Anatomy of the Pass-through of Productivity Shocks

Pablo Mileni Munari   Wenxuan Xu   Andrei Zaloilo *

September 10, 2025

Abstract

Workers experience labor income risk when employers adjust their hourly wages, hours worked, and their separation rates into unemployment in response to productivity shocks. Using French matched employer-employee data, we document how different employers adjust each of these margins for workers with different jobs, thus determining the earnings risk of their employees. We find that high-paying jobs adjust mainly hourly wages in response to changes in the unemployment rate, at 2.6 times that of low-paying ones. At the same time, low-paying jobs adjust primarily hours worked and separation rates. Adjusting hours worked at 40 times the semi-elasticity of high-paying jobs, and 10 times the semi-elasticity for separation rates. We develop an equilibrium labor market search model that incorporates dynamic contracts that allow firms to share risks with their workers through different margins. Firms share risks with workers using margins that are less costly to them, given their heterogeneous cost of creating vacancies and the job mobility of their workers. Consequently, government policies that aim to reduce labor income risk by targeting only one margin (e.g., minimum wage, hours restrictions, firing cost regulation) can be ineffective due to firms offloading risks into other margins.

Keywords: productivity, long-term contracts

JEL codes: D86, H23, J24, J31, J41, J62

---

*Munari, Xu, Zaloilo: Toulouse School of Economics and University of Toulouse Capitole. We are extremely grateful to Christian Hellwig, Fabrice Collard and Eugenia Gonzalez-Aguado for their invaluable advice and continuous guidance. We thank Dirk Krueger, Nicolas Werquin, Iourii Manovskii for generous feedback on our work. We greatly benefited from participants at the macro workshop at the Toulouse School of Economics and at UPenn. We also thank our discussant, Alexander Clymo, at the Junior SAM workshop, and participants of the XXVII Dynamic Macroeconomics workshop for constructive comments and insightful discussions.

1 Introduction

A long-standing literature considers the key role that firms play in determining the labor income risk of their workers (Knight, 1921). While prior work has focused on quantifying the pass-through of productivity shocks to workers' income only through wages (Guiso et al., 2005;Carneiro et al., 2012;Balke and Lamadon, 2022; Souchier, 2022), workers experience labor income risk not just from changes in wages, but also in hours worked and in the risk of separation into unemployment. Employers can choose how to utilize each margin to pass on productivity shocks to their employees, shaping workers' income risk. How firms decide to pass on productivity shocks to workers has clear consequences for risk sharing in the labor market, but also broader implications for job mobility and the design of government insurance.

Using French matched employer-employee data and data from administrative labor surveys, we uncover novel cross-sectional facts about how firms utilize these margins in response to productivity shocks. We find that high-paying jobs mainly adjust hourly wages, in contrast to low-paying ones that adjust working hours and their separation rate with workers. We develop a dynamic contracting framework where (i) firms form long-term contracts with works and are able to adjust all three margins, and (ii) employers differ in their vacancy cost. We then estimate the model parameters by matching it to our novel cross-sectional moments. Our theoretical framework and empirical results highlight the importance of jointly considering hourly wages, working hours, and the separation rate. Consequently, policies targeting only one of these margins can be ineffective due to the endogenous reaction of the others. Given the extensive set of labor market policies that affect these margins (e.g. minimum wage, working-time regulation, layoff restrictions), there can be significant welfare implications from the better design of these policies.

We start by presenting empirical results on firms' responses to productivity shocks. We apply approaches widely used in the empirical literature on wage cyclicality (e.g. Carneiro et al., 2012; Dapi, 2020) to estimate the response of firms to productivity shocks. We estimate the semi-elasticity of individual hourly wages, hours worked, and separation rates with respect to the aggregate unemployment rate. The richness of our administrative data allows us to control for contemporaneous composition effects using worker and job (occupation $\times$ firm $\times$ region) fixed effects (Abowd et al., 1999). We refer to the semi-elasticity of each margin to the unemployment rate as the shock pass-through on that margin. We estimate the heterogeneity of this pass-through across different margins and jobs. We find that high-paying jobs adjust hourly wages and keep working hours and the separation rate stable, in contrast to low-paying ones that keep wages stable and adjust working hours and the separation rate.

Next, we develop a theory of the pass-through of productivity shocks onto hours and separation rates, on top of wages. We leverage on recent work that embeds dynamic contracts into search and matching labor market models with on-the-job search (Menzio and Shi, 2011;Balke and Lamadon, 2022;Souchier, 2022). On one hand, risk-neutral employers find it profitable to insure risk-averse workers against productivity shocks, as this allows them to reduce their expected total payments - the insurance-provision motif. On the other hand, by passing positive shocks to workers, they discourage them from searching for alternative jobs and increase the chances of retaining them when they are the most productive - the incentive-provision motif.

While previous models assume that firms control income risk only through workers' wages, we add to the firm's decision the hours worked and the separation rate. We also introduce heterogeneous vacancy creation costs (hiring costs) across employers, which provides an ex-ante reason for firms to use different margins to insure workers. We provide empirical evidence for these different vacancy costs in the French labor survey. Having heterogeneous vacancy costs also requires moving away from the standard assumption of free entry at a vacancy value of zero. Crucially, this gives firms in our model an option value of waiting to create vacancies (Schaal, 2017), a source of ex-post heterogeneity. Incorporating option values means that a standard free entry condition can no longer characterize firms in our model. We adapt the modeling framework in Schaal (2017): firms discover their vacancy cost upon entering the market, then the tightness function, specific to each hiring cost, ensures that firms are indifferent across all submarkets. This ensures our dynamic contracting framework with non-zero vacancy value remains tractable. The firms' optimal contracts qualitatively generate volatile wages and stable hours and separation rates at high vacancy cost/high-paying employers, as well as stable wages and volatile hours and separation rates at low vacancy cost/low-paying employers.

The heterogeneity in vacancy creation costs alters the firm's decision problem by affecting their trade-off between incentive and insurance provision. For firms with high vacancy creation costs, it becomes more profitable to retain an employee because her employment value for the firm increases since it is more costly to create vacancies for workers that leave (incurring hiring costs). Moreover, the benefit of retaining an employee is more correlated with productivity states for high vacancy cost employers since they lose more when a job is vacant, pushing them to provide the worker with relatively higher utility in times of high productivity. In contrast, this incentive-provision motif is weaker for employers with low vacancy cost, giving them more leeway to smooth workers' utility across states.

To understand how the volatility of a worker's utility maps into the volatility of each margin of the labor contract, we consider first the case of an employer that provides cyclical utility (high vacancy cost). Instead of making all the margins volatile, the optimal way to deliver the variation in utility is to make hourly wages volatile and keep hours and separations stable. Wages act as a direct utility transfer and thus are proportionally cyclical to the utility. When it comes to hours, employers face a conflict between seizing the productivity boost and retaining the worker when she is the most productive. Indeed, if they increased expected working hours in a good state by adjusting working hours upwards, they would have to additionally increase the hourly wage to provide the extra utility needed to increase the probability of retaining the employee. However, this generates an income effect on labor supply which reduces the profitability of increasing expected working hours. Lastly, high vacancy costs reduce the employer's flexibility when it comes to downsizing: during a low productivity state, an employer with high vacancy costs may decide to keep the worker rather than fire them in order to avoid paying the hiring cost once the high productivity state arrives. In contrast, this plays a smaller role for employers with low vacancy cost , who can choose to fire workers when they become unproductive. Moreover, as they are less worried about keeping the workers during high productivity states, they can freely raise working hours during those periods. Finally, as they intend to provide stable utility to their workers, wages move significantly less than in high vacancy cost jobs.

Our theoretical framework together with our empirical results highlights a novel fact, that employers actively use working hours and the separation rate, in addition to wages, to pass through productivity shocks and hence determine labor income risk. Considering the three margins of labor income risk jointly has important policy implications. We show that policies targeting only one of these margins can be ineffective due to the endogenous reaction of the other ones. For example, policies introducing firing costs make it more costly to increase the separation rate. Low-paying employers, for which it is profitable to increase layoffs when hit by a bad productivity shock, would reduce hours or wages to compensate for the decreased profitability induced by the policy. As a result, the policy reduces unemployment risk at the expense of more pronounced volatility in hours and wages. Since most policies act on one of the margins in isolation (e.g. minimum wage, minimum and maximum hours constraints, firing costs regulation), we believe our results have important welfare consequences for designing these policies.

Related literature.— We relate to several strands of the literature investigating fluctuations in labor income. Firstly, we relate to theoretical and empirical studies on long-term contracts that view firms as insurance devices. The first to formalize this idea in an implicit contract model are Baily (1974) and Azariadis (1975). Subsequent work has introduced a lack of commitment on the worker side (Harris and Holmstrom, 1982) and a lack of commitment on the worker and firm side (Thomas and Worrall, 1988). More recent papers have investigated the role of firms in insuring workers in frictional labor markets facing different types of shocks, firm and worker level (Balke and Lamadon, 2022) and firm and sector level (Souchier, 2022). We complement these works by expanding the contracting space to allow for several margins through which the principal can provide utility to the worker. We also introduce heterogeneity in vacancy costs that meaningfully impacts this choice of margins.

Within the literature on the insurance role of firms, Guiso et al. (2005) estimates that insurance within the firm accounts for about 15% of overall earnings variability, suggesting a prominent insurance role of firms. Building on this evidence, Lagakos and Ordoñez (2011) quantifies the heterogeneity in insurance provision across firms. They find that workers in high-skilled sectors receive more insurance than those in low-skilled ones and relate this result to differences in displacement costs. We contribute by extending this empirical analysis of pass-through and its heterogeneity to all the margins of labor income: wages, hours, and separation risk, rather than focusing only on total earnings.

Second, we relate to the empirical literature on the heterogeneity of labor income cyclicality. Guvenen et al. (2017) documents how the aggregate risk exposure of individual earnings to GDP and stock returns varies along the distribution of workers' lifetime earnings, finding a U-shape pattern of this risk over the lifecycle. Relative to these papers, our work highlights the role of employer heterogeneity, specifically job-specific vacancy costs, in determining the degree of labor income cyclicality experienced by workers.

Lastly, there is a literature focused on understanding the statistical properties of idiosyncratic income processes (e.g. Arellano et al., 2017). Among these works, we are most closely related to Halvorsen et al. (2023), which disentangles the role of total hours and wages in generating the observed statistical features of idiosyncratic risk. In our paper, we separately quantify the contribution of the intensive and extensive margins of hours to labor income fluctuations, which allows us to study two different forms of risk: employment risk, and hours risk conditional on working. In addition, we provide a theory that highlights the mechanism underlying the observed cross-sectional heterogeneity of wages and hours fluctuations documented in these papers.

The remainder of the paper is structured as follows. Section 2 describes the data sources, presents the empirical specifications used to estimate the pass-through of shocks onto different margins. Section 3 describes the dynamic contracting framework and the labor market environment, and characterizes the pass-through of shocks for each margin of earnings risk. Section 5 provides a quantitative assessment of the proposed mechanisms and performs policy experiments.

2 Empirical analysis

This section describes our data sources and the main empirical result: high-paying jobs have volatile wages, stable hours, and separation rates, whereas low-paying ones have volatile hours and separation rates, and stable wages. Since we measure response to aggregate productivity shocks, we use the terminology used in the cyclicality literature (e.g. Carneiro et al., 2012 and Stüber, 2017) to describe these findings. Thus, we say that high-paying jobs have pro-cyclical wages, acyclical hours and separation rates whereas low-paying ones have procyclical hours and separation rates and acyclical wages.

2.1 Data

We use administrative tax and survey data from France between 2003 and 2019. The administrative dataset contains information on wages and hours worked, while the survey data includes additional information on workers' employment status, allowing us to better identify separations. We match the two datasets based on worker and firm identifiers, creating a unique dataset containing information about hourly wages (henceforth, "wages"), hours worked, and employment status at the worker-firm level. Importantly, the matched nature of our datasets allows us to relate heterogeneity in the cyclicalities of wages, hours, and separation rates across workers to differences in their respective employers.

Matched Employer-Employee Dataset.—The primary administrative data used in this work is the *Déclarations Annuelles de Données Sociales* (DADS), a matched employer-employee dataset built by the French Statistical Institute (INSEE) from the social contributions declarations of firms. It covers about 85% of all French workers and spans 1976-2019. It provides employment information (salaries, hours worked, occupation, tenure), worker information (age, experience), and firm information (sector, industry, size) on an annual basis. Salaries reported include regular, overtime, and bonus pay. Here, hours worked are remunerated hours and so include regular and overtime hours. The dataset is additionally available in panel form for a subsample (1/12th) of workers, which we also use in our analysis.

Sample Selection.— In the DADS dataset, we limit attention to the period going from January 2003 to December 2019 because of changes in the sampling methodology of DADS in 2002. For computational reasons, we take a subsample of the panel by randomly selecting 20 % of all workers. Since we subsample workers, we observe the whole labor market history of these workers.

In addition to the standard sample selection in the cyclicality literature (see Appendix A.1), we remove observations when the wage is less than 1.03 times the national minimum wage[1]. This prevents our cyclicality measures from capturing the mechanical wage rigidity induced by the minimum wage. Lastly, we remove observations in public sector jobs. We end up with about 27,000 workers and 24,000 firms per year.

Labor Force Survey.— We also employ data from the *Enquête Emploi en continu* (EEC), the French labor force survey. This is a rotating panel of workers who are interviewed for six consecutive quarters. Worker are asked about their current employment status. By comparing their responses across consecutive periods, we can infer whether a separation occurred or not. This allows us to create a dummy variable s_{ijt} which indicates if an individual i working at job j in period t becomes unemployed in $t+1$. In addition, since the firm identifier is provided and DADS contains the universe of French firms, we can assign observations in EEC to their respective earning bin as defined using firm-level data from DADS. We end up with about 24,500 workers and 4,000 firms per quarter[2].

Job Definition and Job Brackets.— Understanding the contribution of firms to the cross-sectional heterogeneity of income cyclicality require granularity on the jobs at these firms. As an example, a manager and a production line worker at Peugeot may have very different wage, hours, and separation rate cyclicalities. To achieve this level of granularity, we define a job as an occupation in a given firm in a given region, e.g. an engineer at Peugeot in Île-de-France (Paris region). For each job, we then compute the average wage that these jobs pay over the entire sample period. We rank them by their average wage percentiles and construct K equally-sized brackets, indexed by $k = 1, ..., K$. In our preferred

[1]We obtain this number by estimating the largest average wage change in a sample where only wages above the minimum wage are considered. To do so, we run the wage cyclicality regression (1) when $k = 3$ on a sample with wages above the minimum wage. We multiply the highest wage cyclicality estimate (-1.845) with the largest change in the unemployment rate over our sample period (1.605, 2008-2009): $-1.845 \times 1.605 \approx 0.03$. Finally, we keep only wages such that $0.97 \times w_{ijt} > minwage_t$, so that if firms wanted to reduce wages as much as 3% they could do so because the would not hit the minimum constraint.

[2]Our sample size is smaller than the full size of EEC, because the firm identifier is provided only for a subset of observations.

specification, $K = 3$, and we show that the qualitative patterns are unaffected when we increase the number of brackets, up to $K = 10$.

We rank jobs according to their average wage, rather than employment size or revenue. We do this to keep in line with the prior literature (see Haltiwanger et al. (2018) for a survey) that has shown that average wage differences across jobs is a key heterogeneity related to labor income risk. Appendix A.2 provides descriptive statistics on the whole sample and for each job bracket.

2.2 Measuring Cyclicality

The empirical strategy that we employ to measure the cyclicality of wages, hours worked, and separation rates is an extension of the worker-level wage regression proposed by Abowd et al. (1999) to study the determinants of wages. We estimate the co-movement of our dependent variables with the unemployment rate, a proxy of the business cycle, controlling for worker observed and unobserved heterogeneity, and unobserved job heterogeneity. This approach is also widely used in the empirical literature on wage cyclicality (e.g. Carneiro et al., 2012, Stüber, 2017, Guiso et al., 2005) to estimate the cyclicality of hours worked and separation rates.

Since we are interested in the heterogeneity of wages, hours and separation rate cyclicality in the cross-section of firms, we allow for an interaction term between the cyclicality coefficient and firm bracket indicators. Formally, the cyclicality of wages, hours worked and separation rates in group k are measured by α_k^y for different outcomes y in separate linear regressions:

$$log(w_{ijt}) = \alpha^w_{k(j)}\mathbf{1}\{j \in k\} \cdot u_t + \alpha^w_0 u_t + \beta^w x_{it} + \mu^w_{k(j)}(t) + FE_i + FE_j + \epsilon_{ijt} \tag{1}$$

$$log(h_{ijt}) = \alpha^h_{k(j)}\mathbf{1}\{j \in k\} \cdot u_t + \alpha^h_0 u_t + \beta^h x_{it} + \mu^h_{k(j)}(t) + FE_i + FE_j + \epsilon_{ijt} \tag{2}$$

$$s_{ijt} = \alpha^s_{k(j)}\mathbf{1}\{j \in k\} \cdot u_t + \alpha^s_0 u_t + \beta^s x_{it} + \mu^s_{k(j)}(t) + \epsilon_{ijt} \tag{3}$$

where w_{ijt} is the hourly wages, and h_{ijt} is the daily hours worked of individual i at job j in period t, and s_{ijt} is a dummy variable equal to 1 if individual i working at job j in period t becomes unemployed in $t+1$, i.e. if she is separated; u_t is the unemployment rate in period t in France as reported by the national statistical agency; x_{it} is a vector of time-varying worker characteristics such as age and experience; $\mu_{k(j)}(t)$ is a group-specific linear time trend; FE_i and FE_j are worker and job fixed effects. Regressions in Equations

(1) and (2) are at the annual frequency, while our data allow Equation (3) to be estimated at the quarterly frequency, which provides a more accurate picture of separations.

Our matched employer-employee dataset allows us to account for compositional effects in the workforce over the business cycle through worker and job fixed effects. This isolates changes in our dependent variable over the business cycle from composition effects associated to differences in the pool of workers and firms between booms and recessions, which have been shown to be significant in previous work[3]. We do not introduce worker and job fixed effects in regression (3) [4]. Regressions (1), (2) and (3) are estimated by OLS. Standard errors are clustered at the regional level.

2.3 Empirical Results

This section details the main empirical finding. We find that high-paying jobs have cyclical wages, acyclical hours and separation rates. Whereas low-paying jobs have cyclical hours and separations rates, and acyclical wages.

2.3.1 High-paying Jobs: Cyclical wages, acyclical hours and separation rates

Figure 1 shows the estimates of coefficient α_k^y of Equations (1), (2) and (3), with $K = 3$. Appendix A.3 provides the full results. In the upper left panel of Figure 1, we observe that wages are the most cyclical for high-paying jobs, while in the upper right panel hours are the most procyclical for low-paying ones. On the other hand, in the bottom panel, we see that low-paying jobs have the most counter-cyclical separation rates.

As expected, the wages and hours worked are procyclical, and separations are counter-cyclical. Our wage cyclicalities measures are also in the range of existing estimates. Gertler et al. (2020) obtains a semi-elasticity ranging from -1.1 to -1.9 for log wage cyclicality in the US using the SIPP panel. As a robustness check, we provide the analogous figure for $K = 10$ job brackets in Appendix **??**, which retains the qualitative pattern above.

[3]See **?** for downward bias of wage cyclicality due to changes in the composition of workers; Hagedorn and Manovskii (2013) for upward bias on wage cyclicality due to changes in the composition of jobs.

[4]While we can introduce worker and job fixed effects in this regression, we are concerned about the effect of the incidental parameters bias from limited mobility in our regression (Andrews et al., 2008;Bonhomme et al., 2023). We believe that this is more pronounced for separations due to the limited number of separations we observe in worker-firm pairs.

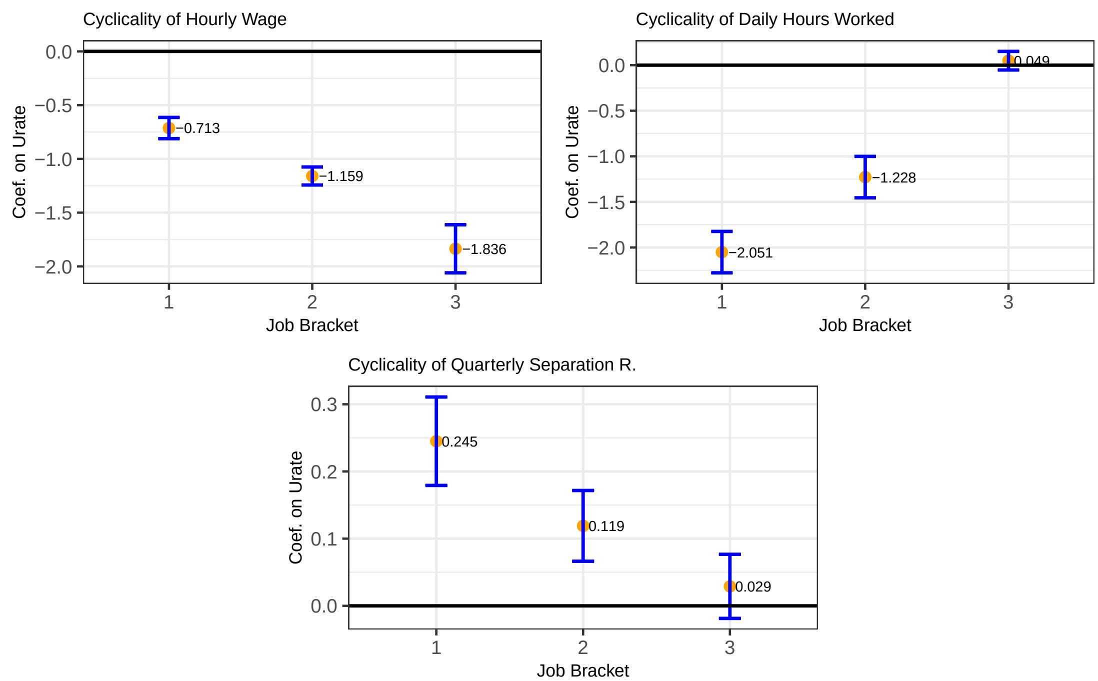

Figure 1: Cyclicalities of the Three Margins of Labor Income Risk by Job Bracket. Point estimates in yellow; 95% C.I. in blue

2.3.2 High-paying jobs: Higher vacancy costs

A main source of ex-ante heterogeneity that differentiates firms in our model is their vacancy creation costs. Here, we provide evidence that jobs that have higher vacancy costs are those that pay higher wages.

Leveraging on the assumption that vacancy costs for jobs in high-skilled occupations are on average higher than for low-skilled ones [5], we compute the share of jobs above the median of the wage distribution that are in high-skilled occupations ("managers of firms with more than 10 employees", "administration and sales executives", "engineers and technical officers" according to PCS2017 classification of INSEE). This is around 14%. Instead, when we focus on jobs below the median, this share drops to 0.5%.

This provides preliminary and indirect evidence that high-paying jobs are jobs with high vacancy costs. We plan to use more precise measures of vacancy costs, such as the time to fill a vacancy or the share of unsuccessful recruiting activity, provided by the French Ministry of Labor, to provide direct evidence of these costs (see Le Barbanchon et al., 2022).

[5]Some reasons can be because of lower relative supply of high-skilled workers, or because higher job-specific human capital is required to accomplish task in high-skilled occupations.

3 Model

This section presents an equilibrium search and matching model with on-the-job search, where firms use long-term labor contracts, drawing from Menzio and Shi (2011) and Balke and Lamadon (2022). This allows us to study how risk-neutral firms pass idiosyncratic and aggregate productivity shocks to risk-averse workers. We augment Balke and Lamadon (2022) by incorporating hours worked and the separation rate as part of the contract, in addition to wages, and introduce heterogeneous vacancy cost.

3.1 Environment

Time is discrete and infinite, and is indexed by t.

Agents.— There is a continuum of ex-ante homogeneous workers that are either be employed or unemployed. When employed, they supply h hours of work in exchange of earnings I. There is no savings technology, so workers consume their current earnings I when employed, and home production b when unemployed [6]. Workers have period utility $u(I, h)$ that is increasing and concave in consumption I (i.e. they are risk-averse) and decreasing and convex in labor supply h. They maximize the expected present value of their utility streams, discounted with a discount factor β.

There is a continuum of heterogeneous firms indexed by j. Firms maximize the present value of their profits (output), discounted at the common discount factor β. Each firm can employ at most one worker. If a firm has a worker, it can produce output, and has to pay wages due to the worker, and if a firm does not have a worker, it can choose to post a job vacancy. Output is produced according to a decreasing returns to scale technology that uses labor h as its sole input:

$$y_{jt} = z_t f(h_{jt}) \tag{4}$$

where z_t is the level of aggregate productivity at time t and we assume $f' > 0, f'' < 0$. We will also write the total payment to the worker for h hours worked as a result of a firm's contracting decisions as $I(h)$.

Firms may only hire if they have zero workers. Attempting to hire a worker requires paying a per-period vacancy posting cost c_j which varies across firms. Each period, they

[6]The ability to save complicates the dynamic contracting problem and have seen recent attempts to characterize the optimal contracts with observable/unobservables savings (see Dilmé and Garrett, 2023). We abstract from this in the paper, but can potentially extend the model in this direction.

can pay this vacancy posting cost to open the vacancy for hiring, or keep the vacancy closed and not be hiring until the next period. This is make our model different from a standard CRS production function. By making hiring costs c_j heterogeneous across firms, we allow for non-zero value of a vacancy. Firms differ in their values of existing without the worker, which makes a firm's intertemporal trade-off of vacancy posting non-trivial, much like Schaal (2017). In Section 4, we show that this heterogeneity can help reconcile our finding of cross-sectional heterogeneity in cyclicality across the three margins across firms.

Labor Market.— There is a continuum of labor markets indexed by the promised value v that is delivered to a worker that finds a job in that submarket. Every period, workers can choose one submarket v to apply to, and firms choose where to post vacancies.

The workers and vacancies that are searching in the same submarket are brought into contact by a meeting technology with constant returns to scale that can be described in terms of the vacancy-to-worker ratio θ (i.e., the tightness). This tightness is specific to each submarket, which means that in submarkets with a high ratio of vacancies to workers it is harder for firms to hire. We denote by $p(v)$ the probability that a worker finds a job in submarket v, and $q(v)$ the probability that a firm fills a vacancy that was posted in submarket v. The job finding probability $p(v) \in [0, 1]$ is assumed to be twice continuously differentiable, strictly increasing and strictly concave in v, where $p(0) = 0$ and $p(\infty) = 1$. Similarly, $q(v) \in [0, 1]$ is a twice continuously differentiable, strictly decreasing and strictly convex function.

As in Schaal (2017), we assume that firms together post a mass and not a discrete number of vacancies. As a result, a law of large numbers applies that allows firms to meet worker in a submarket at the expected rate.

Timing.— In each period t, the timing of events is as follows:

1. Productivity z_t is realized
2. Separations into unemployment occur at rate s_j
3. Job mobility phase: employed and unemployed workers search for jobs; firms post vacancies; new matches are formed; new contracts are signed
4. Firms that have a worker produce and pay their current wages; workers provide labor and consume their wage

Contracts.— When a worker and a firm match, they sign a contract. The contract is subject to limited commitment by the worker. This captures the on-the-job search that workers engage in the is unobservable and non-contractible by firms, and results in job mobility. On the other hand, firms have commitment power in the contracts that they participate in. Additionally, we assume that firms cannot make counteroffers to workers' on-the-job searches, as in Burdett and Mortensen (1998). Following previous work on dynamic contracts (e.g. Spear and Srivastava, 1987), we write the contract recursively in terms of promised values and continuation values. In the dynamic contract, the promised utility to a worker at the start of the period is also denoted by v.

The contract $\mathcal{C}$ specifies the current labor income, the associated amount of current hours of work, a set of future separation rates and promised utilities for each state in the next period. The firm optimal current labor income I and the hours of work h are a function of the current promised utility v as well as the current productivity state z, i.e. $I(v,z)$ and $h(v,z)$. While the set of future state-contingent separation rates and promised utilities are also a function of next period's productivity state, i.e. $s'(v,z;z')$ and $v'(v,z;z')$. Formally a contract $\mathcal{C}$ is then:

$$\mathcal{C} \equiv \left\{ I, h, \{s(z')\}_{z'}, \{v(z')\}_{z'} \right\} \tag{5}$$

where we suppressed the dependence on v and z for notational ease.

3.2 Worker's Problem

Unemployment.— An unemployed worker looks for a job in the submarket that maximizes her expected utility. She faces a trade-off between high future promised utility and low probability of a match. Hence, the value of being unemployed is:

$$U = \sup_{v_0} u(b,0) + \beta \left[p(v_0)v_0 + (1 - p(v_0))\, U' \right]$$

where we abuse notation and denote v_0 as the promised utility offered in a submarket that the unemployed worker searches.

Employment.— An employed worker chooses the submarket in which looking for a job conditional on the future promised value of her current job v'. The value of being employed when the current promised value is v and the current state is z is:

$$W(v) = \sup_{v_1} u(I,h) + \beta E_{z'|z} \left[(1 - s(z')) \left((1 - p(v_1))\, v' + p(v_1)v_1 \right) + s(z')U' \right]$$

3.3 Firm's Problem

There are active and inactive firms. Active firms are already matched with a worker at the beginning of the period, they enjoy the value of a filled job but must decide on the current hours/wage combination of its worker. Inactive firms are not matched to a worker, and so enjoy the value of a vacant job.

Value of a Vacant Job.— When inactive, the firm decides between opening a vacancy today or waiting. If a vacancy is opened in submarket v, the firm pays a cost c and obtains the value of a filled job at probability $q(v)$. If a vacancy is not opened, the firm enjoys the discounted expected value of a vacant job. Hence, we rearrange the value of the two options that determine the current value of a vacant job as:

$$V(z) = \max \Big\{ \sup_{v_2} \big\{ -\frac{c}{q(v_2)} + J(z, v_2) \big\}, \beta E_{z'|z} V(z') \Big\}$$

Value of a Filled Job.— When active, the problem of the firm consists in choosing a contract $\mathcal{C}$, as defined in Equation (5), for its incumbent worker. Its objective is to maximize the stream of profits derived from its current period profits $zF(h) - I$, plus the continuation value. This continuation value is the future value of a filled job if the worker stays (i.e if she does not find a better job and if she is not fired), or it is the future value of a vacant job if the worker does not stay. Hence, the value of a filled job is:

$$J(z, v) = \max_{I,h,\{v',s'\}} zF(h) - I + \beta E_{z'|z} \Big[(1 - p(v'))(1 - s') \, J(z', v') + [1 - (1 - p(v'))(1 - s')] V(z') \Big]$$

$$s.t. \quad u(I, h) + \beta E_{z'|z} \left[(1 - s') \left(p(v') \, \hat{v} + (1 - p(v')) \, v' \right) + s' U(z') \right] \geq v \tag{PK}$$

$$\hat{v}(v') = \arg\max_{v_1} \, (1 - p(v_1))v' + p(v_1)v_1. \tag{IC}$$

Inequality (PK) is a promise-keeping constraint which forces the firm to determine a combination of current wages, hours, and future state-contingent promised values that delivers at least the current promised value v. Equation (IC) is an incentive-compatibility constraint which allows a firm to know which submarket v_1 the worker chooses to search on-the-job as a function of the future promised value v' stipulated by the contract.

3.4 Equilibrium

Following Balke and Lamadon (2022), we define a stationary equilibrium in this economy as a set of symmetric value and policy functions of the firms and workers that are optimal given the firm and worker problems defined in Section 3.2 and Section 3.3. Market clearing requires the market tightness in each submarket generates laws of motions of vacancies and worker flows that are consistent with the policy functions of firms.

4 Characterizing the Optimal Contract

In this section, we analyze the trade-offs faced by firms when setting each margin of the contract. Our focus is on understanding how differences in the vacancy cost across firms make it more profitable to pass shocks through one margin or another. The trade-offs across margins are shaped by a tension between insurance and incentive provision (Balke and Lamadon, 2022). Workers search on-the-job, this search is directed, and firms compete to create vacancies by offering dynamic contracts. While firms can commit to contracts, workers cannot: their choices about which sub-market to search and the level of effort to put into retaining their current jobs are unobservable and hence non-contractible. As a result, firms face a trade-off between fully insuring the worker from shocks and providing higher utility in periods of high productivity, when the benefit of retaining her is larger.

4.1 Income cyclicality

Proposition 1 *(Optimal income growth). For any current state (z, v), the following relationship between income and retention benefit holds:*

$$\eta(v') \left[J(z', v') - V(z')\right] = \frac{1}{u_{I'}(I', h')} - \frac{1}{u_I(I, h)} \tag{6}$$

where $\eta(v') \equiv \left(-\frac{\partial p(v')/\partial v'}{1-p(v')}\right)$ represents the increase in the retention probability induced by a marginally higher promised utility in the next period, $J(z', v') - V(z')$ is the benefit of retaining in a worker in future state (z', v'), I and h are current income and hours, I' and h' are the next period income and hours in state (z', v').

Proposition 1 draws on Proposition 2 in Balke and Lamadon (2022) in our context with non-zero value of a vacant job. It states that income (marginal utility) changes across periods (the right hand-side) are associated with the benefit of retaining a worker in a particular future state z' (the left hand-side). If the retention benefit is positive, then income grows, while if the retention benefit is negative, income decreases.

Equation (6) provides intuition on how firms optimally choose their worker's income cyclicality. Consider a two-state stochastic process for the productivity shock: $z \in \{z_h, z_l\}$, with a high productivity state h and a low productivity state l.

Corollary 1 *(Optimal income cyclicality). Suppose that the increase in the retention probability is the same across states, i.e. $\eta(v_h) = \eta(v_l) = \eta$. Then, the optimal degree of income cyclicality is given by:*

$$\eta \left[(J(z_h, v_h) - V(z_h)) - (J(z_l, v_l) - V(z_l)) \right] = \frac{1}{u_{I_h}(I_h, h_h)} - \frac{1}{u_{I_l}(I_l, h_l)} \tag{7}$$

where $(J(z_h, v_h) - V(z_h)) - (J(z_l, v_l) - V(z_l))$ represents the difference in the retention benefit between the high and low productivity state.

Corollary 1 shows that utility difference across states, i.e. cyclicality, occurs if and only if the benefit of retaining a worker differs between the high and low productivity states. Here, differences in the cyclicality of income across firms is driven by differences in the cyclicality of the retention benefit.

We analyze now how differences in the job-specific cost of replacing a worker - captured by the vacancy cost c - shape the heterogeneity of income cyclicality across jobs. Later, we will examine the consequences for the heterogeneity in the cyclicality of margins.

Lemma 1 *(Heterogeneity in the cyclicality of the retention benefit). The benefit of retaining a worker is higher in the high productivity state than in the low one. The pro-cyclicality of the retention benefit (weakly) increases in the vacancy cost c.*

The intuition for the first part of the lemma is that, if the worker leaves, the firm will become inactive, and the associated foregone revenue is larger in a high productivity state. The intuition for the second part of the lemma is that firms with a high vacancy cost are more inclined to post vacancies in the high productivity state than in the low one because the retention benefit is more cyclical for these firms, i.e. the difference in value of a filled job $J()$ from the value of a vacancy $V()$ changes more with productivity. This implies that the benefit of retaining a worker fluctuates more between states: from not wanting to retain a worker at all in a low productivity state to wanting to retain him at all cost in a high productivity one.

The above intuitions about the optimal contract speaks to the heterogeneous cyclicality of income across jobs.

Proposition 2 *As the job-specific vacancy posting cost c increases, the cyclicality of income rises.*

The result follows from Corollary 1 and Lemma 1.

4.2 Margin Cyclicality

While the above results concern the cyclicality of wages and its heterogeneity in the cross-section of jobs, this can be extended to the cyclicality of other margins of labor income.

Proposition 3 *(Link between optimal income and hours cyclicality). For any contract, as the (pro-)cyclicality of income increases the (pro-)cyclicality of hours decreases. Indeed, in the optimal contract income and wages satisfy the standard efficiency condition between intra-temporal marginal rate of substitution and labor productivity:*

$$zF'(h) = -\frac{u_h(I,h)}{u_I(I,h)} \tag{8}$$

As the vacancy cost c increases, the pro-cyclicality of hours decreases.

Equation (8) indicates that changes in income exert a standard income effect on labor supply: a higher level of I decreases the marginal utility of consumption and hence makes it costlier for the firm to compensate the worker for an additional hour of work. This implies that, for a given productivity state, hours decrease when income increases. Then, as observed in the data, jobs with higher income pro-cyclicality are jobs with lower hours pro-cyclicality, and may even have counter-cyclical hours.

We can now put together the results about income and hours cyclicality to examine the behavior of wage cyclicality in the cross-section. This is easily done since the hourly wage is the ratio between income and hours worked.

Corollary 2 *(Optimal wage cyclicality). In the optimal contract, the cyclicality of wages increases as the vacancy cost c increases.*

The result follows from Proposition 1 and 3 that state that income pro-cyclicality is increasing in c and hours pro-cyclicality is decreasing in c, respectively.

5 Quantitative Analysis

While the theoretical analysis provides intuitions on why firms might have empirically chosen to employ different margins to insure their workers, this is insufficient for us to quantify better designs of labor market policies that improve insurance provision or job mobility. To design better policies, we need to bring the model to the data to see how firms would respond to policy that aim to restrict the way they currently employ different margins.

To bring the model to the data, we need to tackle two important challenges. The first concerns the identification of the inactive vacancies. Firms can hold on to vacancies without posting them in our model, which drives an option value for the firm. These vacancies are not directly observable. We intend to draw from Schaal (2017) in looking at the correlation of hiring, layoffs, and quits within the firm. Intuitively, changes in firm size over the business cycle is informative of how reluctant firms are in downsizing, and how they might be keeping vacancies inactive in downturns. The second concerns the solution to the optimal contracting problem. Here, we draw heavily from Balke and Lamadon (2022) in applying the recursive Lagrangian approach of Marcet and Marimon (2019). While optimizing the promised utilities over all future states is intractable, the recursive Lagrangian approach notes that the first order condition of the contracting problem requires firms to equalize worker marginal utilities over all future states. This allows us to optimize over marginal utilities rather than state-contingent promised utilities, increasing its traceability.

We are in the process of quantification of the model.

6 Conclusion

We study how employers determine the labor income risk faced by their employees by adjusting hourly wages, working hours and the separation rate in response to productivity shocks. We find that high-paying jobs adjust mainly hourly wages, whereas low-paying ones adjust working hours and the separation rate. Additionally, we observe that high-paying jobs have a relatively higher vacancy cost. To explain these novel cross-sectional facts, we augment a dynamic contracting framework to incorporate the three margins in a search and matching labor market model. The model describes the trade-offs across margins in a way that qualitatively replicates the observed cross-sectional heterogeneity in pass-throughs, pointing to the importance of considering the interactions between hourly wages, working hours, and the separation rate. We plan to quantitatively match the relevant cross-sectional moments, and to discuss the design of labor market policies. In particular, we plan to analyze the effectiveness of policies that target only one margin (e.g. minimum wage, minimum and maximum hours constraints). Indeed, we hypothesize that the endogenous reaction of untargeted margins may dampen the direct effects of these policies. Given their pervasiveness in current labor market regulation, we hope that we will add to the policy debates on their effectiveness.

References

Abowd, John M., Francis Kramarz, and David N. Margolis (1999) "High Wage Workers and High Wage Firms," *Econometrica*, 67 (2), 251–333, https://doi.org/10.1111/1468-0262.00020.

Andrews, M. J., L. Gill, T. Schank, and R. Upward (2008) "High Wage Workers and Low Wage Firms: Negative Assortative Matching or Limited Mobility Bias?" *Journal of the Royal Statistical Society Series A: Statistics in Society*, 171 (3), 673–697, 10.1111/j.1467-985X.2007.00533.x.

Arellano, Manuel, Richard Blundell, and Stéphane Bonhomme (2017) "Earnings and Consumption Dynamics: A Nonlinear Panel Data Framework," *Econometrica*, 85 (3), 693–734, https://doi.org/10.3982/ECTA13795.

Azariadis, Costas (1975) "Implicit contracts and underemployment equilibria," *Journal of Political Economy*, 83 (6), 1183–1202.

Baily, Martin Neil (1974) "Wages and employment under uncertain demand," *The Review of Economic Studies*, 41 (1), 37–50.

Balke, Neele and Thibaut Lamadon (2022) "Productivity Shocks, Long-Term Contracts, and Earnings Dynamics," *American Economic Review*, 112 (7), 2139–77, 10.1257/aer.20161622.

Bonhomme, Stéphane, Kerstin Holzheu, Thibaut Lamadon, Elena Manresa, Magne Mogstad, and Bradley Setzler (2023) "How Much Should We Trust Estimates of Firm Effects and Worker Sorting?" *Journal of Labor Economics*, 41 (2), 291–322, 10.1086/720009.

Burdett, Kenneth and Dale T. Mortensen (1998) "Wage Differentials, Employer Size, and Unemployment," *International Economic Review*, 39 (2), 257–273, http://www.jstor.org/stable/2527292.

Carneiro, Anabela, Paulo Guimarães, and Pedro Portugal (2012) "Real wages and the business cycle: Accounting for worker, firm, and job title heterogeneity," *American Economic Journal: Macroeconomics*, 4 (2), 133–152, 10.1257/mac.4.2.133.

Dapi, Bjorn (2020) "Wage Cyclicality and Composition Bias in the Norwegian Economy," *The Scandinavian Journal of Economics*, 122 (4), 1403–1430.

Dilmé, Francesc and Daniel F. Garrett (2023) "Relational Contracts: Public versus Private Savings," *Econometrica*, 91 (3), 1025–1075, https://doi.org/10.3982/ECTA18742.

Gertler, Mark, Christopher Huckfeldt, and Antonella Trigari (2020) "Unemployment Fluctuations, Match Quality, and the Wage Cyclicality of New Hires," *The Review of Economic Studies*, 87 (4), 1876–1914, 10.1093/restud/rdaa004.

Guiso, Luigi, Luigi Pistaferri, and Fabiano Schivardi (2005) "Insurance within the firm," *Journal of Political Economy*, 113 (5), 1054–1087.

Guvenen, Fatih, Sam Schulhofer-Wohl, Jae Song, and Motohiro Yogo (2017) "Worker Betas: Five Facts about Systematic Earnings Risk," *American Economic Review*, 107 (5), 398–403, 10.1257/aer.p20171094.

Hagedorn, Marcus and Iourii Manovskii (2013) "Job Selection and Wages over the Business Cycle," *American Economic Review*, 103 (2), 771–803.

Haltiwanger, John C., Henry R. Hyatt, Lisa B. Kahn, and Erika McEntarfer (2018) "Cyclical Job Ladders by Firm Size and Firm Wage," *American Economic Journal: Macroeconomics*, 10 (2), 52–85, 10.1257/mac.20150245.

Halvorsen, Elin, Hans A Holter, Serdar Ozkan, and Kjetil Storesletten (2023) "Dissecting Idiosyncratic Earnings Risk," *Journal of the European Economic Association*, 22 (2), 617–668, 10.1093/jeea/jvad047.

Harris, Milton and Bengt Holmstrom (1982) "A theory of wage dynamics," *The Review of Economic Studies*, 49 (3), 315–333.

Knight, Frank Hyneman (1921) *Risk, uncertainty and profit*, 31: Houghton Mifflin.

Lagakos, David and Guillermo L. Ordoñez (2011) "Which workers get insurance within the firm?" *Journal of Monetary Economics*, 58 (6), 632–645, https://doi.org/10.1016/j.jmoneco.2011.11.009.

Le Barbanchon, Thomas, Maddalena Ronchi, and Julien Sauvagnat (2022) "Hiring Difficulties and Firms' Growth."

Marcet, Albert and Ramon Marimon (2019) "Recursive Contracts," *Econometrica*, 87 (5), 1589–1631, http://www.jstor.org/stable/45217221.

Menzio, Guido and Shouyong Shi (2011) "Efficient Search on the Job and the Business Cycle," *Journal of Political Economy*, 119 (3), 468–510, 10.1086/660864.

Schaal, Edouard (2017) "Uncertainty and unemployment," *Econometrica*, 85 (6), 1675–1721.

Souchier, Martin (2022) "The Pass-through of Productivity Shocks to Wages and the Cyclical Competition for Workers,"Technical report, Wharton School of the University of Pennsylvania.

Spear, Stephen E. and Sanjay Srivastava (1987) "On Repeated Moral Hazard with Discounting," *The Review of Economic Studies*, 54 (4), 599–617, http://www.jstor.org/stable/2297484.

Stüber, Heiko (2017) "THE REAL WAGE CYCLICALITY OF NEWLY HIRED AND INCUMBENT WORKERS IN GERMANY," *The Economic Journal*, 127 (600), 522–546, http://www.jstor.org/stable/45022972.

Thomas, Jonathan and Tim Worrall (1988) "Self-enforcing wage contracts," *The Review of Economic Studies*, 55 (4), 541–554.

A Data appendix

A.1 Sample selection

We only keep observations of workers in their prime working years from ages 25 to 55. We remove workers who had on average more than 4 different jobs a year. We drop workers who had on average more than 4 different jobs a year. On the job side, we restrict attention to private sector jobs lasting at least a month, with weekly hours between 10 and 100.

A.2 Descriptive statistics

Table 1: Summary Statistics of Jobs in each Job Bracket

Bracket (k)	1	2	3
Avg wage	11.41	14.73	28.28
Avg hours	5.87	6.52	6.60
Median size (firm)	45	139	298
Median size (est)	17	32	77
Share of firms with >10 workers	0.71	0.81	0.87
Share of with >50 workers	0.48	0.61	0.70
Avg worker age	37.86	38.64	40.33
Avg worker exp.	14.17	15.50	16.77
Share of part-timers	0.37	0.19	0.17
Share of permanent contracts	0.52	0.65	0.76

A.3 Additional Cyclicality Results

We report here the regression tables related to the estimation of (1), (2) and (3).

Table 2: Cyclicality Regression Results

	$log(h)$	$log(w)$	s
unemp × Br. 1	-2.051***	-0.7127***	0.245**
	(0.1385)	(0.0598)	(0.0400)
unemp × Br. 2	-1.229***	-1.159***	-0.119*
	(0.1382)	(0.0516)	(0.0324)
unemp × Br. 3	-0.0496	-1.836***	-0.029
	(0.0625)	(0.0292)	
exp	-1.301***	2.093***	
	(0.0662)	(0.1244)	
age	0.5932	1.693	-0.0008**
	(2.986)	(1.896)	(0.00008)
age^2	-0.0043***	-0.0389***	
	(0.0008)	(0.0045)	
Fixed-effects			
Job	Yes	Yes	No
Worker	Yes	Yes	No
Observations	4,048,380	4,048,380	304,396
R^2	0.76460	0.88046	0.0097
Within R^2	0.00224	0.07727	0.0094

Clustered standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

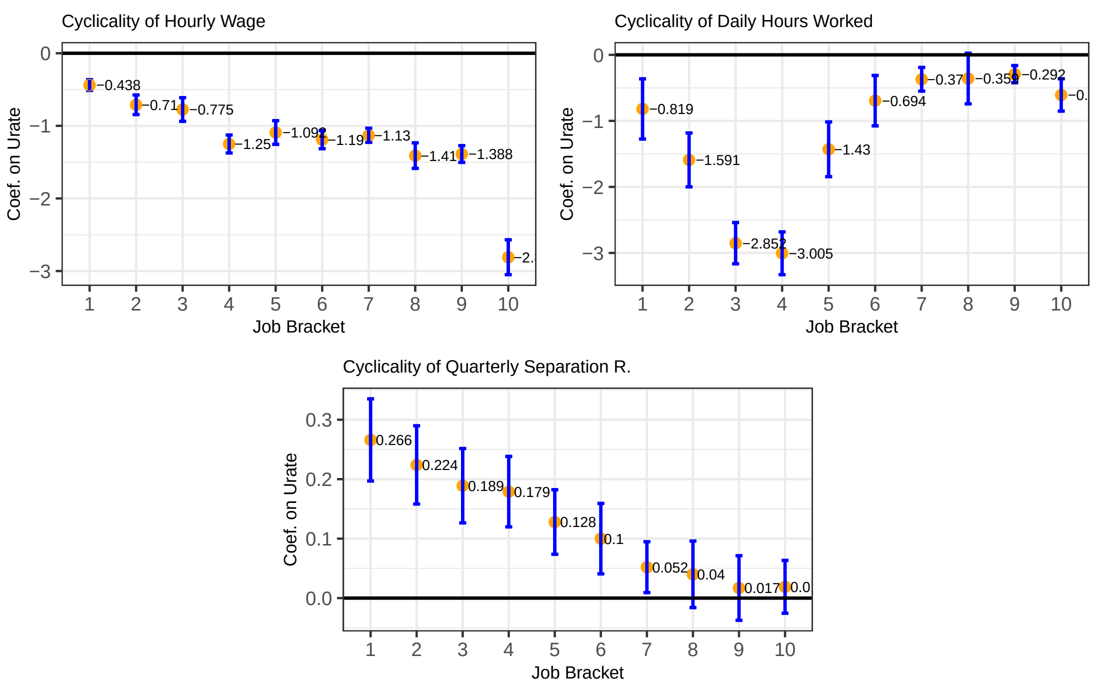

Figure A1: Cyclicality Results for 10 Job Brackets