Labor Supply and Firm Size*

Lin Shao†

Bank of Canada

Faisal Sohail‡

U. Melbourne

Emircan Yurdagul§

U. Carlos III and CEPR

June 2025

Abstract

Larger firms feature i) longer hours worked, ii) higher wages, and iii) smaller (larger) wage penalties for working long (short) hours. We reconcile these patterns in a general equilibrium model, which features the endogenous interaction of hours, wages, and firm size. In the model, workers willing to work longer hours sort into larger firms that offer a wage premium. Complementarities in hours generate wage penalties that increase with the distance from the usual hours. We use the model to argue that variation in average hours across firms contributes significantly to wage inequality.

JEL Codes: E24, J2, J31

Keywords: labor supply, hours, firm size, sorting, inequality, complementarities

*The views expressed in this paper are solely those of the authors and may differ from the official Bank of Canada views. We thank Dirk Krueger (editor) and three anonymous referees for insightful comments and suggestions. We also thank Alexander Bick, Andres Erosa, Belen Jerez, Minjoon Lee, Benjamin Schoefer, Lawrence Uren, Felix Wellschmied, and seminar participants at CEA, Bank of Canada, Singapore Management U. Macro Workshop, Trinity College Dublin, U. Alicante, U. Autonoma de Barcelona, U. Catolica Uruguay, U. Helsinki, U. Houston, U. Melbourne, and VAMS for valuable feedback. Yurdagul gratefully acknowledges the financial support by grants IJC2018-038229-I, PID2019-110779GB-100, RYC2022-035751-I, and CEX2021-001181-M funded by MICIU/AEI /10.13039/501100011033.

†lin.j.shao@gmail.com. Bank of Canada, 234 Wellington Street, Ottawa, ON K1A 0G9, Canada.

‡faisal.sohail@unimelb.edu.au. 111 Barry St, Carlton VIC, 3053, Australia

§emircan.yurdagul@uc3m.es. Calle Madrid 126, Getafe-Spain 28903.

1 Introduction

There exists significant variation in workers' labor supply. This variation, and its interaction with firm-level heterogeneity, can have important implications for inequality. For example, if workers in high-paying firms also work longer, variation in hours amplifies income inequality, while inequality is mitigated if workers in high-paying firms work fewer hours. Thus, understanding how hours vary across firms of different characteristics is crucial for assessing the role of worker and firm heterogeneity on inequality.

In this paper, we study the relationship between firm size, hours and wages. We find that workers in larger firms, which tend to pay higher wages, tend to work longer hours. Moreover, we show that the relationship between hours and wages differs systematically by firm size. Motivated by these findings, we develop a general equilibrium framework with heterogeneous workers and firms to show that incorporating the joint relationship between hours, wages, and firm size has important aggregate implications for wage inequality that are absent in canonical models that abstract from this relationship.

The first part of the paper uses data from the US Current Population Survey (CPS) to document three motivating facts on the relationship between hours and wages by firm size. To begin with, we present a relatively understudied empirical pattern: average hours worked increases with firm size. Secondly, we revisit the well-established size-wage premium, wherein average hourly wages are found to increase with the size of the firm (Brown and Medoff 1989, Oi and Idson 1999). The third fact we present is novel. Our data reveals that the wage penalties associated with relatively long and short working hours vary by firm size. Specifically, we find that workers at larger firms experience smaller wage penalties for working long hours but are subject to greater wage penalties for working shorter hours. Collectively, these three facts build upon the well-documented size-wage premium, offering a more comprehensive perspective of how wages, hours, and their relationship vary across firms of different sizes.

Motivated by our empirical findings, the second part of the paper introduces a theoretical framework to examine the interplay of hours, wages, and firm heterogeneity. In the model, firms differ in their exogenous productivity and decide on their labor input.[1] Aligning with recent findings in Shao et al. (2023) and Kuhn et al. (2023) indicating that working hours are complements and coordinated in production, our model's production function allows for complementarities between workers' hours. Specifically, a firm's labor input is a non-linear aggregate of the hours worked by all its employees, implying workers are more productive if their hours are similar to those of their co-workers. In addition, workers differ in their value of leisure and have additional preferences for working in firms of different productivity levels. Given these sources of heterogeneity, workers decide their labor supply and which firm to work for.

Despite its minimal structure, our model, which is calibrated to match key features of the US economy, successfully replicates all three motivating facts. First, the size-wage premium is generated by an interaction of firm-level heterogeneity in productivity and workers' preferences over working in firms of different characteristics. As in Card et al. (2018), with such idiosyncratic tastes, workers view firms of different productivity as imperfect substitutes. This results in less pronounced increase in firm employment with productivity, allowing the marginal productivity of labor (wages) to rise with firm productivity, and hence size, in equilibrium. We introduce heterogeneity in tastes as a simple and relatively standard feature, aiming to generate a size-wage premium that interacts with other mechanisms in the model to generate our other two motivating facts. Importantly, this modeling choice for generating the size-wage premium does not, on its own, generate increasing hours with firm size, nor does it cause the size-dependent wage penalties for short and long hours. Indeed, as we discuss below, differences in firm productivity and endogenous sorting between workers and firms remain critical for reconciling the data.

Second, the positive correlation between firm size and worker hours can be attributed to the interaction between the size-wage premium and workers' decisions about labor supply. Since income is the product of hourly wage and working hours, a size-wage premium suggests that those who work longer hours will see larger income gains when employed by larger firms. Furthermore, the higher wages offered by large firms can enhance income returns

[1] Several papers document a robust positive correlation between measures of firm productivity and the number of employees in a firm (see, for example, Leung et al., 2008 and Bartelsman et al., 2013). Consistent with this, we use the number of employees in a firm to proxy for productivity.

from working longer hours. In line with this mechanism, workers who prefer longer hours sort into larger firms, and for any given type of worker, they tend to work longer hours in these larger firms. Therefore, average working hours increase with firm size.

Third, differences in short- and long-hour wage penalties across firms result from endogenous worker sorting decisions under the presence of complementarities in working hours. Consistent with the data, complementarities in workers' hours ensure that worker productivity declines as the gap between their hours and a firm's usual working hours widens, and consequently, they suffer more significant wage losses as they deviate from usual hours worked. Since larger firms feature longer average hours worked, long (short) hours are less (more) heavily penalized compared to small firms. Importantly, even without complementarities between workers' hours, the size-wage premium will lead to a positive relationship between hours and firm size. Instead, we introduce complementarities in hours as it is empirically supported, and endogenously converts the size-hour patterns into the size-dependent short- and long-hour penalties.

Our baseline model is kept intentionally simple to illustrate the key mechanisms reconciling our main motivating facts. We also present an extended model that includes factors such as worker efficiency, risk aversion, wealth accumulation, and friction in switching employers. This extended model allows us to show the robustness of our findings in a richer environment and sets the stage for two important exercises. First, we use this model to discuss a testable implication about workers sorting into different-sized firms. Second, we conduct an accounting exercise to quantify the various model mechanisms in shaping wage inequality.

The model has strong implications for worker sorting based on desired hours. Specifically, it predicts that workers who work fewer hours than their co-workers are more likely to sort into smaller firms where their desired hours align more closely with their average coworkers. Conversely, those working longer hours than their co-workers are more likely to sort into larger firms. We find support for this prediction in the data. By tracking CPS respondents over a span of 12 months, we show that workers with relatively shorter hours are more likely to transition into smaller firms and less likely to move to larger firms – just as the theory predicts. This finding complements existing literature on labor market sorting, highlighting the importance of sorting based not only on worker skills and firm heterogeneity but also on workers' desired hours.

We use the extended model to explore how the relationship between firm size, hours and wages – as summarized by our motivating facts – contributes to inequality. In contrast to canonical models, where wage inequality is primarily due to individual and firm-specific factors, our model highlights that a worker's hours relative to their coworkers will also impact wages. We find that variation in relative hours – that is, the ratio of one's own and a measure of usual (or average) coworker hours – accounts for 4.6% of the observed wage variation in the model. Decomposing the contribution of relative hours further into i) variation in average hours across firms, taking individual hours as given and ii) variation in individual hours, taking average hours as given, reveals that variation in average hours across firms *reduces* wage inequality (by 1.8%) while variation in individual hours *raises* wage inequality (by 6.4%).

The overall effect on wage inequality of firm-level differences in average hours across firms masks significant heterogeneity. Differences in average hours across firms have an equalizing effect on inequality among workers who work shorter hours (under 40) while it amplifies inequality for workers working longer hours (over 40). This heterogeneous impact follows directly from our three motivating facts. On its own, the size-wage premium increases wage inequality for workers regardless of hours worked - shifting the wage schedule uniformly upwards for all hours. However, since – as we document – average hours worked are increasing with firm size, the shift in the wage schedule is not uniform across all hours, leading to heterogenous wage-hour schedules across firms (even after accounting for productivity differences), giving rise to the overall and heterogeneous impacts on wage inequality. Indeed, holding all else equal, had the relationship between average hours and firm size been negative, we would predict overall wage inequality to be higher.

Next, we investigate how two key features of our model – firm heterogeneity and complementarities in working hours – shape inequality in i) hours, ii) wages, and iii) income. While firm-level differences have long been recognized as a source of wage dispersion, complementarity in hours introduces a novel channel influencing inequality. Comparing our benchmark model to counterfactual economies that abstract from these features, we find that firm heterogeneity increases inequality in hours, wages, and income. In contrast, complementarities in hours decrease the dispersion in hours by incentivizing workers to align more closely with their coworkers' hours while increasing the dispersion in wages due to penalties associated with deviations from firm-specific reference hours. The compression of the hours distribution dominates quantitatively, so complementarities in hours act to decrease overall income inequality. Lastly, we use our model to conduct a simple quantitative exercise which suggests that changes in the distribution of average hours across firms – driven by trend changes in the composition of firms – can ameliorate the trend increases in earnings inequality.

Taken together, this paper highlights the endogenous interaction between hours, wages, and firm-level heterogeneity. An interaction that is currently under-emphasized in the literature but, we argue, has important implications for earnings inequality.

Related literature. This paper is closely related to several strands of literature studying the interaction of firm characteristics, wages and hours worked. Among our three motivating facts, the size-wage premium has been the most extensively documented and studied. No consensus on the determinants of the size-wage premium exists, and our model generates it through workers' heterogeneous preferences over potential employers, closely following the approach in Card et al. (2018) and Lamadon et al. (2022). Such heterogeneity prevents worker flows from equalizing wages across firms and generates higher wages at larger, more productive firms.

In contrast, the positive relationship between hours worked and firm size that we document is relatively under-studied. Montgomery (1988) and Headd (2000) touch on this relationship by showing that larger firms have a lower fraction of part-time workers. Our analysis moves beyond distinguishing between part-time and full-time workers and considers the entire distribution of hours across firms. Accordingly, our model is rich enough to capture differences between small and large firms that are not driven by discontinuities in part/full-time hiring costs or workers' productivity.

Our third fact, the variation in the long and short hours penalties across size categories of firms, is novel and naturally relates to the literature that documents the presence of such penalties across the economy. Recent work by Yurdagul (2017) and Bick et al. (2022) have documented a hump-shaped relationship between wages and hours in the aggregate while

Shao et al. (2023) document such a relationship within establishments. We build on these findings by providing new evidence showing that the hump-shaped wage-hours profile varies with firm size.

Our paper joins a growing literature studying heterogeneous agent macroeconomics models that aim to be consistent with micro-level evidence on wages and the labor supply of workers. Much of this literature has focused on the response of labor supply to business cycle or life-cycle fluctuations (see for example, Heathcote et al. 2014, Erosa et al. 2016, and Chang et al. 2020). Instead, we focus on the cross-sectional relationship between hours, wages, and firm characteristics. Bick et al. (2022) study the relationship between hours and wages, using an exogeneously specified non-linear wage schedule to replicate the hump-shaped wage-hours profile. In contrast, we generate the non-linear wage schedule endogenously in a general equilibrium model, emphasizing the interaction between hours, wages and firm-level heterogeneity as well as its implications for inequality.

A distinguishing feature of our model is the presence of complementarities in the hours of different workers. Our modeling of complementarity is closely related to Yurdagul (2017) and Shao et al. (2023).[2] The feature aims to capture the necessity for workers to coordinate their work schedule to produce output and is supported by empirical evidence. For instance, recent work by Shao et al. (2023) and Kuhn et al. (2023) use matched employer-employee data to show that workers' hours within the same establishment are gross complements and coordinated. As such, our model relates to the literature exploring the impact of constraints on working hours resulting from coordination (see, for example, Altonji and Paxson 1988, Chetty et al. 2011, Labanca and Pozzoli 2021, and Cubas et al. 2022). We contribute to this strand of literature by explicitly incorporating inflexible working hours (via complementarities) in a theoretical model and studying its interaction with firm heterogeneity and its implications worker sorting across firms and inequality.

Our analysis of inequality is related to recent empirical work such as Song et al. (2019) and Barth et al. (2016) that explore the role of worker and firm heterogeneity in generating wage dispersion. We contribute to this literature by documenting the role of hours heterogeneity

[2]Battisti et al. (2024) also utilize a similar production function to analyze the correlation between estimates of the Frisch elasticity and the extent of hour complementarities.

in wage and income inequality. Blau and Kahn (2011) and Checchi et al. (2016) document the relationship between heterogeneity in hours and income inequality. We relate to this literature by highlighting the importance of complementarity in hours.

An outline of the paper is as follows. Section 2 describes our motivating facts in detail and Section 3 outlines our model. In Section 4 we calibrate the model and discuss how it reconciles the motivating facts. In Section 5 we extend our baseline model with realistic features to make it more suitable for quantitative analysis. Section 6 explores the aggregate implications of our theoretical framework for wage and income inequality. Section 7 concludes.

2 Motivating facts

This section documents three motivating facts about the distribution of hours and wages across firm size. First, we establish a robust positive relationship between firm size and average worker hours. Workers in the smallest firms (1 to 9 employees) work 7% fewer hours per week than those in the larger firms. Second, we confirm the existence of a size-wage premium. Third, we show that workers face penalties for working either short or long hours in all firms. However, the magnitude of these penalties vary systematically across size categories. In particular, we find that the penalty for working long hours decreases with size while the penalty for working short hours increases with size.

Data description To establish these facts, we use data from the Annual Social and Economic Supplement (ASEC) of the CPS covering information from 1991 to 2018. This supplement to the CPS contains detailed information on respondents' economic activity over the previous year. Importantly, it includes information on worker earnings, usual weekly hours worked, and firm size.[3] The partitioning of size bins has varied over time, so for clarity and consistency, we report three categories of firm size; i) small (1 to 9 employees), ii) medium (between 10 and 100 employees), and iii) large firms (over 100 employees). We

[3]Data is extracted from IPUMS and described in Flood et al. (2020). Information on firm size is only available in the ASEC and is reported in bins which record the total number of employees that a worker's employer has at all establishments. For example, the 2019 ASEC, which contains information for the reference year 2018, asks the following: "Counting all locations where this employer operates, what is the total number of persons who work for your employer?". We start our sample in 1991 since, prior to this, the smallest reported firm size category was firms with under 25 employees.

restrict attention to individuals aged between 25 and 64, who report working with a single private employer in the previous year and exclude those who usually work less than 10 hours a week or have implied hourly wages less than half the federal minimum wage. Respondents with imputed values for annual earnings, firm size, hours worked, or weeks worked are also dropped. The final sample includes around 819,000 respondents.

We report a number of additional empirical results in Appendix B. Specifically, Appendix B.1 provides robustness checks for our findings using data from the CPS. Appendix B.2 argues that our findings are not driven by measurement error in hours while Appendix B.3 replicates our empirical analysis at both the establishment and firm level using analogous data from the 1997 to 2018 Canadian Labour Force Surveys (LFS).

Fact 1 Average hours increase with size. We begin by studying the relationship between firm size and hours worked. Figure 1 reports the distribution of usual weekly hours worked by firm size with Panel (a) showing the overall distribution of usual hours worked. While the median weekly hours across all firms is between 40 and 44 hours, there are important differences in the share of short and long hours worked across firm sizes. This is evident in Panels (b) and (c), which report, respectively, the distribution of the right and left tails of the hours distribution. Panel (b) shows that workers in small firms are much more likely to work shorter ($<$40) hours than their counterparts in larger firms.[4] For example, around 3% of employees in larger firms work 30 to 35 hours while the analogous share in small firms is around 6%. Panel (c) shows that employees in medium and large firms are more likely to work between 45-59 hours with a similar likelihood of working very long (≥ 60) hours. For example, just around 8% of employees in small firms (less than 10 employees) work between 50 and 55 hours while the analogous share in larger firms is 10%.

The first column of Table 1 reports the (unconditional) average hours worked by firm size and confirms that, as suggested from Figure 1, there is a positive relationship between

[4] We plot the cumulative distribution of worker shares by firm size and hours worked in Figure A.1 in Appendix A. The cumulative distribution makes clear that the share of hour workers is higher in smaller firms for all levels of shorter hours under 40. Further, for all firm sizes, we do not observe any significant jumps or shifts in the cumlative distribution at or around 35 weekly hours. This is threshold defining part-time and full-time workers – a distinction that has been the focus of the (limited) existing work studying hours across firm size (for example, Montgomery 1988 and Headd 2000).

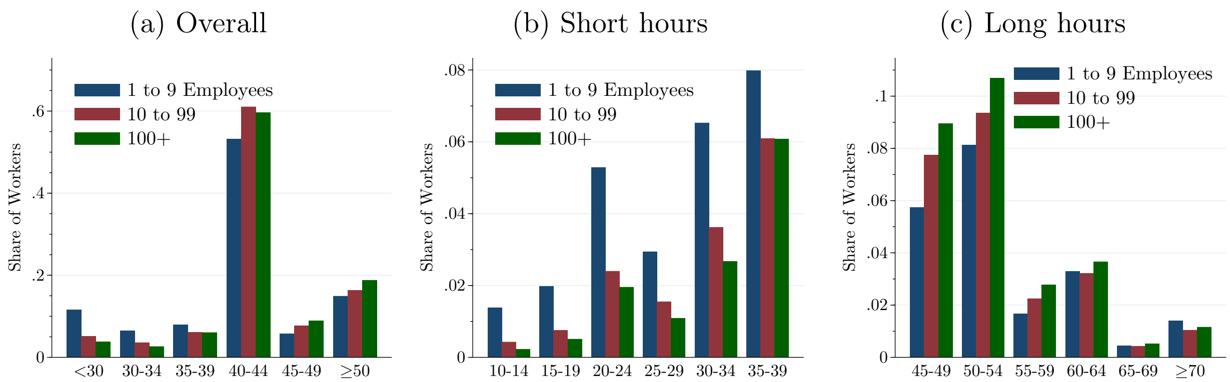

Figure 1: Distribution of working hours by firm size

Notes: The figure reports the share of workers by their usual weekly hours worked and firm size using data from the CPS.

average hours worked and firm size. On average, workers in the largest firms work around 3 hours longer than workers in the smallest firms.

While informative, the cross-sectional average does not control for confounding factors, such as industry, that might impact both firm size and hours worked. To control for such factors, we estimate the following regression,

$$h_i = \alpha + \left(\sum_{f \in F} \beta_f \mathbb{I}_{i,f} \right) + \delta X_i + \epsilon_i \tag{1}$$

where h_i are the usual weekly hours worked. X_i is a vector of individual-level controls which includes a quadratic in years of experience, dummies for the race, education, gender, marital status as well as state, year, and industry fixed effects. The variable $\mathbb{I}_{i,f}$ is an indicator variable which is equal to one if an individual is employed in a firm of size $f \in F$ so that the coefficient β_f captures the additional number of hours worked by firm size.

The last three columns of Table 1 report this coefficient and show that when excluding controls for industry or demographics, workers in the firms with 10 to 99 employees work, on average, around 1.8 and 1.3 hours longer per week than workers in firms with 1 to 9 employees while workers in firms with over 100 employees work between 2.7 and 2.2 hours longer per week. Controlling for industry and demographic characteristics explains some of the differences in hours worked between medium and larger firms and implies that workers in firms with 10 to 99 employees work around 1 hour and 15 minutes longer per week and

Table 1: Firm size and hours worked

	Uncond. Avg.	Conditional Avg. (rel. to small firms) (1)	(2)	(3)
1 to 9 Employees	39.5 hrs	-	-	-
10 to 99 Employees	41.3 hrs	1.874***	1.584***	1.314***
		[+2.4 weeks/yr]	[+2.1 weeks/yr]	[+1.7 weeks/yr]
		(0.046)	(0.044)	(0.044)
100+ Employees	42.2 hrs	2.722***	2.311***	2.180***
		[+3.5 weeks/yr]	[+3.0 weeks/yr]	[+2.8 weeks/yr]
		(0.042)	(0.041)	(0.043)
Year, State FE	-	Y	Y	Y
Demographic Controls	-	N	Y	Y
4-digit Industry FE	-	N	N	Y
N	819,295	819,295	819,295	819,295
R^2	-	0.014	0.106	0.137

Notes: The first column of the table reports the unconditional average of hours worked by firm size. The table reports the coefficient β_f estimated from Equation (1) where the reference size category is firms with 1 to 9 employees. The brackets report the additional number of weeks worked per year implied by the estimated regression coefficient. For example, an additional, relative to small firms, 2 hours worked per week over 52 weeks implies an additional 104 hours worked per year. Given that the median work week consists of 40 hours, this suggests an additional 2.6 (104/40) weeks worked per year. Standard errors are reported in parentheses. *** indicates statistical significance at the 1% level.

workers in firms with over 100 employees work 2 hours and 10 minutes longer per week than workers employed in the smallest firms.

In addition, the differences between medium and large firms are statistically significant confirming that average weekly are positively related with firm size even after controlling for observable characteristics. These differences are also economically significant and could amount to as much as an additional 2 to 3 weeks of work per year in medium and large firms relative to small firms.

We report a number of robustness exercises in Appendix B.1. Specifically, we use the more detailed firm size categories reported in the CPS to show that average weekly hours tend to increase monotonically across finer firm size categories. We also show that hours increase with size when Equation (1) is estimated with additional controls including occupation fixed effects and controls for hourly worker status. Building on Montgomery (1988) and Headd (2000), who show that smaller firms have a higher share of part-time workers, we find a positive relationship between weekly hours and firm size even after controlling for worker's part/full-time status or focusing on workers that work within a narrow window of hours (35 to 45 hours per week).

__Fact 2__ **Average wages increase with size.** The wage premium in large firms has been studied extensively (see, for example, Oi and Idson 1999). We establish the size-wage premium in our data by estimating the following regression,

$$\log(w_i) = \alpha + \left(\sum_{f \in F} \beta_f \mathbb{I}_{i,f}\right) + \left(\sum_{h \in H} \gamma_h \mathbb{I}_{i,h}\right) + \left(\sum_{f \in F} \sum_{h \in H} \theta_{f,h} \left[\mathbb{I}_{i,f} \times \mathbb{I}_{i,h}\right]\right) + \delta X_i + \epsilon_i \qquad (2)$$

where $\log(w_i)$ is the log hourly wages of individual i. Hourly wages are computed as the ratio of annual earnings to the product of usual weekly hours and weeks worked in the year. As in (1), X_i is a vector of individual-level controls which includes demographic controls, state, year, and industry fixed effects. The indicator variable $\mathbb{I}_{i,f}$ is equal to one if an individual is employed in a firm of size $f \in F$. Similarly, $\mathbb{I}_{i,h}$ is equal to one if an individual works h hours.

We partition weekly hours into a set H by grouping hours in 5-hour bins. The partitioned set is $H = \{10-14, 15-19, \ldots, 65-69, 70-99\}$. The final bin $70-99$ is larger as there are relatively few workers working over 70 hours. As most workers work 40 hours, we choose the category $40-44$ as the reference category for hours and omit the coefficients γ_{40} and $\theta_{40,f}$ for all f. The reference size category is firms with 1 to 9 employees.

The regression in (2) extends the specification in Bick et al. (2022) by also controlling for firm size and an interaction term between firm size and usual weekly hour bins. Including these regressors allows us to study i) the size-wage premium (fact 2) and ii) the relationship between hours and wages by firm size (fact 3 discussed below).

The coefficient that captures the firm size wage premium (for workers that work in the 40 hours bin) is β_f. Table 2 reports β_f and shows that it increases monotonically in size.

Indeed, the wage premium between the largest and smallest size categories is around 25%.

Table 2: The size-wage premium

	(1)	(2)	(3)
10 to 99 Employees	0.138***	0.123***	0.114***
	(0.004)	(0.003)	(0.003)
100+ Employees	0.353***	0.290***	0.247***
	(0.003)	(0.003)	(0.003)
Year, State FE	Y	Y	Y
Demographic Controls	N	Y	Y
4-digit Industry FE	N	N	Y
N	819,295	819,295	819,295
R^2	0.149	0.370	0.450

Notes: The table reports the coefficient β_f estimated from Equation (2) where the reference size category is the smallest size firms/establishments. That is, firms with 1 to 9 employees. The reference hours bin is $40 - 44$ hours. Standard errors are reported in parentheses. *** indicates statistical significance at the 1% level.

<u>Fact 3</u> Long-hour (short-hour) penalty decreases (increases) with size. Next, we study the cross-sectional relationship between wages and hours worked by firm size. Panel (a) of Figure 2 plots the unconditional average hourly earnings of workers by hours worked and firm size. This is obtained by estimating Equation (2) without any controls and, in particular, we report the sum of the coefficients γ_h and $\theta_{f,h}$ estimated from (2). This sum captures the wage penalty of working outside of the 40-44 hours bin by size category.[5] This cross-sectional average features a hump-shaped relationship between hours and earnings across all firm sizes. That is, there appears to be a relative wage penalty resulting from working either longer or shorter hours relative to a wage-maximizing level of hours. Importantly, this hump-shape appears to vary by firm size.

Panel (b) reports the same relationship after controlling for observable characteristics. Adding controls significantly shrinks variation in wages but qualitatively implies the same hump-shaped relationship between wages and hours. An aggregate hump-shaped wage-hours

[5]Strictly speaking, the coefficient γ_h captures the wage penalty of working away from the 40 hours bin for small firms. While $\theta_{f,h}$ captures the difference in penalty relative to small firms. Recall that $\theta_{f,h}$ for small firms is zero as small firms are the reference size category.

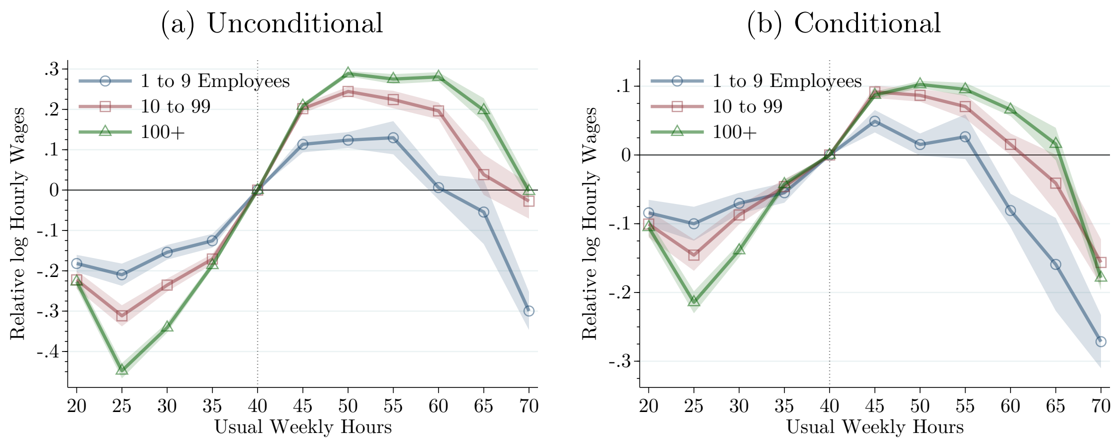

Figure 2: The relationship between wages and hours by firm size

Notes: Panel (a) reports the sum of coefficients $(\gamma_h + \theta_{f,h})$ estimated from Equation 2 with no controls. Panel (b) reports the same sum of coefficients $(\gamma_h + \theta_{f,h})$ when including controls for observable characteristics denoted as X_i in Equation 2. The reference category for hours worked is $40 - 44$ and the reference category for firm size is firms with 1 to 9 employees. The shaded regions are the 95% confidence intervals.

relationship has been documented by Yurdagul (2017) and Bick et al. (2022) in the US.[6] Panel (b) shows that this hump-shaped relationship varies significantly with firm size.

Specifically, the conditional hourly wages in Figure 2 suggest that the penalty for working shorter hours is more severe in larger firms while the penalty for working longer hours is less severe. For example, relative to working 40 hours, working 25 hours in a firm with 1 to 9 employees is associated with a 10% wage penalty, while the analogous penalty in a firm with 10 to 99 (over 100) employees is around 4 (10) percentage points *higher*. On the other hand, working longer hours – say around 60 hours – in a firm with 10 to 99 (over 100) employees results in a wage penalty that is roughly 9 (14) percentage points *lower* relative to working the same hours in a smaller firm.[7] These are significant differences and imply that wage penalties for working 25 hours are 40 (100)% larger in firms with 10 to 99 (over 100) employees while working 60 hours are 75 (116)% smaller. In addition, the peaks of

[6] Earlier work by Hirsch (2005) also explores the aggregate conditional difference in wages based on working hours but focused on part-time workers relative to full-time workers (and selected hours bin).

[7] This difference in relative penalties is captured explicitly by the coefficient $\theta_{f,h}$ in Equation (2). Figure A.2 plots this coefficient and shows that there are indeed statistically significant differences in wage penalties by size. Further, in Appendix B.1, we use the Outgoing Rotation Group of the CPS to control for whether the hourly pay status of workers – that is, whether they earn an hourly wage or are salaried. We find that the composition of hourly earners does not generate qualitatively different implications for wage penalties by firm size. We also argue that our results are robust to controlling for occupations and are not driven by measurement error in hours.

the conditional wage-hour profiles clearly differ by firm size with wages in the largest firms being highest in the 50 hours bin while wages peak in small and medium sized firms at the 45 hours bin.

Together, the three empirical facts highlighted in this section motivate our theoretical analysis. A focus of our framework is the causal link between differential wage penalties and average hours across firms. We generate a link between wage penalties and average hours by allowing for complementarities in hours worked in our theoretical model. A natural consequence of such complementarities is that wage penalties are increasing in the distance from the usual hours worked within a workplace and that wages are maximized at higher hours in larger firms. Since larger firms feature longer average hours, longer hours are penalized less severely compared to smaller firms. Conversely, smaller firms feature shorter hours, and hence shorter hours are penalized less severely compared to larger firms.[8]

Having said this, it should be noted that the cross-sectional wage penalties documented here are not conclusive of complementarities in hours, as they may also reflect unobservable worker or workplace characteristics. For example, Hirsch (2005) exploits the rotating panel design of the CPS and argues that the observed (cross-sectional) wage penalty for working part-time is due to worker heterogeneity. A longitudinal matched employer-employee dataset that includes detailed information on firms and worker characteristics including the distribution of coworker hours would be an ideal dataset to disentangle the role of worker and firm characteristics separately from the role of complementarities in hours. Unfortunately, to our knowledge, no such dataset exists for the US economy.[9] In Appendix B.4, we instead use the short panel of the CPS to conduct longitudinal analysis of the relationship between wages and present evidence consistent with the presence of complementarities in hours while also highlighting a role for worker characteristics in shaping the cross-sectional wage-hours relationship.

[8]This intuition implies that average hours not only affect the wage-hours menu faced by workers but are also affected by these wage-hours menus as they alter workers' labor-supply decisions. This feedback mechanism will also be present in our model.

[9]Shao et al. (2023) use a Canadian matched employer-employee data set which includes information on coworkers hours; the Workplace Enterprise Survey. They document, using both cross-sectional and longitudinal analysis, evidence supporting the presence of complementarities in working hours in Canada. As we discuss below, we will utilize their estimate of the elasticity of substitution of working hours in our quantitative analysis.

Having discussed our primary motivating facts, we now move to describe our theoretical framework.

3 Model

In this section, we present our baseline model. It is a static model that exhibits a minimal framework in order to better highlight the model mechanisms. There are two types of agents each with a unit mass: firms and workers. Below we describe the production function of firms, the preferences of workers, and the market structure. Then we present the maximization of each agent.

3.1 Firms' technology

Firms' only endogenous input is labor. Production of all the firms in the economy can be represented by $Y = zL^{\theta}$, where L denotes the effective labor input. Firms differ in their exogenous productivity, z, a discrete variable which can take J different values with density Λ_j, $j \in \{1, 2, .., J\}$. We think of the productivity term z as broadly capturing all non-labor inputs of the firm as well as its technology. We assume that firms are owned by absentee entrepreneurs and all profits are accrued to them.

As in Yurdagul (2017) and Shao et al. (2023), we allow for complementarities between hours of workers:

$$L = \underbrace{\left(\frac{\int\limits_{i \in \mathcal{N}} l_i^{\rho} di}{\int\limits_{i \in \mathcal{N}} di} \right)^{\frac{1}{\rho}}}_{\text{average}} \underbrace{\left(\int\limits_{i \in \mathcal{N}} di \right)}_{\text{scale}}, \tag{3}$$

where $\mathcal{N}$ is the set of workers, and $\{l_i\}_{i \in N}$ is their hours worked. The first term of labor aggregation is an average hours measure in CES form, where the total weight of each individual is normalized to unity, i.e. a scaling up of the labor unit maintaining the distribution of hours would not alter this term. The second term is meant to scale up the labor aggregation depending on the number of workers involved. Accordingly, the labor aggregation boils down to the standard linear aggregation of hours when $\rho = 1$. In our model, we allow for $\rho < 1$, which captures imperfect substitutability between the working hours of coworkers in a unit.

We abstract from the indices of workers by rewriting the aggregation in terms of the measure of workers employed at each level of hours, and also collect the scale component in one term:

$$L = \left(\sum_{l \in \mathcal{L}} \mu(l) l^{\rho} \right)^{\frac{1}{\rho}} \left(\sum_{l \in \mathcal{L}} \mu(l) \right)^{1-\frac{1}{\rho}}, \tag{4}$$

where $\mu(l)$ is the measure of workers working l hours. Here we assume that the hours of workers are on a discrete grid, i.e. $l \in \mathcal{L} = \{l_1, ..., l_N\}$.[10]

3.2 Workers preferences

There is a continuum of workers with static decisions. Preferences are given by:

$$c_{it} - \nu_{it} \frac{l_{it}^{1+\phi}}{1+\phi}. \tag{5}$$

The value of leisure is log-normally distributed, $\log \nu_{it} \sim N(\log \nu_0, \sigma_\nu)$. We denote the set of value of leisure shocks by B_ν.

Each period, a worker receives shocks to the value obtained in each firm productivity group. Formally, there is a vector $\boldsymbol{\epsilon} = \{\epsilon_1, \epsilon_2, ..., \epsilon_J\}$ with as many components as the number of different firm-level productivity. This vector follows a distribution $G(\boldsymbol{\epsilon})$ and is drawn independently every period. These shocks capture workers' taste for the other factors locating a worker into a type of firm not featured in our model and their role is discussed in more detail below. We interpret both these shocks and the value of leisure as being unobservable characteristics of workers.[11]

Workers sort into labor markets conjecturing a wage outcome $w_j(l)$. Accordingly, their

[10] Discrete hours facilitate establishing the optimality of the demand schedule in equilibrium.

[11] We assume that workers' preferences for firms are uncorrelated with their leisure preferences. Allowing for a correlation in these preference would result in the equilibrium relationship between hours worked and firm size being mechanically influenced by the extent and direction of this correlation. For example, if preferences for firms and leisure are such that workers who wish to work shorter hours also prefer smaller firms, then, in equilibrium, average hours will be shorter in smaller firms even in the absence of the endogenous forces we discuss. Thus, to maintain parsimony and focus on endogenous mechanisms we opt for uncorrelated preferences for firms and leisure.

value function is:

$$V(\nu, \boldsymbol{\epsilon}) = \max_j \{V_j^G(\nu) + \epsilon_j\},$$

where the value conditional on working in a firm of productivity j is:

$$V_j^G(\nu) \quad = \quad \max_{l \in \mathcal{L}} \; w_j(l)l - \nu \frac{l^{1+\phi}}{1+\phi}.$$

The description of workers' problem so far conjectures the outcomes after sorting into a labor market. Below we describe the market structure and the decisions that take place within each market.

3.3 Market structure and the timeline in each market

There are J labor markets, one for each productivity level z_j such that all firms of type z_j participate in labor market j while firms for type z_k participate in a separate labor market k. In other words, labor markets are identified by the productivity of firms and consist of homogenous firms. Upon the realization of their taste shocks (ϵ) and their value of leisure (ν), workers choose which market j to participate in and the hours l to work. We denote the density of workers in market j that wish to work l hours as $\Phi_j(l)$.

The problem of firms involves two sequential stages. First, firms choose the number of workers, M to "interview" for consideration for employment – the *recruitment stage*. We assume labor markets are frictional such that firms cannot target or direct their search towards any particular type of worker and interview participating workers randomly. Accordingly, a firm that chooses to interview M workers receives $M\Phi_j(l)$ workers of each type.

After the recruitment stage, firms decide the number of interviewed workers of each type to hire $\mu_j(l)$, and a wage offer $w_j(l)$ for these workers – the *hiring stage*. If a worker accepts the offer, they work in the current period for the firm. If the offer is rejected, the recruited match breaks up, the firm pays a penalty $\kappa > 0$. Within the same period, the worker is then randomly matched to another firm in the same market j and now makes a take-it-leave-it offer herself with the same hours to that new firm. If this offer is rejected by the new firm, the worker receives no income in the current period and does not work (the new firm does not incur a rejection cost).

3.4 Solution of an equilibrium

A formal definition of the equilibrium is included in Appendix C.1. Informally, an equilibrium consists of policy and wage functions, and a distribution of workers and hours, such that i) policy functions solve the problem for workers and firms, ii) labor markets clear, and iii) the distribution of hours and workers are consistent with the optimal policy and wage functions. We focus on the equilibrium in which all firms with the same productivity have the same optimal policies in both the recruitment and hiring stages. In this *symmetric* equilibrium, firms in a labor market recruit the same number of workers, hire all recruited workers, and make identical wage offers.

Before discussing the policies of firms in equilibrium, it is useful to define the marginal productivity, $f_j(l;\mu)$, of an $l-$hour worker in a firm of type j that has an arbitrary demand schedule of $\mu(l)$,

$$f_j(l;\mu) \equiv \theta z_j L(\mu)^{\theta-1}\left[\frac{1}{\rho}\left(\frac{l}{\tilde{l}_j}\right)^{\rho} + \left(1-\frac{1}{\rho}\right)\right]\tilde{l}_j \tag{6}$$

where $L(\mu)$ is the labor aggregation implied by the demand schedule $\mu(l)$ and

$$\tilde{l}_j \equiv \left[\left(\sum_{l\in\mathcal{L}}\mu(l)l^{\rho}\right) / \left(\sum_{l\in\mathcal{L}}\mu(l)\right)\right]^{\frac{1}{\rho}} \tag{7}$$

is a "weighted average" of l^ρ based on $\mu(l)$.

Under a symmetric equilibrium, the implied labor allocation of firms directly reflects the supply of workers with $\mu_j^*(l) = M_j^*\Phi_j(l)$ where M_j^* denotes the equilibrium number of workers interviewed (and hired). Under such an equilibrium, the marginal product of an l-hour worker can be written as $f_j(l; M_j^*\Phi_j(l))$. Next, we discuss how to solve for a symmetric equilibrium and argue that this equilibrium can be sustained under mild conditions.

We solve for the equilibrium policy and wage functions of firms using backwards induction. Starting from the hiring stage, when a firm chooses to employ a recruited worker, it will offer the lowest wage that leaves the worker indifferent between accepting and rejecting (and is profitable for the firm). To identify this minimal wage, it is useful to consider the case of a possible break-up between a worker and the firm that originally recruited her. In this case, upon a match with a random alternative (representative) firm, the worker will find it optimal to offer an hourly wage that is equal to her marginal productivity per hour in that firm; $f_j(l; M_j^* \Phi_j(l))/l$. This will leave the firm indifferent between accepting or rejecting. We assume that the firm accepts and so the worker becomes employed in the current period.

Hence, the wage that the (initial) recruiting firm should pay to employ a worker at the hiring stage is,

$$w_j(l) = f_j(l; M_j^* \Phi_j)/l. \tag{8}$$

Accordingly, the hiring stage problem (taking as given the total number of interviewed workers, M) becomes,

$$\begin{aligned} \pi_j(M) = \max_{\{\mu(l)\}_{l \in \mathcal{L}}} \; & Y - \sum_{l \in \mathcal{L}} w_j(l) l \mu(l) - \sum_{l \in \mathcal{L}} \kappa(M\Phi_j(l) - \mu(l)) \\ s.t.\; Y = z_j & \left[\left(\sum_{l \in \mathcal{L}} \mu(l) l^{\rho} \right)^{\frac{1}{\rho}} \left(\sum_{l \in \mathcal{L}} \mu(l) \right)^{1-\frac{1}{\rho}} \right]^{\theta} \\ & \mu(l) \in [0, M\Phi_j(l)] \;\; \forall l \in \mathcal{L}. \end{aligned} \tag{9}$$

where the optimal policy of firms at this stage is denoted as $\mu_j^*(l; M)$.

In turn, the problem of the firms at the initial recruiting stage is to choose the number of workers to interview and is given by,

$$\Pi_j = \max_{M \geq 0} \; \pi_j(M), \tag{10}$$

with a corresponding solution denoted M_j^*.

In our model, the firms' problem is not necessarily convex, hence it is not feasible to analytically prove the optimality of the aforementioned policy functions and, in turn, the existence of the symmetric equilibrium. Accordingly, we use a semi-numerical approach, detailed in Appendix C.4, to argue that the policy functions in the symmetric equilibrium are optimal for firms under relatively small values of rejection costs κ.

Briefly, we first characterize – analytically – the set of profit maximizing demand schedules $\mu_j(l; M)$ for any given number of interviewed workers M. We then – numerically – search within this narrow set of alternative to identify the highest profits that firms could achieve (for any M) by deviating from the symmetric equilibrium. We show that for reasonably small values of κ (less than 1% of average wages), the profit maximizing solution from the first (recruitment) stage is the same as in the symmetric equilibrium and that if this is the case, then hiring all interviewed workers achieves higher profits than any other solution that involves rejection of workers.

Then, denoting the solution of (10) as M_j^*, under a symmetric equilibrium, the optimal labor demand and wages are for a firm of productivity z_j are given by,

$$\mu_j^*(l) = M_j^* \Phi_j(l) \tag{11}$$

$$w_j(l) = \theta z_j L(\mu_j^*)^{\theta-1} \left[\frac{1}{\rho} \left(\frac{l}{\tilde{l}_j^*} \right)^{\rho-1} + \left(1 - \frac{1}{\rho}\right) \left(\frac{l}{\tilde{l}_j^*} \right)^{-1} \right], \tag{12}$$

where $\tilde{l}_j^*$ denotes the "weighted average" hours within the firms according to Equation (7) under the symmetric schedule.

Remark on the wage schedule. Equation (12) shows that a worker's wage depends on her hours of work relative to her co-workers in the same firm. Indeed, the highest hourly wage is achieved when $l = \tilde{l}_j^*$ with wages decreasing as working hours move further from $\tilde{l}_j^*$. This characteristic of the wage schedule is crucial in generating not only an aggregate hump-shaped wage-hours relationship but also in making the penalties for short and long hours depend on the usual hours in each firm. Indeed, if usual hours were to differ across firms, so would the wage schedules – an observation highlighted in our motivating facts.

Notice that although our model features a non-standard labor market structure, the implied wage schedule exactly in line with the implications from more standard setups: Workers are paid their marginal product. We discuss this further in the next subsection.

3.5 Discussion

Before discussing our quantitative analysis, we provide more detail on some of the non-standard features of our model and the role they play.

Complementarities in hours. Unlike most macroeconomic models of production, we allow for complementarities in the hours of different workers by aggregating hours in a non-linear manner. Intuitively, such complementarity captures the need for workers to coordinate which can arise naturally when workers' tasks require coordination. The need for such coordination has long been recognized since at least Adam Smith's early discussions of assembly line production. Recently, Bick et al. (2022) and Labanca and Pozzoli (2021) document suggestive evidence supporting the presence of complementarities in hours. More direct evidence is presented in Shao et al. (2023) that uses matched employer-employee data from Canada to estimate the elasticity of substitution between hours worked and find that working hours are indeed gross complements in production.

Given the evidence supporting complementarities in working hours, we explicitly incorporate them into our production function. As we discuss below, these complementarities play a crucial role in generating the observed hump-shaped relationship in wages and hours. However, the presence of complementarities do not, on their own, generate differences in wage penalties by firm size or the patterns of sorting and inequality that we highlight. Indeed, we will specifically discuss how alternative levels of complementarities impact the model's implications in our quantitative analysis below.

Structure of the labor market. In a frictionless environment, the presence of complementarities in hours will push firms towards concentrating their hours at a single point so that all workers in a firm work the same number of hours. However, as highlighted in recent work by Battisti et al. (2024), Labanca and Pozzoli (2021) and Shao et al. (2023), such concentration is not observed in the data with significant dispersion in hours *within* firms. Our two-stage approach to the labor market, where firms first conduct a non-directed search based on the number of workers and then decide their hiring based on hours, is a tractable yet endogenous way of introducing frictions in the labor search process that limits a firm's ability to concentrate at a single level of hours. These frictions include all the costs of looking for particular types of workers that wouldn't be in line with their distributions prevailing in the market.

As we describe in detail in Appendix C.4, the symmetric equilibrium of the quantitative model can be sustained with a only modest level of frictions (that is, cost of rejecting a recruited worker). Firms have little incentive to deviate and concentrate on usual hours (frequently observed hours, such as 40) because these hours receive the highest wages in the market. Alternatively, if they choose to hire only workers of unusual hours (such as 20 or 60), they need to recruit and subsequently reject many workers, resulting in high rejection costs, even if the cost per rejection is low. Accordingly, an endogenous mechanism combined with small rejection costs guarantees the existence of the symmetric equilibrium.

What is essential for our labor market structure is that it generates an equilibrium where there are no incentives for firms and workers to separate once they are randomly matched, and thus, there is dispersion in hours within firms (as in the data). The timing and bargaining protocol that we assume achieves this in one specific manner, but there are alternative setups that could achieve the same outcome. Deviations from this, where workers in a given market have an incentive to separate, would occur only if workers' marginal productivity is higher in another firm of the same type. This, in turn, requires that the distribution of hours varies across firms (since firms in a given market are otherwise identical). In principle, such a scenario could still be consistent with the observed dispersion of hours within firms but would require additional structure to discipline how hours distributions differ between (ex-ante identical) firms while simultaneously avoiding full concentration of hours within firms – an outcome that is counterfactual. Overall, the proposed setup here maintains relative simplicity while incorporating reasonable frictions associated with the labor market.

Our labor market setup also generates, as observed in the data, heterogeneous equilibrium wage schedules between firms. We generate this by assuming that markets are identified by firm type z with competition for labor within each market.[12] A setup where all firms of different z's compete for the same workers would result in a uniform, economy-wide wage

[12] A further split of the same productivity firms into more sub-markets (such as randomly allocating each j-type firm into $K > 1$ different sub-markets), with competition within each sub-market, would not change the results of our analysis.

schedule over hours.

Taste shocks. Workers are assumed to have heterogeneous preferences over their workplaces which are captured by the vector $\boldsymbol{\epsilon}$. In the absence of a consensus explanation for the size-wage premium, we introduce these preferences in order to generate higher wages in larger firms. We interpret this heterogeneity in tastes as capturing a number of (non-wage) factors that affect individuals' sorting into firms of different sizes and productivity. Such factors include differences in non-pecuniary benefits like workplace safety, childcare, or sick leave provision, as well as differences in technology.[13]

Intuitively, the presence of preferences for workplace prevents wages from equalizing across firms (and markets) of different productivity and generates higher wages in more productive firms that desire to be larger. To understand this, consider two extreme scenarios relative to our setup: one where there are no taste shocks and one where these shocks have infinite volatility. Without exogenous preferences over workplaces, worker flows and competition between firms will ensure that wages are equalized across firms, with higher labor input in more productive firms. On the other hand, with an infinite variance of taste shocks, workers' sorting into firms will be fully random and dependent only on taste shocks. Under a common expected value for the taste shock (as in our model), workers will be equally distributed across firm types, and labor input will be equalized across firms while wages will be proportional to productivity. Our current framework featuring taste shocks with finite volatility lies in between these two extremes. Here, workers flow into firms with higher productivity but not as much as in the case without taste shocks. Accordingly, labor input will be increasing in productivity but not so that wages are equalized across firms but rather such that wages are higher in more productive (larger) firms.

Our use of taste shocks to generate the size-wage premium follows the approach of recent work such as Card et al. (2018) and Lamadon et al. (2022). Taste shocks have the benefit of being a simple way to generate a size-wage premium and, importantly, these shocks do not by themselves generate the other empirical facts that we focus on. If anything, the presence of taste shocks adds noise to the sorting decisions of workers and weakens the

[13] We discuss the literature documenting such differences across firms, as well as further details on the role of the taste shocks, in Appendix C.3.

positive relationship between firm size and hours. For instance, in the extreme case of very large taste shocks, workers would sort based primarily on their idiosyncratic preferences for workplace resulting in similar average hours across firm size.

The main difference in our model compared to the more standard approach in Card et al. (2018) and Lamadon et al. (2022) is that these papers assume firm-specific tastes, whereas we assume tastes that are uniform within each productivity level of firms. Our approach simplifies the solution by leaving workers indifferent between any firm in a market while maintaining the implication of higher wages and a bigger size in more productive firms. This is particularly important in our setup which features multiple layers of heterogeneity among workers (including differences in value of leisure) and the production function with complementarities between working hours.

Taste shocks also provide a computational advantage as they effectively 'convexify' the occupational choice decision of workers by introducing randomness. This transforms workers' policy function to a probability between 0 and 1 rather than a binary of 0 or 1 and facilitates convergence in the model solution.

4 Quantitative analysis

This section describes the quantitative implications of the model. We begin by detailing our calibration strategy and then show that the calibrated model can match our three main motivating facts. Following this, we highlight the mechanisms in the model that result in outcomes that are consistent with the data.

4.1 Calibration

We describe the calibration exercise in three parts: the functional forms, parameters calibrated outside the model, and the parameters calibrated targeting features in the data. Since the model is static, we pool the observations both in the data and in the model for our analysis.

Functional forms. Firm productivity is assumed to follow a Pareto distribution with shape parameter λ, and the lowest productivity is normalized to 1. This distribution is approximated using 24 grid points. Our results are robust to increasing this number.

We assume that the ϵ-shocks affecting workers' value in each firm follow a Generalized Extreme Value distribution:

$$G(\boldsymbol{\epsilon}) = \exp\left[-\sum_{j=1}^{J} \exp\left(-\frac{\epsilon_j}{\sigma_\epsilon}\right)\right].$$

The parameter σ_ϵ determines the variance of these shocks.

Parameters calibrated outside the model. The model is calibrated to match key features of the US data.

Our model is intentionally simple, and in particular, it assumes linear utility in consumption, hence no income effects. Only one parameter, ϕ, captures the substitution between consumption and leisure. While the standard values of this parameter in the labor-macro papers are below 5, the corresponding studies typically allow for income effects. In calibrating this parameter for our very simple model, we face the tradeoff between avoiding the hours that increase very steeply with wages (hence with firm productivity and size), and remaining close to the empirical estimates for the Frisch elasticity parameter. Accordingly, we adopt a rather high value of ϕ at 10 in our baseline.[14]

The parameter governing the complementarities between working hours, ρ, is set by using results from Shao et al. (2023). Using matched employer-employee data from Canada, Shao et al. (2023) first provide evidence consistent with the presence of complementarities in hours in production, and then, using a generalized version of the production in this paper, they estimate the substitution parameter ρ to be around -0.46. Based on these results, and with the underlying assumptions that the production technologies adopted by US and Canadian

[14]In our extended model calibration, we will set this parameter at a standard value of 2. Even with ϕ higher than 10, the simple model cannot get to the low elasticities of hours to wages it produces for the extended model, including the steepness of the hours-size profiles. Nevertheless, the main model implications do not change with those values.

firms are similar, we set the value of ρ to be -0.46.[15]

Our model features costs on the firms for rejecting the recruited workers (κ per rejection). In our equilibrium of focus, the firms do not reject any of the recruited workers, hence, do not pay these costs. However, there is a minimum level for these costs needed to make the alternative schedules suboptimal. In Appendix C.4, we use a semi-numerical approach to compute this minimum for our model. We find that the minimum required to have the symmetric solution sustain in the baseline is $\kappa = 0.042$ which is about 0.7 percent of the mean wages in the baseline model.

We set the labor share of output equal to the standard value of 0.67. The working hours grid is assumed to have nine points: $\mathcal{L} = \{20, 25, ..., 60\}$.

Parameters targeting features in the data. The remaining parameters are calibrated to match specific targets in the data. To compare model implications to the data, it is useful to construct model-implied firm size categories to match those in the CPS sample – the source of our motivating facts. These size categories are firms with 1 to 9 employees (small firms), firms with 10 to 99 employees (medium firms), and firms with more than 100 employees (large firms). According to the Business Dynamics Statistics (BDS), the fraction of firms in these three categories in 2015 is 77, 21, and 2 percent, respectively. Firms in our model are categorized as "small", "medium" and "large" so that they replicate this distribution – once firms are sorted by size. We then calibrate the shape parameter of the firm TFP distribution, λ, to match the employment share of the largest firm size category. Table 3 summarizes the statistics for the three size categories in the data and the model. The last two columns compare average employment in the model and the data. The fit of our firm grouping to the data is assuring: not only is our firm grouping similar to the data in terms of the percentage of firms in each size category but also in terms of how average employment increases over the size categories.

For the value of leisure distribution, we target the observed average and dispersion of hours in the data. In particular, we set the level parameter (ν_0) to match the average weekly

[15] We justify the assumption of similarity between the US and Canadian economies by replicating, in Appendix B.3, our motivating facts using data from the Canadian Labour Force Survey (LFS). An alternative calibration targeting moments from Canadian economy gives qualitatively similar implications reconciling these facts.

Table 3: Firm size categorization

	Share of Firms		Share of Employment		Avg. Employment (log diff. from small firms)	
	Data	Model	Data	Model	Data	Model
Small	0.77	0.77	0.11	0.10	-	-
Medium	0.21	0.21	0.23	0.24	2.1	2.1
Large	0.02	0.02	0.66	0.66	5.4	5.5

Notes: The size categories in the data follow the categorization in the CPS. Small firms are firms with 1 to 9 employees, medium firms are those with between 10 and 99 employees, and large firms are those with over 100 employees. Data is from the 2015 Business Dynamics Statistics. The average size of small firms in 2015 was 3.4 employees. We follow the observed fraction of firms in each category to construct the same groupings in our model.

hours of workers. The standard deviation (σ_ν) is calibrated to match the standard deviation in log-hours.

The preference shocks for working at different productivity levels of firms generate noise in workers' choices over firms and prevent sorting from being driven entirely by pecuniary returns. As such, changes in the size of these shocks change the steepness of the wage profiles across firm productivity and size groups. Given this, we set σ_ϵ to match the ratio between the (unconditional) average wages in firms with over 100 employees and firms with under 100 employees.[16]

Table 4: Parameters of the baseline model

Panel A: Outside the Model

Parameter	Basis
$\theta = 0.67$	Labor Share of Output
$\phi = 10$	Low Elasticity
$\rho = -0.46$	Shao et al. (2023)
$\kappa = 0.042$	See Appendix C.4

Panel B: Calibrated

Parameter	Basis	Data	Model
$\lambda = 1.98$	Employment Share (large firms)	0.66	0.66
$\nu_0 = 15.77$	Average Weekly Hours	41.7	41.7
$\sigma_\nu = 2.53$	SD log Weekly Hours	0.22	0.22
$\sigma_\epsilon = 0.37$	Wage gap, large firm to rest	0.29	0.29

Notes: Panel A reports parameters that are set following the literature. Panel B reports the parameters that are calibrated to match specific data features with the model. The last two columns in Panel B report the data target and model implied value. The employment share of firms with over 100 employees is computed from the 2015 Business Dynamics Statistics. Measures of hours and wages are calculated using the pooled CPS sample.

Table 4 presents the parameter values in our calibration. It also reports the implied moments of the model against the data for the targets. The model performs well in matching

[16] The average wage gap between firms with over 100 and under 100 employees has evolved in a non-monotonic manner. It declined from around 0.32 to 0.26 between the 1990s and early 2000s, after which it slowly rose to around 0.29 by 2018. Our calibration target is the average over the entire sample period 1991-2018, which is very similar to the average over the decade before 2018.

the features of the data.

4.2 Model implications

This section compares the model's implications with the data. We begin by showing that the model replicates the motivating facts detailed in Section 2. In the next subsection, we discuss the relevant features of the model that generate these patterns.

Figure 3a reports the relationship of average wage (solid line) and hours worked (dashed line) by firm size category and shows that the model successfully replicates our first two motivating facts. As the wage premium for the largest firms (with respect to the rest) is a target in our calibration, the model quantitatively matches the increase in the average wage along the three size categories.

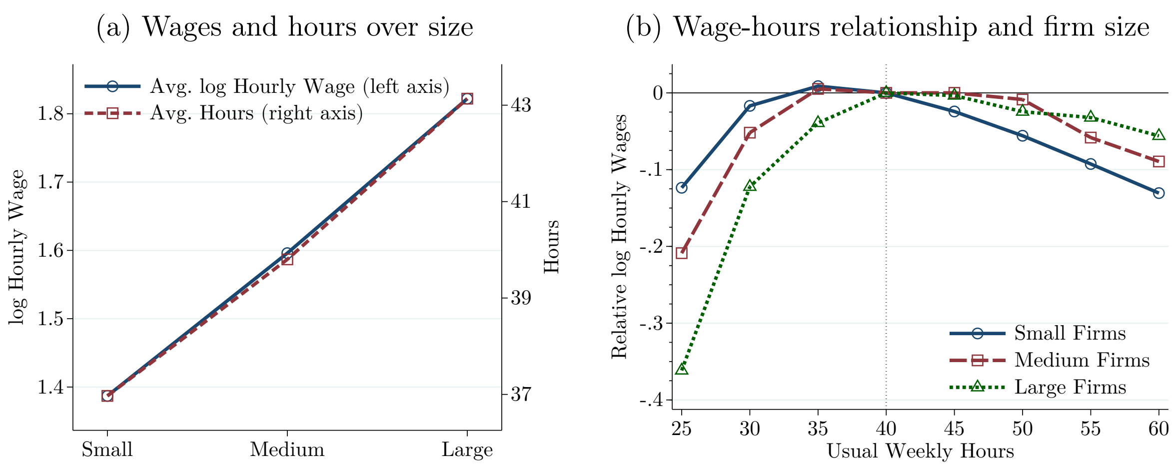

Figure 3: Motivating facts predicted by the baseline model

Notes: The left panel plots the log average wages (solid line, left axis) and average weekly hours worked (dashed line, right axis) for each size group in the model. The right panel plots, for each size category, the sum of coefficients ($\gamma_h + \theta_{e,h}$) estimated from Equation (2) (only controlling for size group, hours bin and their interactions) using simulated data from the model. Section 4.1 describes the construction of size categories in the model.

The model also replicates the (non-targeted) positive relationship between average hours worked and firm size as average hours increase monotonically with firm size. Quantitatively, the model predicts a steeper relationship between hours and size with an increase in average hours from the smallest to largest firms of 6 hours per week in the model compared to a 3 hours increase in the data. In the extended model of Section 5, this fit will improve when allowed for the combination of risk aversion and wealth accumulation.

Finally, the baseline model generates our third motivating fact on the size-dependent long-short hour penalties. To be comparable with our data analysis, we estimate Equation (2) using model simulated data only controlling for size group, hours bin and their interactions.[17] Figure 3b plots the sum of the coefficients of corresponding to size and hours-size interaction bins. Consistent with the data, the model features i) a hump-shaped relationship between hours worked and wages for each size group, and ii) a larger long (short) hour penalties in the smaller (larger) firms.

4.3 Accounting for the motivating facts

Having shown that the model successfully replicates our three motivating facts, we explore, in turn, the mechanisms which generate each fact.

4.3.1 Wage differentials between small and large firms

In the absence of heterogeneous preferences for firms, workers are sorted into firms based on wages alone; therefore they would flow into the highest-paying firm until wages are equalized across firms. Workers' random preference for firms (i.e. taste shock) introduces noise to their sorting and allows the model's equilibrium to sustain a positive wage gap between high- and low-productivity firms. Indeed, the larger the taste shocks, the less important wage differentials are to workers' sorting decisions, and the larger the size-wage premium. Table 5 show that the increase in average wage from small to large firms increases by about 36 percent when we double the standard deviation of the taste shocks.

4.3.2 Increasing hours over firm size

The existence of a size-wage premium results in longer average hours worked in larger firms. This is not because high-productivity firms demand more hours per worker, but rather

[17]Notice, since we interpret taste shocks and the value of leisure as being unobservable to the econometrician, there are no observable controls that we can include when estimating Equation (2). As such, there is no distinction between the unconditional and conditional relationship between wages and hours in this simple model. In the next section, we will introduce (observable) worker heterogeneity and a distinction between the two will arise.

Table 5: Size-wage premium, role of heterogeneity in tastes

	Baseline ($\sigma_\epsilon = 0.37$)	$\sigma_\epsilon = 0.74$
Small	0	0
Medium	0.209	0.281
Large	0.435	0.590

Notes: The table reports log average wages (relative to small firms) for each size group in a model. The first column reports the average under the benchmark calibration where $\sigma_\epsilon = 0.37$ while the second column reports the results for $\sigma_\epsilon = 0.74$ holding all other parameters of the model fixed.

because the high wages in those firms influence worker sorting and their endogenous hours decisions. First, workers with longer desired hours sort into larger firms where wages are higher. To understand this, notice that since income is the product of hours and wages, the income gains from working longer are higher in larger (higher wage) firms. Thus, workers who have longer desired hours will prefer, all else equal, to work in large firms to enjoy the higher total earnings. In the baseline model, where utility is linear in consumption, this is the only force shaping the sorting of workers. More generally, if utility is concave in consumption, a second opposing force pushes in the opposite direction and incentivizes workers with shorter desired hours to work in larger firms. To understand this, notice that employment in smaller (low-wage) firms will result in a lower income and, hence, higher marginal utility for any given level of hours. The marginal utility will be higher for workers with shorter desired hours; hence, these workers' value increases in income more than workers with longer desired hours and will be incentivized to sort into larger firms. In the extended model that we discuss below, both these forces are present, and together, they pin down the sorting of long-hour workers into large firms.

The sorting of workers, which can be considered an extensive margin channel, is not the only channel shaping the overall relationship between firm size and hours. There is another channel, an intensive margin channel, that also plays a role. The intensive margin channel shapes workers' hours regardless of their characteristics across firm size. As with sorting, the intensive margin of hours worked is shaped by two opposing forces, of which, only one is present in the baseline model. First, as wages are higher in larger firms, conditional on sorting, all workers are incentivized to work more to enjoy the higher earnings gain. In the baseline model with linear utility in consumption, this is the only force present and pushes all workers, regardless of characteristics, to work longer hours in larger (high-wage) firm. With non-linear utility in consumption, a second opposing force is present and influences the intensive margin channel. Namely, the marginal utility of the same worker in small firms will be higher due to a lower total income incentivizing them to work longer hours in smaller firms.

In the baseline model, both the extensive and intensive margin channels unambiguously push toward a positive relationship between hours worked and firm size. So, in equilibrium, average hours are increasing with firm size. To explore the relative importance of sorting (the extensive margin channel), we report, in the first two columns of Table 6 the model-implied size-hours relationship to a counterfactual relationship computed by assuming that all workers of the same characteristics work the same hours as their counterparts in small firms. This captures the size-hours relationship due only to differences in the sorting of workers based on their characteristics – effectively holding fixed the intensive margin channel. Since larger firms attract workers with longer desired hours, the counterfactual also features a positive relationship between firm size and average hours. Column (2) suggests that worker sorting accounts for 64 percent of the positive size-hours relationship in the baseline model while the remainder can be attributed to the intensive margin channel, namely that workers of the same characteristics work longer hours in larger firms.

Table 6: Average hours by firm size (log difference from small firms)

	Baseline	Only Sorting	$\gamma = 0.5$	$\rho = 1$
Small	0	0	0	0
Medium	0.074	0.043	0.046	0.066
Large	0.154	0.099	0.102	0.120

Notes: The first column of the table reports the log average hours (in difference to small firms) by firm size in the baseline model. The second column reports a counterfactual relationship that only allows selection between firms. To construct average hours in the counterfactual, we set hours worked for each worker with the same state variable – value of leisure – to be equal to the average hours in small firms. We then compute the average hours in each size group according to the resulting distribution of these states across firms. The third column plots the log of average hours (relative to small firms) across firm size when the utility function is specified as $u(c) = \frac{c^{1-\gamma}}{1-\gamma}$ with $\gamma = 0.5$ and all other parameters are as in the baseline calibration. The last column changes only the substitution parameter ρ to unity (full substitutability).

Though absent in the baseline model (with linear utility), it is also instructive to explore how the size-hours relationship is influenced by non-linear utility in consumption. As discussed above, non-linear utility in consumption introduces opposing forces in both the extensive and intensive margins. In a more general setup, the overall direction of each of these channels will depend on how quickly marginal utility decreases as income increases. To explore the importance of such "income effects", we allow the utility function to be given by $u(c) = \frac{c^{1-\gamma}}{1-\gamma}$ and compare the size-hours relationship when $\gamma > 0$ to the relationship implied by the baseline model (where $\gamma = 0$). Intuitively, higher relative risk aversion makes the opposing forces highlighted above stronger, causing long-hour workers to sort less into high-wage firms and also discouraging workers from significantly extending their working hours in the face of higher wages. The third column of Table 6 reports the relationship between firm size and hours when we set $\gamma = 0.5$, holding all other parameters fixed. In line with intuition, the steepness of the size-hours relationship decreases (to around half of the baseline) when we allow for marginal utility to diminish with consumption.

4.3.3 Size-variant hump-shaped wage schedule

Previously in Figure 3b, we showed that the model gives a similar implication to the data in the observed relationship between hours and wages for each size category. For a given size category, the model generates a hump-shaped relationship between wages and working hours due to complementarities between workers' hours. Such complementarity maximizes an individual worker's marginal productivity when she works the same hours as the rest of her production unit. Any hours worked above the level of her co-workers' in the worker's marginal productivity diminishing for those extra hours. By the same token, when a worker works shorter hours, she pulls her co-workers' productivity down, which is reflected in her marginal productivity, hence in her wages. Accordingly, there is a penalty for working shorter and longer than the usual hours in the firm. Figure 4 clearly illustrates the role of complementarities, by plotting the wage schedule faced by a worker in a firm of an intermediate productivity as ρ changes. As working hours become less complementary (higher ρ), the link between hours – within a production unit – becomes weaker, and penalties become less severe for both short and long hours.

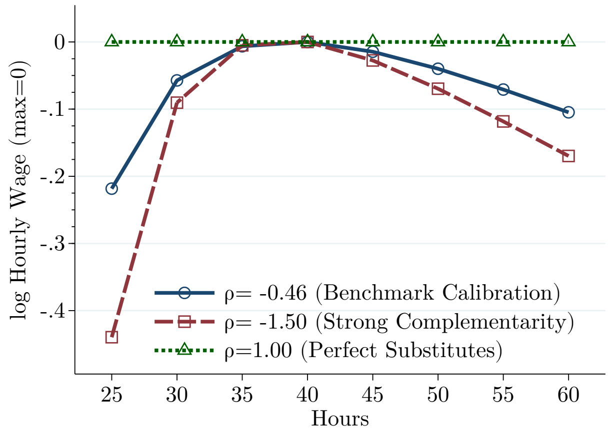

Figure 4: The relationship between wages and hours, role of complementarities

Notes: The figure plots the relative wage schedule by hours worked for different values of ρ, for a middle productivity grid of the medium size category (11th out of 24). To construct the relative wage schedule, we plot the logarithm of $w_j(l)$ in difference to its maximum level, against working hours. The three lines correspond to the benchmark model ($\rho = -0.46$), as well as the alternatives changing the substitutability parameter to values of 1 and -1.5. In the alternative computations, we do not recalibrate any other parameter.

Importantly, our motivating fact 3 is not only about the observation of a hump-shaped relationship for a given size category or overall, but also about how this relationship changes between the size categories. In order to zoom further into the mechanisms behind this feature in our model, we plot in Figure 5a the actual wage schedules faced by workers in different firms as a function of their working hours. For the figure, we picked three productivity levels, 4th, 11th and 20th points of our productivity grid of 24 points. These grids correspond to the mid-points corresponding to small, medium and large firms in our categorization.

As discussed in Section 4.3.2, the more productive firms in our model feature longer usual hours. In the specific levels of z that we use in Figure 5a, the average hours are 37.0, 40.0 and 42.8. The wage maximizing hours, $l_j^* = E_j(l^\rho; M_j^* \Phi_j)^{\frac{1}{\rho}}$, are 35.7, 38.5 and 41.5, respectively. Hence, the wage penalties for longer (shorter) hours decrease (increase) from the lowest productivity point to the highest.[18]

In order to better highlight the implication of the model for the size-specific wage schedules, we next show the average wages faced by workers of different hours within the three

[18] According to the wage equation (12), hourly wages relative to those of the wage-maximizing hours $\tilde{l}_j^*$ are simply functions of the hours relative to $\tilde{l}_j^*$: $\frac{w_j(l)}{w_j(\tilde{l}_j^*)} = \frac{\tilde{l}_j^*}{l}\left[\frac{1}{\rho}\left(\frac{l}{\tilde{l}_j^*}\right)^\rho + (1-\frac{1}{\rho})\right]$. Hence the wage schedules across hours plotted in Figure 5a is the same conditional on the (percentage) distance from the wage-maximizing hours.

size categories of firms. To do that, we first take the average wage for each of the nine hour grids within the three size categories of firms. Here we weight the firm productivities by their density within the size category for each hour grid. We then show the log of this average normalizing the maximum of each line to zero. (For instance, for the wages of 25 hour workers in large firms, we take the average of the 25 hour workers in firms of 15th to 24th productivity grids, i.e. those correspond to large firms in our calibration.) Figure 5b shows the outcome of this exercise, highlighting the wage schedules that penalize more the longer (shorter) hours in smaller (larger) firms. In other words, patterns shown earlier for specific grid points are representative of the wage schedules within the firm groups that these grids belong to.[19]

The wage schedules faced by workers in different firm size groups (Figure 5b) and the observed wage outcomes associated with different hours in each size groups of firms (Figure 3b) can potentially be different. In particular, covariates of wages and hours that are not controlled for in the exercise behind the latter figure would alter the observed wage-hour relationships away from the underlying wage schedules faced by workers. In our baseline model, there are no such covariates, which make the figures 5b and 3b look similar to each other. When we later extend our model to introduce worker heterogeneity in efficiencies, the two patterns will differ when we do not control for the heterogeneities in worker efficiency.

The amplification effect of hours complementarity. Complementarities between workers' hours amplify the positive hour-size relationship. To see this, suppose that, for some reason, a given firm features longer average hours than all other firms. Complementarity in hours implies that, compared to another firm, workers with longer desired hours will be penalized less in this firm while workers with shorter desired hours will be penalized more. As a result, long-hour workers will wish to sort into this firm (an extensive margin effect). This sorting will, in turn, result in similar workers in that firm to work longer (an intensive margin effect). Both these extensive and intensive margin effects, driven by complementarities, will amplify the positive relationship between hours and size.

[19]Notice, the smooth patterns in the wage schedules (Panel (a)) slightly change, when we take a weighted average to map into firm size categories because of our choice to discretized the grid for hours and firm productivities. Indeed, if we were to increase the number of grid points for the levels of hours and productivity, the average wages for larger firms would also be smooth.

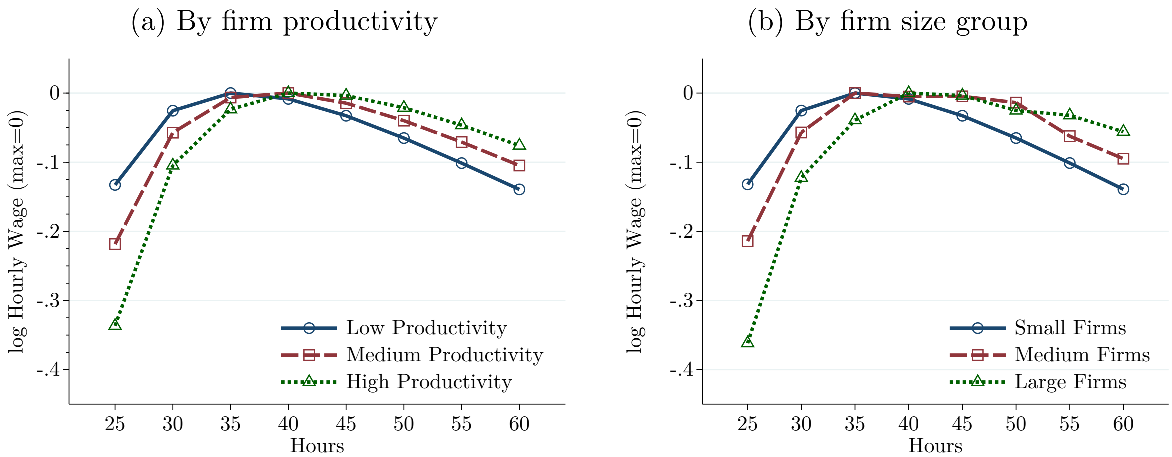

Figure 5: Wage schedules in firms

Notes: Panel (a) plots wages across hours in our baseline model for different productivity levels of firms, in logs and normalized to have the maximum at 0 for each line. The productivity levels are the 4th, 11th and 20th points of the 24-point productivity grid. Panel (b) plots, for each size category, average wages for different hours within the three size categories of firms. To do that, we first take the average wage over productivities for each of the nine hour grids within the three size categories of firms. These averages are weighted by the number of workers in each hours grid. We then show the log of this average normalizing the maximum of each line to zero. Section 4.1 describes the construction of size categories in the model.

We highlight this amplification by comparing the benchmark calibration to an alternative version of the model that does not feature complementarities. In the last column of Table 6 we report the size-hours relationship when we remove complementarities in hours by setting $\rho = 1$. First, even in the absence of complementarities, the sorting of workers is such that average working hours are increasing in firm size. However, in the absence of complementarities, this relationship becomes weaker with the difference in average hours between small and large firms being around one-fourth smaller than the baseline model.

5 Extended model

The baseline model presented in the previous section is rich enough to illustrate the main mechanisms that we argue to be behind the observed motivating facts. However, it abstracts from important channels that can be quantitatively and qualitatively important for these facts, as well as for other important aggregate implications of the model. In this section, we enrich the static baseline framework towards a more realistic and full-fledged dynamic model for two main reasons. First is to show that the main implications of the model are robust to such extensions. Second reason is to later use the extended model to quantify the role of model mechanisms in shaping wage and income inequality, as well as sorting for which the added features will be first-order relevant.

The first important change we introduce in the extended model is heterogeneity in the idiosyncratic worker efficiency. We interpret this efficiency as representing observable characteristics that can influence a worker's earnings and it can be thought to represent factors such as education, experience and industry among others. We model idiosyncratic efficiency as a persistent shock, $\Gamma_x(x'|x)$. This is important because, in addition to generating additional wage heterogeneity, the efficiency of workers can potentially affect the working hours of workers, and their sorting into small and large firms.[20]

The second (duo of) changes with the extended model are to allow for risk aversion in preferences for consumption and, to complement that, allow for wealth accumulation of workers. As we noted earlier, the race between income and substitution effects is key in shaping how hours correlate with wages and firm size; hence, setting up the model to have a more realistic preference structure is important. Moreover, to let the trade-off between income and substitution effects shape endogenously, we also allow for savings.

Finally, we include friction in the choice among different firms. In particular, we assume that only with probability $s < 1$ in each period, a worker can decide on the productivity type z of the firm that she will work for. Otherwise, she needs to remain with the employer type of the last period. With this feature, we aim to capture many different types of rigidities in employer selection that the workers face which can be important in shaping the observed dynamics in terms of transitions between employers, as well as volatility in wages and income.

Ultimately, the firms' labor aggregation now allows for heterogeneity of skills in the production unit in the same way as Shao et al. (2023):

$$L = \left(\sum_{x \in B_x} \sum_{l \in \mathcal{L}} x \mu(l,x) l^{\rho} \right)^{\frac{1}{\rho}} \left(\sum_{x \in B_x} \sum_{l \in \mathcal{L}} x \mu(l,x) \right)^{1-\frac{1}{\rho}}, \tag{13}$$

where $\mu(l, x)$ is the measure of workers of efficiency x, working l hours. Note that the efficiency units are introduced in the labor aggregation such that all workers of different

[20] The value of leisure shocks, v, are also assumed to be persistent shocks with a transition matrix Γ_{ν}.

types have similar complementarity links between each other. The efficiency of each worker only affects (i) the weight of her hours in shaping the "average" hours in the unit, and (ii) the scale of the labor aggregation.[21]

Workers' optimization is now dynamic due to frictions in the choice of firms and the wealth accumulation:

$$V(a,x,\nu,\boldsymbol{\epsilon}) = \max_j \{V_j^G(a,x,\nu)+\epsilon_j\}$$

where:

$$\begin{aligned} V_j^G(a,x,\nu) \quad &= \quad \max_{a'\geq 0, l\in\mathcal{L}} \frac{c^{1-\gamma}}{1-\gamma} - \nu\frac{l^{1+\phi}}{1+\phi} \\ &\quad\quad +\beta E_{x',\nu',\boldsymbol{\epsilon}'|x,\nu}\left[sV(a',x',\nu',\boldsymbol{\epsilon}')+(1-s)\left(V_j^G(a',x',\nu')+\epsilon_j'\right)\right] \\ \text{s.t} \quad &\quad\quad c = w_j(l,x)l + a(1+r) - a'. \end{aligned}$$

The optimal labor demand and wage schedule in the extended model is similar to that in the baseline model and now also capture the heterogeneous skills of workers.[22] These optimal policy functions and wages are given by,

$$\mu_j^*(l,x) \quad = \quad M_j^*\Phi_j(l,x) \tag{14}$$

$$w_j(l,x) \quad = \quad x\theta z_j L(\mu^*)^{\theta-1}\left[\frac{1}{\rho}\left(\frac{l}{\tilde{l}_j^*}\right)^{\rho-1} + \left(1-\frac{1}{\rho}\right)\left(\frac{l}{\tilde{l}_j^*}\right)^{-1}\right], \tag{15}$$

where, as before, M_j^* solves the extended model analog of (10) and the "weighted average"

[21] An alternative would be to have the non-linear aggregation only happen within efficiency units, i.e.

$$L = \sum_{x\in B_x} x\tilde{L}(x)dx, \text{ where } \tilde{L}(x) \equiv \left(\sum_{l\in\mathcal{L}}\mu(l,x)l^\rho\right)^{\frac{1}{\rho}}\left(\sum_{l\in\mathcal{L}}\mu(l,x)\right)^{1-\frac{1}{\rho}}.$$

While this alternative does not significantly affect our results, we choose not to use it for two reasons. First, it requires keeping track of the average hours for each worker efficiency-firm productivity pair (rather than only for each firm productivity). Second, we do not think of the complementarities in hours, or the coordination in tasks, as only happening within efficiency groups but potentially in the whole production unit with workers of different skills, occupations, and hierarchies interacting.

[22] We outline the definition of the stationary general equilibrium for the extended model in Appendix C.2.

Table 7: Parameters of the extended model

Panel A: Outside the Model

Parameter	Basis
$\phi = 2$	Standard
$\gamma = 2$	Standard
$r = 0.04$	Standard
$\theta = 0.67$	Labor Share of Output
$\rho = -0.46$	Shao et al. (2023)
$\kappa = 0.035$	See Appendix C.4

Panel B: Calibrated

Parameter	Basis	Data	Model
$\lambda = 1.97$	Employment Share (large firms)	0.66	0.63
$\nu_0 = 3.78$	Average Weekly Hours	41.7	41.7
$\sigma_\nu = 0.53$	SD log Weekly Hours	0.22	0.22
$\sigma_\epsilon = 0.69$	Wage gap, large firm to rest	0.29	0.29
$\rho_\nu = 0.683$	Autocorrelation of log Hours	0.69	0.69
$\sigma_x = 0.327$	SD log Wages	0.62	0.62
$\rho_x = 0.812$	Autocorrelation of log Wages	0.80	0.77
$s = 0.38$	Prob. switching size category	0.20	0.20
$\beta = 0.903$	Wealth to Income Ratio	2.5	2.5

Notes: Panel A reports parameters that are set outside of the model following standard values in the literature. Panel B reports the calibration of all parameters in the extended model. The last two columns in Panel B compare the data target and model implied value. The employment share of firms with over 100 employees is computed from the 2015 Business Dynamics Statistics. Measures of hours, wages and switching probability of size categories are calculated using the CPS. The data-target for the wealth-income ratio is from Erosa et al. (2016).

term $\tilde{l}_j^*$ incorporates heterogeneity in worker skill and is given by,

$$\tilde{l}_j^* \equiv \left[\left(\sum_{x \in B_x} \sum_{l \in \mathcal{L}} x \mu_j^*(l,x) l^\rho \right) \Big/ \left(\sum_{x \in B_x} \sum_{l \in \mathcal{L}} x \mu_j^*(l,x) \right) \right]^{\frac{1}{\rho}}$$

As with the simpler baseline model, the highest wage is achieved whenever $l = \tilde{l}_j^*$ for all the workers independent of their type x, and wages decrease as hours get further from $\tilde{l}$.

5.1 Calibration of the extended model

We calibrate this model following the same strategy (i.e. same targets) as in the baseline for the parameters that also appeared in that model. For the costs of rejecting the recruited workers, κ, we find that the minimum value we need to assume to sustain the symmetric equilibrium is 0.035, which is about 0.7 percent of the mean wage in this model, which is similar to the baseline.

Among the new parameters in the extended model, we set the risk aversion to the standard value of 2 and the annual interest rate on the savings at 4 percent where one model period is interpreted as one year. Panel A of Table 7 reports the parameters set outside the model. We calibrate all other parameters to match specific targets in the data.

Even though we assume a discrete Markov process for the value of leisure and idiosyncratic efficiency shocks, we parameterize the grids and the evolution of these shocks to resemble AR(1) processes:

$$\log(\nu_{i,t+1}) = (1-\rho_\nu)\log(\nu_0) + \rho_\nu \log(\nu_{i,t}) + \xi_{i,t}, \;\; \xi_{i,t} \sim N(0,\sigma_\nu)$$

$$\log(x_{i,t+1}) = \rho_x \log(x_{i,t}) + \psi_{i,t}, \;\; \psi_{i,t} \sim N(0,\sigma_x)$$

This way, we boil down the corresponding parameters to ρ_ν, σ_ν and ν_0 for the value of leisure, and ρ_x and σ_x for worker efficiency. We follow Tauchen (1986) to map these AR(1) processes to the discrete Markov processes assumed in the model. We calibrate the persistence of the value of leisure, ρ_ν to match that of the log-hours in the data. The dispersion and persistence of the efficiency shock, σ_x and ρ_x are calibrated to match those of the log wages.

The probability of being able to choose the employer type (as opposed to being forced to stay in your previous market), s, directly alters the transition rates between employers and sizes in our model. Accordingly, we calibrate this parameter targeting the probability of a worker changing her size category in our dataset, which is at 20 percent.

As in our baseline calibration, we set the size of the taste shocks, σ_ϵ targeting the wage gap between large firms and the rest. In the (dynamic) extended model, we assume that the taste shocks for the workplace are fully transitory. Tastes that are persistent or permanent would also fulfill their main objective of generating the size-wage premium. Yet, there are two technical reasons for assuming only transitory shocks. First, this allows us to avoid additional state variables in our model, and only operate with ex-ante value and probability functions before the realization of these shocks (see Appendix C.3). Second, the technical benefit of smoothing the policy and value functions of individuals and facilitating the convergence of the model solution is only achieved due to the transitory nature of the shocks. The randomness of the shock in each period from the individual point of view transforms workers' policy functions into probabilities. This would be absent with permanent shocks, and quantitatively less powerful with persistent shocks (conditional on the overall variation of tastes). Finally, the discount factor, β is set to match the average wealth-to-income ratio.

Panel B of Table 7 reports the values of calibrated parameters. It also shows that the model performs well in matching targeted data moments. Though we do not target it, the model also generates a positive covariance between wages and hours as is observed in the data (0.028 in the data and 0.005 in the model).

5.2 Implications of the extended model for the motivating facts

The calibrated version of the extended model matches the motivating facts in a similar fashion to the baseline model. Panel (a) of Figure 6 shows that average wages (as targeted moment) and average hours (an untargeted moment) both increase with firm size. The steepness of average hours over the firms size features a better fit to the data than the baseline model, with an increase from the small to large firms about 5 hours, whereas the data counterpart is around 3 hours.

Panel (b) plots the unconditional relationship between wages and hours by firm size implied by the extended model. Qualitatively, the hump-shaped relationship between hours and wages is still evident. The property that long (short) hour penalties are more severe in small (large) firms still holds and is still explained jointly by the (i) longer average hours in larger firms, and (ii) complementarities between working hours.

Quantitatively, the unconditional wage penalties are more severe relative to those in the data. This is, in part, because the unconditional wage-hours profiles do not control for observed worker heterogeneity, summarized by x, and thus reflect differences in average wages that arise due to workers choice of hours and sorting. Indeed we find that, uniformly across all firm sizes, lower efficiency workers are over-represented among the shortest and longest hours worked contributing to the lower unconditional average wages among these hours. Panel (c) of Figure 6, reports conditional wage-hours profiles which controls for efficiency x and shows relative wage penalties that are closer to those observed in the data though still not a perfect quantitative fit.[23] Similar to the baseline model, the conditional relationship between wages and hours in the model suggests wages that tend to peak at shorter (longer) hours in smaller (larger) firms – a pattern also evident in the data (Panel (b) of Figure 2). We next discuss the intuition behind how the income and substitution effects and their interaction with worker efficiency deliver these results.

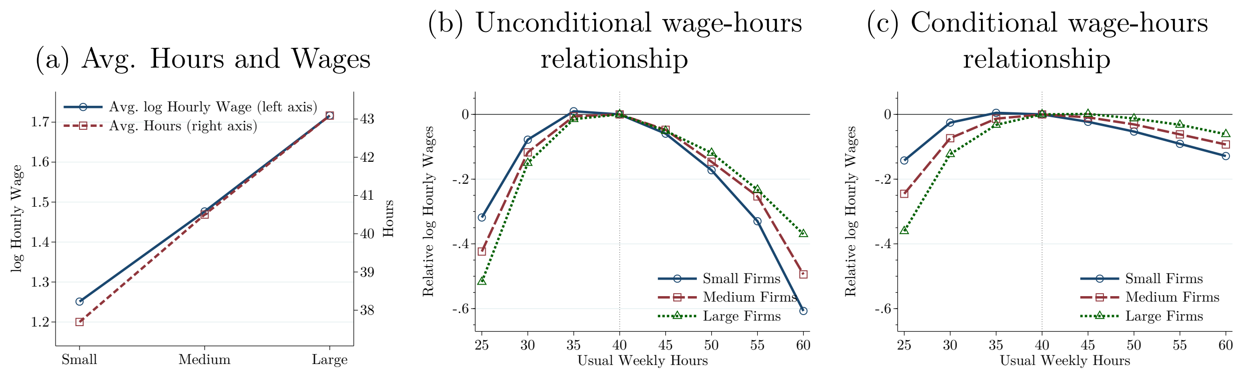

Figure 6: Motivating facts predicted by the extended model

Notes: Panel (a) plots the log average wages (solid line, left axis) and average weekly hours worked (dashed line, right axis) for each size group in the model. Panel (b) plots, for each size category, the sum of coefficients ($\gamma_h + \theta_{e,h}$) from Equation (2) using model simulated data while only controlling for firm size, hours bin and their interactions. Panel (c) reports the same sum of coefficients estimated when also including controls for worker efficiency x. To control for efficiency we include a dummy variable for each level of x similar to the construction of Panel (b) of Figure 2. Section 4.1 describes the construction of size categories in the model.

Income vs. substitution effects, and worker heterogeneity. In our simple model, which features only substitution effects, we noted that long-hour workers sort more into larger, higher-wage firms, and workers work longer hours in these firms than in small firms because the income gains from doing so are larger. In our extended model, where income effects are also present, substitution effects continue to dominate. While the concavity in preferences pushes for income effects, endogenous wealth accumulation brings the consump-

[23] We control for worker efficiency by including a dummy variable for each level of x in the model. This implicitly assumes that x is perfectly observable to the econometrician. In Appendix C.5, we experiment with constructing conditional wage-hours relationships by controlling for a noisy measure of x. We find qualitatively similar results with smaller conditional wage penalties compared to unconditional penalties. Alternative calibrations with higher values of ρ can deliver a closer fit of conditional model-implied wage penalties to the data with little qualitative change to the positive size-hours relationship or differences in wage-hours profiles across firms. We prefer to use the value of ρ estimated in Shao et al. (2023) as it is the outcome of an empirical estimation of a production function aimed at clearly identifying the elasticity of substitution in working hours rather than a value resulting from a model-specific calibration.

tion levels to a sufficiently high level to limit the changes in the marginal utilities from short to long hours. Hence, the model mechanisms highlighted in our baseline model continue to prevail.

A similar intuition also applies to workers' sorting across firms based on their efficiency. In particular, when substitution effects dominate, workers with higher efficiency prefer to sort into larger firms because the income gains from the size-wage premium increase with efficiency. At the same time, working longer hours brings higher income gains for high-efficiency workers. All else equal, having higher efficiency pushes workers to work longer hours and sort into larger firms. However, due to the persistence of efficiency shocks, high-efficiency workers also tend to have higher wealth, reducing their desired hours. Quantitatively, these two opposing forces (due to wealth accumulation and substitution effects) cancel each other out, and as a result, worker efficiency does not play a role in shaping the positive relationship between firm size and average hours.

Having said this, heterogeneity in worker efficiency does play a role in shaping the wage hours profiles in Figure 6b of the extended model. Indeed, we find that low efficiency workers are over-represented among both the shortest and longest hours, contributing to larger unconditional wage penalties for short and long hours. Intuitively, relatively stronger substitution effects lead to low efficiency workers working shorter hours. However, low efficiency workers also tend to have low wealth, and for these workers, substitution effects are dominated by income effects. So, there is a subset of low efficiency workers – that is, those with low wealth – that will work long hours. Ultimately, the combination of the relatively strong substitution effects and wealth accumulation results in an over-representation of low efficiency workers working very short and long hours. This leads to large unconditional wage penalties for working short or long hours. Comparing the conditional and unconditional wage-hours profiles in Figure 6 illustrates the role of worker efficiency in shaping unconditional wage penalties. Importantly, this sorting pattern by worker efficiency is common across firm sizes and does not contribute to differences in the wage-hours profile by firm size.

5.3 Sorting on hours

In our model, worker skills and desired working hours play a role in allocating workers to firms. While the literature has typically focused on sorting based on worker skill and firm productivity (see, for example, Eeckhout, 2018), we argue that sorting based on desired hours is an essential yet understudied factor. In this section, we document that our model's theoretical predictions about sorting based on hours are supported by empirical evidence.

Through the lens of the model, workers sort into firms based on skills x, preferences for workplaces, and desired working hours. Desired hours are influenced by the value of leisure and marginal utility and play a significant role in the allocation process. The relationship between hours and firm size, combined with complementarities in hours, causes workers with longer (shorter) desired working hours to prefer employment in larger (smaller) firms. Given this, a theoretical implication of the model is that workers who work fewer (more) hours than their coworkers will seek to transition to smaller (larger) firms where their hours will be more similar to their coworkers.

We test this implication by comparing the rate of worker transitions in the model to those observed in the data. Specifically, we test whether there are systematic differences in the rates at which workers move between different firm size groups based on their work hours relative to their coworkers. To do this, we use data from the CPS and construct worker transition rates across firm size bins by tracking respondents over 12 months (one period in the model). The unconditional rates of switching to a different size category with this calculation is 21 percent, which is a target in our calibration described in Section 5.1. Beyond that, the conditional switching probabilities that we study here are non-targeted moments of our model.

Before comparing the transition rates by groups, it should be noted that the CPS does not report the hours of coworkers. So, we can not observe how different a worker's hours are relative to their coworkers' – which is the theoretically relevant measure. Ideally, we would like to conduct this exercise using a dataset that includes information on all workers in a firm, such as in matched employer-employee datasets. However, to our knowledge, no matched employer-employee dataset in the US includes information on worker hours. Instead, we utilize the short panel in the CPS and approximate the difference between workers' hours relative to their coworkers by using the average hours in their firm size bin.

We classify workers as working shorter hours if their hours in period t, h_t, are at least 10% less than the average for their firm size group, $\bar{h}_{f,t}$. Similarly, a worker is classified as working longer hours if their hours, h_t, are at least 10% higher than the average for their firm size group, $\bar{h}_{f,t}$. Then, for a short- or long-hour worker, we measure the transition probability across firm size as the share of short- or long-hour workers in firm size f that transition to a firm of size f'.

Table 8: Transitions across firm size based on hours worked

Panel A: Transitions to Small (S) Firms

	Shorter Hours ($h_t < \bar{h}_{f,t}$)		Longer Hours ($h_t > \bar{h}_{f,t}$)		Ratio	
	Data	Model	Data	Model	Data	Model
M-S	0.12	0.06	0.07	0.03	0.54	0.50
	[0.09,0.15]		[0.05,0.08]		[0.34,0.74]	
L-S	0.04	0.05	0.02	0.03	0.47	0.60
	[0.02,0.06]		[0.01,0.02]		[0.09,0.85]	

Panel B: Transitions to Large (L) Firms

	Shorter Hours ($h_t < \bar{h}_{f,t}$)		Longer Hours ($h_t > \bar{h}_{f,t}$)		Ratio	
	Data	Model	Data	Model	Data	Model
S-L	0.09	0.21	0.19	0.26	2.02	1.24
	[0.07,0.11]		[0.16,0.22]		[1.49,2.56]	
M-L	0.22	0.22	0.30	0.26	1.35	1.18
	[0.18,0.26]		[0.26,0.33]		[1.07,1.63]	

Notes: The table compares transition probabilities in the data to those in the benchmark model. Columns 1-2 show workers who work at least 10% shorter hours than the average in their firm size group in period t; columns 3-4 show those who work at least 10% longer hours. The last two columns report the ratio of these transition probabilities. Panel A shows probabilities of transitioning into small firms (S) from medium (M-S) or large firms (L-S). Panel B shows probabilities of transitioning into large firms from small (S-L) or medium firms (M-L). Small firms have under 10 employees, medium firms have 10-99, and large firms have over 100. Transition probabilities in the CPS are based on respondents in the main sample tracked over a 12-month period, identifying their firm size group in periods $t-1$ and t. We match respondents that satisfy the same sample restrictions as in the primary sample over adjacent years resulting in data on around 123,000 unique respondents over two consecutive years. Each probability represents the share of workers moving from the size group in the matrix row in period $t-1$ to the size group in the matrix column in period t. The 95% confidence intervals are in brackets, computed using standard errors from 500 bootstrap simulations.

Table 8 compares (a subset of) these transition probabilities implied by the model to those observed in the data.[24]. Panels A and B report the transition probabilities of workers switching to small (S) and large (L) firms, respectively, where small firms are those with under ten employees and larger firms are those with over 100 employees. The first two columns of each Panel report the transition probabilities of shorter hour workers (those with $h_t < \bar{h}_{f,t}$) that transition, and the following two columns report the transition probabilities of longer hour workers (those with $h_t > \bar{h}_{f,t}$). In the last two columns, we report the ratio of transition probabilities for longer and shorter hour workers.

Focusing first on the data in Panel A, which reports transition rates into small firms, we find that workers are more likely to transition into small firms if they tend to work relatively shorter hours. For example, in the data, 12% of shorter hour workers in medium-sized (M) firms (with between 11 and 99 employees) transition to small firms compared to only 7% of longer hour workers. This difference is evident in the ratio of transition probabilities for longer and shorter hour workers as this ratio is less than one for all firm size bins – transitions into smaller firms are less likely for workers that work longer hours. Sorting based on hours would imply this pattern, predicting fewer transitions towards smaller firms when workers work longer hours. Indeed, the model predicts the same qualitative pattern, implying a ratio of transition probabilities under one for all firm sizes.

In contrast, sorting based on desired hours would predict the opposite pattern for transitions into larger firms with transitions being more likely when workers tend to worker longer hours. Panel B shows that this is what is observed in the data. For example, only 9% of shorter-hour workers in small firms transition to large firms compared with 19% of longer-hour workers in small firms. Consistent with the model, the observed ratio of transition probabilities into larger firms is above one for all firm sizes.

Overall, Table 8 reveals that working shorter (longer) hours is associated with more likely subsequent movements toward small (large) firms. This is consistent with sorting based on hours, and though the model's quantitative fit to (non-targeted) transition probabilities is

[24]For the sake of clarity, we omit elements of the diagonal of transition matrices as well as transitions to medium (M) sized firms with between 11 to 99 employees and the transition probabilities for workers that are classified as working neither shorter nor longer hours (that is, $h_t \approx \bar{h}_{f,t}$. The complete transition matrices in both the data and model are reported in Table A.1 of Appendix A

not exact, it is also qualitatively consistent with the model's predictions.[25] Having said this, the evidence in support of sorting based on hours should be interpreted as suggestive since, in the CPS, we do not observe information on coworkers' hours or other relevant firm characteristics that may shape transitions. Still, we view this as an important first step in establishing the importance of worker sorting based on hours.

6 Implications for inequality

In this section, we use the calibrated model for three exercises. First, we show how dispersion in workers' hours—especially differences in reference hours across firms, a previously overlooked factor—affects wage inequality. Second, we examine how firm heterogeneity and hour complementarities influence inequality in hours, wages, and income. Finally, we estimate how changes in the U.S. firm size distribution may have contributed to trends in earnings inequality.

6.1 The role of dispersion in hours

The theoretical framework we propose implies that hourly wages are determined by the interaction of worker skills, firm size and productivity, and workers' hours relative to their coworkers. These interactions are summarized in the expression for equilibrium wages in Equation (15).

Through the lens of the model, dispersion in worker hours (l) and the firm-level reference hours ($\tilde{l}_j^*$), a summary statistic capturing the usual hours within the firm, can shape the wage structure and overall wage inequality in the economy. Here, we use the model to evaluate the importance of dispersion in worker hours and firm-level reference hours in shaping wage inequality.

As a starting point, we evaluate the role of dispersion in relative hours ($l/\tilde{l}_j^*$) in shaping the wage dispersion, by computing the variance of log-wages when shutting the variation

[25] In Appendix B.4, we test a related implication of the model: We explore how worker's wages change when their hours do not change but they switch to a different firm size category. We show that, in the data, relative to staying in a firm of the same size category, workers that work relatively longer (shorter) hours tend to experience larger (smaller) wage gains when switching to larger firms, while switchers to smaller firms experience larger (smaller) wage losses. This pattern is also predicted by the model.

in this component. From the wage equation (15), this amounts to eliminating variances in the wage-penalty term $\left[\frac{1}{\rho}\left(\frac{l}{\tilde{l}_j^*}\right)^{\rho-1}+\left(1-\frac{1}{\rho}\right)\left(\frac{l}{\tilde{l}_j^*}\right)^{-1}\right]$. We find that variation in relative hours – and thus dispersion in wage-penalties – accounts for 4.6% of overall dispersion in (log) wages. This contribution is significant. For context, we find that variation in firm size and productivity (through the term $z_j L(\mu_j^*)^{\theta-1}$) – which have been studied in previous work, contributes approximately 11.1% to wage inequality.

While this demonstrates the importance of relative hours, it does not enable us to separately assess the role played by dispersion in the firm reference hours ($\tilde{l}$), which arises from the positive relationship between hours and firm size. Indeed, even if firms have the same reference hour, dispersion in worker hours alone would also contribute to wage inequality as a result of hours complementarity. Therefore, to examine the role played by dispersion in average hours across firms, we construct a counterfactual measure of wages assuming all firms have a common reference hour $\tilde{l}$ while leaving all other components in Equation (15) unchanged. Specifically, we set $\tilde{l}$ to be common across firms and equal to the weighted average implied by the population, thus eliminating the role of the positive relationship between average hours and firm size on the wages. We find that inequality in (log) wages is 0.395 under this counterfactual compared to 0.388. This suggests that the dispersion in firm reference hours *dampens* inequality by 1.8%. Further, since the overall impact of relative hours on wage inequality is 4.6%, dispersion in hours l *raises* overall wage inequality by around 6.4%.

The notion that dispersion in worker hours increases wage inequality is intuitive and stems directly from complementarity in hours. On the other hand, the finding that dispersion in firm reference hours ($\tilde{l}_j$) dampens wage inequality is not as straightforward. Indeed, since wage penalties increase with the distance between worker hours and reference hours, changes in reference hours of the firm can have heterogeneous impacts on worker's wages based on the hours they work. We focus on this heterogeneity next.

We find that imposing the common reference hours ($\tilde{l}$) has heterogeneous impacts for short and long-hour workers – increasing wage inequality among relatively short-hour workers while decreasing it for long-hour workers. Specifically, this counterfactual exercise indicates that variation in average hours dampens wage inequality within 25-hour workers by around

7% but raises it by 1.5% within 60-hour workers with little contribution for workers that work around 40 hours. In fact, in Appendix C.6 we show that these effects change monotonically across hours of work, gradually turning from a dampening effect for short hours to an amplification for long hours. More generally, this suggests that a positive size-hours relationship is consequential as it differentially impacts inequality measures among subsets of the population based on their hours. For example, though our model does not feature gender, female workers tend to work shorter hours which can be represented as a relatively high value of leisure (or high utility cost of working). Our results suggest that the differences in average hours across firms compress wage dispersions within this group by allowing workers to mitigate penalties due to their "non-standard" hours.[26]

However, this logic also applies equally to workers that have a relatively low value of leisure, yet in our analysis, wage dispersion amongst this group rises. Our finding that a common reference hour exacerbates inequality among workers with shorter hours while reducing it for those with longer hours hinges critically on our three motivating facts, most notably the positive correlation between hours and firm size. The positive hour-size relationship indicates that larger firms typically have longer reference hours than smaller ones. Therefore, imposing a uniform reference hour across the economy results in an increase in reference hours for small firms and a decrease for large firms. These changes in reference hours lead to wage changes that vary not only by workers' hours but also by the type of firms they are in.

We provide a simple illustration of these differential effects in Figure 7, based on the theoretical characterization of the wage profiles. The dotted and solid lines represent the wage profiles of a high-productivity firm (z_2) and a low-productivity firm ($z_1 < z_2$), respectively, in the benchmark equilibrium. The relative positioning of these two lines is determined by two factors: (i) the size-wage premium, which results in an upward shift from the solid (z_1) to the dotted (z_2) line, and (ii) the positive hour-size relationship, which causes the dotted line

[26] In this exercise, we compute counterfactual wages holding workers' hours and firms fixed. Letting workers instead optimize hours under the counterfactual wage schedules reduces wage inequality by 1.0%. A full model re-run—where agents optimize all choices under the counterfactual with common reference hours—reduces inequality by 1.9%. The differential effects on wage dispersion across hours groups hold in these alternative counterfactuals. See Appendix C.6 for further details on the description and the results of these exercises.

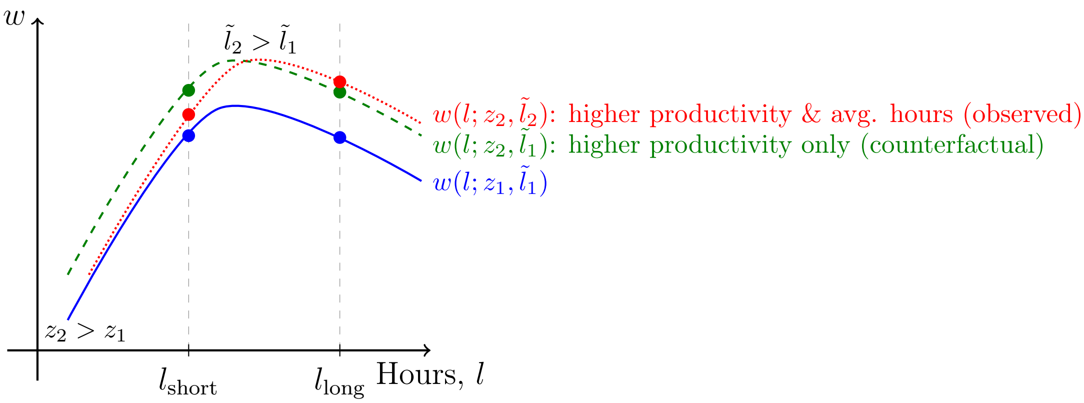

Figure 7: Wage schedules as firm productivity and average firm hours increase

to also shift to the right of the solid line. The differences between these two wage functions, evaluated at l_{short} and l_{long}, represent the wage inequality among workers who work short hours and those who work long hours, respectively.

Now consider a counterfactual exercise where the same reference hours are imposed across all firms. Specifically, we assume that the high-productivity firm will adopt the reference hour of the low-productivity firm, causing the dotted line to shift leftward to align with the dashed line, which represents the counterfactual wage profile.[27] This shift in reference hours results in wage increases for short-hour workers and wage decreases for long-hour workers in high-productivity firms. Consequently, it leads to increased wage inequality within the short-hour groups and reduced inequality within the long-hour groups.

Taken together, our model suggests that variation in relative hours raises overall wage inequality, while the differences in reference hours across firms, particularly the positive size-hours relationship, dampens overall wage inequality and has heterogeneous impacts on inequality based on the hours of workers.

6.2 The role of firm heterogeneity and complementarities

While instructive, the counterfactual exercise above does not clearly delineate the roles of two key features of our model – firm heterogeneity and complementarities in working hours

[27] It is important to note that the underlying intuition would also apply if the low-productivity firm were to adopt the reference hour of the high-productivity firm, or if both firms were to share a common reference hour situated between these extremes.

– in shaping inequality. While the first feature is emphasized in existing work, the second is novel. To quantify the contribution of each feature, as well as the endogenous responses of workers to them, we compare inequality in labor market outcomes across alternative model economies that abstract from one or both features. Specifically, we begin with the simplest version of the model, which features perfect substitutability in hours and one type of firm – the "plain" economy – and then sequentially introduce firm heterogeneity and complementarities in hours, building up to our benchmark model, which includes both.

	Variance in Hours		Variance in Wages		Variance in Income	
	Levels	log diff.	Levels	log diff.	Levels	log diff.
Plain Economy	0.091	-	0.344	-	0.400	-
w. Firm Heterogeneity only	0.096	+5.8%	0.374	+8.8%	0.467	+16.7%
w. Complementarities only	0.044	−51.9%	0.358	+4.1%	0.381	−4.7%
Benchmark Economy	0.048	−47.0%	0.388	+12.7%	0.447	+11.7%

Table 9: Comparison of inequality across alternative model economies

Notes: The models in the first three rows are one or two step deviations from the full extended model calibration given in Section 5, represented in the last row. In perspective from the last row: The first model ("Plain" economy) differs by setting (i) only one firm productivity level, z, hence, one common labor market ($J = 1$), and (ii) a labor aggregation that features full substitutability in hours, that is, $\rho = 1$; the second model differs only by assuming full substitutability in hours; the third model differs only by shutting down the firm heterogeneity. For the first and the third row, we choose the uniform firm productivity to match the average hours worked in the data. All other parameters are kept as in the benchmark calibration of the extended model.

In Table 9, we report the variance in (the log of) hours, wages, and income in each model economy. Starting with the "Plain" economy, we find that there is significant inequality, driven entirely by heterogeneity on the worker side, namely, dispersion in the value of leisure and worker skills – both of which are common across all model economies. To facilitate comparisons across models, we also report the percentage change in the variance of each outcome relative to the levels implied by the "Plain" economy. These percentage differences reflect changes in inequality attributable to model features beyond individual heterogeneity.

Focusing first on the introduction of firm heterogeneity (second row), we find that it increases inequality in hours, wages, and income by approximately 6, 9, and 17%, respectively, relative to the "Plain" economy. The rise in hours dispersion reflects standard income and substitution effects, as workers can now choose from a variety of workplaces that differ in their productivity. Under perfect substitutability of worker hours, wage inequality increases mechanically due to dispersion in firm productivity, while higher income dispersion follows directly from greater dispersion in both hours and wages.

On the other hand, introducing complementarities in hours only (third row) significantly *reduces* inequality in both hours and income, while *increasing* wage inequality. Hours inequality declines by 52%, an effect that is not only in the opposite direction of the impact from firm heterogeneity but also an order of magnitude larger. Intuitively, complementarities incentivize workers to coordinate their hours, as deviations result in wage penalties. As a result, complementarities *compress* dispersion in hours. Despite this compression in hours, wage inequality increases by 4.1% – about half the effect observed under firm heterogeneity – due to the introduction of wage penalties under complementarities. Indeed, when hours are imperfect substitutes, dispersion in relative hours $(l/\tilde{l})$ generates wage penalties that raise wage inequality. Importantly, the reduction in the variance of hours outweighs the increase in variance in hourly wages, leading to *lower* income inequality. The variance in income declines by 4.7% – an effect that moves in the opposite direction of, and is about one-third of the size of, the increase caused by firm heterogeneity.

Comparing the introduction of complementarities and firm heterogeneity in isolation reveals that both features play an important role in shaping inequality, and that abstracting from either mechanism can lead to misleading conclusions about the sources of inequality. While the importance of firm heterogeneity has long been understood in the literature, this analysis, together with the previous exercise, highlights the critical role played by complementarities.

Finally, our benchmark model (fourth row) incorporates both heterogeneous firms and complementarities in hours. Relative to the model without either feature, it exhibits lower dispersion in hours and higher dispersion in wages and income. These outcomes reflect the combined effects observed in the models with only one feature (second and third rows of Table 9). For instance, while firm heterogeneity puts upward pressure on the variance of hours, complementarities exert stronger downward pressure, resulting in a level of hours dispersion that is lower than in the "Plain" economy but higher than in the model without firm heterogeneity. Similarly, wage inequality rises when both firm heterogeneity and complementarities are present, as each independently increases wage dispersion. Income inequality also increases, though less than in the model without complementarities, since the compressive effect of complementarities partially offsets the inequality induced by firm heterogeneity.

The benchmark model also reveals modest interaction effects that act to *dampen* inequality in hours, wages and incomes. To understand this, note that under complementarities in working hours, firm heterogeneity not only raise dispersion in hours (through standard income and substitution effects), but also introduces heterogeneity in reference hours $\tilde{l}_j$, as workers can now choose among multiple firms. Compared to the model with only one firm, this dispersion in $\tilde{l}_j$ allows workers to endogenously sort into firms where their preferred hours are closer to the firm-specific reference hours $\tilde{l}_j$, thereby minimizing wage penalties. This sorting channel – absent in the model without complementarities – leads to smaller overall dispersion in hours. By a similar token, the same endogenous sorting mechanism dampens wage inequality, as workers select into firms where their hour choices are penalized less. As a result, this interaction also contributes to a reduction in income inequality.

Taken together, the model comparisons summarized in Table 9 highlight that both complementarities in working hours and firm heterogeneity are key mechanisms in shaping inequality in hours, wages, and income. Complementarities in hours are an order of magnitude more important than firm heterogeneity for shaping inequality in hours, while they are approximately half as important for wage inequality and one-third as important for income inequality.

6.3 A remark on changes in inequality and the composition of firms

Wage (and income) inequality in the US has trended upwards since the 1970s, while dispersion in hours has remained relatively stable. Over the same period, there has been a trend shift in the composition of firms towards older and larger firms, effectively shifting the firm size distribution over time. In light of these changes, the literature studying wage (and income) inequality has attributed little role to dispersion in hours in shaping trend changes in inequality while attributing much more to changes in the dispersion of firm productivity.[28]

However, our model-based results suggest that dispersion in average hours across firms – which can change naturally as the firm size distribution shifts – will also shape wage inequality even if the underlying distribution of worker hours remains unchanged. Using data on average hours across firm size bins in the CPS and information on the firm size distribution in the Business Dynamics Statistics (BDS), we find that, between 1995 and 2018, the standard deviation in average (log) hours across firms has declined by around 25%.[29]

Through the lens of our model, a decline in the dispersion of average hours across firms—holding everything else constant—reduces wage inequality by lowering overall dispersion in relative hours. However, since this decline is driven by a shift toward more productive (and larger) firms, it also increases wage inequality through greater productivity dispersion. We use the model to disentangle the relative contributions of these two forces to changes in wage inequality.

By comparing inequality changes across the two model economies, we quantify the effects of (i) the shift in firm size distribution—pushing inequality up—and (ii) the decline in average hours dispersion—which we hypothesize reduces inequality.

In the model without complementarities, shifting the firm size distribution (while holding hours dispersion fixed) leads to a 9.0% increase in wage inequality and a 7.2% increase in income inequality. Although firm-level average hours dispersion declines, it has no effect on wages in this setting, so inequality rises entirely due to greater firm productivity dispersion. In contrast, the benchmark model with complementarities—which allows average hours dispersion to influence wages—sees only a 2.8% rise in wage inequality and a 2.4% rise

[28] Heathcote et al. (2023) show that, since the late 1960s, the variance of (log) hourly wages has roughly doubled for both men and women while the variance in (log) annual hours has remained stable for men and declined by around one-third for women. Pugsley and Sahin (2019), among others, document changes in the firm size and age distribution over time in the U.S. For example, in the BDS, the employment share in firms with over 100 employees has increased from 59% in 1978 to 66% in 2015. Song et al. (2019) argue that most of the rise in earnings inequality is due to earnings at the firm level.

[29] To measure the standard deviation of average hours across firm size, we use the detailed firm-size bin categories reported in the CPS to construct average hours by size bin $\bar{l}_i$ where i denotes the size bin. We then match the size bins reported in the BDS to those in the CPS to get measures of the share of firms, s_i, in each size bin. The variance is then given by $(\sum_i s_i \bar{l}_i^2) - (\sum_i s_i \bar{l}_i)^2$. The observed decline in this variance is *not* driven by changes in average hours across firm size bins but rather due to a shift in the size distribution towards larger firms while differences in average hours across firm sizes have remained stable.

in income inequality. Since both economies share the same shift in firm size distribution, the difference in outcomes shows that falling dispersion in average hours helped offset rising inequality—consistent with the model's logic.

We interpret the results of this simple exercise as highlighting the importance of the empirical patterns and mechanisms we have highlighted here and believe that the development of more comprehensive employer-employee data sources would allow for a quantitatively rigorous assessment of this importance.

7 Conclusion

This paper studies the relationship between hours, wages and firm-level heterogeneity – specifically firm size. Using micro-data from the US, we document that workers' average wages and average hours increase with firm size, and, novel to the literature, that wage penalties for long (short) hours are larger in smaller (larger) firms.

Motivated by this evidence, we develop a general equilibrium model of heterogeneous firms and workers. Our framework generates a size-wage premium through heterogeneity in workers' preferences for the workplace. The size-wage premium leads workers willing to work longer hours to endogenously sort into larger (more productive) firms, as well as making similar workers work longer hours in larger firms. The existence of complementarities between workers' hours combined with the longer hours in larger firms results in less severe long-hour wage penalties and more severe short-hour wage penalties in larger firms – as observed in the data.

We use our model to understand how variation in the hours of workers and the average hours of firms shapes inequality. We argue that variation in workers' hours relative to a measure of a firm's usual hours significantly impacts inequality, raising wage dispersion to a similar extent as variation in firm productivity. Through the lens of our model, the positive size-hours relationship that we document dampens overall wage inequality and has heterogeneous impacts based on workers' hours. Our results underscore the importance of a previously neglected factor shaping the wage structure and inequality: dispersion in workers' hours and especially the differences in reference hours across firms.

References

Agell, J. (2004): "Why are Small Firms Different? Managers' Views," *Scandinavian Journal of Economics*, 106, 437–452.

Altonji, J. G. and C. H. Paxson (1988): "Labor Supply Preferences, Hours Constraints, and Hours-Wage Trade-Offs," *Journal of Labor Economics*, 6, 254–76.

Artuç, E., S. Chaudhuri, and J. McLaren (2010): "Trade shocks and labor adjustment: A structural empirical approach," *American economic review*, 100, 1008–45.

Autor, D. H. and D. Dorn (2013): "The growth of low-skill service jobs and the polarization of the US labor market," *American economic review*, 103, 1553–1597.

Bartelsman, E., J. Haltiwanger, and S. Scarpetta (2013): "Cross-country differences in productivity: The role of allocation and selection," *American Economic Review*, 103, 305–34.

Barth, E., A. Bryson, J. C. Davis, and R. Freeman (2016): "It's Where You Work: Increases in the Dispersion of Earnings across Establishments and Individuals in the United States," *Journal of Labor Economics*, 34, S67–S97.

Battisti, M., R. Michaels, and C. Park (2024): "Labor supply within the firm," *Journal of Labor Economics*, 42, 511–548.

Bick, A., A. Blandin, and R. Rogerson (2022): "Hours and wages," *The Quarterly Journal of Economics*, 137, 1901–1962.

Blau, F. D. and L. M. Kahn (2011): "Inequality and Earnings Distribution," ISBN: 9780199606061.

Brown, C. and J. Medoff (1989): "The Employer Size-Wage Effect," *Journal of Political Economy*, 97, 1027–1059.

Caliendo, L., M. Dvorkin, and F. Parro (2019): "Trade and labor market dynamics: General equilibrium analysis of the china trade shock," *Econometrica*, 87, 741–835.

Card, D., A. R. Cardoso, J. Heining, and P. Kline (2018): "Firms and labor market inequality: Evidence and some theory," *Journal of Labor Economics*, 36, S13–S70.

Chang, Y., S.-B. Kim, K. Kwon, and R. Rogerson (2020): "Cross-sectional and aggregate labor supply," *European Economic Review*, 126, 103457.

Checchi, D., C. García-Peñalosa, and L. Vivian (2016): "Are changes in the dispersion of hours worked a cause of increased earnings inequality?" *IZA Journal of European Labor Studies*, 5, 15.

Chetty, R., J. N. Friedman, T. Olsen, and L. Pistaferri (2011): "Adjustment costs, firm responses, and micro vs. macro labor supply elasticities: Evidence from Danish tax records," *The Quarterly Journal of Economics*, 126, 749–804.

Cubas, G., C. Juhn, and P. Silos (2022): "Coordinated Work Schedules and the Gender Wage Gap," *The Economic Journal*, 133, 1036–1066.

Eeckhout, J. (2018): "Sorting in the labor market," *Annual Review of Economics*, 10, 1–29.

Erosa, A., L. Fuster, and G. Kambourov (2016): "Towards a Micro-Founded Theory of Aggregate Labour Supply," *Review of Economic Studies*, 83, 1001–1039.

Flood, S., M. King, R. Rodgers, S. Ruggles, and J. R. Warren (2020): "Integrated Public Use Microdata Series, Current Population Survey: Version 7.0 [dataset]," *Minneapolis, MN: IPUMS*.

Headd, B. (2000): "The characteristics of small-business employees," *Monthly Labor Review*.

Heathcote, J., F. Perri, G. L. Violante, and L. Zhang (2023): "More unequal we stand? Inequality dynamics in the United States, 1967–2021," *Review of Economic Dynamics*, 50, 235–266.

Heathcote, J., K. Storesletten, and G. L. Violante (2014): "Consumption and labor supply with partial insurance: An analytical framework," *American Economic Review*, 104, 2075–2126.

Hirsch, B. T. (2005): "Why Do Part-Time Workers Earn Less? The Role of Worker and Job Skills," *ILR Review*, 58, 525–551.

Hofferth, L. S., M. S. Flood, and M. Sobek (2020): "American Time Use Survey Data Extract Builder: Version 2.8 [dataset]," *College Park, MD: University of Maryland and Minneapolis, MN: IPUMS*.

Idson, T. L. (1990): "Establishment size, job satisfaction and thestructure of work," *Applied Economics*, 22, 1007–1018.

Kuhn, M., J. Luo, I. Manovskii, and X. Qiu (2023): "Coordinated Firm-Level Work Processes and Macroeconomic Resilience," *Journal of Monetary Economics*, 137, 107–127.

Labanca, C. and D. Pozzoli (2021): "Constraints on hours within the firm," *Journal of Labor Economics*.

Lamadon, T., M. Mogstad, and B. Setzler (2022): "Imperfect Competition, Compensating Differentials, and Rent Sharing in the US Labor Market," *American Economic Review*, 112, 169–212.

Leung, D., C. Meh, and Y. Terajima (2008): "Firm size and productivity," Tech. rep., Bank of Canada.

McFadden, D. L. (1978): "Modeling the Choice of Residential Location," *Spatial Interaction Theory and Residential Location*.

Melo, P. C., D. J. Graham, and R. B. Noland (2009): "A meta-analysis of estimates of urban agglomeration economies," *Regional Science and Urban Economics*, 39, 332–342.

Montgomery, M. (1988): "On the Determinants of Employer Demand for Part-Time Workers," *The Review of Economics and Statistics*, 70, 112–17.

Morchio, I. and C. Moser (2018): "The Gender Pay Gap: Micro Sources and Macro Consequences," MPRA Paper 99276, University Library of Munich, Germany.

Oi, W. Y. (1974): "On the Economics of Industrial Safety," *Law and Contemporary Problems*, 38, 669–699.

Oi, W. Y. and T. L. Idson (1999): "Firm size and wages," in *Handbook of Labor Economics*, ed. by O. Ashenfelter and D. Card, Elsevier, vol. 3 of *Handbook of Labor Economics*, chap. 33, 2165–2214.

Pugsley, B. W. and A. Sahin (2019): "Grown-up business cycles," *The Review of Financial Studies*, 32, 1102–1147.

Rosen, S. (1986): "The theory of equalizing differences," in *Handbook of Labor Economics*, ed. by O. Ashenfelter and R. Layard, Elsevier, vol. 1 of *Handbook of Labor Economics*, chap. 12, 641–692.

Shao, L., F. Sohail, and E. Yurdagul (2023): "Are working hours complements in production?" *Journal of Economic Dynamics and Control*, 153, 104696.

Song, J., D. J. Price, F. Guvenen, N. Bloom, and T. von Wachter (2019): "Firming Up Inequality," *The Quarterly Journal of Economics*, 134, 1–50.

Tauchen, G. (1986): "Finite state markov-chain approximations to univariate and vector autoregressions," *Economics Letters*, 20, 177–181.

Wolpin, K. I. (1984): "An estimable dynamic stochastic model of fertility and child mortality," *Journal of Political Economy*, 92, 852–874.

Yurdagul, E. (2017): "Production Complementarities and Flexibility in a Model of Entrepreneurship," *Journal of Monetary Economics*, 86, 36 – 51.

Online Appendix for:

Labor Supply and Firm Size

Lin Shao Faisal Sohail Emircan Yurdagul

A Additional figures and tables

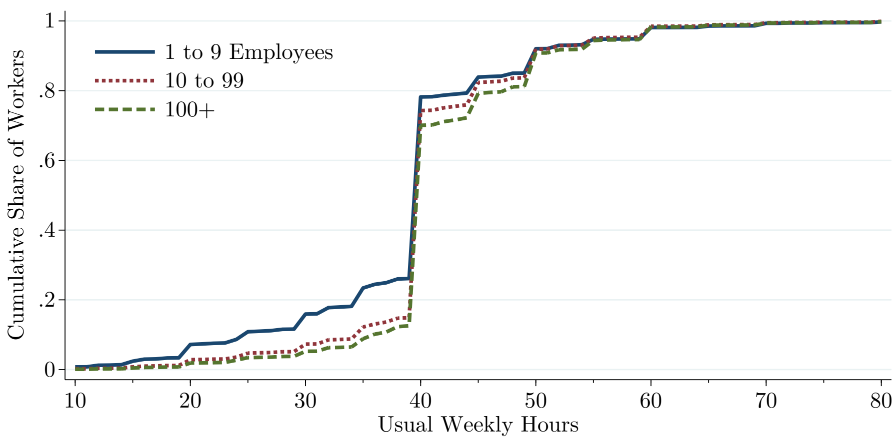

Figure A.1: Cumulative distribution of weekly hours worked by firm size

Notes: The figure plots the cumulative share of workers by their usual weekly hours worked and firm size. The dotted, vertical line marks 35 weekly hours worked.

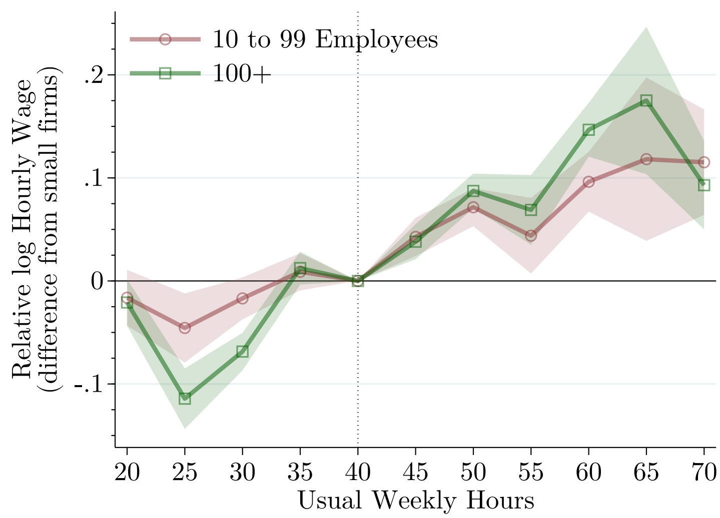

Figure A.2: Wage penalty relative to small firms, $\theta_{e,h}$

Notes: The figure reports the coefficient $\theta_{e,h}$ as estimated from Equation (2). The reference group for usual hours worked in the regression is workers that work $40-49$ hours. The reference group for size is firms with 1 to 9 employees. The shaded regions are the 95% confidence intervals.

Table A.1: All Transition Probabilities across firm size based on hours worked

	Data			Model		
	$h_t < \bar{h}_{f,t}$	$h_t \approx \bar{h}_{f,t}$	$h_t > \bar{h}_{f,t}$	$h_t < \bar{h}_{f,t}$	$h_t \approx \bar{h}_{f,t}$	$h_t > \bar{h}_{f,t}$
S-S	0.75	0.58	0.59	0.68	0.67	0.65
	[0.72, 0.78]	[0.56, 0.60]	[0.56, 0.62]			
S-M	0.16	0.24	0.22	0.11	0.10	0.09
	[0.13, 0.18]	[0.22, 0.26]	[0.19, 0.25]			
S-L	0.09	0.17	0.19	0.21	0.23	0.26
	[0.07, 0.11]	[0.16, 0.19]	[0.16, 0.22]			
M-S	0.12	0.08	0.07	0.06	0.04	0.03
	[0.09, 0.15]	[0.07, 0.09]	[0.05, 0.08]			
M-M	0.65	0.64	0.64	0.73	0.72	0.71
	[0.61, 0.70]	[0.62, 0.66]	[0.60, 0.67]			
M-L	0.22	0.28	0.30	0.22	0.24	0.26
	[0.18, 0.26]	[0.26, 0.29]	[0.26, 0.33]			
L-S	0.04	0.02	0.02	0.05	0.04	0.03
	[0.02, 0.06]	[0.02, 0.03]	[0.01, 0.02]			
L-M	0.09	0.09	0.08	0.11	0.10	0.09
	[0.06, 0.12]	[0.08, 0.10]	[0.06, 0.10]			
L-L	0.87	0.89	0.90	0.84	0.86	0.88
	[0.84, 0.90]	[0.87, 0.90]	[0.88, 0.92]			

Notes: The table reports transition probabilities in the data and in the benchmark model for workers that work at least ten percent shorter, similar and longer hours than the average hours in their firm size group in period t. Transition probabilities in the CPS are based on respondents in the main sample tracked over a 12-month period, identifying their firm size group in periods $t-1$ and t. We match respondents that satisfy the same sample restrictions as in the primary sample over adjacent years resulting in data on around 123,000 unique respondents over two consecutive years. Workers are grouped based on their firm size group and hours worked in the first period t. "S" indicates small firms with under 10 employees, "L" large firms with over 100 employees and "M" indicates medium-sized firms with between 11 and 99 employees. For example, "S-L" indicates transition rates from small to large firms. The 95% confidence intervals are reported in brackets and are computed using bootstrapped standard errors from 500 bootstrap simulations with full replacement.

B Data appendix

In this appendix, we report supplementary empirical results including robustness exercises using the CPS. We also explore the role of measurement error in hours for our empirical findings related to hours and replicate our motivating facts (at both the establishment and firm levels) using the Canadian Labour Force Surveys (LFS).

B.1 Additional evidence from the CPS

In this section we use the CPS to provide additional empirical results and robustness to our main motivating facts using the CPS.

Average hours by detailed firm size categories

For expositional clarity, in the main text we explores the distribution of by firm size using only three size categories. Here, we report the same results using the more detailed firm size categories included in the CPS. Indeed, for most of our sample, CPS respondents have reported their firm size in one of seven categories. These are, firms with i) 1 to 9 (under 10), ii) 10 to 24, iii) 25 to 99, iv) 50 to 99, v) 100 to 499, vi) 500 to 999, and vii) 1000+ employees. Between 2010 and 2017, the size categories of 10 to 24 and 25 to 99 were instead reported as 10 to 49 and 50 to 99.

Panel (a) of Figure B.3 reports the distribution of hours worked using these detailed firm size categories. As with the more course firm size categories, Panel (a) makes clear that even with more detailed firm size categories employees in smaller firms tend to work shorter hours compared to employees in larger firms. Focusing on the share of workers working the modal number of hours (40-44), we find that this share is lowest for firms with under 1 to 9 employees at 53% and tends to increase in firm size with 62% of employees in firm with 100 to 499 and 500 to 499 working 40 to 44 hours. The share is slightly lower for the largest firms with over 1000 employees at around 59%. Panels (b) and (c) focus on the distribution of short and long hours worked, respectively and show that the entire distribution of hours is shifted to the left for smaller firms and to the right for larger firms. Indeed, the share of workers that work under 25 hours is around 4% for firms with over 1000 employees and 12% for firms with 1 to 9 employees while the analogous shares of workers that work over 55 hours for these firms is 9% and 7%.

Consistent with the difference in hours distributions by firm size, the first column of Table B.2 reports the unconditional average of weekly hours worked and shows that workers in larger firms tend to work longer. Indeed, employees in firms with over 1000 employees work, on average, 42.2 hours per week while employees in firms with 1 to 9 employees work 39.6 hours, 10 to 24 employees work 41.1 hours and 25 to 99 employees work 42.0 hours. Overall, using more detailed firm size categories reveals that average hours worked are increasing, in a concave manner, with firm size.

Finally, we consider the conditional average of hours worked by detailed firm size categories by estimating the coefficient β_f in Equation (1) where f now comprises the finer firm size categories in the CPS. Table B.2 reports the estimates of β_f. The first three columns report this estimate when using all available data. To accommodate the change in firm size

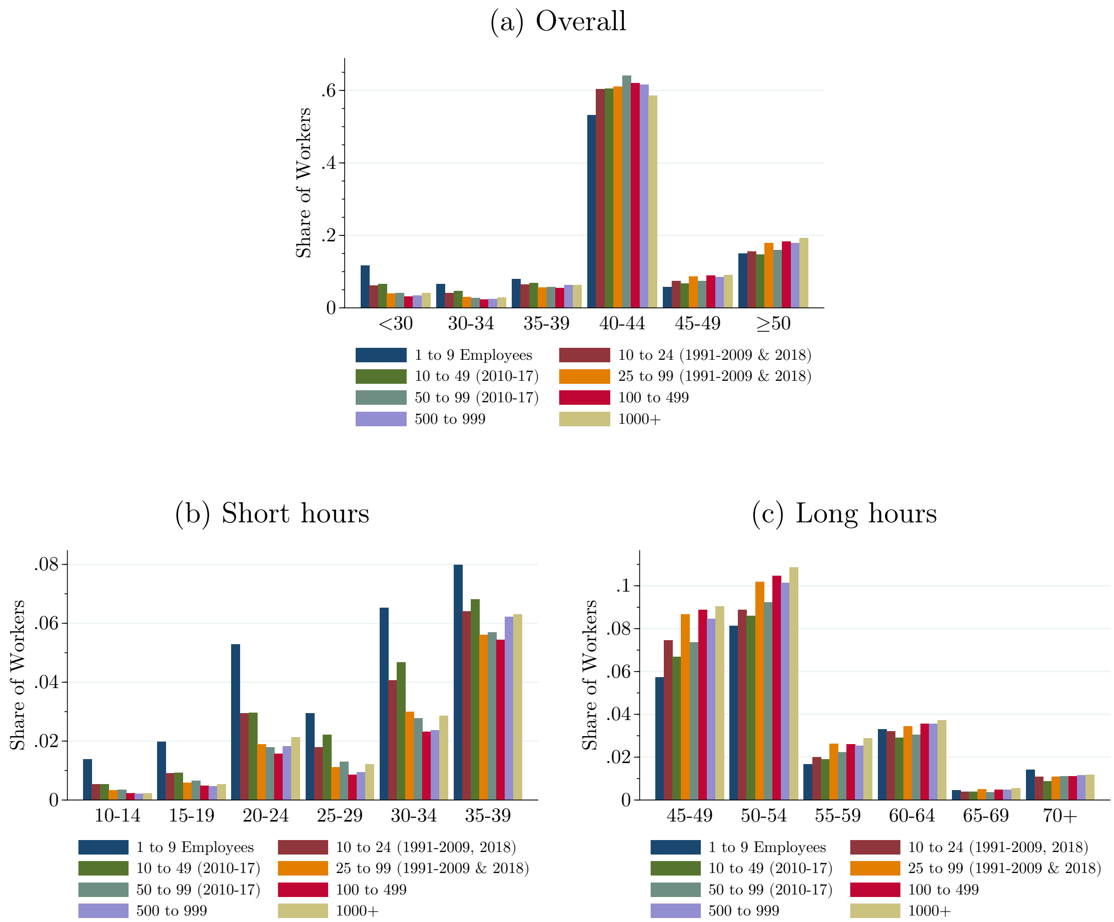

Figure B.3: Distribution of working hours by detailed firm size categories

Notes: The figure reports the share of workers by their usual weekly hours worked and detailed firm size categories reported in the CPS. The firm size categories 10 to 49 and 50 to 99 are only available in the CPS from 2010 to 2017 and replace the categories 10 to 24 and 25 to 99 in these years.

categories over CPS years, the fourth column reports β_f using the 2010 to 2017 CPS and the last column uses the 1991 to 2009 and 2018 CPS. Regardless of the size categories reported, we find that the coefficient β_f increases with firm size. For example, relative to firms with 1 to 9 employees, column (3) shows that employees in firms with 10 to 24 employees work almost 1 hour longer per week, 25 to 99 employees work 1 hour and 38 minutes longer, 100 to 499 employees work 2 hours 8 minutes longer, 500 to 999 employees work 2 hours and 11 minutes longer while workers in firms with over 1000 employees worker 2 hours and 13 minutes longer. These weekly differences are economically significant and amount to workers in firms with over 1000 employees working an additional 2.9 weeks longer per year than in firms with 1 to 9 employees.

It is important to note that though the increase in hours with firm size is monotonic, it is not linear. This can also be seen in the unconditional averages reported in first column of B.2. Instead, the coefficient β_f increases linearly between smaller firm size categories of 10 to 24 employees, and 100 to 499 but increases relatively modestly for larger firms. Indeed, while the estimates of β_f are statistically significantly different from each other for size categories 10 to 24 and 25 to 99 as well as 25 to 99 and 100 to 499, the estimates of β_f for 100 to 499 and larger firm sizes are not statistically significant from each other.

Table B.2: Detailed firm size categories and hours worked

	Uncond. Avg.	Conditional Avg. (rel. to small firms)				
		(1)	(2)	(3)	(4)	(5)
1 to 9 employees	39.5 hrs	-	-	-	-	-
10 to 24	41.0 hrs	1.412***	1.184***	0.986***	-	0.877***
		(0.058)	(0.055)	(0.055)	-	(0.060)
10 to 49	40.6 hrs	1.359***	1.151***	0.991***	1.185***	-
		(0.069)	(0.066)	(0.065)	(0.085)	-
25 to 99	41.9 hrs	2.354***	1.991***	1.662***	-	1.574***
		(0.053)	(0.050)	(0.051)	-	(0.057)
50 to 99	41.6 hrs	2.321***	1.980***	1.672***	1.827***	-
		(0.086)	(0.082)	(0.082)	(0.098)	-
100 to 499	42.2 hrs	2.779***	2.438***	2.155***	2.354***	2.051***
		(0.048)	(0.046)	(0.047)	(0.087)	(0.056)
500 to 999	42.1 hrs	2.649***	2.349***	2.224***	2.434***	2.116***
		(0.058)	(0.055)	(0.057)	(0.106)	(0.067)
1000+	42.2 hrs	2.712***	2.259***	2.235***	2.475***	2.113***
		(0.044)	(0.042)	(0.045)	(0.082)	(0.054)
Year, State FE	-	Y	Y	Y	Y	Y
Demographic Controls	-	N	Y	Y	Y	Y
4-digit Industry FE	-	N	N	Y	Y	Y
CPS Samples	-	All	All	All	2010-2017	1991-2009, 2018
N	819,295	819,295	819,295	819,295	247,476	571,819
R^2	-	0.014	0.107	0.137	0.133	0.142

Notes: The first column of the table reports the unconditional average of hours worked by firm size. The remaining columns report the coefficient β_f estimated from Equation (1) where the reference size category is firms with 1 to 9 employees. Standard errors are reported in parentheses. *** indicates statistical significance at the 1% level.

Hourly and salaried workers

As highlighted in Bick et al. (2022), workers that are paid by the hour experience a relatively stable penalty when working over 60 weekly hours. In contrast, salaried workers experience much larger penalties when working long hours above 60. Given this, our empirical finding that the long (and short) hour penalty varies with firm size could follow simply due to differences in the compositions of workers across firms. For example, if larger firms feature a higher share of hourly workers working longer hours than smaller firms then this could generate the relatively flatter long hours penalty. Figure B.4 tests whether this is the case by plotting the share of workers that are paid hourly by firm size and usual hours worked bins. The figure shows that the share of hourly workers declines as hours worked increase across all firm sizes. Further, the share of hourly workers is relatively similar across firm size bins for usual hours above 40, suggesting that the composition of workers is likely not the primary driver of the flatter long-hours penalty in larger firms.

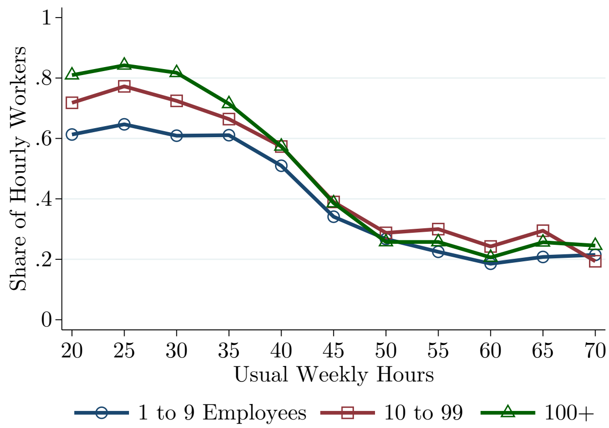

Figure B.4: Share of hourly workers by firm size and hours worked

Notes: The figure plots the share of workers that are paid by the hour, by firm size and usual hours worked. Data is from the outgoing rotation group (ORG) sub-sample in the pooled CPS sample. The ORG sub-sample make up around 25% of the pooled CPS sample and contains information on whether respondents are paid by the hour.

Having said this, larger firms feature a relatively higher share of short-hour, hourly workers than smaller firms. To concretely test whether differences in composition drive the differences in the short and long hour penalties by firm size, we re-estimate the regression in Equation 2 while also including an indicator for whether workers are salaried or paid by the hour. Figure B.5 reports the sum $(\gamma_h + \theta_{e,h})$ (Panel (a)) and the coefficient $\theta_{e,h}$ (Panel (b)) as estimated from this regression. Due to the smaller sample size when restricting attention to respondents with information on hourly or salaried status, we group usual hours worked into 10-hour bins and include 2-digit industry fixed effects. The reference group for usual hours worked in the regression is workers that work $40 - 49$ hours.

Panel (a) shows that the hump-shaped nature of the wage-hours profile remains unchanged when controlling for the salaried status of workers. Also persisting are apparent differences in the wage penalties between the smallest firm size categories and larger firms. This can be seen more clearly in Panel (b), which shows that medium and large firms exhibit more severe short-hours and less severe long-hour penalties compared to small firms. However, the difference in penalties between medium and large firms becomes much smaller when controlling for whether workers are paid by the hour – particularly for low levels of usual hours worked.

Taken together, this evidence suggests that differences in the composition of workers are not likely drivers of the differences in wage penalties observed across firm size bins.

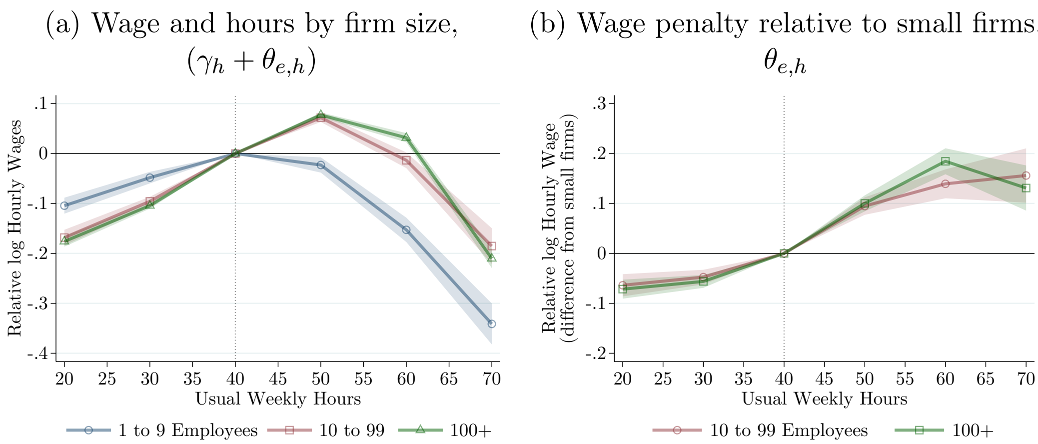

Figure B.5: Wage profiles by firm size and hours worked, controlling for hourly workers

Notes: The figure reports the coefficient $(\gamma_h + \theta_{e,h})$ in Panel (a) and $\theta_{e,h}$ in Panel (b) as estimated from Equation 2 with an additional indicator variable for whether or not a worker is paid by the hour. The reference group for usual hours worked in the regression is workers that work $40 - 49$ hours. The reference group for size is firms with 1 to 9 employees. The shaded regions are the 95% confidence intervals. Data is from the outgoing rotation group sub-sample in the pooled CPS sample.

Evidence by Industry and Occupation

Our primary empirical evidence focuses on evidence in the aggregate economy, here we report additional evidence on the relationship between firm size and hours worked as well as the wage-hours profiles by occupation and industry.

We begin by estimating Equation (1) on a restricted sample for a given occupation or industry. We find that worker hours are increasing in firm size even within occupations. Panel (a) of Figure B.6 reports the estimated coefficient β_f when restricting the sample to one of sample of each of one of 11 broad occupational categories. We construct these categories using 3-digit CPS occupation codes following Autor and Dorn (2013). Panel (a) shows that, with the exception of agricultural occupations, the estimated value of β_f is increasing in firm size for all others occupations. For example, compared to managers in firms with 1 to 9 employees, managers in firms with between 10 to 99 and over 100 employees work, respectively, 1.5 and 2.5 hours longer. Having said this, the difference in hours by firm size is not statistically significant for all occupations including protective services, technicians and food preparation and cleaning.

Panel (b) of Figure B.6 reports the coefficient β_f when we estimate Equation (1) for workers in a single industry. As with occupations, the hours worked in all industries are increasing with firm size.

Next, we estimate a variant of (2), which does not control for firm size, and restricts attention to a given occupation or industry. Specifically we estimate,

$$\log(w_i) = \alpha + \left(\sum_{h \in H} \gamma_h \mathbb{I}_{i,h} \right) + \delta X_i + \epsilon_i \tag{B.1}$$

We exclude controls for firm size to ensure large sample sizes when restricting smaller occu-

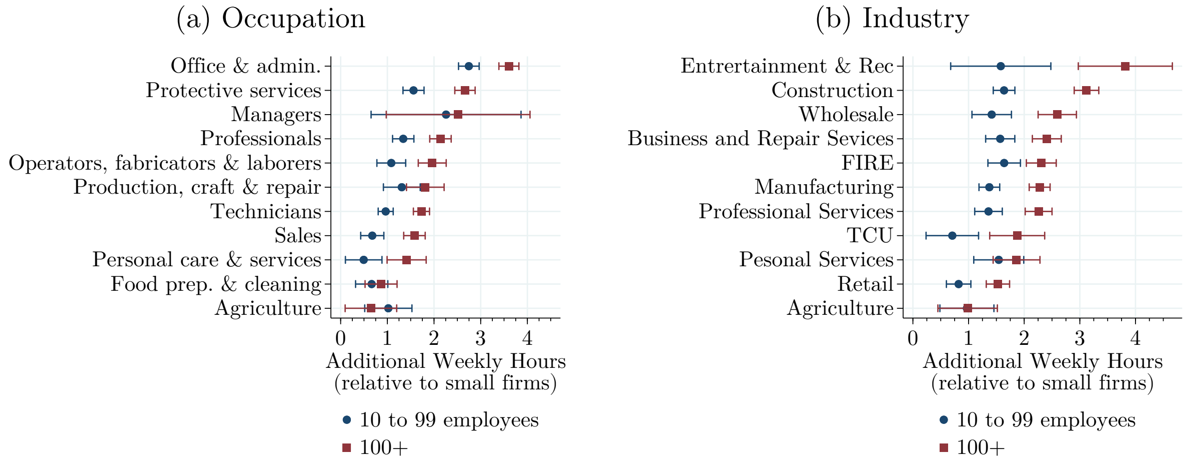

Figure B.6: Firm size and hours worked, by industry and occupation

Notes: The Figure reports the coefficient β_f from estimating equation (1) on a sample of workers in a given occupation (Panel (a)) and workers in a given industry (Panel (b)). The reference group for firm size is firms with 1 to 9 employees. The vertical markers indicate the 95% confidence interval for the estimated coefficient. Occupations were grouped into broad categories by following Autor and Dorn (2013).

pation or industry sub-samples. We also group hours in 10 hours bins for the same reason. The coefficient of interest is γ_h and it provide an estimate of the wage penalty for working relatively short or longer hours. Panel (a) of Figure B.7 reports this coefficient across occupations. We find significant variation in the short-hours penalty across occupations with technicians exhibiting a short-hours premium (of around 5%) and sales occupation exhibiting the largest penalty (around 30%) relative to workers that work 40 to 49 hours. Importantly, we document wage penalties in all but one occupation. There is substantially less variation in the penalty for working longer hours (60 to 69) with two occupations (sales and office & admin.) exhibiting long-hour penalties.

Panel (b) reports results across industries and, as with occupations, suggests significant variation in the penalty for working shorter hours. The long hours penalty across industry is less robust with many industries featuring modest penalties or a modest wage premium in the case of Retail and Finance, Insurance, and Real Estate (FIRE). While instructive, it should be noted that variation in either short or long hour penalties across industries may reflect differences in the hours that maximize wages. Indeed, we find that industries with a large penalty for working 20 to 29 hours also feature modest penalties for working 60 to 69 hours. This could simply reflect that the hours at which wages are maximized in these industries are relatively high.

Taken together, Figures B.6 and B.7 suggest that the new facts we emphasize in the aggregate are qualitatively consistent with evidence observed across most occupations and industries.

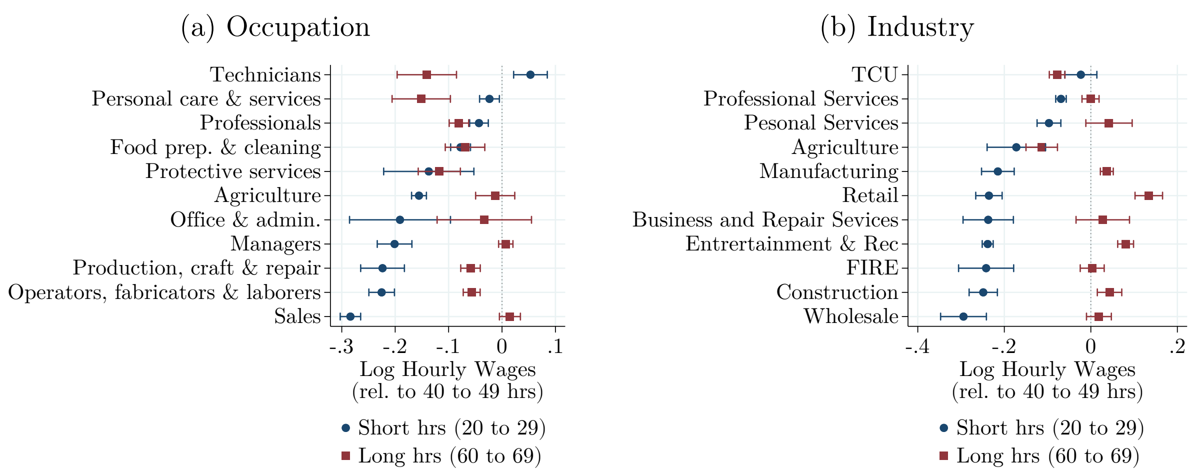

Figure B.7: Wage penalties for working short and long hours, by industry and occupation

Notes: The Figure reports the coefficient γ_h from estimating equation (B.1) on a sample of workers in a given occupation (Panel (a)) and workers in a given industry (Panel (b)). The reference group for hours is workers who worked between 40 and 49 hours. The vertical markers indicate the 95% confidence interval for the estimated coefficient. Occupations were grouped into broad categories by following Autor and Dorn (2013).

Controlling for occupations

Our primary empirical analysis does not include controls for worker occupations. We make this choice to capture the idea that production involves the interaction of workers employed in different types of occupations. Specifically, the complementarity in hours may be particularly salient between occupations rather than within occupations. Here we show that our findings from Section 2 are robust to controlling for occupation.

To do this, we estimate Equations (1) and (2) while also including an additional regressor that includes dummies for 3-digit occupations as recorded in the IPUMS variable `occ90ly`. The first and second columns of Table B.3 reports the coefficient β_f from estimating, respectively, a version of Equations (1) and (2) which controls for occupations. Controlling for occupations has little impact on the positive relationship between hours and firm size (first column) or on wages and firm size (second column).

Table B.3: The size-wage premium and the hours-size relationship, controlling for occupations

	Weekly Hours	log wages
10 to 99 Employees	1.217***	0.106***
	(0.043)	(0.003)
100+ Employees	2.044***	0.236***
	(0.043)	(0.003)
Year, State FE	Y	Y
Demographic Controls	Y	Y
4-digit Industry FE	Y	Y
3-digit Occupation FE	Y	Y
N	819,295	819,295
R^2	0.203	0.529

Notes: The first and second columns of the table report the coefficient β_f from estimating Equations (1) and (2), respectively, while also including controls for occupations. The reference size category is the smallest size firms. That is, firms with 1 to 9 employees. The reference hours bin is $40-44$ hours. Data is from the pooled CPS sample. Standard errors are reported in parentheses. *** indicates statistical significance at the 1% level.

Figure B.8 plots the relationship between hours and wages by firm size in the CPS as estimated from Equation (2) while also controlling for occupation. In particular, Panel (a) reports the sum of the coefficients γ_h and $\theta_{f,h}$ which captures the wage penalty of working outside of the 40-45 hours bin by firm size. Panel (b) reports the coefficient $\theta_{f,h}$ estimated from the same regression. The figure shows that controlling for occupation does not significantly alter the wage-hours relationship across firms. Panel (b) shows that the difference in relative penalties continue to be statistically significant.

(a) Wage-hours relationship by firm size (b) Wage penalties relative to small firms

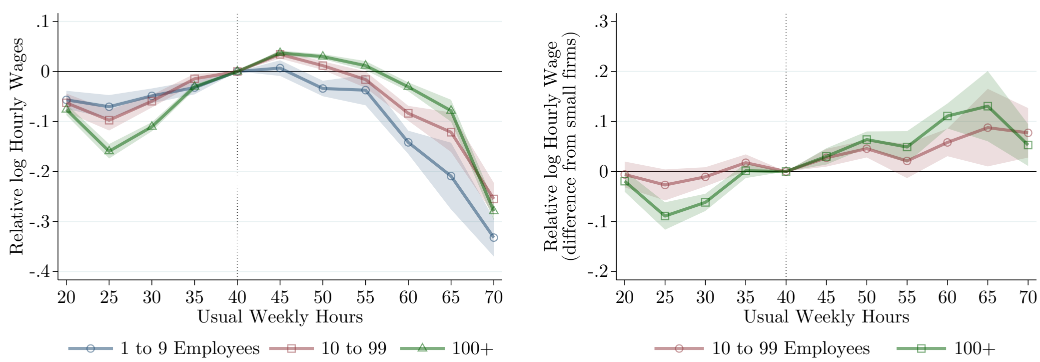

Figure B.8: The relationship between wages and hours, controlling for occupations

Note: Panel (a) reports the the sum of coefficients $(\gamma_h + \theta_{f,h})$ estimated from a version of Equation (2) which also includes controls for occupation. The reference group for usual hours worked in the regression is workers who work $40 - 44.9$ hours. The reference group for size is the smallest size category. Panel (b) reports the coefficient $\theta_{f,h}$ estimated from the same regression. The shaded regions are the 95% confidence intervals.

Average hours and firm size, additional controls

As an additional robustness check on our finding that average weekly hours worked are increasing in firm size, we estimate Equation (1) using additional sets of controls and restrictions. Table B.4 reports the results of these exercises. The first column reports the coefficient β_f when we include a dummy varibale indicating whether a worker works full-time, that is works at least 35 hours per week. Unsurprisingly, the coefficient on full-time status is larger with full-time workers working over 17 hours longer per week. Importantly, the coefficient β_f is strictly positive and increasing in firm size despite controls for full-time worker status suggesting that both full-time and part-time workers tend to work longer in larger firms.

The second column of Table B.4 estimates Equation (1) while also including a dummy variable for whether a worker is paid hourly or not. It shows that hourly workers tend to work longer but even after controlling for workers who are paid hourly, the coefficient β_f is increasing in firm size.

Finally, the third column estimates β_f while restricting the sample to only those sets of workers atht work between 35 and 45 hours (inclusive). We find that even within this narrow window of hours worked, average weekly hours are increasing with firm size with workers in mid-sized firms working 12 minutes longer and workers in larger firms working 18 minute longer per week than workers in small firms.

Table B.4: Firm size and hours worked, additional controls

	(1)	(2)	(3)
10 to 99 employees	0.168***	1.565***	0.210***
	(0.033)	(0.088)	(0.012)
100+	0.695***	2.405***	0.306***
	(0.033)	(0.087)	(0.012)
Full-Time employee	16.994***	-	-
	(0.032)	-	-
Paid Hourly	-	3.480***	-
	-	(0.050)	-
Year, State FE	Y	Y	Y
Demographic Controls	Y	Y	Y
4-digit Industry FE	Y	Y	Y
Hours Range	All Hours	All Hours	[35,45]
N	819,295	196,729	590,911
R^2	0.410	0.172	0.072

Notes: The table reports the coefficient β_f estimated from Equation (1) where the reference size category is firms with 1 to 9 employees. The first column includes an additional indicator variable which is equal to 1 if workers are full-time workers and work at least 35 hours, 0 otherwise. The second columns includes an indicator variable which is 1 if a worker is paid hourly and 0 otherwise. The last column estimates the equation by restricting attention to only those sets of workers that work between 35 and 45 hours (both inclusive). Standard errors are reported in parentheses. *** indicates statistical significance at the 1% level.

Actual Hours Worked

Our primary empirical evidence is based on reported usual weekly hours worked in the ASEC data. Here, we explore the robustness of our findings if we instead use the reported actual hours worked. Before reporting our results, it should be noted that a worker's firm size (in the previous year) is only available at an annual frequency in the CPS ASEC. In contrast, actual hours worked are reported in the CPS basic monthly surveys at a monthly frequency. Thus, it is important to ensure that the reported actual hours in the CPS Basic monthly data correspond to employment with the same employer referenced in the annual ASEC data. To ensure this, we merge the ASEC and Basic Monthly CPS surveys with information from the Job Tenure Supplement (JTS) of the CPS. The JTS has been conducted biennially since 1996 and includes information on a worker's tenure with their employer. This information allows us to restrict attention to workers who have been employed with the employer for at least one year to limit the cases where a reported employer in the basic monthly survey of the CPS may be different from the ASEC. In addition to this restriction, we apply the same sample restrictions as in our benchmark analysis and remove workers who report working fewer hours due to being on vacation. Due to the biennial nature of the JTS, the final sample is around one-eighth the size of our benchmark sample with around 100,000 workers with a correlation between usual and actual hours worked of 0.75.

We begin by exploring the relationship between firm size and actual hours worked. The first column of Table B.5 reports the unconditional average, by firm size, of actual hours worked. As with usual weekly hours, actual hours tend to increase with firm size. The remaining columns of the table confirm this by reporting the coefficient β_f from estimating Equation (1) using actual hours worked as the measure of hours. Workers in larger firms tend to report working 1.4 and 2.3 longer (actual) hours than workers in smaller firms.

Next, we investigate the relationship between wages, firm size and actual hours by estimating Equation (2) using actual instead of usual hours worked. First, we find that the firm size wage premium (for workers that work in the 40-hour bin) is qualitatively similar when using actual instead of usual hours worked. Specifically, when including all controls, we find that wages in medium and large firms are 10 and 23% higher, respectively than in small firms. Importantly, using actual instead of usual hours worked results in qualitatively similar differences in the wage-hours relationship by firm size. Panel (a) of Figure B.9 shows that, as with usual hours worked, the profile of wages across actual hours worked also varies with firm size. In particular, we observe a qualitatively similar hump-shaped relationship between hourly wage and actual hours worked for all firm sizes, with penalties for working relatively shorter and longer hours. Across firm size, we find that the penalties for working relatively shorter hours are not statistically significantly different across firm size bins – likely due to the much smaller sample size. However, the penalties for working longer hours in the largest firms do differ significantly across firm sizes, with the penalties for working longer actual hours in larger firms being less severe compared to medium and small firms. This can be seen more clearly in Panel (b), which plots wage penalties in medium and large firms relative to small firms.

Overall, using actual instead of usual hours worked has limited qualitative impact on our motivating facts. Given that we use annual earnings to derive hourly wages (together with weeks worked), the most relevant measure of worker hours is the usual weekly hours rather

Table B.5: Firm size and actual hours worked

	Uncond. Avg.	Conditional Avg. (rel. to small firms) (1)	(2)	(3)
1 to 9 Employees	39.8 hrs	-	-	-
10 to 99 Employees	41.8 hrs	2.159***	1.781***	1.502***
		[+2.8 weeks/yr]	[+2.3 weeks/yr]	[+2.0 weeks/yr]
		(0.149)	(0.142)	(0.146)
100+ Employees	42.6 hrs	3.214***	2.722***	2.578***
		[+4.2 weeks/yr]	[+3.5 weeks/yr]	[+3.4 weeks/yr]
		(0.136)	(0.131)	(0.143)
Year, State FE	-	Y	Y	Y
Demographic Controls	-	N	Y	Y
4-digit Industry FE	-	N	N	Y
N	98,879	98,879	98,879	98,879
R^2	-	0.015	0.100	0.127

Notes: The first column of the table reports the unconditional average of actual hours worked by firm size. The remaining columns report the coefficient β_f estimated from Equation (1) where the reference size category is firms with 1 to 9 employees and the measure of hours is actual weekly hours worked. The brackets report the additional number of weeks worked per year implied by the estimated regression coefficient. For example, an additional, relative to small firms, 2 hours worked per week over 52 weeks implies an additional 104 hours worked per year. Given that the median work week consists of 40 hours, this suggests an additional 2.6 (104/40) weeks worked per year. Standard errors are reported in parentheses. *** indicates statistical significance at the 1% level.

than actual hours worked, which may fluctuate more frequently. In the next section, we discuss potential measurement error in usual weekly hours worked.

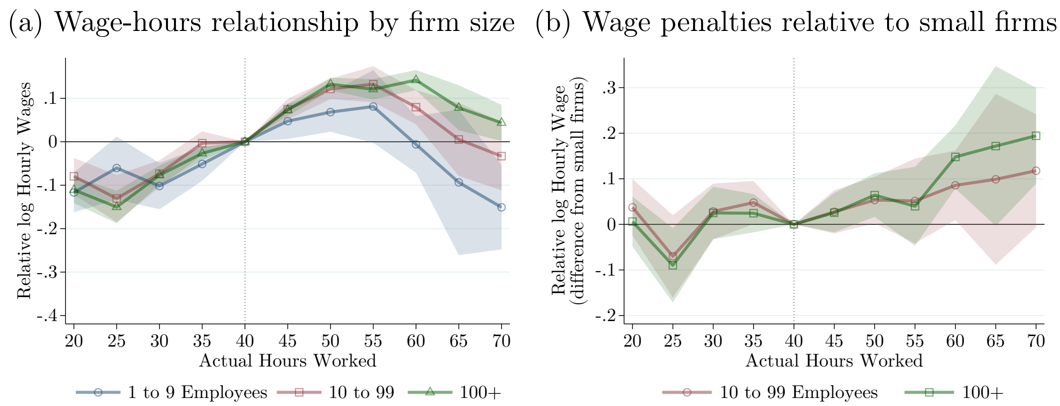

Figure B.9: The relationship between wages and actual hours worked

Note: Panel (a) reports the the sum of coefficients $\left(\gamma_h + \theta_{f,h}\right)$ estimated from Equation (2) where the measure of hours worked is actual hours worked. The reference group for actual hours worked in the regression is workers that work $40 - 44.9$ hours. The reference group for size is the smallest size category. Panel (b) reports the coefficient $\theta_{f,h}$ estimated from the same regression. The shaded regions are the 95% confidence intervals.

B.2 Measurement error in hours

To address concerns of measurement error in reported hours in the CPS and their impact on our motivating facts, we follow Bick et al. (2022), and merge the CPS data with data from the American Time Use Surveys (ATUS). The ATUS is a survey conducted since 2003 of a sub-sample of CPS respondents that asks respondents to complete a time diary detailing how their time use over a 24 hour period. Importantly, respondents of the ATUS are asked how long they have worked, which provides a high-quality measure of hours worked in a day. Our approach is to use firm size information from the CPS and the higher quality measure of hours worked in the ATUS to re-evaluate the relationship between average hours worked and firm size. Further, by using the difference between hours reported in the CPS and the ATUS we can create alternative wage-hours profiles by firm size to evaluate the extent to which mis-measurement in hours may be driving the differences in wage-hours profile across firms of different sizes.

We extract ATUS data from IPUMS, (Hofferth et al. 2020) and merge it with our main March CPS sample. Before describing our results, it is important to note that the ATUS does not elicit information on firm size and is conducted 2 to 5 months after a respondent's final (eighth) interview in the CPS. Given this, we must take information on firm size as reported in the CPS. A concern with this is that respondents may switch employment across firm size categories between the CPS and ATUS. To minimize this possibility, we only merge data from respondents' final CPS interview and remove all respondents who report a different 2-digit industry in the ATUS compared to the CPS. We further restrict attention to those who report working a single job at the time of the ATUS interview and those who have high-quality time use diaries (as perceived by the interviewer). Finally, we also removed all respondents who had completed their time diary on a public holiday (such as New Year's Day, Easter, Memorial Day, 4th of July, Thanksgiving, and Christmas). All other sample restrictions are the same as in the CPS. The resulting merged sample consists of around 2500 respondents – a significantly smaller sample size than our primary CPS sample.

Average hours by firm size We being by exploring whether the average reported hours in the ATUS are increasing with firm size as in the CPS. Since the ATUS includes hours worked for different days of the week, we construct a measure of weekly hours in the ATUS by multiplying the average daily hours reported in weekdays by five and multiplying the average daily hours reported during weekends by two and summing these together. The first two columns of Table B.6 compare the unconditional average weekly hours in the CPS and this measure of average weekly hours constructed using the ATUS. Compared to the CPS, average hours in the ATUS are higher for small firms with under ten employees and lower for larger firms with over 10 employees with only a modest difference in the (unconditional) average hours reported by employees in firms of 10 to 99 employees and firms with over 100 employees. This result is consistent with Bick et al. (2022) who show that respondents that report working shorter hours in the CPS tend to report longer hours in the ATUS, and those that report longer hours in the CPS report shorter hours in the ATUS. Given that larger firms tend to have a higher share of long hour workers, a natural consequence of this is that average hours in large (small) firms will be lower (higher) in the ATUS.

We conduct a more rigorous analysis of the relationship between hours worked in the

ATUS and firm size by estimating Equation (1) using daily hours reported in the ATUS. In addition to the controls described in the main text, we also control for whether respondents' time diary was completed on a weekday or a weekend (that is, Saturday or Sunday). The last three columns of Table B.6 report the estimate of β_f, which indicates the additional hours worked *per day* by workers in a firm of size f relative to workers in firms with 1 to 9 employees.

Table B.6: Firm size and hours worked in the CPS-ATUS

	Unconditional Avg.		Conditional Avg. (daily hours rel. to small firms)		
	CPS	CPS-ATUS	(1)	(2)	(3)
1 to 9 Employees	39.5 hrs	39.6 hrs	-	-	-
10 to 99 Employees	41.3 hrs	40.4 hrs	0.056	0.084	0.123
			[+0.4 weeks/yr]	[+0.5 weeks/yr]	[+0.8 weeks/yr]
			(0.280)	(0.281)	(0.295)
100+ Employees	42.2 hrs	40.7 hrs	0.142	0.201	0.279
			[+0.9 weeks/yr]	[+1.3 weeks/yr]	[+1.8 weeks/yr]
			(0.261)	(0.268)	(0.295)
Year, State FE	-	-	Y	Y	Y
Demographic Controls	-	-	N	Y	Y
4-digit Industry FE	-	-	N	N	Y
N	819,295	2,559	2,559	2,559	2,559
R^2	-	-	0.505	0.513	0.559

Notes: The first two columns of the table report, respectively, the unconditional average hours in the CPS and CPS-ATUS samples. The remaining columns report the coefficient β_f estimated from Equation (1) where the reference size category is firms with 1 to 9 employees and the dependent variable is daily hours worked in the ATUS-CPS merged sample. The brackets report the additional number of weeks worked per year implied by the estimated regression coefficient. For example, an additional, relative to small firms, 12 minutes (0.2 hours) worked per day over 5 working days and 52 working weeks implies an additional 52 hours worked per year. Given that the median work week consists of 40 hours, this suggests an additional 1.3 (52/40) weeks worked per year. Standard errors are reported in parentheses. ** indicates statistical significance at the 5% level.

When including all demographic and industry controls, the coefficient on firm size suggests that workers in medium-sized firms work around 7 minutes longer per day, which, assuming workers work five days a week, amounts to an additional 37 minutes of work per week (or an additional 0.8 weeks per year assuming the modal hours worked of 40). Employees in large firms (with over 100 employees) work around 17 minutes longer per day than those in small firms. This amounts to an additional hour and twenty-four minutes of work per week or almost 2 additional weeks worked per year. Qualitatively, the conditional averages in the merged sample imply an increasing relationship between firm size and hours work – consistent with the CPS – although the relationship is slightly flatter. Having said this, the coefficients in Table B.6 are not statistically significant. This is likely due to the much smaller sample size in the merged CPS-ATUS sample.

Taken together, it is encouraging that the higher quality time-use data is qualitatively consistent with an increasing relationship between firm size and hours.

Wages and hours by firm size Next, we revisit the relationship between wages and hours by firm size using time use data from ATUS. With the underlying assumption that working hours reported in the ATUS are reported without measurement error, we compare hours reported in the ATUS to hours reported in the CPS to quantify the degree of measurement error in the CPS. We then use this measure of measurement error to adjust the reported hours (and wages) of workers in the CPS and recompute wage-hours profiles by firm size.

Before conducting this exercise, it is important to note that classical measurement error, which would effect employees of all firms identically is not likely to change our finding that wage-hours profiles differ by firm size. Instead, if measurement error is not classical and is, for example, correlated with firm size then this might be a driver of the differential profiles we document in the main text. Bick et al. (2022) found that measurement error in hours, as proxied by comparing ATUS and CPS hours, was correlated with reported CPS hours such that those that reported shorter (longer) hours in the CPS tended to report longer (shorter) hours in the ATUS. Given our finding that workers in larger firms tend to work longer, such correlations of measurement error with reported would also lead to correlated measurement errors with firm size.

Table B.7 confirms this. In it, we report the difference between average weekly hours in the CPS with average weekly hours in the ATUS by hours worked bins (as reported in the CPS) and firm size. Given the low number of observations in the merged CPS-ATUS sample when grouping by hours bin and firm size, we group together workers that work between 10 to 24, 25 to 35, 36 to 44, 45 to 55 and greater than 56 hours. For workers in small firms, we group all workers working 45 hours and above into one bin as there are only 2 respondents that work above 56 in firms with 1 to 9 employees on a weekend.[1] As in Bick et al. (2022), we find that, across all firm sizes, workers that report shorter (longer) hours in the CPS report longer (shorter) hours in the ATUS. Importantly, the difference in ATUS and CPS hours does indeed differ by firm size with larger differences, particularly for workers that work over 55 hours, in larger firms. We also observe that employees of medium sized firms that report shorter hours in the CPS tend to understate their hours substantially.

Taking the differences reported in Figure B.7 as a measure of the degree of measurement error in the CPS, we adjust the reported usual weekly hours of workers by adding in this measure to the reported hours worked. Since reported hours tend to be bunched around multiples of five, we round the difference between ATUS and CPS to the nearest multiple of 5. For example, a worker in a firm with 10 to 99 employees that reported working 45 hours in the CPS has their hours adjusted down by 10. Similarly, a worker in a firm with under 10 employees that reports working 15 hours will have their hours adjusted upwards by 10. With these adjusted measure of hours, we also recompute hourly wages by taking annual income and dividing it by the product of weeks worked and adjusted weekly hours.

With these adjusted measures of hours and wages, we re-estimate Equation 2. The resulting wage-hours profile by firm size is reported in Panel (a) of Figure B.10. Consistent

[1]Recall, to construct a measure of average weekly hours in the ATUS we must utilize information on time use diaries over weekday and weekends.

Table B.7: Average difference in ATUS and CPS weekly hours by firm size and reported hours in the CPS

	1 to 9 Employees		10 to 99 Employees		100+ Employees	
	Diff.	N	Diff.	N	Diff.	N
10 to 24 hrs	8.8	17	7.8	21	1.7	60
25 to 35 hrs	-2.3	37	4.7	49	-0.3	109
36 to 44 hrs	-2.0	133	0.7	389	-0.5	1,094
45 to 55 hrs	-1.2	45	-8.3	110	-3.2	389
$>$ 56 hrs			-2.0	19	-10.8	87

Notes: The figure reports the difference in the merged CPS-ATUS sample between average weekly hours as reported in time use diaries and average usual weekly hours worked in the March CPS by firm size and by usual weekly hours bins in the CPS.

with Bick et al. (2022) correcting for measurement error in hours in the CPS does not change the overall hump shape of the wage-hours profile. Further, we find that the adjusted wage penalties for working shorter hours (under 35 hours) are more severe in large firms, as in the baseline estimates in Figure 2 of the main text. Having said this, the penalties work working shorter hours in medium sized firms tend to be are not consistently larger than those in small firms.

(a) Wages and hours by firm size, $(\gamma_h + \theta_{e,h})$

(b) Wage penalty relative to small firms, $\theta_{e,h}$

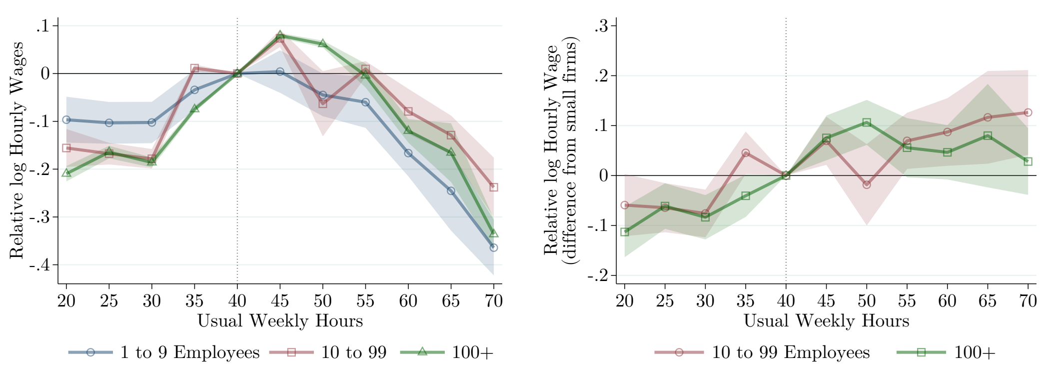

Figure B.10: Wage profiles by firm size and hours worked, adjusting for measurement error

Notes: The figure reports the coefficient $(\gamma_h + \theta_{e,h})$ in Panel (a) and $\theta_{e,h}$ in Panel (b) as estimated from Equation 2 by adjusted hours and wages reported in the CPS using the CPS-ATUS merged sample. The reference group for usual hours worked in the regression is workers that work $40 - 44$ hours. The reference group for size is firms with 1 to 9 employees. The shaded regions are the 95% confidence intervals.

Penalties for working longer hours also continue to differ by firm size, particularly for workers working between 40 and 60 hours. However, for even longer hours worked, there is little difference between the long hour wage penalties of workers in medium and large firms. Though, due to the nature of the adjustment of hours, there are much fewer workers that have hours adjusted above 56 and the estimated penalties for this region are much noisier for all firm sizes. Panel (b) shows this clearly by only plotting the wage penalties in medium and large firms, relative to small firms. Focusing on workers that work between 35 and 45 hours, the region of the hours distribution where the vast majority of workers work, we can see that adjusting for measurement error does not change our baseline empirical finding; penalties for working shorter (longer) hours are more severe in larger (smaller) firms.

Taken together, the results in this section confirm that measurement error in hours is not a likely driver of our empirical findings related to hours worked and firm size.

B.3 Evidence from the Canadian LFS

Our primary empirical analysis utitlizes data from the US. In this appendix, we document our three motivating facts by firm size using data from the Canadian Labor Force Surveys (LFS) between 1998 and 2018. Similar to the CPS, the LFS is a nationally representative survey containing detailed information on respondents' economic activity for the month they are interviewed such hourly earnings, usual weekly hours worked, and firm size. Firm size is recorded in one of four bins and for clarity, we combine the larger two size bins into one and report size in three categories; i) small (under 20 employees), ii) medium (between 20 and 100 employees), and iii) large firms (over 100 employees). Our sample starts in 1998 as this is the first year that information on establishment size is available in the LFS. Our treatment of the LFS data remains identical to that of the CPS. In particular, we restrict attention to respondents aged 25 and 64 who worked for a single private employer during the reference month. We exclude workers who usually work fewer than 10 hours per week and those that earned less than half the minimum wage.[2] Since 1997, the LFS has also reported establishment size and we conduct our empirical analysis below at both the firm and establishment level.

<u>Fact 1</u> Average hours increase with size.

We begin by showing that workers in larger establishments work longer hours than workers in smaller ones. To do this, we estimate Equation (1) using LFS data and report the coefficient β_f in Table B.8 where f represents establishment size (first three columns) and firm size (last three columns). As with the CPS data, usual hours worked are longer in larger firms however the difference in hours across establishment sizes is smaller in the LFS than in the CPS. For example, workers in both medium to large sized establishments work around 54 minutes longer per week than similar workers in the smallest establishments. Interestingly, average hours in medium sized *firms* tend to be higher than those in large firms though both are higher than average hours of workers in small firms.

<u>Fact 2</u> Average wages increase with size.

Next, we estimate Equation (2) using LFS data and report the coefficient β_f, which captures the size-wage premium, in Table B.9. Consistent with the data from the US, wages in the largest establishments and firms are indeed higher than those in the smallest ones. The LFS data indicates a wage premium of around 19% for workers in firms or establishments with over 100 employees compared to those with under 20.

<u>Fact 3</u> Long-hour (short-hour) penalty decreases (increases) with size.

Finally, we show that, consistent with data from the CPS, the short and long hours penalties also vary systematically by firm and establishment size in the LFS. Figure B.11 plots the relationship between hours and wages in the LFS as estimated from Equation (2). In particular, Panels (a) and (b) report the sum of the coefficients γ_h and $\theta_{f,h}$ which captures

[2] The minimum wage in Canada is taken to be an employment-weighted average of the minimum wage across provinces as reported by Statistics Canada.

Table B.8: Size and hours worked

	Uncond. Avg.	Establishment Size Conditional Avg. (rel. to small estabs.) (1)	(2)	(3)	Uncond. Avg.	Firm Size Conditional Avg. (rel. to small firms) (4)	(5)	(6)
1 to 19 Employees	37.2 hrs	-	-	-	37.0 hrs	-	-	-
20 to 99 Employees	38.5 hrs	1.347***	0.992***	0.896***	38.8 hrs	1.815***	1.393***	1.304***
		[+1.8 weeks/yr]	[+1.3 weeks/yr]	[+1.2 weeks/yr]		[+2.4 weeks/yr]	[+1.8 weeks/yr]	[+1.7 weeks/yr]
		(0.010)	(0.009)	(0.009)		(0.013)	(0.012)	(0.012)
100+ Employees	39.1 hrs	1.934***	1.357***	0.911***	38.5 hrs	1.520***	1.131***	0.979***
		[+2.5 weeks/yr]	[+1.8 weeks/yr]	[+1.2 weeks/yr]		[+2.0 weeks/yr]	[+1.5 weeks/yr]	[+1.3 weeks/yr]
		(0.009)	(0.009)	(0.009)		(0.011)	(0.010)	(0.011)
Year, Province FE	-	Y	Y	Y	-	Y	Y	Y
Demographic Controls	-	N	Y	Y	-	N	Y	Y
4-digit Industry FE	-	N	N	Y	-	N	N	Y
N	6,552,536	6,846,599	6,846,599	6,846,599	6,552,536	6,552,536	6,552,536	6,552,536
R^2	-	0.018	0.110	0.144	-	0.015	0.109	0.145

Notes: The table reports the coefficient β_f estimated from Equation (1) where the reference size category is the smallest size category. The first three columns report results where f represents establishment size categories. The last three columns report results where f represents firm size categories. All data is from the pooled LFS sample and firm size data is available starting 1998 while establishment size data is available starting 1997. Standard errors are reported in parentheses. *** indicates, respectively, statistical significance at 1% level.

Table B.9: The size-wage premium

	Establishment Size (1)	(2)	(3)	Firm Size (4)	(5)	(6)
10 to 99 Employees	0.094***	0.078***	0.082***	0.097***	0.078***	0.080***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
100+ Employees	0.242***	0.202***	0.192***	0.212***	0.176***	0.185***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Year, Province FE	Y	Y	Y	Y	Y	Y
Demographic Controls	N	Y	Y	N	Y	Y
4-digit Industry FE	N	N	Y	N	N	Y
N	6,846,599	6,846,599	6,846,599	6,552,536	6,552,536	6,552,536
R^2	0.159	0.283	0.346	0.146	0.273	0.342

Notes: The table reports the coefficient β_f estimated from Equation 2 where the reference size category is the smallest size category. The reference hours bin is $40 - 44.9$ hours. The first three columns report results where f represents firm size categories. The last three columns report results where f represents establishment size categories. All data is from the pooled LFS sample and firm size data is available starting 1998 while establishment size data is available starting 1997. Standard errors are reported in parentheses. *** indicates, respectively, statistical significance at 1% level.

the wage penalty of working outside of the 40-44.9 hours bin by firm and establishment size, respectively. The panels show that for both firms and establishment size categories, there exist hump-shaped relationships between hours and wages. Importantly, as with the US data, the short hours wage penalty in larger firms is much more severe than the penalty in smaller firms and establishments. For example, relative to working 40 hours, working 25 hours in the smallest establishments is associated with a 12% penalty while the analogous penalty in the largest establishments is 22%. Conversely, the penalty for working longer hours tends to be more severe in smaller establishments than in larger ones. For example, working 60 hours in small establishments (with under 20 employees) results in 10% lower wages (relative to working 40 hours) and the analogous measure for establishments with over 100 employees is only around 7%.

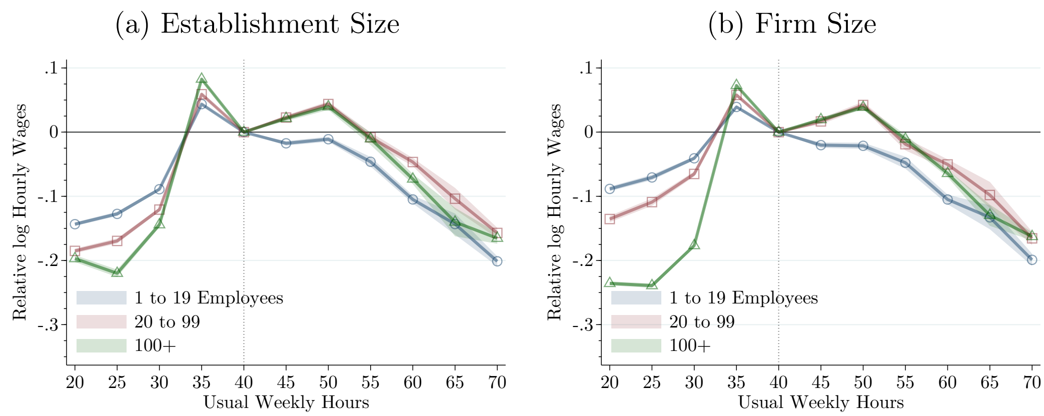

Figure B.11: The relationship between wages and hours, Canada

Note: Panels (a) and (b) report the the sum of coefficients $(\gamma_h + \theta_{f,h})$ estimated from Equation (2) using LFS data where f represents, respectively, firm size and establishment size. The reference group for usual hours worked in the regression is workers that work $40 - 44.9$ hours. The reference group for size is the smallest size category that is firms or establishments with under 20 employees. The shaded regions are the 95% confidence intervals. Data is from the pooled LFS sample.

Taken together, the analysis with the LFS data is encouraging as it confirms that our main empirical findings are not simply an artifact of the US data or driven by our use of firm size. Additionally, replicating our motivating facts using Canadian data is encouraging as it suggests that the US and Canadian economies are similar and may share similar fundamentals such as the substitution parameter ρ.

B.4 Longitudinal analysis

Our primary empirical evidence relies on cross-sectional relationships. However, our model has several implications and assumptions that can be tested using longitudinal data. As discussed in the main text (in Section 5.3), the ideal longitudinal dataset would be a matched employer-employee data set and would include detailed information on firms and worker characteristics including the distribution of coworker hours. Unfortunately, to our knowledge, no such data set exists for the US.

In the absence of such ideal data, in this section, we exploit the short panel dimension of the CPS and conduct two longitudinal analysis that aim to test the model's implications and assumptions. First, we construct measures of longitudinal wage penalties to argue that both complementarities and worker characteristics shape the cross-sectional wage penalties highlighted in this paper and in Bick et al. (2022). Second, related to our analysis of sorting on hours, we test a prediction of our model by exploring, in both the data and the model, how wages vary when workers change firm size but keeps hours fixed. Throughout, we also discuss the limitations of these longitudinal exercises using the CPS.

Longitudinal Wage Penalties

We interpret the hump-shaped relationship between hours and wages as evidence supporting the presence of complementarities in working hours. While this interpretation is shared with existing empirical work including Bick et al. (2022) and Shao et al. (2023), the (cross-sectional) hump-shaped relationship between hours and wages alone is not conclusive evidence for complementary in coworker hours. Indeed, such non-monotonicity could simply reflect worker sorting on unobservable individual or firm characteristics. For instance, using the Outgoing Rotation Group (ORG) of the CPS, Hirsch (2005) finds that the cross-sectional wage penalty for part-time workers (relative to full-time) workers can be explained by individual characteristics suggesting that the penalty for working shorter hours documented by Bick et al. (2022) and in this work may be driven by worker characteristics. To control for (fixed) unobservable characteristics, we extend the setup in Hirsch (2005) and study how a given worker's wages change when their hours change, that is, we estimate longitudinal wage penalties.

We begin by presenting results using the CPS ORG, which, though it does not contain information on firm size, has the advantage of being a larger sample size than the March CPS. We prepare the CPS ORG data in the same manner as in the March CPS. In particular, we restrict attention to those respondents aged between 25 and 64 that work at least 10 hours per week. We also drop any respondents that have imputed earnings or imputed hours in either the first or second year. For workers that are paid at a hourly rate (with no usual overtime), we define hourly wages as the reported wage. For all other workers, we define hourly wages as the ratio of weekly earnings and weekly hours. For workers whose hours vary we use actual weekly hours worked while for all others we use usual weekly hours worked. In addition to these restrictions, we also remove those that have implied hourly wages that is less than half of the minimum wage, those whose weekly earnings are top-coded and we trim top and bottom 5% of wage changes. Our final sample covers the period from September 1995 – the first month in which we can determine whether earnings were imputed – to December 2019 and includes almost 1 million unique respondents tracked over a calendar year. Using information on respondent's hours and wages over adjacent years we estimate the following variant of Equation (2),

$$\Delta \log(w_i) = \alpha + \left(\sum_{h' \in H} \gamma_{h'} \mathbb{I}_{i,h'} \right) + \delta X_i + \epsilon_i \tag{B.2}$$

where $\Delta \log(w_i)$ is the log difference in hourly wages of individual i between year t and $t-1$. h' is the new level of working hours in period t, and the indicator variable $\mathbb{I}_{h'}$ is equal to one if an individual's hours in period t are $h' \in H$ where H is the same partitioning of hours as in Equation (2). X_i is a vector of individual-level controls, which includes demographic controls, state, year, and industry fixed effects in both periods t and $t-1$. In addition to these controls, we also control for occupation in periods t and $t-1$. This specification is closely related to Hirsch (2005) but moves beyond the binary division of part-time and full-time work and allows for detailed hours changes. In particular, we can estimate this regression for each initial level of hours worked $h \in H$ so that the coefficient γ'_h captures the effect of moving from working h hours to h' hours (or equivalently, of changing their hours by $h' - h$). Since most workers work 40 hours, we choose the category $h' = 40$ as the reference category for hours changes and omit the coefficient γ_{40}.

Figure B.12 plots the coefficient $\gamma_{h'}$ as estimated for workers that are initially working 40 to 44 hours – the level of hours worked by most workers. For completeness, we report the estimates using both the ORG and March CPS samples. We find that workers that switch to working longer hours tend to experience wage penalties of similar size to those in our cross-sectional estimates suggesting that the cross-sectional wage penalty for working longer hours is likely not driven by unobserved worker characteristics. On the other hand, working shorter hours is generally associated with either modest increases or little change in wages, suggesting that the cross-sectional penalty for working shorter hours is likely to be driven by worker characteristics. This result is consistent with Hirsch (2005) who found that transitions from full-time to part-time work were associated with modest wage gains and also with the individual fixed effects estimates reported in the Appendix of Bick et al. (2022) which are most closely related to Figure B.12.

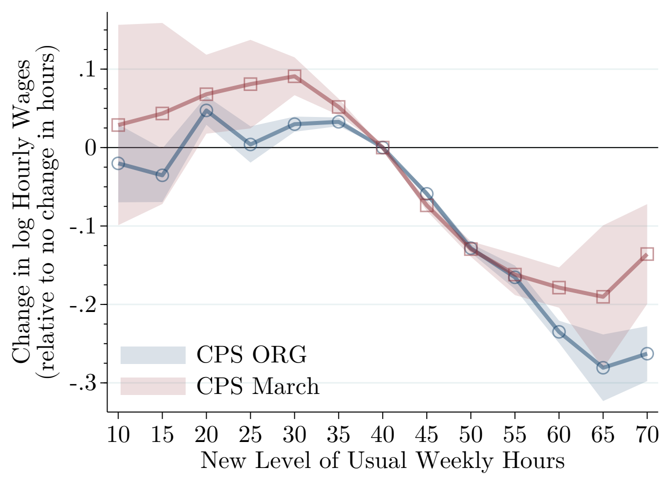

Figure B.12: The longitudinal relationship between wage changes and hours

Notes: The figure reports the estimated coefficient $\gamma_{h'}$ from Equation (B.2) for workers that were initially working 40-44 hours in the CPS ORG and CPS March samples. The shaded regions are the 95% confidence intervals.

At first, these findings suggest that complementarities in hours may not be the main driver of wage penalties for working shorter hours with individual characteristics also being important. However, in order to test for the presence of complementarities, simply considering changes in a given worker's hours is not sufficient. The wage penalties implied by complementarities depend on *deviations* from one's hours relative to their coworkers and not simply the absolute levels of hours. Thus, the theoretically relevant measure to test for complementarities is the change in a worker's hours relative to the hours of their co-workers. Complementarities would suggest that workers who move further away from their firm's reference hours will experience lower wages, while those that move closer should experience wage gains. As stated above, to our knowledge, no dataset available for the US would allow us to test this prediction. Shao et al. (2023) use Canadian data to test this prediction of complementarities and find strong evidence supporting complementarities.

In the absence of the ideal data, we can attempt to make some progress by using information on firm size in the March CPS data and the fact that average hours worked across firm size bins differs to test a related prediction of complementarities in hours. Specifically, given that average hours are longer in larger firms, the presence of complementarities would predict that the wage-maximizing reference hours are longer in larger firms compared to smaller firms. Thus, under complementarities, a worker that decreases their hours, say, from 40 a week to 30 but switches to a larger firm will experience relatively smaller wage gains (or equivalently more severe wage penalties) compared to a worker that changes their hours in the same manner but instead switches to a smaller firm. The idea behind this is that when a worker reduces their hours and switches to a smaller firm, they are more likely to be moving closer to the reference level of hours – reducing any penalties – while if they decrease their hours but switch to a larger firm, they are more likely to be moving further away – amplifying any penalties. Under the same intuition, workers that increase their hours, say from 40 to 50, should experience larger wage gains (or equivalently smaller wage penalties) when switching to a larger firm.

To test this prediction, we can re-estimate Equation B.2 using data from the March CPS separately for workers who switch to working for a smaller firm and workers who switch to working for a larger firm. We consider a worker to have switched to working for a smaller firm if they are i) initially working for a medium-sized firm (10 to 99 employees) and switch to a small firm (1 to 9 employees) or ii) initially working for a large firm (100+ employees) and switch to either a medium or small firm. Similarly, we consider a worker to have switched to a larger firm if they are i) initially working for a medium firm and switch to a large firm or ii) initially working for a small firm and switch to either a medium or large firm. By comparing the coefficient γ_h' for workers that switch to smaller firms and those that switch to larger firms, we can test whether an implication of complementarities in hours is borne out in the data. More specifically, is it the case that γ_h' is smaller (larger) for switchers to smaller firms when h' is low (high) relative to switchers to larger firms?

Figure B.13 reports the results from estimating Equation B.2 separately for workers that were working 40-49 hours in period $t-1$ and switched to either smaller or larger firms. Due to the smaller sample size in the March CPS generally and of the group of firm-size switchers, we group hours in period t into bins of 10 hours. First, regardless of whether a worker switches to a larger or smaller firm, we find that the estimated longitudinal wage penalties for firm-size switchers feature a more prominent hump-shaped relationship, implying penalties for switching to shorter or longer hours. This differs from the estimates in Figure B.12 and highlights one of several challenges with longitudinal analysis in the CPS. Namely, workers that switch hours – particularly those that also switch firm size and thus employers – are likely to be strongly selected. For example, workers who switch to working shorter hours may draw particularly good match qualities and thus experience wage increases or little wage

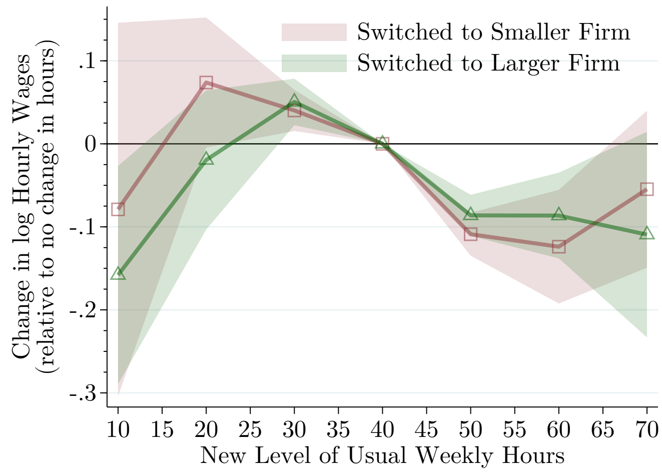

Figure B.13: The longitudinal relationship between wage changes and hours, by changes in firm size

Notes: The figure reports the coefficient $\gamma_{h'}$ from Equation (B.2) for workers that were initially working 40-44 hours in the CPS March sample estimated separately on the sample of workers that switched to a larger and smaller firm between adjacent years. The shaded regions are the 95% confidence intervals.

changes. By the same token, workers that switch firm size categories (and thus employers) may have done so involuntarily and thus may be more likely to experience wage losses.

Having said this, the difference between workers switching to smaller or larger firms is most relevant for testing the presence of complementarities. Our estimates suggest a systematic difference in the estimates of γ'_h based on whether a worker switches to a larger or smaller firm. We find that workers that *decrease* their hours and switch to a larger firm experience *smaller* wage gains relative to workers that switch to smaller firms. On the other hand, workers that *increase* their hours and switch to a larger firm experience *smaller* wage losses relative to workers that switch to smaller firms. Though not statistically significant, qualitatively, this pattern is consistent with the presence of complementarities in hours when hours are increasing with firm size. However, these results should be interpreted with caution, not only due to issues of selection and statistical significance but also since this exercise cannot account for heterogeneity in the reference hours within broad firm size categories and assumes that all large (small) firms have higher (smaller) levels of wage-maximizing hours. Despite these limitations, we find it encouraging that Figure B.13 supports predictions consistent with complementarities and interpret this as suggestive evidence in support of the presence of complementarities in working hours.

Taken together, Figure B.12 suggests that the cross-sectional penalty for working shorter hours may be driven by unobservable worker characteristics – as suggested by Hirsch (2005). While Figure B.13 suggests evidence consistent with complementarities for both long and short-hour penalties. Disentangling the role of complementarities and worker characteristics is not feasible using the CPS nor, to our knowledge, any US dataset. In the absence of the ideal dataset, we interpret the results in Figures B.12 and B.13 as highlighting the importance of both complementarities and worker characteristics in shaping the cross-sectional wage penalties we highlight with "true" wage penalties due to complementarities likely lying somewhere in between the cross-sectional and longitudinal estimates.

Wage changes with fixed hours

Here, we consider how wages change when a worker changes firm sizes but keeps hours fixed. This exercise parallells the analysis in Figure B.13 and tests a central implication of our model. Intuitively, our model predicts that workers who keep their hours fixed but change firm sizes should experience wages changes through two channels: i) changes in firm productivity—a standard mechanism—and ii) changes in the worker's hours relative to the the firm-level reference hours l^*—a mechanism specific to our model, driven by complementarities in working hours.

For instance, a worker who consistently works long hours and moves to a larger firm should experience wage gains due to both working in a more productive firm and being closer to that firm's reference hours. Conversely, a worker who consistently works short hours and moves to a larger firm would receive a similar wage boost through the productivity channel, but—unlike the long-hour worker—they would incur an additional wage penalty from being farther from the firm's reference hours. This intuition gives us a testable implication: long-hour workers should experience larger wage gains than short-hour or typical-hour workers when moving to a large firm, due to the second channel—penalties that depend on the distance from the firm's reference hours. A parallel implication applies to workers moving to smaller firms: short-hour workers should experience the largest wage gains, as their hours are more closely aligned with the reference hours in smaller firms.

To explore whether this pattern is observed in March CPS data, we estimate the following regression,

$$\Delta \log(w_i) = \alpha + \left(\sum_{\Delta f \in \mathcal{F}} \gamma_{\Delta f} \mathbb{I}_{i,\Delta f} \right) + \delta X_i + \epsilon_i, \tag{B.3}$$

where $\Delta \log(w_i)$ denotes the log difference in hourly wages of individual i between year t and $t-1$. The variable Δf indicates whether the worker moves to a firm of the same, smaller, or larger size in period t compared to period $t-1$. The indicator variable $\mathbb{I}_{\Delta f}$ equals one if a worker's transition corresponds to a size change in the set $F =$ No Change, Smaller, Larger, which partitions all possible firm size transitions into three categories. The vector X_i contains individual-level controls, as in Equation (B.2). Since most workers do not change firm size categories, we use this group (no change) as the reference. The coefficient $\gamma_{\Delta f}$ therefore captures the differential change in wages for workers who switch to a smaller or larger firm, relative to those who remain in the same firm size category.

We estimate this regression separately for three groups of workers who, in adjacent years, work (i) 40 hours per week, (ii) between 25 and 34 hours, and (iii) between 46 and 55 hours. The first group represents the typical worker, as 40 hours is the modal value in the data, and serves as a benchmark for evaluating wage changes associated with a "typical" transition across firm size categories. The second and third groups capture short-hour and long-hour workers, respectively.

Table B.10 reports the estimated coefficients $\gamma_{\Delta f}$ for each of these three worker groups. Column 1 reports the results for the workers who work the modal number of hours (40). We

Table B.10: Wage changes for modal, shorter and longer hour workers, relative to no change in firm size

	Modal Hours $h = h' = 40$	Shorter Hours $h, h' \in [25, 34]$	Longer Hours $h, h' \in [46, 55]$
Switched to Smaller Firm	-0.010*	-0.001	-0.015
	(0.004)	(0.020)	(0.011)
Switched to Larger Firm	0.015***	-0.004	0.023*
	(0.004)	(0.017)	(0.010)
Year, State FE	Y	Y	Y
Demographic Controls	Y	Y	Y
4-digit Industry FE	Y	Y	Y
N	50,465	2,476	8,427
R^2	0.012	0.062	0.037

Notes: The table reports the coefficient $\gamma_{\Delta f}$ estimated from Equation (B.3) using the March CPS where the reference firm size change group is no change in firm size. The first column reports the estimates among the group of workers that work 40 hours in adjacent years. The second and third columns report, respectively, results for workers that work between $[25, 34]$ and $[46, 55]$ hours in adjacent years. Standard errors are reported in parentheses. * and *** indicates statistical significance at the 10% and 1% level, respectively.

find that, relative to no change in firm size, switching to a smaller firm is associated with a 1% decline in wages, while switching to a larger firm is associated with a 1.5% increase. This "typical" transition aligns with the predictions from the firm productivity channel.

Column 2 reports $\gamma_{\Delta f}$ for short-hour workers. Relative to modal-hour workers who switch to smaller firms, short-hour workers making the same transition experience much smaller wage losses (0.1% vs. 1%). In contrast, compared to modal-hour workers who switch to larger firms, short-hour workers see significantly smaller wage gains (-0.4% vs. 1.5%). These results are consistent with the model's intuition: short-hour workers move closer to their firm's wage-maximizing hours when switching to a smaller firm, but farther away when switching to a larger firm.

The last column reports results for long-hour workers. Compared to modal-hour workers who switch to smaller firms, this group experiences somewhat larger wage losses (1.5% vs. 1%), though the difference is not statistically significant. In contrast, their gains from switching to larger firms are noticeably higher (2.3% vs. 1.5%). These findings are again consistent with the model's predictions: for long-hour workers, moving to a larger firm aligns their hours more closely with the firm's reference level, resulting in higher wage gains. Conversely, moving to a smaller firm increases the distance from the reference hours, leading to a larger wage loss.

Table B.11 reports the model-implied values of $\gamma_{\Delta f}$ and shows that the model reproduces the qualitative pattern observed in the data. Although the model is not disciplined to match earnings growth, it nonetheless captures the same directional patterns: long-hour (short-hour) workers experience larger (smaller) wage gains than modal workers when switching to larger firms, and larger (smaller) wage losses when switching to smaller firms. Quantitatively, the model predicts significantly larger wage changes and, by construction, can fully account for all sources of variation in wage changes. This difference is, in part, because our model abstracts from many factors that shape wage changes such as downward wage rigidity and (possibly heterogeneous) life-cycle earnings profiles.

Table B.11: Wage changes for modal, shorter and longer hour workers in the model, relative to no change in firm size

	Modal Hours $h = h' = 40$	Shorter Hours $h, h' \in [25, 34]$	Longer Hours $h, h' \in [46, 55]$
Switched to Smaller Firm	-0.328	-0.303	-0.343
Switched to Larger Firm	0.330	0.297	0.359

Notes: The table reports the coefficient $\gamma_{\Delta f}$ estimated from Equation (B.3) using simulated data from the model.

Overall, we find it encouraging that the data bears out additional (un-targeted) predictions of the data and interpret this as suggestive evidence in support of complementarities in hours. However, as with the evidence on longitudinal wage penalties, data such as the CPS has several limitations in this context including being unable to identify the wage-maximizing level of hours and small sample sizes.

C Model appendix

In this appendix, we provide additional details and discussions related to the model and quantitative analysis. We also conduct sensitivity analysis on the parameter ρ governing the degree of elasticity of substitution in working hours.

C.1 Definition of the equilibrium in the baseline model

We start by introducing the notation for the policy functions, distributions and wages. Without loss of generality, one can think of the ϵ-shocks as being realized after the value of leisure shock of a worker. The probability of a worker choosing a firm productivity level z_j is denoted by $\mathbf{o}_j(\nu) \in \{1, ..., J\}$. Meanwhile, $\mathbf{l}_j(\nu)$ denotes the labor supply policy function. We denote the workers' policy for accepting the offer post-recruitment by $A_j^*(w, l, \nu)$, for the offer made to the alternative firm in case of rejecting the recruiting firm by $Q_j^*(l, \nu)$.

We denote the policy functions of the firms for recruitment (first stage) by M_j^*, employment and wage offers (second stage) by $\mu_j^*(l; M)$ and $W_j^*(l; M)$. The policy of a firm for accepting the offer of a worker that rejected its recruiting firm by $B_j^*(w, l, \mu)$.

Recall that $\Phi_j(l)$ is the density of l hours supplied in market j. $\varphi_j(\nu)$ is the distribution of workers over employer types $j \in \{1, .., J\}$, the value of leisure $\nu \in B_\nu$.

Equilibrium. The equilibrium consists of a set of policy functions: M_j^*, $\mu_j^*(l; M)$, $W_j^*(l; M)$, and $B_j^*(w, l, \mu)$ for firms $j \in \{1, .., J\}$; $\mathbf{l}_j(\nu)$, $\mathbf{o}_j(\nu)$, $Q_j^*(l, \nu)$ and $A_j^*(w, l, \nu)$ for workers, wage functions $w_j(l)$ (to be conjectured by workers before entry into the market), and a distribution of workers $\varphi_j(\nu)$ and hours $\Phi_j(l)$ defined for each employer type $j \in \{1, .., J\}$, the value of leisure $\nu \in B_\nu$, hours $l \in \mathcal{L}$ such that:

(i) The policy functions solve the problems of workers and firms given the wages.

(ii) Labor markets clear. The total number of workers demanded by the firms in the recruitment stage is equal to the total supply:

$$M_j^* \Lambda_j = \sum_{v \in B_\nu} \varphi_j(\nu), \ \ \forall j \in \{1, 2, .., J\}$$

and the total measure of workers demanded by all firms for each level of firm productivity z_j and working hours $l \in \mathcal{L}$ is equal to the corresponding labor supply:

$$\mu_j^*(l) \Lambda_j = \sum_{v \in B_\nu} \varphi_j(\nu) \mathbb{1}[\mathbf{l}_j(\nu) = l], \ \ \forall j \in \{1, 2, .., J\}$$

(iii) The density of hours in the recruited stacks is consistent with the distribution of workers over states and their policy functions:

$$\Phi_j(l) = \frac{\sum\limits_{v \in B_\nu} \varphi_j(\nu) \mathbb{1}[\mathbf{l}_j(\nu) = l]}{\sum\limits_{\tilde{l} \in \mathcal{L}} \sum\limits_{v \in B_\nu} \varphi_j(\nu) \mathbb{1}[\mathbf{l}_j(\nu) = \tilde{l}]}, \ \ \forall j \in \{1, 2, .., J\}, \ l \in \mathcal{L}$$

C.2 Definition of the equilibrium in the extended model

The extended model differs from the baseline in being dynamic, and having the additional heterogeneity in worker efficiency and wealth. Here, we provide the definition of the stationary general equilibrium modified for the extended model. It is necessary to update the notation of some of the policy and density functions. The probability of a worker choosing a firm productivity level z_j is now denoted by $\mathbf{o}_j(a, x, \nu) \in \{1, ..., J\}$. $\mathbf{l}_j(a, x, \nu)$ denotes the labor supply policy function, and $\mathbf{a}_j(a, x, \nu)$ denotes the savings policy function. Workers' policy for accepting the offer post-recruitment is, $A_j^*(w, l, a, x, \nu)$, the offer made to the alternative firm in case of rejecting the recruiting firm is $Q_j^*(l, a, x, \nu)$. Finally

The distribution of workers over employer types $j \in \{1, .., J\}$, assets $a \geq 0$, efficiency $x \in B_x$, and the value of leisure $\nu \in B_\nu$ is denoted by $\varphi_j(a, x, \nu)$,

Stationary general equilibrium of the extended model. A stationary general equilibrium consists of a set of policy functions: M_j^*, $\mu_j^*(l, x; M)$, $W_j^*(l, x; M)$, and $B_j^*(w, l, x, \mu)$ for firms $j \in \{1, .., J\}$; $\mathbf{l}_j(a, x, \nu)$, $\mathbf{o}_j(a, x, \nu)$, $\mathbf{a}_j(a, x, \nu)$, $Q_j^*(l, a, x, \nu)$, and $A_j^*(w, l, a, x, \nu)$ for workers, wage functions $w_j(l, x)$ (to be conjectured by workers before entry into the market), a time-invariant distribution of workers $\varphi_j(a, x, \nu)$ and hours $\Phi_j(l, x)$ defined for each employer type $j \in \{1, .., J\}$, wealth $a \geq 0$, efficiency level $x \in B_x$, the value of leisure $\nu \in B_\nu$, hours $l \in \mathcal{L}$ such that:

(i) The policy functions solve the problems of workers and firms given the wages.

(ii) Labor markets clear. The total number of workers demanded by the firms in the recruitment stage is equal to the total supply:

$$M_j^* \Lambda_j = \int\limits_{a=0} \sum_{x \in B_x} \sum_{v \in B_\nu} \varphi_j(a, x, \nu) da, \ \ \forall j \in \{1, 2, .., J\}$$

and the total measure of workers demanded by all firms for each level of firm productivity z_j and working hours $l \in \mathcal{L}$ is equal to the corresponding labor supply:

$$\mu_j^*(l, x)\Lambda_j = \int\limits_{a=0} \sum_{x \in B_x} \sum_{v \in B_\nu} \varphi_j(a, x, \nu) \mathbb{1}[\mathbf{l}_j(a, x, \nu) = l] da, \quad \forall x \in B_x \, \forall j \in \{1, 2, .., J\}$$

(iii) The density of hours in the recruited stacks is consistent with the distribution of workers over states and their policy functions:

$$\Phi_j(l, x) = \frac{\int\limits_{a=0} \sum\limits_{x \in B_x} \sum\limits_{\nu \in B_\nu} \varphi_j(a, x, \nu) \mathbb{1}[\mathbf{l}_j(a, x, \nu) = l] da}{\sum\limits_{\tilde{l} \in \mathcal{L}} \int\limits_{a=0} \sum\limits_{x \in B_x} \sum\limits_{v \in B_\nu} \varphi_j(a, x, \nu) \mathbb{1}[\mathbf{l}_j(a, x, \nu) = \tilde{l}] da}, \ \ \forall j \in \{1, 2, .., J\}, \ l \in \mathcal{L}$$

(iv) The evolution of the distribution across workers satisfies, for each $a \geq 0$, $x \in B_x$,

$\nu \in B_\nu$, and $j \in \{1, 2, .., J\}$:

$$\varphi_j(a, x, \nu) = \int_{\tilde{a}=0} \sum_{\tilde{x} \in B_x} \sum_{\tilde{\nu} \in B_\nu} \Gamma_x(\tilde{x}, x) \Gamma_\nu(\tilde{\nu}, \nu) \times$$

$$\sum_{\tilde{j}=1}^{J} \varphi_{\tilde{j}}(\tilde{a}, \tilde{x}, \tilde{\nu}) \left(s\mathbf{o}_j(a, x, \nu) + (1-s)\mathbb{1}[\tilde{j} = j] \right) \mathbb{1}[\mathbf{a}_{\tilde{j}}(\tilde{a}, \tilde{x}, \tilde{\nu}) = a] d\tilde{a}.$$

C.3 The role of taste shocks

Described in Section 3, our model features shocks to the return of workers to working for firms of differing productivity. The computational advantage of these shocks is that they help 'convexify' the occupational choice of workers by introducing additional randomness in their decision.

In particular, by assuming a Generalized Extreme Value Distribution for these shocks, the occupational choice of workers can be considered as a probability which is given by the value obtained in each occupation – net of the ϵ-shocks, relative to the aggregation of values in all other firm productivity levels,

$$H_j(\nu) = \frac{\exp\left(V_j^G(\nu)\right)^{\frac{1}{\sigma_\epsilon}}}{\sum_{k=1}^{J} \exp\left(V_k^G(\nu)\right)^{\frac{1}{\sigma_\epsilon}}}.$$

Having a probability as the policy function, instead of an binary indicator of 0 or 1 for choosing each occupation, smooths out the value function of workers, and help with convergence.

Existing literature has used similar "tastes" in different models of discrete choice, such as McFadden (1978) for households' location choice and Wolpin (1984) in a model of fertility. The role of taste heterogeneity in shaping wage heterogeneity between employers has previously been highlighted and modeled in Card et al. (2018). Other papers that use similar shocks in the context of occupational choice of workers are Artuç et al. (2010) and Caliendo et al. (2019).

As with Card et al. (2018), these taste shocks play a crucial role in generating a size-wage premium in our model. We interpret these taste shocks as capturing a number of data features that affect individuals' sorting into firms of different productivity and size, which are not accounted for by wages. Below we discuss several features of the labor market that motivate introducing the taste shocks in the model.

Small and large firms differ in the multi-dimensional non-pecuniary benefits they offer to workers. Studies showed that workers' heterogeneous preferences over these non-pecuniary characteristics are important in generating earning inequality (Rosen, 1986, Morchio and Moser, 2018 and Lamadon et al., 2022). Importantly, we emphasize that these non-pecuniary characteristics might differ across firm size groups. While small firms have a more friendly and less rigid work environment (Agell, 2004 and Idson, 1990), large firms might excel in some other dimensions such as safety in the work environment (Oi, 1974).

Further, there may exist logistical and technological reasons why different workers do not find it equally feasible to work in larger firms. For example, several studies have documented firms in urban areas are more productive and larger than firms in rural areas (Headd, 2000 and Melo et al., 2009). If workers are tied to a specific location, moving cost or commuting costs could contribute to the taste shocks we build into our model. There also exists cross-sector difference in firm size, with smaller firms more likely to be in the construction, services, and agriculture sector and large firms more likely to be in the manufacturing, retail, transportation, and finance sector (Headd, 2000). In this regard, the taste shock could also be interpreted as the limited transferability of sector-specific human capital.

The taste shocks we introduce are a reduced-form representation of the real-world features we discuss here. Having said this, much of the worker side's heterogeneity above can be argued to be persistent, yet our taste shocks are independently drawn each period. The reason for this is tractability and simplicity. We could generate similar implications with our model in a static model with static heterogeneity in workplace preferences.

Notice also that our taste shocks are over firms' productivity levels, even though the discussion above corresponds to the preferences or ability to work in firms of different sizes. However, firm size and productivity will be isomorphic in our model and are strongly positively correlated in the data (see, for example, Leung et al. 2008 and Bartelsman et al. 2013).

C.4 Details of the model solution

Here we establish the optimality of firms' policy functions (firms' labor demand schedule) in the symmetric equilibrium. We do this for the extended model described in Section 5, since it is a more general form of the model in Section 3. First, we derive the necessary condition of the firm's optimization problem. The policy function in the symmetric equilibrium meets this condition. Next, we prove key analytical properties of all firm demand schedules that satisfy this condition, narrowing down the set of alternatives to the symmetric solution. Third, we show that even with a small cost κ, firms do not benefit from deviating from the symmetric solution.

We start from the second (*hiring*) stage problem. That is, the problem of choosing the type-specific demand schedule $\mu(l,x)$ given its stack from the first (*recruitment*) stage. This problem is represented by:

$$\pi_j(M) = \max_{\{\mu(l,x)\}_{l\in\mathcal{L},x\in B_x}} Y - \sum_{x\in B_x}\sum_{l\in\mathcal{L}} w_j(l,x)l\mu(l,x) - \sum_{x\in B_x}\sum_{l\in\mathcal{L}} \kappa(M\Phi_j(l,x)-\mu(l,x)) \tag{C.1}$$

$$s.t.\ Y = z_j\left[\left(\sum_{x\in B_x}\sum_{l\in\mathcal{L}}\mu(l,x)l^\rho\right)^{\frac{1}{\rho}}\left(\sum_{x\in B_x}\sum_{l\in\mathcal{L}}\mu(l,x)\right)^{1-\frac{1}{\rho}}\right]^\theta$$

$$\mu(l,x)\in[0,M\Phi_j(l,x)]\ \forall x\in B_x, l\in\mathcal{L}.$$

This is the same as the problem represented by equation (9) in Section 3, generalized to the extended model in Section (5). It is also convenient to generalize here the notation for the marginal productivity of a worker (introduced for the baseline model in equation (6)) with l hours and efficiency x:

$$f_j(l,x;\mu) \equiv \theta x z_j L(\mu)^{\theta-1}\left[\frac{1}{\rho}\left(\frac{l}{\tilde{l}_j}\right)^\rho + \left(1-\frac{1}{\rho}\right)\right]\tilde{l}_j \tag{C.2}$$

where

$$\tilde{l}_j \equiv \left[\left(\sum_{x\in B_x}\sum_{l\in\mathcal{L}}\mu(l,x)l^\rho\right)\Big/\left(\sum_{x\in B_x}\sum_{l\in\mathcal{L}}\mu(l,x)\right)\right]^{\frac{1}{\rho}}. \tag{C.3}$$

The necessary conditions for the optimality of a demand schedule μ in this stage are:

$$f_j(l,x;\mu)-w_j(l,x)\begin{cases}\geq 0 & \text{if } \mu(l,x)=M\Phi_j(l,x)\\ =-\kappa & \text{if } \mu(l,x)\in(0,M\Phi_j(l,x))\\ <-\kappa & \text{if } \mu(l,x)=0\end{cases}$$

for any $l\in\{l_1,..,l_N\}$ and $x\in B_x$, where $\bar{\mu}_j(l,x)$ is the density of worker (l,x) in market j. Replacing the wage equation that the firms take as given, expressed by our Equation (15), the necessary conditions boil down to:

$$f_j(l,x;\mu) - f_j(l,x;M_j^*\Phi_j) \begin{cases} \geq 0 & \text{if } \mu(l,x) = M\Phi_j(l,x) \\ = -\kappa & \text{if } \mu(l,x) \in (0, M\Phi_j(l,x)) \\ < -\kappa & \text{if } \mu(l,x) = 0 \end{cases}$$

The demand schedule involving the employment of all the recruited workers, $\mu(l,x) = M\Phi_j(l,x)$, satisfies the necessary condition (by satisfying the first condition above). We cannot analytically show the sufficiency conditions of optimality for this solution, hence, we need to rule out other potential choices involving the rejection of certain hours and efficiency types. We do that using a semi-numerical approach. Since there are infinitely many alternatives in principle, we first theoretically narrow down the set of potential profit-enhancing alternatives. We will then use numerical methods to search within this set to ensure that such an alternative does not exist in our calibrated models. Below suppose there exists a demand schedule $\mu(l,x) = \tilde{\mu}_j(l,x)$ that delivers a higher level of profits than $\mu(l,x) = M\bar{\mu}_j(l,x)$. Denote by $L(\tilde{\mu})$ the effective labor, and by $E_j(l^\rho;\tilde{\mu})$ the weighted l^ρ across the workers in a firm with that alternative.

The next claim argues that for each efficiency level x, there is at most one hour level $\hat{l}$, for which the alternative features rejection of some, but not all, of the recruited workers, i.e. $\tilde{\mu}(\tilde{l},\tilde{x}) \in (0, \Phi_j(\tilde{l},\tilde{x}))$.

Claim 1. Suppose there exists a demand schedule $\mu(l,x) = \tilde{\mu}(l,x)$ that delivers a higher level of profits than $\mu(l,x) = M\bar{\mu}_j(l,x)$. For some $\tilde{x} \in B_x$, if there is one $\tilde{l} \in \{l_1,..,l_N\}$ that gives $\tilde{\mu}(\tilde{l},\tilde{x}) \in (0, \Phi_j(\tilde{l},\tilde{x}))$, then one of the following is true,

- **(Case 1) $\tilde{\mu}(l,\tilde{x}) = 0$ for all $l > \tilde{l}$, and $\tilde{\mu}(l,\tilde{x}) = M\Phi_j(l,\tilde{x})$ for all $l < \tilde{l}$.**
- **(Case 2) $\tilde{\mu}(l,\tilde{x}) = M\Phi_j(l,\tilde{x})$ for all $l > \tilde{l}$, and $\tilde{\mu}(l,\tilde{x}) = 0$ for all $l < \tilde{l}$.**

Proof. Suppose that for some $\tilde{x} \in B_x$, there is one $\tilde{l} \in \{l_1,..,l_N\}$, such that a profit enhancing alternative demand satisfies $\tilde{\mu}_j(\tilde{l},\tilde{x}) \in (0, M\Phi_j(\tilde{l},\tilde{x}))$. The necessary condition for this to be optimal is:

$$f_j(\tilde{l},\tilde{x};\tilde{\mu}) - f_j(\tilde{l},\tilde{x};M_j^*\Phi_j) + \kappa = 0.$$

Define:

$$g(\hat{l},\hat{x};\tilde{\mu},M\Phi_j) \equiv f_j(\hat{l},\hat{x};\tilde{\mu}) - f_j(\hat{l},\hat{x};M_j^*\Phi_j) \tag{C.4}$$

If $\tilde{\mu}$ features an interior solution for $(\tilde{l},\tilde{x})$, then the necessary condition is:

$$g(\tilde{l},\tilde{x};\tilde{\mu},M\Phi_j) = -\kappa \tag{C.5}$$

We will first show that the function $g(\hat{l},\hat{x};\tilde{\mu},M\Phi_j)$ is strictly monotone in $\hat{l}$. To do so, we will show that the derivative of the g function with respect to $\hat{l}$ is non-zero for any $\hat{l}$. Taking derivative of $g(\hat{l},\hat{x};\tilde{\mu},M\Phi_j)$ with respect to $\hat{l}$, we get:

$$g_1(\hat{l},\hat{x};\tilde{\mu},M_j^*\Phi_j) = \theta\hat{x}z_j\hat{l}^{\rho-1}\left[L(\tilde{\mu})^{\theta-1}E_j(l^\rho;\tilde{\mu})^{\frac{1}{\rho}-1} - L(M_j^*\Phi_j)^{\theta-1}E(l^\rho;M_j^*\Phi_j)^{\frac{1}{\rho}-1}\right] \tag{C.6}$$

The term in brackets is invariant to $\hat{l}$. Hence, for this derivative to be zero for some l, it has to be equal to zero for all $l \in \{l_1, .., l_N\}$, and it also requires that the term in the bracket to be zero:

$$\left(\frac{L(\tilde{\mu})}{L(M_j^*\Phi_j)}\right)^{\theta-1} = \left(\frac{E_j(l^\rho;\tilde{\mu})}{E(l^\rho;M_j^*\Phi_j)}\right)^{1-\frac{1}{\rho}} \implies \frac{\tilde{N}}{M_j^*} = \frac{E(l^\rho;M_j^*\Phi_j)}{E_j(l^\rho;\tilde{\mu})}, \tag{C.7}$$

where the second equation follows from the definition of the labor aggregation L.

If the bracket in Equation (C.6) is equal to zero, then Equation (C.5) implies:

$$(1-\frac{1}{\rho})\theta\tilde{x}z_j\left[L(\tilde{\mu})^{\theta-1}E_j(l^\rho;\tilde{\mu})^{\frac{1}{\rho}} - L(M_j^*\Phi_j)^{\theta-1}E(l^\rho;M_j^*\Phi_j)^{\frac{1}{\rho}}\right] = -\kappa \implies$$

$$(1-\frac{1}{\rho})\theta\tilde{x}z_jL(M_j^*\Phi_j)^{\theta-1}E(l^\rho;M_j^*\Phi_j)^{\frac{1}{\rho}}\left[\frac{E_j(l^\rho;\tilde{\mu})}{E(l^\rho;M_j^*\Phi_j)}-1\right] = -\kappa \implies \frac{E_j(l^\rho;\tilde{\mu})}{E(l^\rho;M_j^*\Phi_j)} < 1.$$

Together with Equation (C.7), this requires $\tilde{N} > M_j^*$, i.e. that the alternative schedule hires more than the number of recruitment from the previous stage, which cannot happen. This therefore proves our conjecture that function $g(\hat{l},\hat{x};\tilde{\mu},M\Phi_j)$ is strictly monotone in $\hat{l}$.

If the term in brackets in Equation (C.6) is negative (positive), it is negative (positive) for all l for the given $\hat{x}$, i.e. function g is strictly monotone in l. Since $g(\tilde{l},\tilde{x};\tilde{\mu},M_j^*\Phi_j) = -\kappa$, we have either,

- $g(l,\tilde{x};\tilde{\mu},M_j^*\Phi_j) < -\kappa$ **for all** $l > \tilde{l}$**, and** $g(l,\tilde{x};\tilde{\mu},M_j^*\Phi_j) > -\kappa$ **for all** $l < \tilde{l}$**, or**
- $g(l,\tilde{x};\tilde{\mu},M_j^*\Phi_j) > -\kappa$ **for all** $l > \tilde{l}$**, and** $g(l,\tilde{x};\tilde{\mu},M_j^*\Phi_j) < -\kappa$ **for all** $l < \tilde{l}$**.**

A case where for some l and $\tilde{x}$, $g(l,\tilde{x};\tilde{\mu},M_j^*\Phi_j)$ is strictly between $-\kappa$ and 0 does not satisfy the necessary conditions of optimality. Accordingly, above conditions imply:

- $\tilde{\mu}(l,\tilde{x}) = 0$ **for all** $l > \tilde{l}$**, and** $\tilde{\mu}(l,\tilde{x}) = M\Phi_j(l,\tilde{x})$ **for all** $l < \tilde{l}$**, or**
- $\tilde{\mu}(l,\tilde{x}) = M\Phi_j(l,\tilde{x})$ **for all** $l > \tilde{l}$**, and** $\tilde{\mu}(l,\tilde{x}) = 0$ **for all** $l < \tilde{l}$**.**

This proves the claim.

Taking stocks, we proved that if for some $\tilde{x} \in B_x$, there is one level of hours for which the alternative demand features positive demand but less than the full-recruitment of that type, this demand schedule either features hitting the upper bound for all hours higher than l, and 0 for all hours below, which we refer to as Case 1. The other possibility is that this demand schedule features no demand for all hours higher than l, and the upper bound for all hours below, which we refer to as Case 2.

Since $\rho < 1$, we are in Case 1 if $L(\tilde{\mu})^{\theta-1}E_j(l^\rho;\tilde{\mu})^{\frac{1}{\rho}-1} > L(M_j^*\Phi_j)^{\theta-1}E(l^\rho;M_j^*\Phi_j)^{\frac{1}{\rho}-1}$. We are in Case 2 if $L(\tilde{\mu})^{\theta-1}E_j(l^\rho;\tilde{\mu})^{\frac{1}{\rho}-1} < L(M_j^*\Phi_j)^{\theta-1}E(l^\rho;M_j^*\Phi_j)^{\frac{1}{\rho}-1}$.

Next, in Claims 2a and 2b, we approach the problem for a given hours level, trying to rule out alternatives for varying levels of efficiency.

Claim 2a. Suppose there exists a demand schedule $\mu(l,x) = \tilde{\mu}_j(l,x)$ **that delivers a higher level of profits than** $\mu(l,x) = M\bar{\mu}_j(l,x)$**. For a given** $\tilde{l} \in \{l_1, .., l_N\}$**, if there**

is one $\tilde{x} \in B_x$ that gives $\tilde{\mu}(\tilde{l}, \tilde{x}) = M\Phi_j(\tilde{l}, \tilde{x})$, then for all other $x \in B_x$, we should also have $\tilde{\mu}(\tilde{l}, x) = M\Phi_j(\tilde{l}, x)$.

Proof. Suppose that for some $\tilde{l} \in \{l_1, .., l_N\}$, there is one $\tilde{x} \in B_x$ such that a profit enhancing alternative demand satisfies $\tilde{\mu}_j(\tilde{l}, \tilde{x}) = \Phi_j(\tilde{l}, \tilde{x})$. The necessary condition for this to be optimal is:

$$g(\tilde{l}, \tilde{x}; \tilde{\mu}, M\Phi_j(\tilde{l}, \tilde{x})) = f_j(\tilde{l}, \tilde{x}; \tilde{\mu}) - f_j(\tilde{l}, \tilde{x}; M_j^*\Phi_j) \geq 0.$$

Since the function $g(\hat{l}, \hat{x}; \tilde{\mu}, M\Phi_j(\tilde{l}, \tilde{x}))$ is homogeneous of degree 1 with respect to $\hat{x}$ (for a given $\tilde{l}$), this implies that for any $x \in B_x$,

$$g(\tilde{l}, x; \tilde{\mu}, M\Phi_j) \geq 0$$

From the necessary conditions of optimality, it follows that $\tilde{\mu}(\tilde{l}, x) = M\Phi_j(\tilde{l}, x)$ for any $x \in B_x$.

Claim 2b. Suppose there exists a demand schedule $\mu(l, x) = \tilde{\mu}_j(l, x)$ that delivers a higher level of profits than $\mu(l, x) = M\bar{\mu}_j(l, x)$. For a given $\tilde{l} \in \{l_1, .., l_N\}$, if there is one $\tilde{x} \in B_x$ that gives $\tilde{\mu}(\tilde{l}, \tilde{x}) \in (0, M\Phi_j(\tilde{l}, \tilde{x}))$, then this $\tilde{x}$ is the minimum in the set of efficiencies, and for all higher efficiencies, the schedule should give zero demand, i.e. $\tilde{x} = \min\{x : x \in B_x\}$, and $\tilde{\mu}(\tilde{l}, x) = 0\ \forall x > \tilde{x}$.

Proof. Suppose that for some $\tilde{l} \in \{l_1, .., l_N\}$, there is one $\tilde{x} \in B_x$ such that a profit enhancing alternative demand satisfies $\tilde{\mu}_j(\tilde{l}, \tilde{x}) \in (0, \Phi_j(\tilde{l}, \tilde{x}))$. The necessary condition for this to be optimal is:

$$g(\tilde{l}, \tilde{x}; \tilde{\mu}, M\Phi_j(\tilde{l}, \tilde{x})) = f_j(\tilde{l}, \tilde{x}; \tilde{\mu}) - f_j(\tilde{l}, \tilde{x}; M_j^*\Phi_j) = -\kappa.$$

The function $g(\hat{l}, \hat{x}; \tilde{\mu}, M\Phi_j(\tilde{l}, \tilde{x}))$ is homogeneous of degree 1 with respect to $\hat{x}$. Moreover, necessary conditions for optimality imply that $g(l, x; \tilde{\mu}, M\Phi_j(\tilde{l}, \tilde{x})) \notin (-\kappa, 0)$, for any (l, x) pair. This condition is violated for any $x < \tilde{x}$, hence, $\tilde{x}$ has to be the minimum in the possible efficiency levels that can be demanded by the firm. Moreover, it follows that for any $x > \tilde{x}$,

$$g(\tilde{l}, x; \tilde{\mu}, M\Phi_j) < -\kappa.$$

From the necessary conditions of optimality, it follows that $\tilde{\mu}(\tilde{l}, x) = 0$ for any $x > \tilde{x}$.

Claim 3. Suppose there exists a demand schedule $\mu(l, x) = \tilde{\mu}_j(l, x)$ that delivers a higher level of profits than $\mu(l, x) = M\bar{\mu}_j(l, x)$. Suppose there is a combination $\tilde{l} \in \{l_1, .., l_N\}$, and $\tilde{x} \in B_x$ for which $\tilde{\mu}(\tilde{l}, \tilde{x}) \in (0, M\Phi_j(\tilde{l}, \tilde{x}))$. Then, $\tilde{x} = \min\{x : x \in B_x\}$ and $\tilde{\mu}(\tilde{l}, x) = 0$ for all $x > \tilde{x}$. Moreover, one of the two cases below has to be true:

- **(Case A) $\tilde{\mu}(l, x) = 0$ for all $l > \tilde{l}$ and $x \in B_x$, and $\tilde{\mu}(l, \tilde{x}) = M\Phi_j(l, \tilde{x})$ for all $l < \tilde{l}$ and $x \in B_x$.**

- **(Case B) $\tilde{\mu}(l, \tilde{x}) = M\Phi_j(l, \tilde{x})$ for all $l > \tilde{l}$ and $x \in B_x$, and $\tilde{\mu}(l, \tilde{x}) = 0$ for all $l < \tilde{l}$ and $x \in B_x$.**

Proof. Suppose there exists a demand schedule $\mu(l,x) = \tilde{\mu}_j(l,x)$ that delivers a higher level of profits than $\mu(l,x) = M\bar{\mu}_j(l,x)$. Moreover, suppose that in this schedule, there is a combination $\tilde{l} \in \{l_1, .., l_N\}$, and $\tilde{x} \in B_x$ for which $\tilde{\mu}(\tilde{l},\tilde{x}) \in (0, M\Phi_j(\tilde{l},\tilde{x}))$.

Case A. In the first part of the proof, we assume that

$$L(\tilde{\mu})^{\theta-1} E_j(l^\rho; \tilde{\mu})^{\frac{1}{\rho}-1} > L(M_j^*\Phi_j)^{\theta-1} E(l^\rho; M_j^*\Phi_j)^{\frac{1}{\rho}-1}.$$

We know from Claim 1 that in this case, for the same $\tilde{x}$, all the higher level of hours should have zero demand, and all the lower levels of hours should have full demand, i.e. we are in Case 1 highlighted in Claim 1: $\tilde{\mu}(l,\tilde{x}) = 0$ for all $l > \tilde{l}$, and $\tilde{\mu}(l,\tilde{x}) = M\Phi_j(l,\tilde{x})$ for all $l < \tilde{l}$. From Claim 2a, this gives that $\tilde{\mu}(l,x) = M\Phi_j(l,x)$ for all $l < \tilde{l}$ and any $x \in B_x$. In other words, for any l lower, regardless of the x, there is full demand. From Claim 2b, we know that the combination for which the demand is interior should have the lowest efficiency level, and all higher efficiency for the same hours $\tilde{l}$ should have zero demand: $\tilde{x} = \min\{x : x \in B_x\}$ and $\tilde{\mu}(\tilde{l},x) = 0$ for all $x > \tilde{x}$.

The only combinations to still check are those with efficiency higher than $\tilde{x}$, and hours higher than $\tilde{l}$. Suppose there is positive demand for such a point, i.e. $\tilde{\mu}(l,x) > 0$, for some $\hat{l} > \tilde{l}$ and $\hat{x} > \tilde{x}$. Since we are in Case 1 represented in Claim 1, i.e. the bracket in Equation (C.6) is negative, this would require full demand for all hours lower than $\hat{l}$ and any $\hat{x} > \tilde{x}$. Since, $\tilde{l} < \hat{l}$, this implies full demand for our initial reference hours $\tilde{l}$, with $\hat{x} > \tilde{x}$, which contradicts Claim 2b.

Case B. Now assume that

$$L(\tilde{\mu})^{\theta-1} E_j(l^\rho; \tilde{\mu})^{\frac{1}{\rho}-1} < L(M_j^*\Phi_j)^{\theta-1} E(l^\rho; M_j^*\Phi_j)^{\frac{1}{\rho}-1}.$$

Then we are in Case 2 of Claim 1: $\tilde{\mu}(l,\tilde{x}) = M\Phi_j(l,\tilde{x})$ for all $l > \tilde{l}$, and $\tilde{\mu}(l,\tilde{x}) = 0$ for all $l < \tilde{l}$. This implies $\tilde{\mu}(l,x) = M\Phi_j(l,x)$ for all $l > \tilde{l}$ and any $x \in B_x$. We also know from Claim 2b, that $\tilde{x}$ has to be the lower possible efficiency level, and all higher efficiency for the same hours $\tilde{l}$ should have zero demand: $\tilde{x} = \min\{x : x \in B_x\}$ and $\tilde{\mu}(\tilde{l},x) = 0$ for all $x > \tilde{x}$.

Next turn to points with a lower l than $\tilde{l}$, and an efficiency higher than $\tilde{x}$. Suppose there is positive demand for such a point, i.e. $\tilde{\mu}(l,x) > 0$, for some $\hat{l} < \tilde{l}$ and $\hat{x} > \tilde{x}$. Since we are in Case 2, this requires full demand for any $\hat{x} > \tilde{x}$ and hours higher than $\hat{l}$, including $\tilde{l}$, which contradicts Claim 2b and completes the proof.

Taking stocks. Taken together, the claims 1 to 3 prove that there cannot be multiple points for which a profit-enhancing alternative features rejection of a strict subset of the initially recruited workers, i.e. $\mu(l,x) \in (0, M\Phi_j(l,x))$. If there is one point for which this happens, it has to be with the lowest efficiency level, and either

- all the longer hours feature zero demand and hours shorter or equal to feature full demand for all efficiency levels, or
- all the shorter hours feature zero demand and hours longer or equal to feature full demand for all efficiency levels.

This reduces substantially the space of alternative demand schedules that we need to compare our symmetric candidate solution $N\Phi_j$. We numerically evaluate the possibility of a profit-enhancing alternative demand schedule as follows. For each level hours $l \in \{l_1, .., l_{N_l}\}$, we have two possibilities, (i) zero demand for all the hours weakly below, and full demand for all the hours above, or (ii) zero demand for all the hours weakly above, and full demand for all the hours below. For each of these possibilities, we also allow the demand for l to be a fraction of the upper bound, over a grid of fractions of size N_{frac}. We only need to pursue this possibility for the pairs with the lowest efficiency x. Accordingly, we have in total $N_l \times N_{\text{frac}} \times 2$ alternative demand schedules to evaluate for a given size of initial recruitment.

For each initial size of recruitment M, and an alternative schedule $\tilde{\mu}_j(l,x)$, that is feasible (i.e. $\tilde{\mu}_j(l,x) \leq N\Phi_j(l,x)$ for any (l,x)), the implied profits would be:

$$\tilde{\pi}_j(M, \tilde{\mu}_j) = z_j \left[\left(\sum_{x \in B_x} \sum_{l_1}^{l_N} \tilde{\mu}_j(l,x) l^\rho \right)^{\frac{1}{\rho}} \left(\sum_{x \in B_x} \sum_{l_1}^{l_N} \tilde{\mu}_j(l,x) \right)^{1-\frac{1}{\rho}} \right]^\theta \tag{C.8}$$

$$- \sum_{x \in B_x} \sum_{l_1}^{l_N} w_j(l,x) l \tilde{\mu}_j(l,x) - \kappa \sum_{x \in B_x} \sum_{l_1}^{l_N} (M\Phi_j(l,x) - \tilde{\mu}_j(l,x)) \tag{C.9}$$

Since the total costs of rejection are proportional to κ, we can find the minimum κ required, if any, to make each alternative less profitable than the symmetric schedule $(M_j^* \Phi_j)$ as, $\underline{\kappa}(M, \tilde{\mu}_j; M_j^* \Phi_j) =$

$$\frac{z_j \left[\left(\sum_{x \in B_x} \sum_{l_1}^{l_N} \tilde{\mu}_j(l,x) l^\rho \right)^{\frac{1}{\rho}} \left(\sum_{x \in B_x} \sum_{l_1}^{l_N} \tilde{\mu}_j(l,x) \right)^{1-\frac{1}{\rho}} \right]^\theta - \sum_{x \in B_x} \sum_{l_1}^{l_N} w_j(l,x) l \tilde{\mu}_j(l,x) - \tilde{\pi}_j(M_j^*, M_j^* \Phi_j)}{\sum_{x \in B_x} \sum_{l_1}^{l_N} (M\Phi_j(l,x) - \tilde{\mu}_j(l,x))}$$

There are two scenarios that would challenge the optimality of the symmetric demand schedule. The first is if a firm that chooses the recruitment size of the symmetric schedule in the first stage, M_j^*, chooses to reject some recruited workers in the second stage. Figure C.1a shows the rejection costs needed to rule out such deviations in our baseline model,[3]

$$\max_{\tilde{\mu}_j} \ \underline{\kappa}(M_j^*, \tilde{\mu}_j; M_j^* \Phi_j), \tag{C.10}$$

$$\text{s.t } \tilde{\mu}_j(l,x) \leq M_j^* \Phi_j(l,x) \, \forall l \in \mathcal{L}, x \in B_x \tag{C.11}$$

We find that with a fixed cost per rejection of 0.016, which is around 0.3 percent of the mean wage in equilibrium, the firms find it optimal to hire all the recruited workers, i.e. not reject anyone of their recruited stack, if they recruited the symmetric size (M_j^*) in the first stage.

The second challenge to the optimality of the symmetric demand schedule is if the firms

[3]The description here allows for efficiency heterogeneity (x) to have it applicable to the extended model. When the baseline model is concerned, one should treat the corresponding arguments of the measures and wages as equal to a constant.

recruited more or less than the symmetric size M_j^* in the first stage. Note that even with the costs computed above for the second stage, the firms might still find it optimal to reject some recruited workers if the recruitment size inherited from the first stage differs from M_j^*. In particular, if the recruitment size is too large compared to the symmetric size, then the firms would clearly reject some of the workers. Accordingly, the optimal size of the first stage does depend on the costs of rejection, κ. When we evaluate the possibility of profit-enhancing alternatives in terms of the recruitment size of the first stage, we need to take into account the possibility of rejection for off-the-equilibrium recruitment sizes. Figure C.1b shows the rejection costs required to leave the symmetric recruitment size, M_j^*, optimal at the first stage in our baseline model (conjecturing the hiring of all the recruited workers in that case). In particular, it plots:

$$\max_{N,\tilde{\mu}_j} \underline{\kappa}(M, \tilde{\mu}_j; M_j^* \Phi_j), \tag{C.12}$$

$$\text{s.t } \tilde{\mu}_j(l,x) \leq M\Phi_j(l,x)\, \forall l \in \mathcal{L}, x \in B_x \tag{C.13}$$

It documents that the highest required rejection cost to maintain the symmetric recruitment is equal to 0.042. This is about 0.7 percent of the mean wages in the symmetric equilibrium in the baseline, and 1.1 percent of the bottom percentile of the wage distribution. The required costs that are higher for the first stage than the second stage implies guaranteeing the optimality of the first stage also suffices to guarantee the optimality in the second stage.

The figures in the lower panel computes the corresponding numbers for the extended model. From figures C.1c and C.1d we conclude that assuming a value of rejection costs at 0.035 is sufficient to sustain the symmetric equilibrium in the extended model of Section 5. In relative terms, this is also about 0.7 percent of the mean wage in the respective model as we found in the baseline. In addition, the minimum κ's to assume in the alternative models shown in the paper, with respect to the corresponding mean wage in equilibrium are 0.9 percent in the large taste shocks model ($\sigma_\epsilon = 0.5$, used in Table 5), 0.8 percent in the high risk aversion model ($\gamma = 0.5$, used in Table 6), and 1.2 percent in the high complementarities model ($\rho = -1.5$, used in Figure 4).

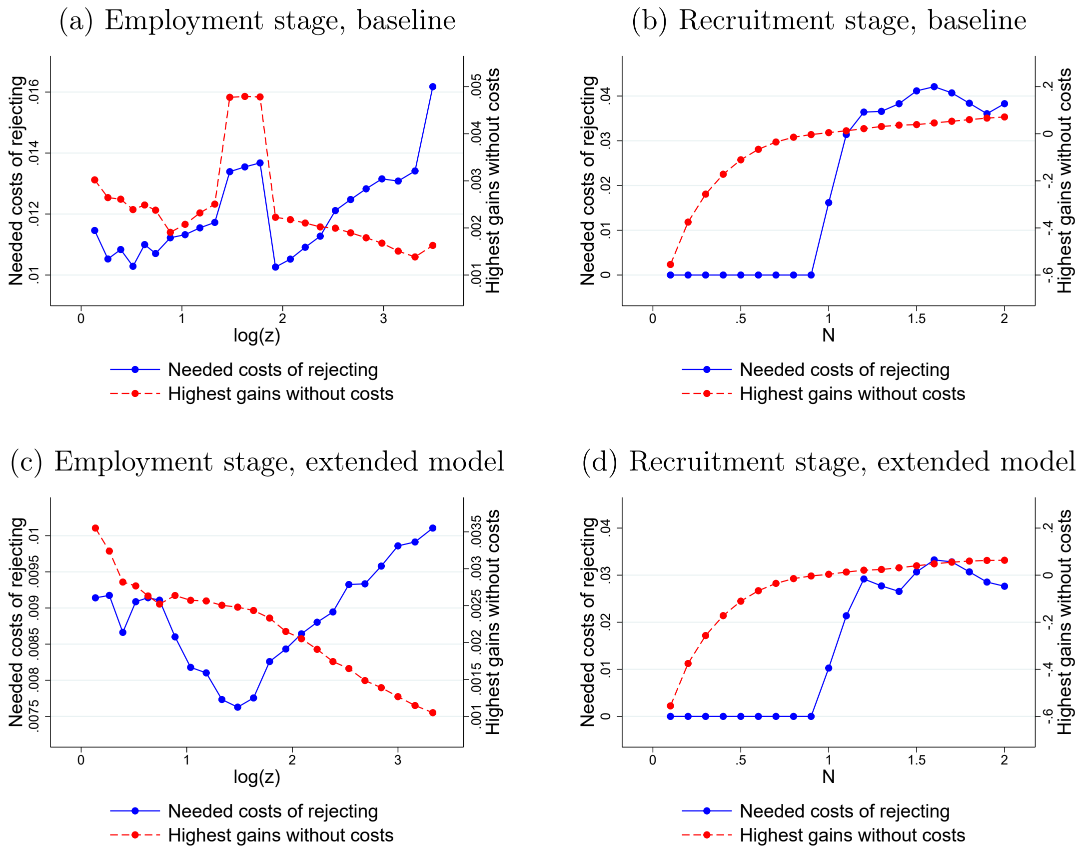

Figure C.1: Rejection costs to sustain symmetric equilibrium

Notes: The blue line in Figure (a) plot the minimum rejection costs κ needed to sustain the optimality of hiring all the recruited workers for the firms that recruited the symmetric share M_j^* in the benchmark model. The formula for such κ is given in Equation (C.10). The x-axis in these figures is the log of the productivity level z_j of the firms. The red line in Figure (a) is the highest profit gain among the alternatives (relative to the equilibrium demand with hiring all the recruitments, $M_j^* \Phi_j$) that would be attained when ignoring the rejection costs. The blue line in Figure (b) show for the benchmark model the minimum required rejection costs (as given in Equation (C.12)) to sustain the optimality of recruiting the symmetric share M_j^* in the first stage. The red line is the largest profit gains among the alternative M when ignoring the rejection costs. Figures (c) and (d) repeat figures (a) and (b) for the extended model.

C.5 Alternative construction of conditional wage-hours profiles.

To construct conditional wage-hours profiles using model simulated data, our preferred approach is to directly control for the level of worker efficiency x. This implicitly assumes that worker efficiency is observable to the econometrician. In this section, we show that even using noisy controls for worker efficiency implies smaller conditional wage penalties relative to unconditional wage penalties.

Specifically, we first generate an intermediate variable $\hat{x}$ that depends on the actual level of efficiency x and a noisy signal, $m \sim N(0, \sigma_m^2)$ such that,

$$\log \hat{x} = \log x + m$$

We assume that $\hat{x}$ is only observable in the form of deciles denoting, the decile d as $\hat{x}^d$. Then, in the model counterpart estimation of Equation (2) we include dummies for these deciles similar to our controls X in the data.

We choose, σ_m so that the residuals from the following simple OLS regression on log wages,

$$\log(w_i) = \alpha + \sum_{d \in [1,10]} \delta_d \mathbb{I}_{i,\hat{x}^d} + \epsilon_i,$$

accounts for the same share of variation in observed log-wages in the model as the controls X do in the data – in the data, we find this share to be around 25% and we choose σ_m to match this share in the data.

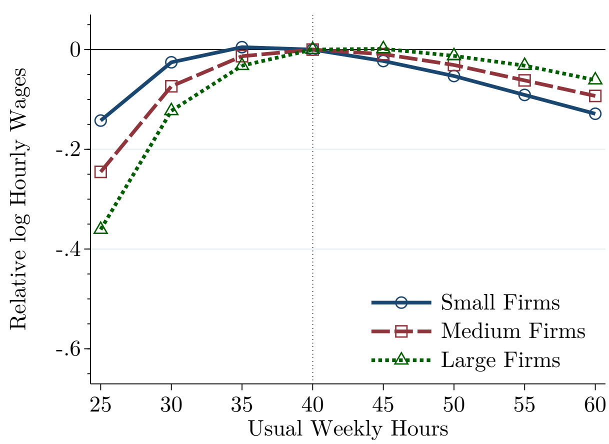

Figure C.2: Conditional wage-hours relationship by firm size, controlling for a noisy measure of x

Notes: The figure plots the sum of the coefficients $(\gamma_h + \theta_{e,h})$ from Equation (2) using model simulated data that controls for deciles of a noisy measure of x. Section 4.1 describes the construction of size categories in the model.

In Figure C.2, we compare the wage-hours profiles under alternative controls. In Panel (a), we do not use any controls, in Panel (b), we control for deciles of the noisy measure of x constructed as described above and in Panel (c) we fully control for x. Comparing these panels shows that qualitatively, the main message does not change regardless of our choice of controls. That is, wage-hours profiles are hump-shaped across all size categories and differ systematically across firm size as in the data. Quantitatively, when using noisy controls for x, wage penalties for short and long hours are roughly in between the results with no controls (Panel (a)) and with exact controls for each x level (Panel (c)). For instance, the 25-hour penalties in small firms is around 22% – in between 35% in Panel (a) and 15% percent Panel (b). The analogous measures in large firms are, respectively, 50, 60 and 40 %. The same pattern is evident for long-hour penalties.

Overall, we prefer to utilize the full controls for x since they are the simplest to implement and do not require us to take a stand on the nature of the noise with which worker efficiency is observed. Importantly, (and intuitively), noisy controls for x do not change the qualitative implications of our model and simply imply wage penalties that are systematically in between the unconditional and our preferred conditional levels.

C.6 Details on the effects of firm-level variation in hours given in Section 6.1

In this appendix we give the details on the computation and the results of the exercise on shutting down the firm-level variation in the reference hours in our extended benchmark model. In our main exercise reported in Section 6.1, we take the model's baseline distribution of workers across firms, with their hours as in the baseline, and compute the counterfactual wages they would have earned –with the same hours in the same firms– had the reference hours been computed across all workers in the economy (rather than only across co-workers). Formally, we compute wages still using the wage equation (15), but computing the reference hours taking the average across all workers over all firms:[4]

$$\tilde{l} \equiv \left[\left(\sum_{j=1}^{J} \sum_{x \in B_x} \sum_{l \in \mathcal{L}} x \Lambda_j \mu_j^*(l,x) l^{\rho} \right) \Big/ \left(\sum_{j=1}^{J} \sum_{x \in B_x} \sum_{l \in \mathcal{L}} x \Lambda_j \mu_j^*(l,x) \right) \right]^{\frac{1}{\rho}}$$

The variance of these counterfactual (log) wages in the economy is 1.8% higher than that of the actual wages. As we highlight in the paper, the dampening effect is more positive for short-hour workers, around 7% for 25-hours. This decreases over the hours spectrum, eventually turning into a negative effect for longer hours. In particular, having firm-level variation in the reference hours increases the dispersion within the 60-hour workers by 1.5%. The green short-dashed line in Figure C.3 documents the pattern of these effects for each level of hours, highlighting the gradual change from a dampening to an amplifying effect of firm-level variation in reference hours from short to long hours.

For simplicity, the exercise described above does not allow the workers to reoptimize their hours or the sorting into firms. In the first additional exercise to support the results above, we allow the workers to choose their hours optimally when they face the counterfactual wages with economy-wide reference hours. We set these reference hours as the same as in the previous exercise, coming from the average over all workers in our baseline. The variance

[4]These wages are not consistent with the production function that we assume, and imposed as a counterfactual. One can interpret such wages as a "mistake" that the firms make when they compute the marginal productivities in the economy.

of the (log) wages with the counterfactual wages, but with optimized hours, is 1% higher than the baseline. The pattern across hours with this alternative are illustrated by the red solid line in Figure C.3, showing virtually the same feature as in the first exercise.

Finally, we re-run the model, allowing the workers to optimize on everything (firms and hours to work, savings). Here, we also allow the reference hours to adjust, computed as the average in this counterfactual economy. This economy features variance of (log) wages that is 1.9% higher than the baseline, again highlighting a dampening effect of firm-level variation in wage dispersion. The blue long-dashed line in Figure C.3 shows that the dampening effect is as high as 10% for short hours, and becoming an amplification as large as 3% for long hours.

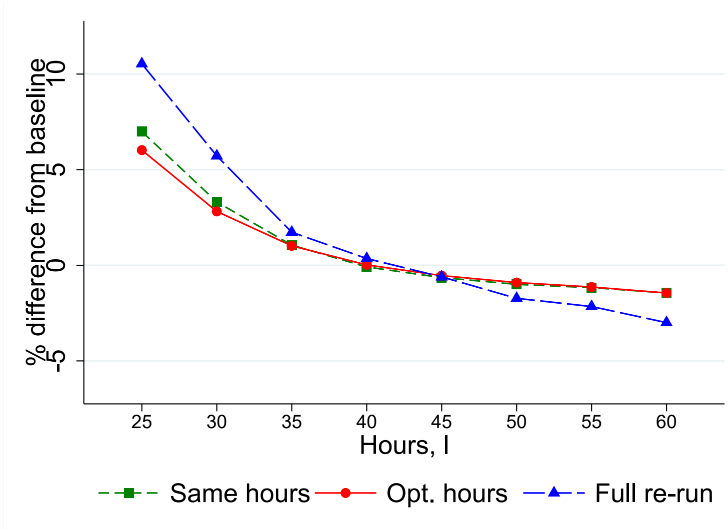

Figure C.3: Percent changes in the variance of log hourly wage, by hours worked

Notes: The figure displays the percent changes in the variance of log hourly wages within various hours, after imposing a common average hour that equals the weighted average determined by the population. The short-dashed green line gives the results of the main counterfactual, where we keep the hours and the employed firms of workers as in the benchmark. The solid red line allows the workers to choose their hours optimally, given their employers and wealth as in the benchmark. The long-dashed blue line gives the results in case of re-running the model allowing the workers to optimize in all dimensions.