



A decomposition of labor earnings growth: Recovering Gaussianity?

Pierre Pora, Lionel Wilner

► To cite this version:

Pierre Pora, Lionel Wilner. A decomposition of labor earnings growth: Recovering Gaussianity?. Labour Economics, 2020, 63, pp.101807. <10.1016/j.labeco.2020.101807>. <hal-04799388>

HAL Id: hal-04799388

<https://hal.science/hal-04799388v1>

Submitted on 22 Nov 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

A Decomposition of Labor Earnings Growth: Recovering Gaussianity?*

Pierre Pora[†] Lionel Wilner[‡]

January 11, 2019

Abstract

Recent works have concluded to non-Gaussian features of labor earnings growth. We argue in this paper that it is mainly due to working hours' volatility. Using the non-parametric approach developed by [Guvenen et al. \(2016\)](#), we find on French data that labor earnings changes exhibit strong asymmetry as well as high peakedness. However, after decomposing labor earnings growth into wage and working time growth, the log-normality of hourly wages remains a quite plausible assumption since deviations from Gaussianity stem mainly from working time changes. The joint dynamics of hourly wages and working time help explain those deviations which relate most likely to labor supply decisions at the extensive margin.

Keywords: Labor earnings growth; non-Gaussian distributions; skewness; kurtosis.

JEL Classification: E24, J22, J31.

*We thank Stéphane Bonhomme, Nila Céci-Renaud, Élise Coudin, Laurent Davezies, Thierry Magnac, David Margolis, Sébastien Roux and Gregory Verdugo for useful suggestions, as well as attendees at EALE (Ghent, 2016), ESPE (Glasgow, 2017), IAAE (Milan, 2016), INSEE seminar (2017), IPDC (Thessaloniki, 2017) and JMA (Bordeaux, 2018) for their comments. All errors and opinions are ours.

[†]INSEE-CREST. Corresponding author. Address: 88 avenue Verdier, 92120 Montrouge (France). Phone: (+33)187695944. Email: pierre.pora@insee.fr

[‡]INSEE-CREST.

1 Introduction

Due to the scarcity of detailed longitudinal data on labor earnings, much is still ignored about the distribution of earnings changes conditional on earnings history. To overcome this issue, an original approach has been recently proposed by [Guvenen et al. \(2016\)](#) which questions the assumptions that earnings changes are Gaussian and that earnings dynamics are linear. They estimate nonparametrically the first moments of the distribution of earnings changes and relate them to the location in the distribution of recent earnings. They also resort to impulse response functions in order to characterize the transient nature of shocks, and more generally labor earnings dynamics. This method contrasts with the literature that has modeled earnings dynamics thanks to sophisticated specifications combining both observed and unobserved heterogeneity as well as uncertainty (often approximated by random walks, autoregressive processes, ARMA, etc.). Their main empirical findings are the following: (i) log-earnings changes are not Gaussian, (ii) their distribution is leptokurtic and negatively skewed, and (iii) individual labor earnings dynamics are nonlinear and heterogeneous. They also examine how omitting to account for these deviations from normality as well as for nonlinearity might underestimate the welfare cost of earnings fluctuations.

Our contribution is twofold. First, we decompose labor earnings growth into wage and working time growth at both the extensive and intensive margins of employment, which sounds like a natural exercise to perform given that a large part of the literature has focused on modeling hourly wages rather than annual earnings. On top of previous stylized facts, we find that (i) most deviations from Gaussianity, including the negative skewness and the fat tail, are mainly due to working time changes at the extensive margin; (ii) though one rejects the log-normality assumption as far as wage growth is concerned, maintaining this assumption might be quite reasonable in practice since wage changes are less asymmetric and exhibit a much lower, homogeneous kurtosis. Second, we analyze the dynamics of labor earnings, hourly wages and working time; we find that (iii) much nonlinearity and heterogeneity of earnings dynamics are driven by working time dynamics at the extensive margin; (iv) large and positive hourly wage growth coincide with higher probability to leave salaried employment few years after. This suggest

that transitions from (and to) employment explain much of positive (negative) large hourly wage changes, and could be consistent with the countervailing role of severance payments. Overall, deviations from normality are primarily the result of unemployment risk. It follows that the welfare cost of non-Gaussian earnings fluctuations depends on the level of insurance provided by unemployment benefits.

We rely on the DADS panel, a French longitudinal administrative database, the filling of which is mandatory for payroll taxes, and which contains information on individuals' labor earnings. We focus on men working in the private sector from 1995 to 2015. We adopt the same method as [Guvenen et al. \(2016\)](#) to measure labor earnings growth. Importantly, we dispose of information on working hours and employment spells duration, which enables us to use this approach on both working time and hourly wages.

The rest of the paper is organized as follows. Next section is devoted to a brief literature review. Section 3 presents our data. In section 4, we describe our empirical approach to disentangle wage from working time changes in labor earnings growth. Section 5 presents our results, section 6 is devoted to some robustness checks and section 7 concludes.

2 Literature review

Numerous papers in the earnings dynamics literature model volatility thanks to parametric models, including, e.g., [Moffitt and Gottschalk \(2002, 2011\)](#), [Baker and Solon \(2003\)](#), [Low, Meghir, and Pistaferri \(2010\)](#), [Altonji, Smith, and Vidangos \(2013\)](#) as well as [Magnac, Pistoiesi, and Roux \(2018\)](#). Most papers belonging to this stream rely on the assumption of Gaussian shocks. Few depart from this hypothesis, but noticeable exceptions include (i) [Bonhomme and Robin \(2009\)](#) who model the transitory component of log-earnings as a first-order Markov process, the marginal distribution of which is a mixture of normals, and the transition probability of which is ruled by a copula; (ii) [Arellano, Blundell, and Bonhomme \(2017\)](#) who model earnings dynamics as the sum of some general Markovian, persistent component and some transitory innovation.

In a recent work, [Guvenen et al. \(2016\)](#) have an agnostic look at individual-level

earnings data. They adopt a descriptive approach which involves a non-parametric estimation of earnings changes. They resort to comprehensive data issued from the US Social Security Administration (SSA) over the 1978-2011 period. The Master Earnings File they use is derived from the W-2 form and presents at least three advantages: (i) it has a large sample size, (ii) with low measurement error, and (iii) no top-coding of annual labor earnings. Their main result is that labor earnings changes are not log-normal. These annual variations are negatively skewed and display a very high kurtosis, which is not consistent with usual features of normal distributions. The asymmetry stems from the fact that large, upwards changes are less likely than large, downwards changes (also called "disaster shocks"). The excess kurtosis means that most individuals experience very small earnings changes, but that a small, non-negligible number of individuals face very large changes. Moreover, positive (negative) changes tend to be transitory (persistent) for high income individuals while the reverse holds for low income individuals. Finally, large changes tend to be more transitory than small changes.

Two important *caveats* due to data limitations are worth being mentioned: (i) they restrict their attention to men in the private sector, and (ii) they do not dispose of any information regarding working time, which does not enable them to determine whether labor earnings growth comes from working time or from wage dynamics. In theory, our dataset enables us to deal with women and with the public sector as well, but to ensure comparison with [Guvenen et al. \(2016\)](#) for instance, we also focus on men working in the private sector. More importantly, worked hours have been available since 1995, which enables us to decompose labor earnings growth.

In another contemporaneous project, [Busch, Fialho, and Guvenen \(2018\)](#) propose to decompose labor earnings into the product of hourly wages and annual working hours. Also, closely related to ours is a recent paper by [Hoffmann and Malacrino \(forthcoming\)](#): based on Italian administrative data, they first show that both the negative skewness of annual labor earnings growth and its fluctuations over the business are driven by employment time changes, i.e., changes in the number of weeks an individual works during a year. Conversely, they find the distribution of weekly earnings growth to be symmetric and stable over the

business cycle. As a result, they praise that modelling earnings processes over the business cycle disentangles carefully employment and wages. They also provide with a process calibrated on the CPS which succeeds in generating the fluctuations of the earnings growth distribution in the US. By contrast, [Busch et al. \(2018\)](#) resort to a quantile-based measure of skewness of hourly wage shocks; and find that it is negative on German data, concluding therefore to an asymmetric distribution in that country.

Why do we care about nonlinear features of labor earnings dynamics? First, because they may lead to rare, yet considerable shocks having first-order consequences. Besides, log-normality assumptions on wage growth might result in a dramatic underestimation of the welfare costs of idiosyncratic earnings fluctuations ([Guvenen et al., 2016](#)). Moreover, because deviations from Gaussianity may lead to a wrong understanding of the relationship between the business cycle and labor earnings dynamics. In particular, the skewness is cyclical, i.e. downward asymmetry is stronger during recessions, which is even more crucial than counter-cyclical variance, as [Guvenen, Ozkan, and Song \(2014\)](#) show.

Why does labor earnings uncertainty matter? Essentially because it has a direct impact on consumption and saving behaviors, as pointed out by [Blundell and Preston \(1998\)](#), [Parker and Preston \(2005\)](#), [Blundell, Pistaferri, and Preston \(2008\)](#) and [Arellano, Blundell, and Bonhomme \(2017\)](#), not pretending to be exhaustive. A better understanding of income uncertainty – especially its impact on household decisions – would help fill in the gap between income and consumption inequalities ([Pistolesi, 2014](#)). In particular, a high volatility prevents individuals from smoothing their consumption: inequality results therefore in inefficient allocations. More precisely, uncertainty on future earnings influences labor supply at both extensive and intensive margins when markets are incomplete, i.e. when there are no state contingent markets for household-specific shocks. Agents insure themselves against earnings fluctuations, either by saving more (precautionary saving) if they are prudent, and by working longer hours (precautionary labor supply) if they are risk-averse. [Pijoan-Mas \(2006\)](#) explains that the relative importance of the two channels depends on the persistence of the non-deterministic component of the wage process (in a standard permanent-transitory model, the permanent

component is often assumed to be a random walk while the transitory component follows some MA process, for instance); the inability of markets to insure workers against idiosyncratic labor market risk has consequent macroeconomic implications and would be responsible for large welfare losses. [Jessen, Rostam-Afschar, and Schmitz \(2018\)](#) estimate on German data (namely, the GSOEP) that the precautionary labor supply amounts to 2.8% of worked hours and that it is higher for self-employed. Overall, the idea that labor supply acts as a smoothing mechanism is widely spread. Finally, [Bozio, Breda, and Guillot \(2016\)](#) provide indirect evidence on how taxation¹ can reduce wage inequalities in the long run, which suggests a mechanism through which the policymaker might deal with these inequalities.

3 Data

3.1 The DADS panel

Our analysis is based on a large panel of French salaried employees working in the private sector, the longitudinal version of the *Déclaration Annuelle de Données Sociales* (DADS). By law,² French firms have to fill in the DADS – an annual form which is the analogue of the W-2 form in the US – for every employee subject to payroll taxes. This panel contains information about individuals born on October of even-numbered years; it is therefore a representative sample of the French salaried population at rate 1/24. Since filling in the form is mandatory, and because of the comprehensiveness of the panel with respect to individuals’ careers, the data is of exceptional quality and has low measurement error in comparison with survey data; it has thus this desirable feature, on top of a large sample size and no top-coding.

The database contains detailed information about gross and net wages, work days, working hours,³ other job characteristics (the beginning and the end of an employment’s spell, seniority, a dummy for part-time employment), firm characteristics (industry, size, region) and individual characteristics (age, gender). Our

¹More precisely, employer social security contributions as payroll taxes.

²The absence of a DADS as well as incorrect or missing answers are punished with fines.

³This information has been available since 1995 only.

variables of interest are: (i) real annual earnings defined as the sum of all salaried earnings, (ii) work days, comprised between 0 and 360, (iii) working time measured in working hours, and (iv) hourly wages defined as the ratio of annual earnings over working time. We decompose working time into its extensive and intensive margins by considering working time measured in hours as the product of work days (extensive margin) and working hours per day (intensive margin).

Our working sample is composed of male salaried employees⁴ working in metropolitan France between 1995 and 2015, aged 20 to 60, at the exclusion of agricultural workers and household employees.

The empirical analysis described in Section 4 requires to select individuals with a strong attachment to the labor market. We rely on “relatively stable” workers to describe changes in earnings between year t and year $t+5$. We impose in particular that these individuals be present at least two years between $t-5$ and $t-2$, on top of being present in $t-1$ and in t . To deal with very low working time or very low hourly wages, we focus on individuals for which work days exceed 45 days a year, working hours per day exceed $1/8$ of the annual legal duration of work divided by 360 and whose hourly wages exceed 90% of the minimum hourly wage. Our results are nevertheless robust to different choices of censoring thresholds (see Section 6).

Table 1 provides with some descriptive statistics on the selection of “relatively stable” workers. First comes the censoring related to working time and hourly wages. Second, we restrict our attention to individuals who were present at least two years between $t-5$ and $t-2$, on top of being present in $t-1$ and in t . The first step leaves us with a higher share of workers in the manufacturing industry and changes slightly the age distribution of the sample by selecting out some young workers. The second step implements the employment stability criterion, which amplifies slightly the previous distortion of the age distribution.

3.2 Some background on the French wage distribution

In the US and in the UK at least, the rise in earnings inequality in the 80s-90s has been driven by skill-biased technical change (SBTC), which has been extensively

⁴This is in line with Guvenen et al. (2015) where the analysis is based on salaried employees. In a more recent version, (Guvenen et al., 2016) include self-employed workers as well.

Table 1 – Descriptive statistics on the selection process

	Base sample		Censoring		Final sample	
N	8 218 738		7 594 734		5 388 366	
	Frequency (in %)	Average earnings (2015 €)	Frequency (in %)	Average earnings (2015 €)	Frequency (in %)	Average earnings (2015 €)
Industry						
Manufacturing	25,1	25 500	26,6	26 400	29,2	28 000
Construction	11,4	19 800	12	20 600	12,3	22 900
Trade	15,4	21 600	15,8	22 800	15,9	25 600
Services	48,1	20 200	45,7	23 100	42,7	27 700
Age						
23-24	6,3	10 600	5,6	12 300	2,9	14 900
25-29	16,2	15 800	15,9	17 200	13,7	19 000
30-34	15,5	20 300	15,6	21 500	15,8	23 100
35-39	15,1	23 300	15,3	24 700	15,7	26 300
40-44	14	25 700	14,3	27 100	15,3	28 700
45-49	13,3	27 500	13,5	29 000	14,3	30 400
50-54	11,4	29 000	11,6	30 500	13	31 800
55-59	8,2	29 300	8,2	31 300	9,3	32 900

documented since the seminal work by [Katz and Murphy \(1992\)](#).

By contrast, in France, most studies conclude that wage dispersion has not increased over the last 30 years.⁵ According to [Verdugo \(2014\)](#) who relates changes in the wage structure to changes in education levels, wage inequality has decreased continuously from 1969 to 2008. The rise of high-skilled labor supply might have hidden the effects of SBTC. The increase in educational attainment after 1980 would have resulted in a large decline of the skill premium, which would account for between 1/3 and 1/2 of the decrease in inequality at the top of the distribution. Inequality decreased also at the bottom thanks to the minimum wage. Both phenomena resulted in a compression of the wage distribution. Moreover, these changes – at least at the top – cannot be fully explained by selection into employment.

⁵Yet, [Bozio, Breda, and Guillot \(2016\)](#) do find some evidence of SBTC in France provided that the measure wage inequality is expressed in terms of labour cost.

4 Empirical analysis

4.1 Methodology

We follow the same descriptive approach as [Guvenen et al. \(2016\)](#); contrary to them, and thanks to the information on the duration of employment spells and working time provided by our data, we do not restrict ourselves to the sole dynamics of labor earnings, but also decompose earnings in three components: hourly wages and working time at both the intensive and extensive margins of employment. We rely therefore on nonparametric estimations of the distribution of individual labor earnings growth, and of each its three components, which enables us not to posit any parametric assumption on the distribution of changes, contrary to many papers in the literature devoted to earnings dynamics.

Another difference with [Guvenen et al. \(2016\)](#) is that we rank individuals according to their recent hourly wages rather than according to their recent labor earnings. This choice emphasizes the role played by the heterogeneity along the hourly wage (the usual proxy for productivity) distribution.⁶

Let denote the logarithm of labor earnings for individual i on year $t = 1, \dots, T$ by \tilde{e}_{it} . We propose the following decomposition:

$$\tilde{e}_{it} = \widetilde{w}_{it} + \widetilde{d}_{it} + \widetilde{h}_{it}, \quad (1)$$

where \widetilde{w}_{it} is the logarithm of hourly wages, \widetilde{d}_{it} accounts for the logarithm of total duration of employment spells (comprised between 1 and 360) and \widetilde{h}_{it} is the logarithm of working hours per day. Hence \widetilde{d}_{it} represents the extensive margin of employment (individuals leaving and getting back to employment) whereas \widetilde{h}_{it} is related to the intensive margin of employment (how many hours individuals work given that we observe them during an employment spell).

We aim at measuring changes at the individual level and within a 5-year horizon, which [Guvenen et al. \(2016\)](#) call "permanent changes". We consider a normalized version of log earnings (resp. hourly wages, employment duration or hours

⁶For the sake of comparison, section 6 provides with the results obtained by ranking according to recent earnings.

per day), net of age effects. Let \tilde{y}_{it} denote generically a labor outcome, either earnings or any of its three components. We start by regressing \tilde{y}_{it} on a set of age, period and cohort dummies:

$$\tilde{y}_{it} = y_0 + \sum_c \alpha_c \mathbb{1}_{cohort_i=c} + \sum_a \beta_a \mathbb{1}_{age_{it}=a} + \sum_j \gamma_j \mathbb{1}_{t=j} + \epsilon_{it}. \quad (2)$$

The inclusion of year dummies is another slight difference with [Guvenen et al. \(2016\)](#). Sampling issues lead us to introduce them (i) since we want to control for any disruption caused by minor, methodological changes in the production of the DADS panel which occurred in 2002, 2009 and 2013, and (ii) because our sample includes even-year born individuals only: even (odd) ages are thus observed in even (odd) years.⁷

The identification of age-period-cohort (APC) models can be achieved at the cost of some normalization. The major threat to the simultaneous identification of α , β and γ stems from the collinearity between age, cohort and period: age is equal to current period minus year-of-birth. Several solutions have been investigated in the sociology literature, e.g. [Mason et al. \(1973\)](#) who propose to assume that any two ages, periods or cohorts have the same effect, on top of removing one dummy in each dimension. [Deaton and Paxson \(1994\)](#) and [Deaton \(1997\)](#) suggest a transformation⁸ of period effects in order to meet two requirements: (i) these time effects sum to zero, and (ii) they are orthogonal to a time trend, so that age and cohort effects capture growth while year dummies account for cyclical fluctuations or business-cycle effects that average to zero over the long run. To sum up, the parameters of the model (α, β, γ) are identified provided that $\alpha_{\underline{c}} = 0$, $\beta_{\underline{a}} = 0$, $\sum_{t=1}^T \gamma_t = 0$ and $\sum_{t=1}^T (t-1)\gamma_t = 0$. The corresponding transformation of time dummies $d_j = \mathbb{1}_{t=j}$ writes as follows:

$$d_t^* = d_t - [(t-1)d_2 - (t-2)d_1],$$

with $d_1^* = d_2^* = 0$. In practice, it is convenient to include all age dummies but the first, all cohort dummies but the first and all transformed dummies d_t^* but d_1^* and

⁷Including these year dummies does not influence our results regarding labor earnings growth and dynamics: see Section 6.

⁸An insightful presentation of this method is provided by [Afsa and Buffeteau \(2006\)](#).

d_2^* in the regression.

We adapt this transformation to our dataset with even-year born individuals only. This limitation leads us to impose further that both odd-year and even-year time effects sum to zero, i.e. to consider two restrictions: $\sum_j \gamma_{2j} = 0$ and $\sum_j \gamma_{2j+1} = 0$, instead of the sole constraint $\sum_t \gamma_t = 0$, on top of imposing $d_1^* = d_3^* = 0$, on the one hand, and $d_2^* = 0$, on the other hand. Corresponding transformations depend now on the parity of t :

$$d_{2j+1}^* = d_{2j+1} - [jd_3 - (j-1)d_1]$$

$$d_{2j}^* = d_{2j} - [(j-1)d_3 + d_2 - (j-1)d_1].$$

We estimate those four (one for labor earnings and one for each of its component) age-period-cohort models independently. Accounting decomposition 1 ensures that $\alpha^e = \alpha^w + \alpha^d + \alpha^h$; similar equalities hold for β^y and γ^y .

As far as annual earnings are concerned, our variable of interest is $e_{it} = \tilde{e}_{it} - \hat{\alpha}_{cohort_i} - \hat{\beta}_{age_{it}} - \hat{\gamma}_t$, which can be interpreted as log-labor earnings, net of age, period and cohort effects. We do the same for each of labor earnings components, Equation (1) guarantees that $e_{it} = w_{it} + d_{it} + h_{it}$. The 5-year change in normalized log-earnings $\delta^5 e_{it} = e_{i,t+5} - e_{it}$ accounts for the relative change in individual i 's earnings between t and $t+5$ with respect to her analogues of the same age and cohort; it is once again decomposed between hourly wages and working time at both the intensive and extensive margins of employment.

Figure 1 provides a synthetic view of the current approach. We use hourly wages between $t-5$ and $t-1$ to depict heterogeneity along the hourly wages distribution, while focusing more specifically on the distribution of changes between t and $t+5$, and its relationship to past changes, i.e. changes between $t-1$ and t .

In the rest of the paper, we distinguish between labor earnings *growth*, which corresponds to individual earnings changes, from earnings *dynamics*, which refers to the relationship between past and future earnings changes. These two definitions must be understood as being conditional on recent hourly wages. Impulse response functions turn out to be an efficient tool to describe the latter.

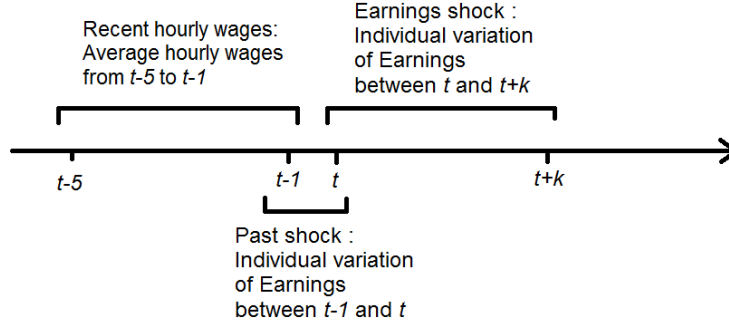


Figure 1 – Labor earnings changes

4.2 Cross-sectional distributions of labor earnings growth

We aim at comparing workers with similar histories in terms of hourly wages. We introduce therefore a measure of recent hourly wages W_{it} similar to the one used by [Guvenen et al. \(2016\)](#) with respect to recent earnings. This measure of recent hourly wages approximates average hourly wages between $t - 5$ and $t - 1$, net of age, period and cohort effects:

$$W_{it} = \frac{\sum_{\tau=t-5}^{t-1} \exp(\widetilde{w}_{it})}{\sum_{\tau=t-5}^{t-1} \exp(\widehat{\alpha}_{cohort_i}^w + \widehat{\beta}_{age_{i\tau}}^w + \widehat{\gamma}_{\tau}^w)} \quad (3)$$

We divide workers into 8 age groups: 23-24, 25-29, 30-34, 35-39, 40-44, 45-49, 50-54 and 55-59. For each year t and each age group, we rank workers according to their recent hourly wages W_{it} , and consider 100 percentile groups.

We estimate various features – moments and quantiles – of the distributions of labor earnings (resp. hourly wages, employment duration and hours per day) changes $\delta^5 y_{it}$ for each (percentile group, age group, year t) and we average these features across all years and over all age groups. This procedure enables us to characterize the distribution of 5-year growth $\delta^5 y_{it}$ conditional on the rank in the distribution of recent hourly wages W_{it} .

We resort to local statistical indicators in order to describe the heterogeneity of annual individual changes in normalized labor earnings and any of its three components.

The variance, i.e. the second moment of the standardized variable, describes

the dispersion of 5-year changes.⁹ An alternative, quantile-based measure of dispersion is the difference between the 90th and 10th percentiles ($P90$ and $P10$, respectively) of 5-year changes.

The skewness, i.e. the third moment of the standardized variable, accounts for the degree of asymmetry. A related quantile-based measure is Kelley’s measure of skewness (Kelley, 1947) defined as the relative share of $P90-P10$ that can be explained by $P90-P50$ and $P50-P10$:

$$\text{Kelley's Skewness} = \frac{(P90 - P50) - (P50 - P10)}{P90 - P10} \quad (4)$$

It is constant and equal to 0 for Gaussian distributions.

The kurtosis, i.e., the fourth moment of the standardized variable, measures the peakedness of the tails of the distribution of those changes. We consider the normalized kurtosis which is constant and equal to 0 for Gaussian distributions. A quantile-based measure of the heaviness of tails is Crow-Siddiqui’s measure of kurtosis (Crow and Siddiqui, 1967). It is defined as:

$$\text{Crow-Siddiqui's kurtosis} = \frac{P97.5 - P2.5}{P75 - P25} \quad (5)$$

It is constant and equal to roughly 2.91 for Gaussian distributions.

We prefer quantile-based over moment-based measures of dispersion, asymmetry and heaviness of tails because the former are more robust to the presence of outliers. Such measures are also praised by Arellano (2014) who refers to Kim and White (2004). However, because Hoffmann and Malacrino (forthcoming) attribute their divergence with Busch et al. (2018) to the choice of moment-based as opposed to quantile-based measures of asymmetry, we provide with both measures (see Section 6). Importantly, in our setting, they yield qualitatively the same results.

4.3 Impulse response functions

Dynamics is described here thanks to impulse response functions that relate future to past changes. From this point of view, labor earnings, hourly wages or working

⁹One could also think of *volatility*. By contrast, in standard life-cycle models, *wage risk*, or *wage uncertainty*, refers to the conditional distribution of future wages given current wages.

time dynamics can be described in a similar fashion once again.

Impulse response functions are estimated non-parametrically in order to allow for non-linearities. To account for heterogeneity in labor earnings dynamics across the earnings distribution, we estimate separate impulse response functions for 21 subgroups (P0-P5, P5-P10, P10-P15,..., P90-P95, P95-P99 and P99-P100) of the distribution of recent earnings. We compute labor earnings, hourly wages, working time changes at both the intensive and extensive margins between $t - 1$ and t as $\delta^1 y_{i,t-1} = y_t - y_{t-1}$.

Within each subgroup, we rank workers according to $\delta^1 y_{i,t-1}$ and create 20 past change groups of the same size ($P0^\delta - P5^\delta$, $P5^\delta - P10^\delta$, ..., $P90^\delta - P95^\delta$, $P95^\delta - P100^\delta$).¹⁰ This gives us $21 \times 20 = 420$ recent earnings \times past change groups denoted as g_{mn} with $(m, n) \in \{P0 - P5, \dots, P95 - P99, P99 - P100\} \times \{P0^\delta - P5^\delta, \dots, P95^\delta - P100^\delta\}$.

For each of these groups, we estimate the average past change $\mathbb{E}[\delta^1 y_{i,t-1} \mid g_{mn}]$ and the average future change $\mathbb{E}[\delta^k y_{i,t} \mid g_{mn}]$, with $k = 1, \dots, 5$. To the extent that within each of these recent earnings \times past change groups, workers have similar recent earnings and experience similar past changes, $\mathbb{E}[\delta^k y_{i,t} \mid g_{mn}]$ approximates a nonparametric estimation of $\mathbb{E}[\delta^k y_{it} \mid \delta^1 y_{i,t-1}, Y_{it}]$. Hence, plotting $\mathbb{E}[\delta^k y_{i,t} \mid g_{mn}]$ against $\mathbb{E}[\delta^1 y_{i,t-1} \mid g_{mn}]$ for various percentile groups corresponds to an impulse response function that may be nonlinear, heterogeneous across the distribution of earnings and which exhibit some asymmetry between positive and negative changes.

5 Results


5.1 Lifecycle profiles

Figure 2 depicts the age profiles – namely, the estimated coefficients $\widehat{\beta}_a$ – of labor earnings, hourly wages, working hours per day and employment duration. Both labor earnings and employment duration exhibit a hump-shaped pattern which peaks at age 55 and 53 respectively, when they are 111% (resp. 19%) higher than

¹⁰Within each recent earnings subgroup, cells are defined by the rank in the distribution of past change conditional on age group (25-34 and 35-50) and year t .

those at age 25. On the contrary, wages keep on increasing after age 55, because of a selection effect: workers who are actually observed with positive earnings after 55 are also the more productive. For instant, hourly wages of workers aged 60 (55) are 121% (90%) higher than those of workers aged 25. Overall, hourly wages' growth is slower than earnings ones. The concave profile observed for hourly wages until age 55 is consistent with diminishing marginal returns of age. The slight decrease of earnings from 55 to 60 must be related to the gradual exit of labor workforce, and in particular to a decrease in employment duration after 55, which is in line with usual findings in the literature devoted to older workers' productivity. Interestingly, while employment duration tends to increase over the lifecycle, working hours per day decrease after 30, when they are only 4% higher than those at 25 or at 40. This empirical evidence suggests that, as they grow of age, workers adjust their labor supply by working more days, but less intensively within each day.


Figure 2 – Earnings, hourly wages and working time lifecycle profiles



5.2 Earnings growth dispersion

We now turn to the dispersion of growth rates. Figure 3 plots the P90-P10 difference of 5-year earnings, wages, hours changes and employment duration against recent hourly wages.

Figure 3 – P90-P10 of 5-year earnings, hourly wages and working time changes




The volatility of future labor earnings turns out to have a U-shape along the distribution of recent hourly wages; besides, individuals with low hourly wages are much more exposed to earnings volatility. However, the main lesson of the decomposition of labor earnings growth into hourly wages and working time growth is the following: much of its dispersion stems from the volatility of working time, except perhaps for top-earners who are subject to significant uncertainty on their hourly wage rate. While the dispersion in hourly wage changes increases along the distribution, especially at the top, the dispersion of working time changes at both margins decreases when moving to better paid workers, with the exception of top earners for which the dispersion of hours per day changes increases once again.

5.3 Asymmetry of changes

Figure 4 displays Kelley’s measure of skewness of 5-year earnings, wages, employment duration and hours changes along the distribution of recent hourly wages. Kelley’s skewness of earnings varies between $-.20$ and $-.05$: earnings changes are negatively skewed, which means that large downward changes are more frequent than large upward changes. As a result, the log-normality assumption is not likely since a skewness of zero would be expected in that case: in this respect, this result is rather consistent with previous findings by [Guvenen et al. \(2016\)](#) in the US. Formally, we are not able to reject log-normality: computing confidence intervals by bootstrap requires to reestimate age-period-cohort models and redefine ranking among ”relatively stable” workers at each step, which is a computationally challenging task. However, we obtain a lower bound of such confidence intervals by fixing this ranking once-and-for-all; this lower bound never includes the horizontal axis (see Figure 10 in Appendix).

Figure 4 – Kelley’s measure of skewness of 5-year earnings, hourly wages and working time changes



More interestingly, most of the asymmetry stems from working time, especially from employment duration since changes in work hours per day do look symmetric: their Kelley's skewness is not significantly different from 0, except in the highest quintile of the hourly wage distribution. By contrast, changes in employment duration appear highly asymmetric, in particular at the bottom of the distribution where large, negative changes are more frequent, the Kelley's skewness being less than -0.1 in the lowest quintile. This empirical evidence emphasizes the role played by labor supply decisions at the extensive margin like dismissals or firings.

Most strikingly, hourly wages changes exhibit a positive asymmetry at the bottom of the distribution, but almost no asymmetry at all elsewhere. It could be partly due to measurement error in working hours. However, if it were the case, then working hours changes would exhibit a negative asymmetry. We do find labor supply changes at the extensive margin, which are measured independently of working hours, to display such asymmetry, but not changes at the intensive margin that would incorporate such an error.

This *caveat* aside, in the first quintile, wage drops are constrained by the minimum wage while wage rises remain unbounded, which help explain the positive asymmetry there –and more generally, the downward-sloping pattern of wages' skewness. An important exception concerns the top 3% of earners who are the only ones to experience more often large negative changes than large positive changes: these workers' wage might be composed of both fixed and variable parts, the latter being perhaps more dependent on aggregate fluctuations, and possibly subject to downward movements. These workers are also more stable: the growth of their annual earnings follows closely the growth of their wages. Mechanically, there is more room to fall at these compensation levels.

Overall, the symmetry of hourly wage changes is rejected on the data since estimated confidence intervals do not include the horizontal axis (see Figure 11 in Appendix); the true confidence intervals might include the horizontal axis, especially in the upper half of the distribution. Making this parametric assumption in a structural model sounds reasonable for the asymmetry on hourly wages' growth rates is rather small and homogeneous along the distribution. This assumption is all the more convincing than individuals have a strong attachment to the labor

workforce.

Stressing that negative asymmetry in labor earnings changes is more likely to stem from labor supply changes at the extensive margin, rather than from changes in the wage rate concurs with the empirical evidence by [Hoffmann and Malacrino \(forthcoming\)](#) in Italy, but contrasts with [Busch et al. \(2018\)](#)’s results for Germany: according to the latter, full-time wage changes drive cyclicalities in the skewness of labor earnings changes. [Hoffmann and Malacrino \(forthcoming\)](#) suggests that this follows from the use of quantile-based rather than moment-based measures of skewness, but since we rely also on quantile-based measures, our results suggest that this methodological choice is not the sole explanation. Moreover, Figure 13 in Appendix provides with moment-based measures and also confirms this diagnosis.


5.4 Peakedness of the distribution and heaviness of tails

Figure 5 displays the Crow-Siddiqui’s measure of kurtosis of earnings, wages and hours changes. Since Gaussian distributions have a constant Crow-Siddiqui of roughly 2.91, and because we focus on deviations from normal distributions, we plot the normalized Crow-Siddiqui defined as the Crow-Siddiqui minus 2.91. Earnings exhibit a higher peakedness than the Gaussian reference, but this peakedness is rather homogeneous along the recent hourly wage distribution. Moreover, hourly wages have a normalized Crow-Siddiqui of slightly less than 2, and completely homogeneous along the hourly wage distribution, which makes them even closer to a normal distribution. On the whole, fat tails arise in earnings growth rather because of working time changes.

The Crow-Siddiqui’s kurtosis of employment duration is not displayed on Figure 5 because its values are mechanically far higher from the ones related to the other labor outcomes: most workers’ employment duration remains stable at 360 days, hence the denominator of (5) can be very low.

To sum up, the tails of the distribution of wage changes are not particularly heavier at either the bottom, the middle or the top of the distribution; moreover, the asymmetry looked rather small and homogeneous along that distribution, with the noticeable exception of the very top of the distribution. Hence we conclude

Figure 5 – Crow-Siddiqui’s measure of kurtosis of 5-year earnings, hourly wages and working time changes



to the following empirical evidence: (i) wage volatility is much less heterogeneous than working time instability, (ii) the Gaussian approximation for wage growth looks quite reasonable, and (iii) most of non-Gaussian features of log-earnings changes stem from working time.


5.5 Labor earnings dynamics

We now report impulse response functions. Figure 6 displays the growth of earnings, wages, working hours and employment duration for various time horizons $\delta^k y_{it} = y_{i,t+k} - y_{it}$, $k = 1, \dots, 5$ and for various values of corresponding past changes $y_{it} - y_{i,t-1}$.


First, labor earnings changes display some kind of mean reversion: the higher the past, the lower the future change. Downward past changes tend to come along with positive future changes, while upward past changes are rather associated with negative future changes. Furthermore, the higher the magnitude of past changes,

Figure 6 – Impulse-response functions of earnings, hourly wages and working time


(a) Earnings




(b) Hourly wages



(c) Employment duration



(d) Hours per day



the higher the magnitude of future changes: workers who experience large changes between $t - 1$ and t tend to experience larger changes between t and $t + k$ than those who were submitted to smaller changes.

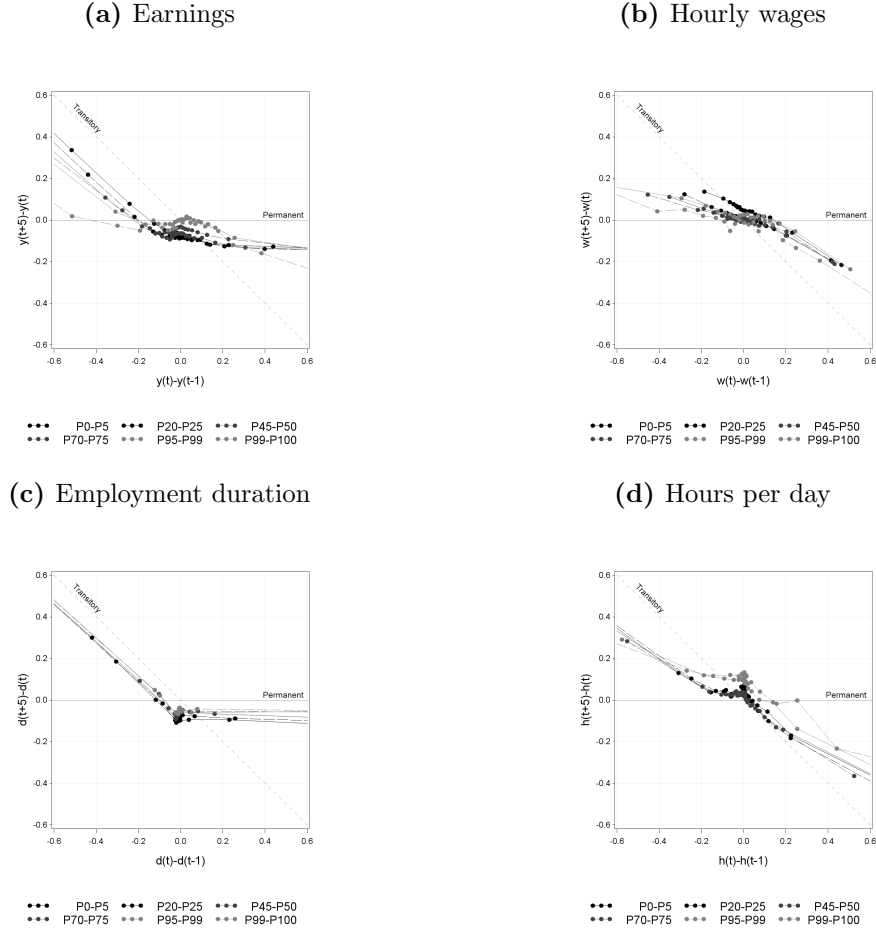
Second, impulse response functions tell us about the persistence of changes: past changes are all the more transitory than $\delta^k y_{it} \approx -(y_{it} - y_{i,t-1})$, i.e. than their effect has vanished after k years. Conversely, when $\delta^k y_{it} = 0$, past changes are more persistent. As far as annual earnings are concerned, large, negative past changes tend to be transitory while large, positive ones are more persistent. For instance, individuals who experienced a 60 log-points loss between $t - 1$ and t with respect to the age trend recover more than half of this loss within 5 years.

Third, the impulse response function related to employment duration shows unambiguously that negative shocks are transitory while positive shocks are strikingly permanent. On the contrary, hourly wages' impulse response functions are almost linear with a negative slope, which is consistent with some mean reversion (up to some pass-through). The same holds, more or less, for working hours. Finally, annual earnings' impulse function looks like a mix of those three functions, but its shape looks qualitatively similar to the employment duration's one, which emphasizes once again the role played by labor supply at the extensive margin.

Figure 7 plots 5-year growth of earnings, wages, hours and employment duration against past changes in earnings, wages, hours and employment duration for various locations in the distribution of recent hourly wages. Like [Guvenen et al. \(2016\)](#), we find labor earnings changes to exhibit a "butterfly" pattern: highly negative (positive) changes are more transitory (persistent) for individuals that earned low hourly wages, but less so for their better paid counterparts.¹¹ The "butterfly" pattern for earnings contrasts with the patterns corresponding to each of the three earnings components, for which heterogeneity in the dynamics is much more limited. This suggests that heterogeneity in labor earnings dynamics is not driven by any of its components displaying substantial heterogeneity in its dynamics along the wages distribution, but rather from each component explaining a different share of labor earnings changes along the wages distribution.

¹¹A more complete description of the heterogeneity of these individual dynamics can be found in [Pora and Wilner \(2018b\)](#).

Figure 7 – Impulse-response functions of earnings, hourly wage and working time changes at a 5-year horizon



5.5.1 Decomposition of labor earnings changes

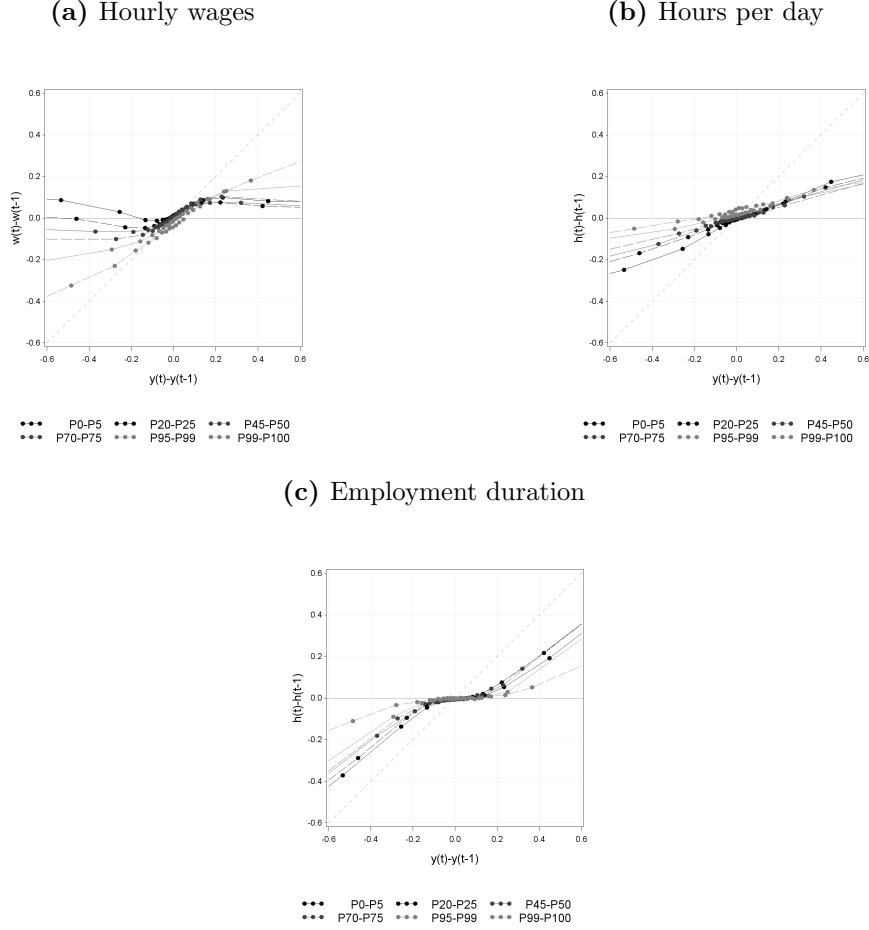
We decompose now labor earnings changes into their simultaneous wage, working hours and employment duration components. We obtain:

$$e_{i,t} - e_{i,t-1} = \mathbb{E}[w_{i,t} - w_{i,t-1} \mid e_{i,t} - e_{i,t-1}] + \mathbb{E}[h_{i,t} - h_{i,t-1} \mid e_{i,t} - e_{i,t-1}] + \mathbb{E}[d_{i,t} - d_{i,t-1} \mid e_{i,t} - e_{i,t-1}] \quad (6)$$

The non-parametric estimations of $\mathbb{E}[w_{i,t} - w_{i,t-1} \mid e_{i,t} - e_{i,t-1}] = e_w(e_{i,t} - e_{i,t-1})$, $\mathbb{E}[h_{i,t} - h_{i,t-1} \mid e_{i,t} - e_{i,t-1}] = e_h(e_{i,t} - e_{i,t-1})$ and $\mathbb{E}[d_{i,t} - d_{i,t-1} \mid e_{i,t} - e_{i,t-1}] = e_d(e_{i,t} - e_{i,t-1})$ yield an exact decomposition of labor earnings changes. Figure 8 plots our estimates for various positions in the distribution of recent hourly wages. The figures display the estimated functions \hat{e}_w , \hat{e}_h and \hat{e}_d given that $\forall x, e_w(x) + e_h(x) +$

$e_d(x) = x$. Earnings, wages, working hours and employment duration are net of the systematic age component so that these changes represent deviations from the average earnings, wages, working hours and employment duration lifecycle profiles.

Figure 8 – Earnings changes decomposition



At the very bottom of the distribution, labor earnings changes are roughly pure changes in working time: they are hardly related to hourly wage changes. As one gets higher in the distribution, the share of labor earnings changes that stems from wage changes increases; the slope of \hat{e}_w is steeper for high earnings. However, even among the highest earnings group P99-P100, working time changes still explain 1/2 of 60 log-points labor earnings changes. Moreover, for those individuals, the slope of \hat{e}_w is steeper for earnings changes smaller than 20 log-points, regardless of their sign: large (small) annual earnings changes correspond rather to substantial

working time changes (wage changes).

These stylised facts are consistent with an increasing share of labor earnings volatility due to wage volatility along the earnings distribution. Additionally, we show large labor earnings changes to be primarily driven by working time changes, rather than by hourly wages changes at the extensive margin, even at the top of the wage distribution. As a consequence, wages changes are unlikely to generate non-Gaussian features which relate to large earnings changes like downwards asymmetry and fatness of the tails, even at among top-earners.

5.5.2 Wage growth and probability of employment


While the latter accounting decomposition is suggestive, it does not tell the whole story, because hourly wage changes are likely to be correlated with working time changes. We investigate therefore how past hourly wages changes depend on subsequent labor supply decisions at the extensive margin by estimating on Figure 9 the probability of being out of employment at time $t + k$ as a function of past hourly wage changes $w_{i,t} - w_{i,t-1}$.

We find that the probability of being out of employment at time $t + k$ has a U-shape with respect to past hourly wages changes: workers who experienced large hourly wages changes between $t - 1$ and t , both positive and negative, are less likely to remain in employment in the future. The difference with respect to past wage growth can be quite large: workers that experience 40 log-points (resp. -40 log-points) hourly wages changes between $t - 1$ and t are 12 (resp. 8) percentage points more likely to be without employment in $t + 1$.

Interpretation of this pattern can prove tricky, because causality flows either way. Namely, changes in the wage rate are likely to trigger working time responses both on the supply and demand sides. Reversely, changes in the wage rate may depend on future employment through institutional setting, in particular because our measure of earnings incorporates some part of severance payments.¹² Disentangling these channels is a challenging task that is beyond the scope of this paper. However, workers being much more likely to leave employment after they experience massive positive hourly wages growth is puzzling and suggests sever-

¹²It may also be that changes in the wage rate reflect anticipations of future employment, in the case where agents are forward-looking.

Figure 9 – Non-employment response to hourly wage changes



ance payment to be at play here. As a consequence, non-Gaussian features that are related to the tails of the hourly wages growth distribution might be due to employment decisions at the extensive margin. This stresses even more the key role of labor supply decisions in deviations from normality when it comes to the distribution of labor earnings changes.

6 Robustness checks

This section describes some robustness checks performed along several dimensions.

First, we rely on moments of the distribution of changes in labor outcomes (standard deviation, skewness and kurtosis) as opposed to quantile-based indicators. Though this methodological choice has a certain quantitative impact, it does not affect our findings from a qualitative point of view (see Appendix B.1).

Second, we perform the whole analysis conditional on the ranking in the distribution of annual earnings, as opposed to the distribution of hourly wages. The main advantage of the former relies in a better approximation of a productivity, but it is perhaps more dependent on the quality of the "working hours" variable. Once again, qualitatively speaking, our results are robust to this change

(see Appendix B.2).

Third, from an econometric perspective, we removed period effects from equation (2). Although this yields different lifecycle profiles, when it comes to earnings changes our results hardly vary from a qualitative perspective (see Appendix B.3).

Fourth, we decompose our analysis by age groups as Guvenen et al. (2016) do. Most stylized facts documented in the US are also observed in France. However, for the sake of clarity, since our goal here is to disentangle carefully hourly wage from working time effects, we present results obtained by pooling over these age groups (see Appendix B.4).

Fifth, our findings are robust to a bunch of issues related to sample composition and to data issues such as left-censoring and right-censoring (winsorization) of annual earnings, imputation, definition of working time and hourly wages. Changing the censoring threshold from 45 days a year to 30 days a year, from 1/8 to 1/12 of the annual legal duration of work and from 90% to 60% of the minimum hourly wage does not impact our findings (see Appendix B.5). Neither does the trimming of very high annual earnings. Defining working time in full-time units (FTU) instead of hours has also a very minor impact on results.¹³

7 Conclusion

This investigation of labor earnings dynamics builds on a nonparametric estimation of individual earnings, wages and working time changes. We find the same striking results as in the US: labor earnings exhibit several non-Gaussian features in France too, including a negative skewness and a high kurtosis. More interestingly, the availability of working hours and work days in our dataset enables us to disentangle hourly wage from working time growth at both the intensive and extensive margins of employment. Major deviations from normality stem rather from working time at the extensive margin, i.e. employment duration and to a

¹³We also replicated our approach on a sample of women working in the private sector. The main, stylized facts documented here still hold. However, labor supply decisions may be more difficult to disentangle from hourly wage growth in that case. We also believe that they deserve specific attention: in an on-going research project Pora and Wilner (2018a), we investigate gender differences in labor earnings growth as well as in labor earnings dynamics along the lifecycle.

lesser extent at the intensive margin, i.e. hours worked per day, which is conform to the rationale and consistent with assumptions made by the literature on earnings dynamics *à la* [Mincer \(1958\)](#). Large changes in annual earnings reflect mostly in working time changes while small variations in earnings correspond to wage changes. This is in line with the job ladder hypothesis: the unemployment risk may be sufficient to generate non-Gaussian labor earnings volatility, without need of wage changes themselves deviating from normality. Furthermore, a joint analysis of both employment and wage dynamics shows large and positive hourly wage changes to correlate with subsequent departure from employment. A likely explanation seems to be the inclusion of severance payments in our measure of earnings, which would imply that some part of the slight deviations from normality in the hourly wage growth distribution is itself related to working time dynamics at the extensive margin. To sum up, transitions from and to employment are key drivers of non-Gaussian, nonlinear labor earnings dynamics.

A challenging task would consist in quantifying the role played by institutional setting on deviations from Gaussianity. Moreover, institutional constraints determine the level of insurance provided by unemployment benefits, which in turn might alleviate or increase the welfare cost of earnings uncertainty. In particular, extending the methodology at stake by including other sources of income (e.g., self-employment and other non-salaried earnings like welfare benefits), studying household income as opposed to individual earnings (in order to enlighten within-household insurance with respect to labor earnings risk) are left for further research.

Finally, this descriptive framework has focused on cross-sectional distributions of labor earnings, hourly wages and working time changes, on the one hand, and on the correlation between two subsequent changes, on the other hand. Yet, from a truly dynamic perspective at the individual level, current changes are likely to be a response to previous changes: a more complete, possibly structural model of both hourly wages and working time is needed to capture better this phenomenon.

References

- Afsa, C., and S. Buffeteau. 2006. “L’activité féminine en France : quelles évolutions récentes, quelles tendances pour l’avenir ?” *Économie et Statistique* 398-399:85–97.
- Altonji, J.G., A.A. Smith, and I. Vidangos. 2013. “Modeling earnings dynamics.” *Econometrica* 81:1395–1454.
- Arellano, M. 2014. “Uncertainty, Persistence, and Heterogeneity: A Panel Data Perspective.” *Journal of the European Economic Association* 12:1127–1153.
- Arellano, M., R. Blundell, and S. Bonhomme. 2017. “Earnings and Consumption Dynamics: a Nonlinear Panel Data Framework.” *Econometrica* 85:693–734.
- Baker, M., and G. Solon. 2003. “Earnings Dynamics and Inequality among Canadian Men, 1976–1992: Evidence from Longitudinal Income Tax Records.” *Journal of Labor Economics* 21:289–321.
- Blundell, R., L. Pistaferri, and I. Preston. 2008. “Consumption Inequality and Partial Insurance.” *The American Economic Review* 98:1887–1921.
- Blundell, R., and I. Preston. 1998. “Consumption inequality and income uncertainty.” *The Quarterly Journal of Economics* 113:603–640.
- Bonhomme, S., and J.M. Robin. 2009. “Assessing the Equalizing Force of Mobility Using Short Panels: France, 1990-2000.” *The Review of Economic Studies* 76:63–92.
- Bozio, A., T. Breda, and M. Guillot. 2016. “Taxes and Technological Determinants of Wage Inequalities: France 1976-2010.” Working Paper No. 2016-05, PSE.
- Busch, C., D. Domeij, F. Guvenen, and R. Madera. 2018. “Asymmetric business-cycle risk and social insurance.” NBER Working Paper 24569.
- Busch, C., P. Fialho, and F. Guvenen. 2018. “Higher-Order Wage and Hours dynamics over the Life-Cycle: Evidence from France and Germany.” Work in progress.

- Crow, E.L., and M.M. Siddiqui. 1967. "Robust Estimation of Location." *Journal of the American Statistical Association* 62:353–389.
- Deaton, A.S. 1997. "Econometric Issues for Survey Data." In T. W. Bank, ed. *The Analysis of Household Surveys: A Microeconomic Approach to Development Policy*. The Johns Hopkins University Press, pp. 123–127.
- Deaton, A.S., and C.H. Paxson. 1994. "Saving, Growth, and Aging in Taiwan." In D. A. Wise, ed. *Studies in the Economics of Aging*. University of Chicago Press, pp. 331–362.
- Guvenen, F., F. Karahan, S. Ozkan, and J. Song. 2015. "What do data on millions of US workers reveal about life-cycle earnings risk?" NBER Working Paper 20913.
- . 2016. "What do data on millions of US workers reveal about life-cycle earnings risk?" mimeo.
- Guvenen, F., S. Ozkan, and J. Song. 2014. "The Nature of Countercyclical Income Risk." *The Journal of Political Economy* 122:621–660.
- Hoffmann, E.B., and D. Malacrino. forthcoming. "Employment Time and the Cyclicity of Earnings Growth." *Journal of Public Economics*.
- Jessen, R., D. Rostam-Afschar, and S. Schmitz. 2018. "How important is precautionary labour supply?" *Oxford Economic Papers* 70:868–891.
- Katz, L.F., and K.M. Murphy. 1992. "Changes in Relative Wages, 1963-1987: Supply and Demand Factors." *The Quarterly Journal of Economics* 107:35–78.
- Kelley, T.L. 1947. *Fundamentals of Statistics*. Harvard University Press.
- Kim, T.H., and H. White. 2004. "On More Robust Estimation of Skewness and Kurtosis." *Finance Research Letters* 1:56–73.
- Low, H., C. Meghir, and L. Pistaferri. 2010. "Wage Risk and Employment Risk over the Life Cycle." *The American Economic Review* 100:1432–1467.

- Magnac, T., N. Pistolesi, and S. Roux. 2018. "Post Schooling Human Capital Investments and the Life-Cycle of Earnings." *The Journal of Political Economy* 126:1219–1249.
- Mason, K.O., W.M. Mason, H.H. Winsborough, and W. Poole. 1973. "Some Methodological Issues in Cohort Analysis of Archival Data." *American Sociological Review* 38:85–97.
- Mincer, J. 1958. "Investment in human capital and personal income distribution." *The Journal of Political Economy* 66:281–302.
- Moffitt, R.A., and P. Gottschalk. 2011. "Trends in the covariance structure of earnings in the US: 1969–1987." *The Journal of Economic Inequality* 9:439–459.
- . 2002. "Trends in the transitory variance of earnings in the United States." *The Economic Journal* 112:C68–C73.
- Parker, J.A., and B. Preston. 2005. "Precautionary Saving and Consumption Fluctuations." *The American Economic Review* 95:1119–1143.
- Pijoan-Mas, J. 2006. "Precautionary savings or working longer hours?" *Review of Economic Dynamics* 9:326 – 352.
- Pistolesi, N. 2014. "Income and Consumption Risk: Evidence from France." *Annals of Economics and Statistics* 113-114:347–377.
- Pora, P., and L. Wilner. 2018a. "Ceilings and Floors: The Gender Pay Gap over the Lifecycle, France 1995-2012." Work in progress.
- . 2018b. "Heterogeneous Exposure to Labor Earnings Risk." Working Paper INSEE F1802.
- Verdugo, G. 2014. "The great compression of the French wage structure, 1969–2008." *Labour Economics* 28:131–144.

A Testing for symmetry of annual changes

We compute confidence intervals for our measures of asymmetry and peakedness of tails of the distribution of 5-year earnings (resp. hourly wages and working time) changes. Bootstrapping the entire procedure would be computationally burdensome, for it would require to redefine ranks in the distribution of recent wages. We rely on an approximation that provides us with a lower bound of confidence intervals: we assume these ranks are fixed. Hence we bootstrap only within age-year-recent wages cells. The difference between our estimated confidence intervals and the true underlying value depends on how precise our measure of ranks is. It is a decreasing function of the density of recent hourly wages.


It is important to account for any correlation between changes observed for the same individual over different years, especially given that these changes may overlap, for instance if this individual is observed both at time t and for some time $t + k$ with $k < 5$. In order to deal with this issue, we cluster our bootstrap procedure at the individual level, i.e. we resample over individuals rather than over individual-year observations.

Figures 10 and 11 display the results, firstly for labor earnings changes and secondly for hourly wages changes. They show that (i) the asymmetry of 5-year earnings changes is always significantly negative; (ii) the asymmetry of 5-year hourly wages changes is significantly positive in the lower half of the recent wages distribution.

Figure 10 – Testing the symmetry of annual earnings changes



Figure 11 – Testing the symmetry of hourly wage changes



B Robustness checks

B.1 Moment-based measures

Figure 12 – Standard deviations of 5-year earnings, hourly wages and working time changes




Figure 13 – Skewness of 5-year earnings, hourly wages and working time changes





Figure 14 – Kurtosis of 5-year earnings, hourly wages and working time changes



●●● Log Earnings ○○ Log Hourly Wages
 +++ Log Working Hours ■■ Log Employment Duration


B.2 Earnings ranks

Figure 15 – Kelley's Skewness of 5-year earnings, hourly wages and working time changes



B.3 Omission of year fixed effects

Figure 16 – Kelley's Skewness of 5-year earnings, hourly wages and working time changes



B.4 Age patterns

Figure 17 – P90-P10 of 5-year earnings changes: by age group




Figure 18 – P90-P10 of 5-year hourly wages changes: by age group




Figure 19 – P90-P10 of 5-year working hours changes: by age group





Figure 20 – P90-P10 of 5-year employment duration changes: by age group



B.5 Censoring issues

Figure 21 – Kelley's Skewness of 5-year earnings, hourly wages and working time changes

