Employment, Firm FE, and Residual




(a) Hourly wage growth




(b) Employment probability



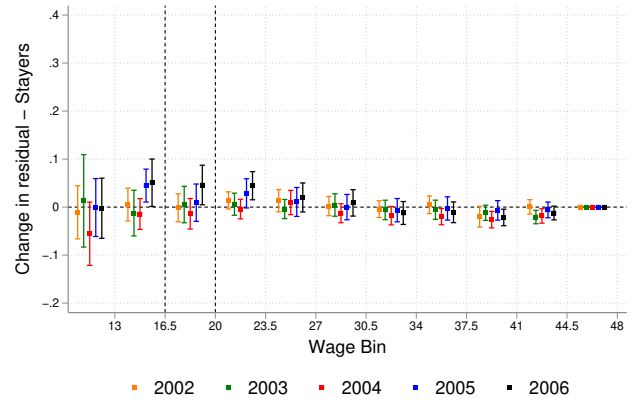
(c) Change in F-FE (switchers)



(d) Change in F-FE (stayers)



(e) Change in TV-AKM residual (switchers)



(f) Change in TV-AKM residual (stayers)

Notes: This figure plots the estimated event study coefficients β_{kt} of the worker-level equation (14) with their corresponding 95% confidence intervals using the processed yearly panel described in Section 3 (Panel (c) of Table 1). Within each panel, the horizontal axis accounts for different wage bins (k), and each series represents a different year (t). Each panel uses a different dependent variable. Panel (a) uses the 2-year hourly wage growth. Panel (b) uses the 2-year change in employment status. Panels (c) and (d) use the change in the firm fixed effect (F-FEs) estimated with the TV-AKM model for populations of job switchers and stayers, respectively. Panels (e) and (f) use the change in the TV-AKM residual for populations of job switchers and stayers, respectively. All regressions control for a vector of baseline characteristics, including gender, age, tenure, firm size, and firm industry (5-digit codes). Standard errors are clustered at the (lagged) industry level.

Table 1: Descriptive Statistics

	Observations	Workers	Firms	log Wage		Age		Male
				Mean	SD	Mean	SD	Mean
Panel (a) Yearly Panel								
1997-2001	2,696,801	850,906	65,604	4.17	0.87	36.06	13.12	0.61
2001-2005	2,544,589	846,076	67,625	3.97	0.87	37.07	12.93	0.59
2005-2009	3,371,364	1,060,687	82,805	4.13	0.76	36.57	12.80	0.60
2009-2013	4,129,817	1,246,366	91,389	4.46	0.73	36.56	12.70	0.58
All years	10,901,288	1,720,159	155,329	4.23	0.82	36.54	12.86	0.59
Panel (b) Yearly Panel + Restrictions								
1997-2001	1,704,472	544,473	19,050	4.34	0.80	36.22	10.49	0.60
2001-2005	1,627,119	537,998	19,446	4.12	0.82	36.79	10.47	0.58
2005-2009	2,094,946	683,081	23,013	4.25	0.71	36.55	10.54	0.59
2009-2013	2,597,104	821,254	25,766	4.56	0.64	36.61	10.57	0.58
All years	5,693,415	1,053,848	30,290	4.39	0.75	36.29	10.16	0.58
Panel (c) Yearly Panel + Restrictions + LCS								
1997-2001	1,652,323	528,189	16,880	4.36	0.80	36.16	10.49	0.60
2001-2005	1,569,502	518,296	16,799	4.13	0.82	36.72	10.47	0.58
2005-2009	2,062,616	671,234	21,370	4.25	0.71	36.49	10.54	0.59
2009-2013	2,572,405	812,510	24,458	4.56	0.64	36.57	10.56	0.58
All years	5,687,957	1,051,564	29,900	4.39	0.75	36.29	10.16	0.58

Notes: This table presents descriptive statistics of the yearly worker-level panel built from the raw monthly administrative records of the SSA. Panel (a) displays statistics for the baseline yearly panel, where each observation reports annual outcomes for the primary employer of the worker, defined as the employer that had larger earnings. Panel (b) displays statistics for the yearly panel after imposing restrictions on age, labor market attachment, outliers, and firm size. Panel (c) further restricts the sample to firms belonging to the largest connected set (LCS). SD accounts for Standard Deviation. See Section 3 and Appendix A.3 for additional details on the panel construction.

Table 2: Bias-Corrected Variance Decomposition: Log Hourly Wages

	1997-2001		2001-2005		2005-2009		2009-2013	
	Level	%	Level	%	Level	%	Level	%
Variance of Log Hourly Wages	0.628	100	0.666	100	0.509	100	0.408	100
Std. Dev.	0.792		0.816		0.714		0.639	
Panel (a) KSS (Bias corrected)								
$V(\alpha_i)$	0.227	36.1	0.236	35.4	0.200	39.3	0.179	43.8
Std. Dev.	0.476		0.485		0.447		0.423	
$V(\psi_j)$	0.189	30.0	0.196	29.4	0.114	22.3	0.073	17.8
Std. Dev.	0.434		0.443		0.337		0.270	
$2 \cdot C(\alpha_i, \psi_j)$	0.112	17.9	0.105	15.8	0.086	16.9	0.062	15.1
Share explained		84.0		80.6		78.5		76.7
Panel (b) Plug-in (Naive)								
$V(\alpha_i)$	0.285	45.4	0.288	43.3	0.246	48.4	0.214	52.5
Std. Dev.	0.534		0.537		0.496		0.463	
$V(\psi_j)$	0.224	35.7	0.226	34.0	0.136	26.6	0.086	21.1
Std. Dev.	0.473		0.476		0.368		0.294	
$2 \cdot C(\alpha_i, \psi_j)$	0.046	7.3	0.048	7.2	0.046	9.0	0.037	9.1
Share explained		88.4		84.4		84.0		82.7
N Movers	104,837		96,900		168,534		219,736	
N Firms	14,286		14,009		19,197		22,351	
Movers/Firms	7.34		6.92		8.78		9.83	

Notes: This table presents AKM variance decompositions of log hourly wages. The AKM model is estimated separately for 5-year overlapping intervals –1997-2001, 2001-2005, 2005-2009, and 2009-2013– using the processed yearly panel described in Section 3 (Panel (c) of Table 1). Panel (a) shows bias-corrected variance decompositions using the leave-one-out method of [Kline et al. \(2020\)](#). Panel (b) shows naive plug-in decompositions that are contaminated by limited mobility bias. Each panel shows results in terms of variances, standard deviations, and share explained. N Movers is the number of movers used to identify the firm fixed effects. N Firms is the number of firms. Movers/Firms is the ratio between the two.

Table 3: GAP Design: Results

	$\overline{\log(w)}$	ψ	$\bar{\alpha}$	$\overline{\log(w)}$	ψ	$\bar{\alpha}$	$\overline{\log(w)}$	ψ	$\bar{\alpha}$	$\overline{\log(w)}$	ψ	$\bar{\alpha}$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
$\hat{\beta}$	3.144 (0.368)	2.693 (0.296)	0.381 (0.096)	2.855 (0.302)	2.511 (0.247)	0.295 (0.087)	3.084 (0.356)	2.637 (0.290)	0.397 (0.092)	2.865 (0.298)	2.501 (0.246)	0.326 (0.084)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	No	No	No	No	No	No	No	No	No
Year x Size FE	No	No	No	Yes	Yes	Yes	No	No	No	No	No	No
Year x Age FE	No	No	No	No	No	No	Yes	Yes	Yes	No	No	No
Year x Size x Age FE	No	No	No	No	No	No	No	No	No	Yes	Yes	Yes
Observations	70,686	70,686	70,686	70,686	70,686	70,686	70,227	70,227	70,227	70,227	70,227	70,227

Notes: This table presents the estimated DID coefficient β of the firm-level equation (13). The regression considers a balanced panel of firms observed in our data in all 17 years between 1997 and 2013. Each column corresponds to a different regression with a different dependent variable. Columns (1), (4), (7), and (10) use the hours-weighted average log real hourly wage at the firm level, computed using the processed yearly panel described in Section 3 (Panel (c) of Table 1). Columns (2), (5), (8), and (11) use the firm fixed effect (F-FEs) estimated with the TV-AKM model. Columns (3), (6), (9), and (12) use the hours-weighted firm-level average worker fixed effect estimated with the TV-AKM model. Columns (1)-(3) include firm and year fixed effects. Columns (4)-(6) include firm and year-by-size fixed effects, where size corresponds to size quintiles based on the average size in the pre-reform period. Columns (7)-(9) include firm and year-by-age fixed effects, where age corresponds to age quintiles. Columns (10)-(12) include firm and year-by-size-by-age fixed effects, thus interacting the year indicator with all combinations of size-age quintiles. Standard errors (in parentheses) are clustered at the 2-digit industry level.

Table 4: Worker-Level Analysis: Wage Growth Effect Decomposition

	2005			2006		
	Reallocation	Causal	Residual	Reallocation	Causal	Residual
Wage bin 1	19.5%	83.9%	-3.4%	9.9%	90.3%	-0.2%
Wage bin 2	3.2%	77.6%	19.2%	1.4%	81.3%	17.3%
Wage bin 3	-3.5%	112.4%	-8.9%	-2.9%	87.8%	15.0%
Wage bin 4	-3.2%	78.8%	24.5%	-3.1%	75.1%	28.0%

Notes: This table presents the decomposition of the estimated effect on changes in the two-year real hourly wage by wage bin based on equation (15). Reallocation is the share of the wage effect explained by workers switching to firms with higher TV-AKM fixed effects. Causal is the share of the wage effect explained by the change in the TV-AKM firm fixed effect for job stayers. Residual is the share of the wage effect explained by the change in the TV-AKM residual for both switchers and stayers. The decomposition is done separately for each of the bottom four wage bins and for 2005 and 2006.

Minimum Wages and the Distribution of Firm Wage Premia

Online Appendix

A Data Appendix	1
A.1 Quality tests of the administrative records on work hours	1
A.2 Building balance sheet variables from corporate tax data	3
A.3 From raw data to annual panel	4
B Additional Figures and Tables	7
C CCHK details	17
Appendix Bibliography	19

A Data Appendix

A.1 Quality tests of the administrative records on work hours

We implement a series of tests proposed in [Lachowska et al. \(2022\)](#) –for administrative records in the State of Washington– to examine the quality of the administrative records of Uruguay’s SSA on hours. We present results on our estimation sample using two years as a reference period: 2003 (pre-NMW reform) and 2006 (post-NMW reform). As a reminder, our measure of annual hours is the total hours worked in a year. When employers report both weekly hours worked and total days worked within a month, monthly hours can be computed without any imputation. This case represents roughly 75% of our baseline yearly panel observations. When employers report weekly hours but days worked are missing, the worker is assumed to have worked the complete month in the firm. This case represents roughly 5% of our baseline yearly panel observations. Finally, when employers report days worked but weekly hours worked are missing, the worker is assumed to work the standard labor contract of 8 hours per day. This case represents roughly 20% of our baseline yearly panel observations. We present our diagnostics separately for the raw measure of hours (no imputation) and the final measure of hours.

First, we analyze the extent to which employers report that a worker had positive earnings but zero or missing hours in the same period. Panel (a) of Figure [A.1](#) shows the proportion of records that have zero (blue bars) or missing (red bars) working hours. In both years, about 20% of the raw SSA records have zero or missing hours worked despite showing positive values in the earnings records. In most cases, however, employers report missing hours worked but report the monthly days worked by the employee. A potential explanation for this pattern is that in Uruguay, workers employed in sectors such as industry

and services are subject to a legal regime of 8 hours per day; thus, employers may only report the number of days worked to the SSA. Panel (b) of Figure A.1 reproduces Panel (a) but using the final measure of hours worked variable with imputed values. After the imputation, less than 1% of the SSA records showed zero or missing hours worked.

Second, we compare the distribution of hours based on the SSA records with the distribution of working hours based on the Uruguayan national household survey (ECH) for the sample of wage and salary workers in the private sector. For example, if the administrative records display a more significant spike at 20 or 40 hours than the ECH, it may suggest that employers are unable to accurately track the actual hours worked by a substantial number of employees and instead report standard working hour packages (e.g., either part-time (20 hours) or full-time (40 hours)). Figure A.2 shows the distribution of weekly hours based on administrative records for the years 2003 and 2006, with and without the imputation procedure. Both distributions look similar, except for the 56-hour spike observed when applying the imputation procedure. This spike is due to records that report 7 days worked per week. In Figure A.3, we show the distribution of hours worked according to surveyors' responses to the ECH in the corresponding years. Similar to the US CPS, the ECH asks respondents to report "usual" weekly hours at their main job, whereas the administrative data report workers' weekly hours worked for each employer. In the administrative records, we observe some spread in hours reported, with some "heaping" at around 20, 30, 40, 44, 48, and 56 hours per week. In particular, about 15% of the employer reports cluster at the 44-hour-per-week spike for both analyzed years, which corresponds to a standard 8-hour workday from Monday to Friday plus a half-day of work on Saturday. Although the distribution based on the ECH responses also shows a "heaping" pattern at those round numbers, the gaps between spikes are much more prominent than those observed in the administrative hours distribution. For example, in the ECH, almost 36% of wage and salary private sector workers reported 40 hours per week in 2003, while only 8% in the administrative reports clustered at that same hour-per-week spike. These results imply that employers in Uruguay do not merely artificially report a standard working week. [Lachowska et al. \(2022\)](#) observed qualitatively similar findings when comparing CPS data with administrative records in the State of Washington.

Third, we regress the yearly changes in log earnings on the changes in log hours. If hours are measured with substantial error, we would expect the slope coefficient in the regression to be attenuated, while it would be close to one if hours are measured accurately. Panel (a) of Table A.1 presents the estimated coefficient of such regressions using observations of workers who remained with the same employer in successive years (i.e., "stayers"). For estimates in years 2003 and 2006, the underlying samples correspond to stayers observed in years 2002–2003 and 2005–2006, respectively. Regressions use the hours worked variable constructed using the imputation procedure when the records show missing values for daily hours. Following [Lachowska et al. \(2022\)](#), Columns (1) through (4) show slope coefficients from models with and

without employer fixed effects and with and without clustered standard errors. The results in Panel (a) of Table A.1 show estimated slope coefficients close to 1, which is consistent with what we expect from data with little measurement error in hours, given that most of our sample consists of salaried workers who presumably have proportional changes in earnings and hours from year to year. Panel (b) of Table A.1 presents estimates with the weekly hours variable constructed without the imputation procedure and shows similar results.

A.2 Building balance sheet variables from corporate tax data

The Corporate Income Tax Data (CIT) covers the universe of medium-sized and large private firms in Uruguay. The data contains information on balance sheets and income statements, allowing us to measure firms' income and cost structure. We have access to forms 2148 and 2149 for the period 2009-2017. The two forms are different but have the same cell numbering to be filled out. In the following, we list the definitions of our key variables and some details about their construction.

Value of Production. Corresponds to revenue plus operational grants. This is reported by firms in cell 100.

Total gross wages. Corresponds to total gross wages paid by the firm (before taxes and social security contributions). It comprises the sum of all wages, salaries, and bonuses paid to workers. The construction of this variable is based on the data from the SSA records before imposing our sample restrictions. We merge this variable with the CIT at the firm level using the procedure described in Section 3.

Production costs. Corresponds to all operating expenses (excluding labor costs). It includes net purchases, administrative costs, and other minor costs (e.g., services and third-party costs). Production and administrative costs are netted out of wage costs, as these cells in the CIT forms usually include employee compensation. This is reported by firms in cell 177.

Value added. Corresponds to the value a firm can create from inputs during production. It is computed as follows: value of production minus production costs.

Profits. Corresponds to the value of the fiscal year's profits before taxes. It is computed as the following: value of production plus other operating income (cell 184) minus all operating expenses (including production costs and other operating costs) plus the net financial income (cells 108 and 109).

Employment. To normalize value added, we build a full-time equivalent employment measure as follows. From the SSA records, we identify at the firm-by-month level the number of workers employed at the firm. Then we sum all employed workers each month, weighting by 0.5 the ones that report working less than 40 hours a week. Then, the firm-by-year full-time equivalent employment is the average across the twelve months of the monthly measure.


Value added per worker. Corresponds to value added divided by employment.

Labor Share. We build two measures of the firm-level labor share. The first is total gross wages over total value added. The second is total gross wages over total gross wages plus pre-tax profits.

A.3 From raw data to annual panel


We impose the following restrictions when transitioning from the raw monthly data to a unique worker-year observation. First, we drop observations of jobs in the public sector (on average, 28% of the number of initial observations each year). Second, we drop duplicate observations ($<0.5\%$) and observations for which monthly earnings or monthly hours are zero or missing (7%). Third, we only keep observations that pertain to job matches recorded as salaried workers or employees (that is, we drop observations associated with shareholders, internships, and explicitly short-termed temporary work arrangements, 12%). After this basic cleaning, workers may appear more than one time in a firm-month pair ($<0.5\%$). In those cases, we keep the observation with maximum monthly earnings, leaving us with a panel with unique observations at the worker-firm-month level. Then, we drop worker-year combinations that exhibit 10 or more different employers within a year, where the total number of employers is recorded before applying the previous restrictions ($<0.1\%$). From this panel, we identify the primary annual employer as the employer that yielded the maximum annual earnings and drop the observations related to secondary employers (4%). After collapsing the data at the worker-year level, we also drop the observations for which the share of total annual earnings pertaining to monthly taxable income is zero ($<0.5\%$) and the observations where the firm has a missing industry indicator (4%).

Figure A.1: Percentage of Observations with Zero or Missing Work Hours (SSA Records)




Notes: This figure plots the percentage of observations with positive earnings but zero or missing working hours. Panel (a) shows raw reported hours, while Panel (b) shows our final measure of hours with imputations based on days worked. The data comes from SSA administrative records for the years 2003 and 2006.


Figure A.2: Distribution of Weekly Hours Worked (SSA Records)




(a) Year 2003 (without imputations)



(b) Year 2006 (without imputations)




(c) Year 2003 (with imputations)




(d) Year 2006 (with imputations)

Notes: This figure plots the distribution of weekly working hours in the administrative data. Panels (a) and (b) show raw reported hours, while Panels (c) and (d) show our final measure of hours with imputations based on days worked. The data come from SSA administrative records for the years 2003 and 2006, respectively. Histograms are trimmed at 100 hours per week.

Figure A.3: Distribution of Weekly Hours Worked, (Household Survey Data, ECH)



(a) Year 2003



(b) Year 2006

Notes: This figure plots the distribution of usual weekly working hours reported in Uruguay's national household survey. The data come from ECH recorded by the National Institute of Statistics (*Instituto Nacional de Estadísticas*, INE) for the years 2003 and 2006, respectively. To make the household survey sample comparable to administrative records, we select workers in the ECH who are private-sector salaried employees aged 20-60. Histograms are trimmed at 100 hours per week.


Table A.1: Estimates of the Change in Log Earnings on the Change of Log Hours (SSA Registers)

Panel (a) With imputations				
	Change in Log Monthly Earnings			
	(1)	(2)	(3)	(4)
Change in Log Monthly Hours	1.055	1.055	1.126	1.126
	(0.008)	(0.008)	(0.008)	(0.008)
Constant	0.160	0.160	0.160	0.160
	(0.001)	(0.001)	(0.001)	(0.001)
Observations	397,656	397,656	381,029	381,029
R-sq	0.178	0.178	0.355	0.355
Employer fixed effects	No	No	Yes	Yes
Standard errors	Standard	Clustered by worker	Standard	Clustered by worker
Panel (b) Without imputations				
	Change in Log Monthly Earnings			
	(1)	(2)	(3)	(4)
Change in Log Monthly Hours	1.131	1.131	1.218	1.218
	(0.009)	(0.009)	(0.010)	(0.010)
Constant	0.161	0.161	0.161	0.161
	(0.001)	(0.001)	(0.001)	(0.001)
Observations	300,050	300,050	286,719	286,719
R-sq	0.196	0.196	0.380	0.380
Employer fixed effects	No	No	Yes	Yes
Standard errors	Standard	Clustered by worker	Standard	Clustered by worker

Notes: This table shows the estimated slope coefficient from a regression of the change in log monthly earnings on the change of log monthly hours. It includes annualized data based on years 2002, 2003, 2005, and 2005 from the SSA administrative records. Panel (a) shows our final measure of hours with imputations based on days worked, while Panel (b) shows raw reported hours.

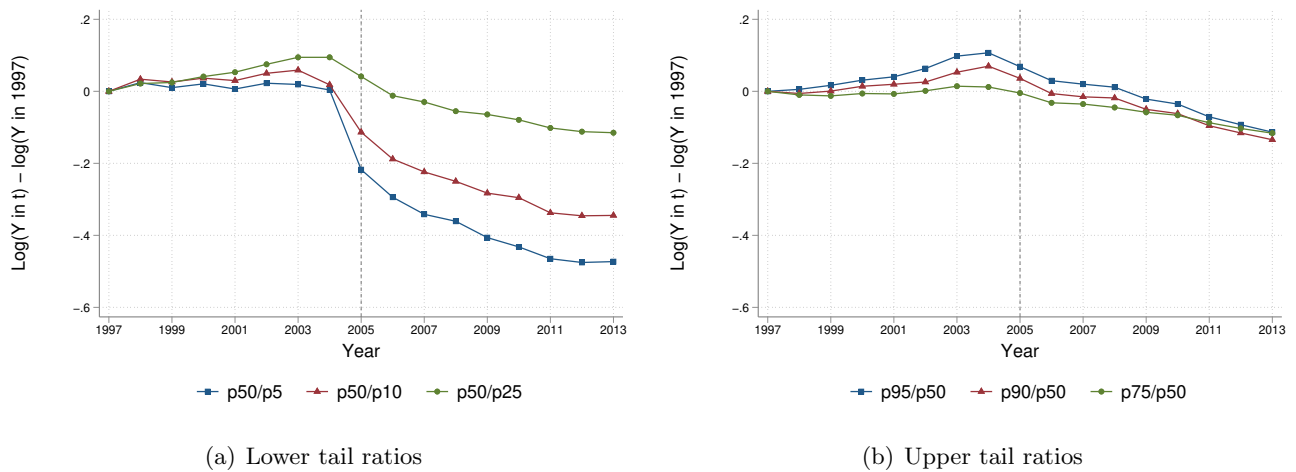
B Additional Figures and Tables

Figure B.1: Wage, Earnings, and Hours Inequality




Notes: This figure uses the processed yearly panel described in Section 3 (Panel (c) of Table 1) to extend Panel (a) of Figure 2. SD accounts for Standard Deviation. The blue series plots the standard deviation of the log monthly earnings. The red series plots the standard deviation of the log hourly wages. The green series plots the standard deviation of the log monthly hours.

Figure B.2: Evolution of Percentiles Ratios




Notes: This figure uses the processed yearly panel described in Section 3 (Panel (c) of Table 1) to extend Panel (b) of Figure 2. Panel (a) shows lower tail ratios of percentiles. The blue series plots the p_{50}/p_5 ratio. The red series plots the p_{50}/p_{10} ratio. The green series plots the p_{50}/p_{25} ratio. Panel (b) shows upper tail ratios of percentiles. The blue series plots the p_{95}/p_{50} ratio. The red series plots the p_{90}/p_{50} ratio. The green series plots the p_{75}/p_{50} ratio.

Figure B.3: Age-Wage Pattern




Notes: This figure plots the age-wage pattern observed in the processed yearly panel described in Section 3 (Panel (c) of Table 1). This figure is computed as follows. We first average log-hourly wages by age and cohort and then residualize the average log-hourly wages against cohort fixed effects. The residual is then plotted against age. The coefficients that emerge from a regression on the residual against age and age squared reveal that the curve flattens when age is equal to 52.


Figure B.4: AKM Event Studies for Job Switchers



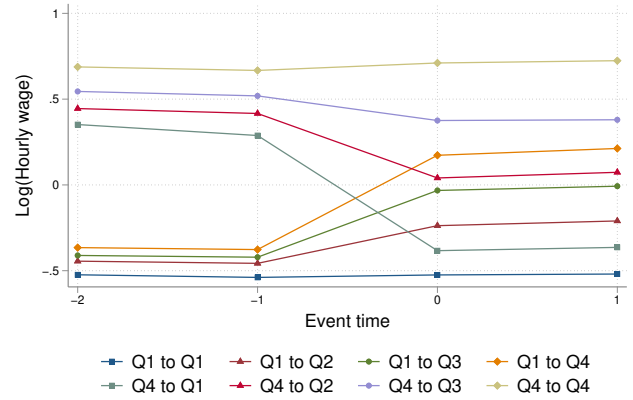
(a) 1997-2001



(b) 2001-2005




(c) 2005-2009



(d) 2009-2013


Notes: This figure plots “event studies” of log hourly wages of switchers around job transitions as in [Card et al. \(2013\)](#) and [Card et al. \(2018\)](#). We define a switcher within a 5-year period as a worker that stays for (at least) 2 consecutive years in a firm and switches to a different firm and stays for (at least) 2 consecutive years. We compute employment-weighted quartiles of firms based on the hourly wages of switchers’ coworkers and plot the residualized (against year fixed effects and age controls) log hourly wage of switchers that experience transitions between firms in different quartiles.

Figure B.5: Projections Between AKM F-FEs and TV-AKM F-FEs




Notes: This figure plots binned scatter plots of the firm fixed effects estimated with the TV-AKM model against against the firm fixed effects estimated with the AKM model (Panel (a)) and vice-versa (Panel (b)).

Figure B.6: Correlation Between TV-AKM F-FEs and Alternative Measure of Labor Share (2009-2013)



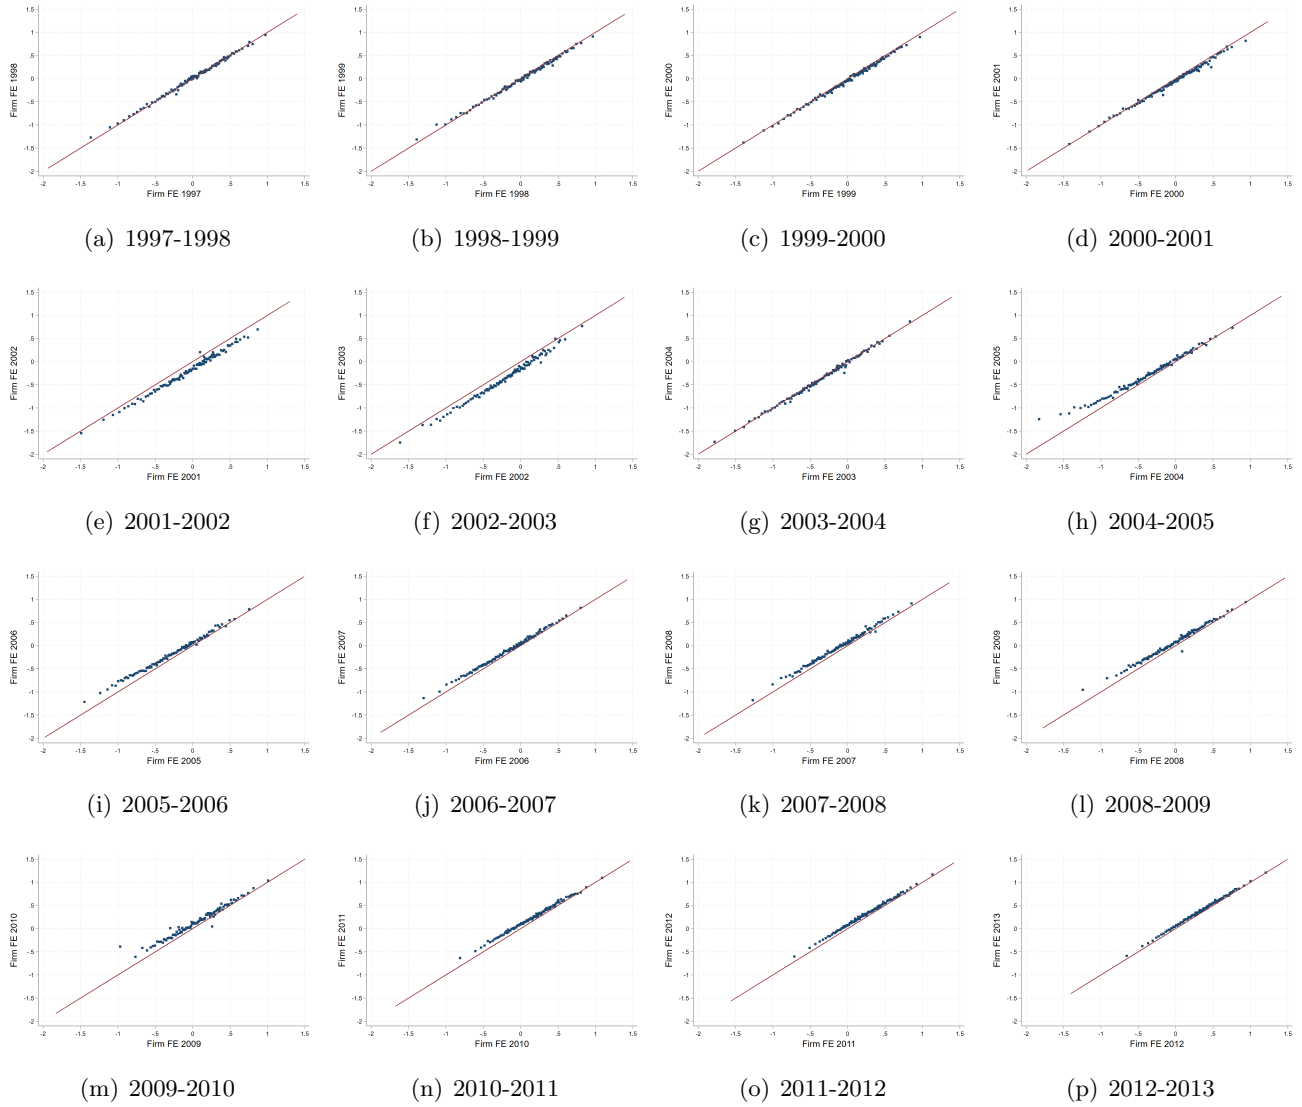
Notes: This figure plots a binned scatter plot of the firm fixed effects estimated with the TV-AKM model against log labor share conditional on value added per worker. Value added is computed using corporate income tax records and is normalized by the full-time equivalent workers in each firm, which is computed using the SSA records. Firm-specific labor shares are computed by dividing total gross wages (computed using the SSA data) by total gross wages plus pre-tax profits. Details on the computation of these variables can be found in Appendix A.2. Given the limited availability of the tax data, these plots only consider the years 2009-2013 and focus on a balanced panel of firms observed in the 5 years. The blue series controls by value added per worker, the red series controls by value added per worker and industry indicators, and the green series controls by value added per worker and firm indicators. Projection slopes and adjusted R-squares are recovered from OLS regressions. Panels winsorize at the percentiles 1 and 99 of the labor share distributions, respectively.

Figure B.7: Standard Deviation of TV-AKM F-FEs




Notes: This figure plots the year-by-year standard deviation of the firm fixed effects estimated with the TV-AKM model. All four series are normalized, so the standard deviation in 2004 is equal to one. The blue series plots the employment-weighted annual standard deviation considering the full (unbalanced) sample of firms. The red series plots the employment-weighted annual standard deviation considering the balanced sample of firms that are observed in each of the 17 years. The green series plots the unweighted annual standard deviation considering the full (unbalanced) sample of firms. The yellow series plots the unweighted annual standard deviation considering the balanced sample of firms that are observed in each of the 17 years.

Figure B.8: Within-Firm Short-Run Persistence in F-FEs: All Years




Notes: This figure presents binned scatter plots of the firm fixed effects (F-FEs) estimated with the TV-AKM model in year t against the F-FEs estimates in year $t - 1$ for all pairs of consecutive years between 1997 and 2013. Observations are weighted by employment in year $t = 1$.


Figure B.9: Firm-Level Design: Different Samples



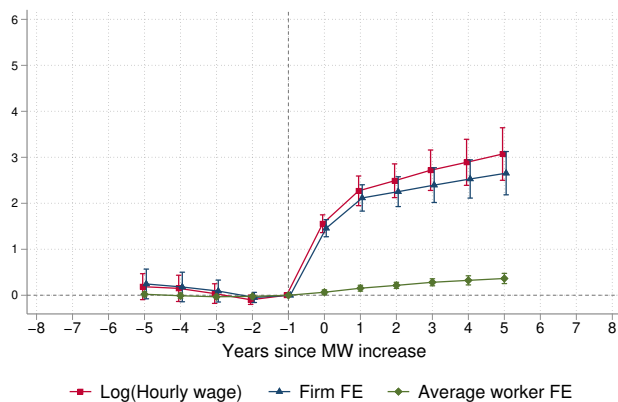
(a) Balanced panel 1997-2013



(b) Balanced panel 1998-2012



(c) Balanced panel 1999-2011



(d) Balanced panel 2000-2010

Notes: This figure plots the estimated event study coefficients β_t of the firm-level equation (12) with their corresponding 95% confidence intervals, for different balanced panels. Panel (a) considers a balanced panel of firms observed in our data in all 17 years between 1997 and 2013. Panel (b) considers a balanced panel of firms observed in our data in all 15 years between 1998 and 2012. Panel (c) considers a balanced panel of firms observed in our data in all 13 years between 1999 and 2011. Panel (d) considers a balanced panel of firms observed in our data in all 11 years between 2000 and 2010. Within panel, each series corresponds to a different regression with a different dependent variable. Panel (a) uses the hours-weighted average log real hourly wage at the firm level, computed using the processed yearly panel described in Section 3 (Panel (c) of Table 1). The blue series uses the firm fixed effect (F-FEs) estimated with the TV-AKM model. The green series uses the hours-weighted firm-level average worker fixed effect estimated with the TV-AKM model. Standard errors are clustered at the 2-digit industry level.

Table B.1: Comparison of the Largest Connected Set and the Leave-One-Out Largest Connected Set

	Observations	Workers	Firms	log Wage		Age		Male
				Mean	SD	Mean	SD	Mean
Panel (a) Yearly Panel + Restrictions + LCS								
1997-2001	1,652,323	528,189	16,880	4.36	0.80	36.16	10.49	0.60
2001-2005	1,569,502	518,296	16,799	4.13	0.82	36.72	10.47	0.58
2005-2009	2,062,616	671,234	21,370	4.25	0.71	36.49	10.54	0.59
2009-2013	2,572,405	812,510	24,458	4.56	0.64	36.57	10.56	0.58
Panel (b) Yearly Panel + Restrictions + Leave-one-out LCS								
1997-2001	1,460,818	384,407	14,286	4.42	0.79	36.43	10.39	0.59
2001-2005	1,365,097	361,526	14,009	4.19	0.82	37.07	10.36	0.58
2005-2009	1,859,788	498,579	19,197	4.29	0.71	36.76	10.43	0.59
2009-2013	2,344,894	614,978	22,351	4.59	0.64	36.79	10.45	0.58

Notes: This table presents descriptive statistics of the yearly worker-level panel built from the raw monthly administrative records of the SSA. Panel (a) displays statistics for the baseline yearly panel, where each observation reports annual outcomes for the primary employer of the worker, defined as the employer that had larger earnings, after imposing restrictions on age, labor market attachment, outliers, firm size, and further restricting the sample to firms belonging to the largest connected set (LCS). Panel (b) furthermore restricts the sample to firms belonging to the leave-one-out connected set as defined in [Kline et al. \(2020\)](#). SD accounts for Standard Deviation. See Section 3 and Appendix A.3 for additional details on the panel construction.

Table B.2: Bias-Corrected Variance Decomposition: Monthly Earnings

	1997-2001		2001-2005		2005-2009		2009-2013	
	Level	%	Level	%	Level	%	Level	%
Variance of Log Monthly Earnings	0.636	100	0.688	100	0.546	100	0.458	100
Std. Dev.	0.798		0.830		0.739		0.677	
Panel (a) KSS (Bias corrected)								
$\mathbb{V}(\alpha_i)$	0.211	33.2	0.234	33.9	0.190	34.8	0.184	40.2
Std. Dev.	0.459		0.483		0.436		0.429	
$\mathbb{V}(\psi_j)$	0.142	22.4	0.179	26.0	0.096	17.6	0.062	13.6
Std. Dev.	0.377		0.423		0.310		0.250	
$2 \cdot \mathbb{C}(\alpha_i, \psi_j)$	0.177	27.8	0.146	21.2	0.139	25.5	0.103	22.4
Share explained		83.3		81.1		77.8		76.2
Panel (b) Plug-in (Naive)								
$\mathbb{V}(\alpha_i)$	0.296	46.6	0.298	43.2	0.262	48.0	0.242	52.9
Std. Dev.	0.544		0.545		0.512		0.492	
$\mathbb{V}(\psi_j)$	0.199	31.3	0.218	31.7	0.134	24.5	0.087	19.1
Std. Dev.	0.446		0.467		0.366		0.295	
$2 \cdot \mathbb{C}(\alpha_i, \psi_j)$	0.069	10.8	0.071	10.4	0.069	12.6	0.057	12.5
Share explained		88.7		85.3		85.1		84.5
N Movers	104,837		96,900		168,534		219,736	
N Firms	14,286		14,009		19,197		22,351	
Movers/Firms	7.34		6.92		8.78		9.83	

Notes: This table presents AKM variance decompositions of log monthly earnings. The AKM model is estimated separately for 5-year overlapping intervals –1997-2001, 2001-2005, 2005-2009, and 2009-2013– using the processed yearly panel described in Section 3 (Panel (c) of Table 1). Panel (a) shows bias-corrected variance decompositions using the leave-one-out method of [Kline et al. \(2020\)](#). Panel (b) shows naive plug-in decompositions that are contaminated by limited mobility bias. Each panel shows results in terms of variances, standard deviations, and share explained. N Movers is the number of movers used to identify the firm fixed effects. N Firms is the number of firms. Movers/Firms is the ratio between the two.

Table B.3: Bias-Corrected Variance Decomposition: Monthly Hours

	1997-2001		2001-2005		2005-2009		2009-2013	
	Level	%	Level	%	Level	%	Level	%
Variance of Log Monthly Hours	0.087	100	0.089	100	0.083	100	0.081	100
Std. Dev.	0.294		0.298		0.288		0.285	
Panel (a) KSS (Bias corrected)								
$\mathbb{V}(\alpha_i)$	0.027	31.7	0.031	34.9	0.025	29.9	0.026	31.7
Std. Dev.	0.166		0.176		0.158		0.161	
$\mathbb{V}(\psi_j)$	0.030	35.0	0.032	36.0	0.026	31.2	0.026	31.7
Std. Dev.	0.174		0.179		0.161		0.161	
$2 \cdot \mathbb{C}(\alpha_i, \psi_j)$	0.006	7.0	0.002	1.9	0.005	6.0	0.005	5.6
Share explained		73.6		72.8		67.1		69.0
Panel (b) Plug-in (Naive)								
$\mathbb{V}(\alpha_i)$	0.049	56.8	0.052	59.0	0.046	55.5	0.044	53.7
Std. Dev.	0.222		0.229		0.215		0.209	
$\mathbb{V}(\psi_j)$	0.044	50.9	0.045	50.7	0.037	44.7	0.034	41.6
Std. Dev.	0.210		0.212		0.193		0.184	
$2 \cdot \mathbb{C}(\alpha_i, \psi_j)$	-0.020	-22.8	-0.023	-25.5	-0.015	-18.6	-0.010	-12.0
Share explained		84.8		84.2		81.5		83.3
N Movers	104,837		96,900		168,534		219,736	
N Firms	14,286		14,009		19,197		22,351	
Movers/Firms	7.34		6.92		8.78		9.83	

Notes: This table presents AKM variance decompositions of log monthly hours. The AKM model is estimated separately for 5-year overlapping intervals –1997-2001, 2001-2005, 2005-2009, and 2009-2013– using the processed yearly panel described in Section 3 (Panel (c) of Table 1). Panel (a) shows bias-corrected variance decompositions using the leave-one-out method of [Kline et al. \(2020\)](#). Panel (b) shows naive plug-in decompositions that are contaminated by limited mobility bias. Each panel shows results in terms of variances, standard deviations, and share explained. N Movers is the number of movers used to identify the firm fixed effects. N Firms is the number of firms. Movers/Firms is the ratio between the two.

C CCHK details

This appendix reproduces the main derivations in [Card et al. \(2018\)](#).

Setup. Consider J firms indexed by $j \in \{1, \dots, J\}$. There are two types of workers heterogeneous in skill, $s \in \{l, h\}$, indexed by i . The total number of low- and high-skill workers is given by $\{\mathcal{L}, \mathcal{H}\}$. Firms post wages (w_{lj}, w_{hj}) , which are observed by workers.

Workers' problem. Workers have preferences for firms, with ϵ_{isj} denoting the preference shock of worker i of skill s for firm j . When employed in firm j , worker i of skill s gets indirect utility:

$$u_{isj} = \eta_s \log(w_{sj} - b_s) + \nu_{sj} + \epsilon_{isj}, \quad (\text{C.1})$$

where η_s mediates the skill-specific labor supply elasticity, b_s is an outside option equal to all workers within skill s , and ν_{sj} is a firm-specific amenity common to all workers of skill s .

Assuming that the $\{\epsilon_{isj}\}$ are independent draws from a type I extreme value distribution, the firm choice probabilities are given by:

$$p_{sj} = \Pr\left(\arg \max_k \{u_{isk}\} = j\right) = \frac{\exp(\eta_s \log(w_{sj} - b_s) + \nu_{sj})}{\sum_k \exp(\eta_s \log(w_{sk} - b_s) + \nu_{sk})}. \quad (\text{C.2})$$

Assume away strategic interactions in wage setting between firms and J large. In that case:

$$p_{sj} \approx \lambda_s \exp(\eta_s \log(w_{sj} - b_s) + \nu_{sj}), \quad (\text{C.3})$$

where λ_s is a constant common to all firms. Then, the firm-specific labor supply elasticities are given by:

$$\log L_j(w_{lj}) = \log(\mathcal{L}\lambda_l) + \eta_l \log(w_{lj} - b_l) + \nu_{lj}, \quad (\text{C.4})$$

$$\log H_j(w_{hj}) = \log(\mathcal{H}\lambda_h) + \eta_h \log(w_{hj} - b_h) + \nu_{hj}. \quad (\text{C.5})$$

Firms' problem. Firms have technology $Y_j = R_j f(L_j, H_j)$, with R_j a firm-specific productivity shifter and (L_j, H_j) the number of low- and high-skill workers hired by firm j . We assume $Y_j = R_j (\theta_l L_j + \theta_h H_j)$. The firms' problem is to set (w_{lj}, w_{hj}) to minimize costs internalizing the labor supply elasticities:

$$\min_{w_{lj}, w_{hj}} w_{lj} L_j(w_{lj}) + w_{hj} H_j(w_{hj}) \quad s.t. \quad R_j (\theta_l L_j + \theta_h H_j) \geq Y. \quad (\text{C.6})$$

Firms don't observe $\{\epsilon_{isj}\}$, so they cannot perfectly price-discriminate. The FOCs are given by:

$$L_j + w_{lj} \frac{\partial L_j}{\partial w_{lj}} - \mu_j R_j \theta_l \frac{\partial L_j}{\partial w_{lj}} = 0, \quad (\text{C.7})$$

$$H_j + w_{hj} \frac{\partial H_j}{\partial w_{hj}} - \mu_j R_j \theta_h \frac{\partial H_j}{\partial w_{hj}} = 0, \quad (\text{C.8})$$

where μ_j is the constraint multiplier (marginal cost of production). Grouping terms yields:

$$w_{lj} \left(\frac{1 + e_{lj}}{e_{lj}} \right) = \mu_j R_j \theta_l, \quad (\text{C.9})$$

$$w_{hj} \left(\frac{1 + e_{hj}}{e_{hj}} \right) = \mu_j R_j \theta_h, \quad (\text{C.10})$$

where e_{sj} are the labor supply elasticities. From equations (C.4) and (C.5), these elasticities equal:

$$e_{Lj} = \frac{\eta_l w_{lj}}{w_{lj} - b_l}, \quad (\text{C.11})$$

$$e_{Hj} = \frac{\eta_h w_{hj}}{w_{hj} - b_h}, \quad (\text{C.12})$$

which together with the FOCs imply that:

$$w_{lj} = \frac{b_l}{1 + \eta_l} + \eta_l \frac{\mu_j R_j \theta_l}{1 + \eta_l}, \quad (\text{C.13})$$

$$w_{hj} = \frac{b_h}{1 + \eta_h} + \eta_h \frac{\mu_j R_j \theta_h}{1 + \eta_h}. \quad (\text{C.14})$$

When $b_s = 0$, the labor supply elasticity is equal to η_s and the equilibrium wage is equal to a fraction of the value of the marginal product (wage markdown), which is a function of the labor supply elasticity.

Assume that the reference wages, b_s , are proportional to the relative productivities, so $b_l = \theta_l b$ and $b_h = \theta_h b$. Let $T_j = \mu_j R_j / b$ be the relative productivity of firm j relative to the “outside sector”. Then

$$\log w_{lj} = \log \frac{\theta_l b}{1 + \eta_l} + \log(1 + \eta_l T_j), \quad (\text{C.15})$$

$$\log w_{hj} = \log \frac{\theta_h b}{1 + \eta_h} + \log(1 + \eta_h T_j). \quad (\text{C.16})$$

When $\eta_l \approx \eta_h$, then $\log w_{sj} \approx \alpha_s + \psi_j$, which coincides with an AKM decomposition. Normalizing $b = \eta_l = \eta_h = \mu_j = 1$ and ignoring the constant $\log(0.5)$ yields equation (5).

Appendix Bibliography

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

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

Kline, Patrick, Raffaele Saggio, and Mikkel Sølvsten, “Leave-out estimation of variance components,” *Econometrica*, 2020, *88* (5), 1859–1898.

Lachowska, Marta, Alexandre Mas, and Stephen A. Woodbury, “How reliable are administrative reports of paid work hours?,” *Labour Economics*, 2022, *75*, 102131.