Heterogeneous Exposure to Labor Earnings Risk*

Pierre Pora† Lionel Wilner‡

Preliminary draft, please do not quote

Abstract

We analyze labor earnings risk based on a comprehensive panel of French individuals issued from administrative records. We use the approach of Guvenen, Karahan, Ozkan, and Song (2015). After documenting that earnings shocks are not gaussian but exhibit both strong asymmetry and high peakedness, we investigate further the heterogeneity of exposure to labor earnings risk. We disentangle the risk related to annual working time from the risk specific to hourly wage. This decomposition proves fruitful in understanding why women or individuals working in the private sector are exposed to a higher annual earnings volatility.

Keywords: Labor earnings risk, non-gaussian shocks, nonparametric estimation, skewness, kurtosis.

JEL Classification: E24, J24, J31.

*We are grateful to Thierry Magnac for his suggestions that are at the origin of this project. We thank Nila Céci-Renaud and Élise Coudin for useful comments. All errors and opinions are ours.

†CREST (INSEE).

‡CREST (ENSAE). Please address correspondence to lionel.wilner@ensae.fr

1 Introduction

Due to the scarcity of exhaustive data on labor earnings, much is still ignored about earnings risk. An original approach to learn more about it has been recently proposed by Guvenen et al. (2015): in an agnostic way they explore what data tell about earnings dynamics. Based on nonparametric estimation of the first moments of the distribution of earnings changes (or shocks, which can be either transitory when they correspond to one-year changes, or permanent when they correspond to five-year changes), the authors put those statistics in close relation to the location in the distribution of recent earnings. The method contrasts with previous literature that has modelled earning dynamics by resorting to sophisticated specifications combining both observed and unobserved heterogeneity as well as uncertainty (often approximated by random walks, autoregressive processes, ARMAs, etc.). Instead of enriching the model to obtain a better fit of the data, this reverse engineering technique at stake leads to rather unexpected results. The main surprise is the following: shocks on (log) earnings are not gaussian, but they rather exhibit a strong asymmetry and a high peakedness, which questions log-normal models *à la* Mincer (1958).

First, we replicate this approach on a French comprehensive dataset. Thanks to administrative forms, the filling of which is mandatory for payroll taxes, exhaustive, longitudinal information on individuals' labor earnings is available. Focusing on men working in the private sector, we find similar results which are consistent with intrinsic features of labor earnings data; such empirical properties would not be country-specific but they might constitute stylized facts that apply to labor earnings in general. For instance, the bottom 90% face shocks that keep on being less and less favorable on average as one gets richer; but the reverse holds for the top 10%. Extremities of the distribution are much more exposed to earnings instability (or risk, or volatility), this dispersion being however lower in France than in the US. We also find that annual earnings changes are negatively skewed (a notable exception concerns the bottom 5%) and exhibit a very high kurtosis.

Second, we investigate further dimensions of heterogeneity in those earnings

changes. In particular, we pay attention to four of them: age, gender, sector and education. The main findings are the following. As in the US, young workers are most exposed to labor earnings uncertainty. Women tend to experience slightly less favorable earnings shocks than men do, which is consistent with a glass ceiling. Individuals working in the private sector are exposed to less favorable earnings shocks and they are subject to a higher risk. Finally, though earnings tend to grow more rapidly for graduates, the effect of education on the dispersion of annual earnings changes varies much with the position in the earnings hierarchy.

In a last step, we propose a decomposition of the labor earnings volatility into a risk that is specific to the working time instability and into another risk that depends only on hourly wage shocks. This approach is meaningful to disentangle both risks and to analyze where the volatility comes from. Moreover, it sheds light on the reasons why subpopulations are exposed differently to annual earnings risk. The main lessons that can be drawn from this exercise are: the hourly wage keeps on increasing on average, regardless of the location in the distribution; the skewness of shocks on hourly wages is everywhere positive (more favorable changes on hourly wages are more frequent, which is consistent with the rationale); after taking into account the fact that women are more likely to experience a lower working time than men do, the difference in terms of hourly wage prospects vanishes in the bottom of the distribution, but they persist at the top, as they would in presence of a glass ceiling; the risk faced by workers in terms of hourly wage is roughly the same in both private and public sectors.

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 the measure of labor earnings risk. Section 5 provides our results and Section 6 furnishes some robustness checks.

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), Magnac, Pistolesi, and Roux (2013) and Ceci-Renaud, Charnoz, and Gaini (2014). Most of this literature relies on the assumption of Gaussian shocks. Few exceptions depart from this hypothesis, but they include Bonhomme and Robin (2009) who rely on copulas to model the transition probability of an AR(1) transitory component of log earnings. In a recent work, Guvenen et al. (2015) propose to have an agnostic look at individual-level earnings data. In particular, 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 period 1978-2011. The Master Earnings File they use is derived from the W-2 form and presents at least three advantages: it has a large sample size, with low measurement error and no top-coding of annual labor earnings. Their main result is that labor earnings shocks are not log-normal, contrary to what most models *à la* Mincer posit. These annual variations are negatively skewed and display a very high kurtosis, which does not coincide with usual features of normal distributions. The asymmetry stems from the fact that upwards shocks are less likely than (large or disaster) downwards shocks. The high kurtosis means that most individuals experience very small earnings shocks, but also that a small and non-negligible number of individuals face very large shocks. Another lesson is that positive (negative) shocks tend to be transitory (permanent) for high income individuals; the reverse holds for low income individuals; large shocks tend to be more transitory than small shocks. However, two important *caveats* (that are due to data limitations) are worth being mentioned: they restrict their attention to men in the private sector and they do not dispose of any information regarding working time, which does not enable them to decompose labor earnings risk.

In the US and in the UK at least, the rise in earnings inequality has been driven by skill-biased technical change, which has been extensively documented at least 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. Using an approach relying on quantile regressions, Charnoz, Coudin, and Gaini (2014) show that inequality has remained rather stable over the 1976-1992 period

and has even diminished over the 1995-2004 period. These stylised facts may be due to the existence of the minimum wage which plays a role in compressing the earnings distribution, to unemployment and to the natural replacement of older cohorts by younger ones. According to Verdugo (2014), wage inequality decreased continuously from 1969 to 2008. Interestingly, he relates changes in the wage structure to changes in education levels. 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 of inequalities at the top of the distribution. Since inequalities also decreased in the bottom with the rise of the minimum wage, this conjunction of phenomena resulted in a compression of the French wage structure. Moreover, these changes cannot be fully explained by selection into employment (at least at the top) as well as by changes in education and experience composition of the labor force.

Why do we care about labor earnings uncertainty? Essentially because it has a direct impact on consumption inequality, as pointed out by Blundell and Preston (1998). In particular, a high volatility prevents individuals from smoothing their consumption. Hence inequality results in inefficient allocations.

3 Data

Our analysis is based on a comprehensive panel of French salaried employees, which merges the longitudinal versions of the *Déclaration Annuelle de Données Sociales* (DADS) in the private sector and of the *Système d'Information sur les Agents du Service Public* (SIASP) in the public sector. By law,[1] French firms have to fill in the DADS (an annual form that is the analogue of the W-2 form in the US) for every employee subject to payroll taxes; the same holds in the public sector with SIASP. Every firm (more precisely, every establishment) has a unique identifier, the SIRET,[2] a 14-digit number, while individuals are identified by their NIR, a social security number with 13 digits. From 1967 to 1975 however, the DADS

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

[2] The SIRET is a concatenation of the SIREN, a firm identifier, and of an establishment identifier.

panel does not contain any information on firms while from 1976 to 2011, the data is available at the individual-firm level. The panel contains information about individuals born on October of even-numbered years –a representative sample of the French salaried population at rate 1/24. Since filling in the form for payroll taxes is mandatory in both private and public sectors, 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. Some years are missing (1981, 1983 and 1990) because there was no data collection by Insee during the 1982 and 1990 censuses. In 1994, the quality is nevertheless questionable. Furthermore, the data on the public sector has been available since 1988 only, which leads us to restrict our attention to the 1988-2012 period. Finally, our dataset contains detailed information about gross and net wages, work days, work hours,[3] other job characteristics from 1976 onwards (like 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 variable of interest will be annual earnings defined as the sum of salaried earnings.

Our working sample is composed of salaried employees working in metropolitan France and aged between 20 and 60. Two years are missing during our period of interest: 1990 and 1994. We then impute annual earnings to individuals who have a positive wage both preceding and following years with the average annual earnings of those years.

The empirical analysis described in section 4 requires to select individuals with a relatively strong attachment to the labor workforce, namely with stable trajectories. We rely on "relatively stable" workers to describe annual changes in earnings between year t and year $t+1$ (see *infra*). We impose in particular that those individuals were also present at least two years between $t-5$ and $t-2$, on top of being present in $t-1$, in t and in $t+1$. Finally, to deal with very low annual earnings, we focus only on individuals earning more than 1/8 of the annual minimum wage $\underline{w}$.

In Table 1, we give a few descriptive statistics on the successive steps of the

[3] in the DADS (private sector) and since 1995 only.

selection of "relatively stable" workers. First comes the censoring at $\frac{1}{8}\underline{w}$. Second comes the restriction to individuals that were present two years between $t-5$ and $t-2$, on top of being present in $t-1$ and t. Third, we focus on individuals that were also present in $t+1$.

The first step of the selection process mostly increases the average earnings both for men and women, and in the public and the private sectors, specifically at age 25. Second and third steps deal with stability criterions. They eliminate most of the younger workers, whose career just started, and increase the proportion of individuals working in the public sector. This is consistent with a lower unemployement risk among workers of the public sector.

4 Empirical analysis

We follow the same descriptive approach as Guvenen et al. (2015). In particular, we rely on nonparametric estimations of shocks on annual labor earnings. A decisive advantage of this methodology is not to posit any parametric assumption on the form of shocks, contrary to most of the literature devoted to earnings dynamics.

Let denote annual earnings of individual i on year t by y_{it}. We aim at measuring individual-level log earnings changes at a one-year horizon, i.e., annual growth. We consider a normalized version of log earnings net of age and cohort effects. Hence we regress individual log earnings on a full set of age and (year-of-birth) cohort dummies:

$$\log y_{it} = \sum_c \alpha_c \mathbb{1}[\text{cohort}_i = c] + \sum_a \beta_a \mathbb{1}[\text{age}_{it} = a] + \epsilon_{it}. \tag{1}$$

This model is estimated on a sample of workers aged 20 to 60 in metropolitan France, provided that their annual earnings exceed 1/8 of the annual minimum wage $\underline{w}$. The estimates of (1) enable us to define our variable of interest, which we note $\log \tilde{y}_{it} = \log y_{it} - \beta_a$ and which can be interpreted as (log) earnings net of age effects. Therefore the annual change in normalized (log) earnings $\log\left(\frac{\tilde{y}_{i,t+1}}{\tilde{y}_{it}}\right)$ accounts for the relative evolution of individual i's annual wages with respect to

Table 1: Descriptive statistics on the selection process

		Initial data	Censoring at 1/8 minimal wage	Past recent earnings selection	Final selection
Total number of observations (individual*year)		24 161 331	22 097 285	14 844 143	13 011 106
Under 25		15,9 %	12,8 %	3,0 %	2,9 %
25-29		13,7 %	13,9 %	13,5 %	13,6 %
30-39		26,5 %	27,6 %	30,5 %	31,0 %
40-49		24,6 %	25,8 %	30,3 %	31,0 %
50-59		16,8 %	17,6 %	21,6 %	20,9 %
Over 59		2,6 %	2,3 %	1,1 %	0,5 %
Men		53,4 %	54,1 %	54,6 %	54,8 %
Women		46,6 %	45,9 %	45,4 %	45,2 %
Private sector		78,1 %	77,5 %	75,9 %	75,5 %
Public sector		21,9 %	22,5 %	24,1 %	24,5 %
Average earnings at 25 (adjusted euros 2012)	Men	13 206 €	14 324 €	15 676 €	16 124 €
	Women	11 685 €	12 909 €	14 068 €	14 611 €
	Private sector	12 335 €	13 543 €	14 768 €	15 282 €
	Public sector	13 222 €	14 214 €	15 821 €	16 210 €
Average earnings at 45 (adjusted euros 2012)	Men	27 284 €	28 087 €	29 065 €	29 525 €
	Women	19 050 €	20 062 €	21 491 €	21 895 €
	Private sector	23 081 €	24 135 €	25436 €	25 930 €
	Public sector	24 307 €	24 941 €	25 939 €	26 196 €

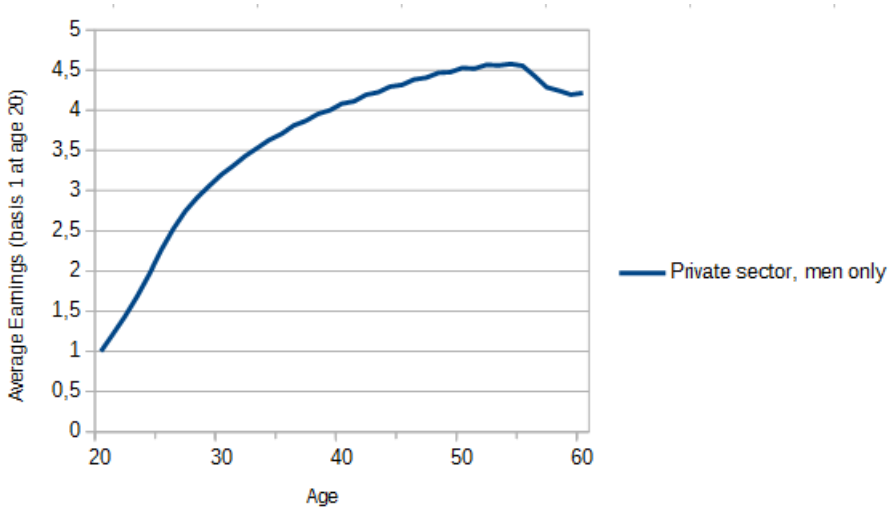

Figure 1: Age profile of average annual earnings (men in the private sector)

his analogues of the same age.

Figure 1 displays the (exponential of) estimated age coefficients $\exp(\beta_a)$ and therefore represents the average life-cycle profile of log earnings. From 20 to 27, average earnings almost grow linearly: average earnings are as high as 2.7 times what they were at 20. This corresponds mostly to the gradual entrance of youth in the labor market, which comprises non full-time work periods; in other words, the rise in the annual job duration helps explain the linear growth. On top of cumulating more experience, young salaried workers may also see their hourly wage increase rapidly due to an iterative improvement of the worker-firm matching, which results in a better recognition of their working abilities (Topel and Ward, 1992). Then we observe an hump-shaped pattern that peaks at age 54 where individual earnings are about 66% higher than those at 27; this pattern is consistent with diminishing marginal returns of experience and seniority (Becker, 1964). Guvenen et al. (2015) found on US data that the rise in earnings between 25 and 55 was about 127%; on French data, the corresponding figure amounts to 101%[4]. This may result of our concept of labor earnings being perhaps narrower since it covers only salaried wages. Earnings decrease slightly from 54 to 60, which stems from a gradual exit of the labor workforce that may occur during the year, and hence

[4]For men working in the private sector. Within the whole population (men and women, public and private sectors, the estimated rise in earnings between 25 and 55 is 94%

lowers the working time. We analyze then earnings changes between t and $t+1$ according to the location of individuals in the distribution of recent earnings $\overline{y}$. The definition of the latter relies also on normalized earnings that are averaged over the past five years, and more precisely over the years when an individual is recorded in the panel between $t-5$ and $t-1$:

$$\overline{y}_{it} = \frac{\sum_{\tau=t-5}^{t-1} y_{i\tau} \mathbb{1}[y_{i\tau} > \frac{w}{8}]}{\sum_{\tau=t-5}^{t-1} \exp(\beta_{\text{age}_{i\tau}}) \mathbb{1}[y_{i\tau} > \frac{w}{8}]}. \tag{2}$$

Workers are ranked according to their average normalized recent labor income. They are next divided into 100 percentile groups or bins. Hence we resort to local statistical indicators in order to describe the heterogeneity of annual individual changes in normalized earnings, in the sense that they are specific to the location in the distribution of recent earnings. Those statistical indicators correspond to first moments (mean, variance, skewness and kurtosis) of the distribution of evolutions. The mean indicates how much normalized earnings evolve at some point of the distribution of past earnings while the variance describes the dispersion (or volatility, instability, uncertainty, risk, etc.). The skewness accounts for the degree of asymmetry in these shocks while the kurtosis measures the peakedness and the heaviness of the tails of the distribution of those shocks. In practice, these indicators are computed for each quantile of rank α from the subset of workers whose average past normalized earnings is comprised between the $100(\alpha-1)$th percentile and the 100αth percentile. In variants, we also do the computations between the $1000(\alpha-1)$th thousandth and the 1000αth thousandth, but also between the $20(\alpha-1)$th twentieth and the 20αth twentieth.

Since the data on working time has been available in the DADS panel since 1995, the decomposition of earnings as the product of the hourly wage and the working time can be done only from that date onwards. As will be explained below, to isolate hourly wage risk from working time risk we propose to compute preceding indicators for individuals working full-time during a whole year.

5 Results

5.1 Stylized facts: a comparison with the US

We present in this subsection some kind of replication of the paper by Guvenen et al. (2015) on French data. For the sake of a meaningful comparison, we have to restrict our attention to men only, since women are absent from the analysis in the US, and to the private sector. We find roughly similar results concerning the distribution of earnings shocks, namely normalized earnings growth $\log\left(\frac{\tilde{y}_{i,t+1}}{\tilde{y}_{it}}\right)$ or log earnings net of the age effect.

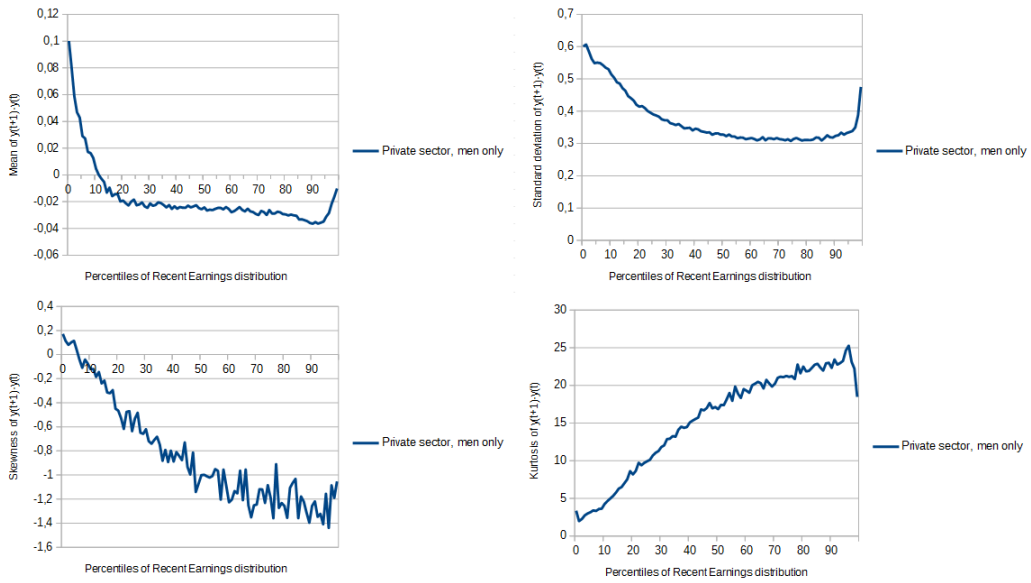

Figure 2: Distribution of annual earnings changes (first moments, men in the private sector)

Having a look at the first moment of earnings growth as a function of the location in the distribution of recent earnings (Figure 2) leads to a somehow unexpected diagnosis. While the bottom 90% of the distribution experience a mechanical adjustment effect (shocks are on average higher for lower normalized recent earnings), the top 10% face shocks which tend to be increasingly favorable with past earnings. For instance, earnings growth of individuals whose recent earnings lie between the 10th and the 11th percentile of the distribution is on average about 1.2pp higher than the average one of comparable workers (that is, having the same

age, this effect being neutralized thanks to the normalization of earnings). It is 2.1pp (2.8pp) lower for individuals between the 50th (90th) and the 51th (91th) percentile but only .4pp lower for individuals belonging to the top 1%.

Examining now the second moment (or rather the standard deviation) of earnings growth (Figure 2), we obtain the same U-shaped pattern as Guvenen et al. (2015) do. Individuals most exposed to volatility belong to extremities of the distribution. In particular, the very bottom of the distribution is most exposed to shocks with standard deviations being as large as .6, which is a slightly higher dispersion than in the very top of the distribution (.58). Defining labor earnings risk as the uncertainty faced by individual earnings changes, this risk is first higher at the very bottom, then at the very top, and finally in the middle of the distribution. Interestingly, in absolute terms the dispersion of shocks is much higher in the US (between .4 and .9 depending on the location in the distribution) than in France (between .3 and .6). This striking difference stems most likely from a higher job protection of the French labor market which moderates the risk of unemployment for insiders.

Another slight interesting difference provided with this international comparison comes perhaps from the asymmetry of shocks, as measured by the skewness of the distribution of earnings growth (Figure 2). While this third moment is systematically negative along the whole distribution in the US, this is not the case in France: in the lowest part of the distribution (namely, the first 5 percentiles), there is evidence of positive skewness. This exception aside, we conclude however rather to a similar diagnosis: earnings shocks are negatively skewed in France (-1 at the median), even though the asymmetry is less pronounced than it is in the US (-2 at the median). The log-normality assumption is not likely here since it would imply with a skewness of zero. A plausible explanation has been provided in Guvenen et al. (2015): the higher an individual's current earnings, the more room they have to fall and the less room left they have to move up.

The profile of the kurtosis of shocks looks similar to the American one (Figure 2). The fourth moment of the distribution of annual earnings changes is much higher than 0 (i.e., the kurtosis' value of a normal distribution). It increases mono-

tonically (and linearly) up to the 80th percentile until it reaches its maximal value (nearly 27); it then decreases to 18. Moreover, and by contrast to the US, the kurtosis may reach much more extreme values (nearly 50) for specific age groups (40-49 for instance). As a result, the French distribution of earnings shocks sounds even further from a log-normal distribution than the American one.

5.2 An heterogeneous exposure to labor earnings risk

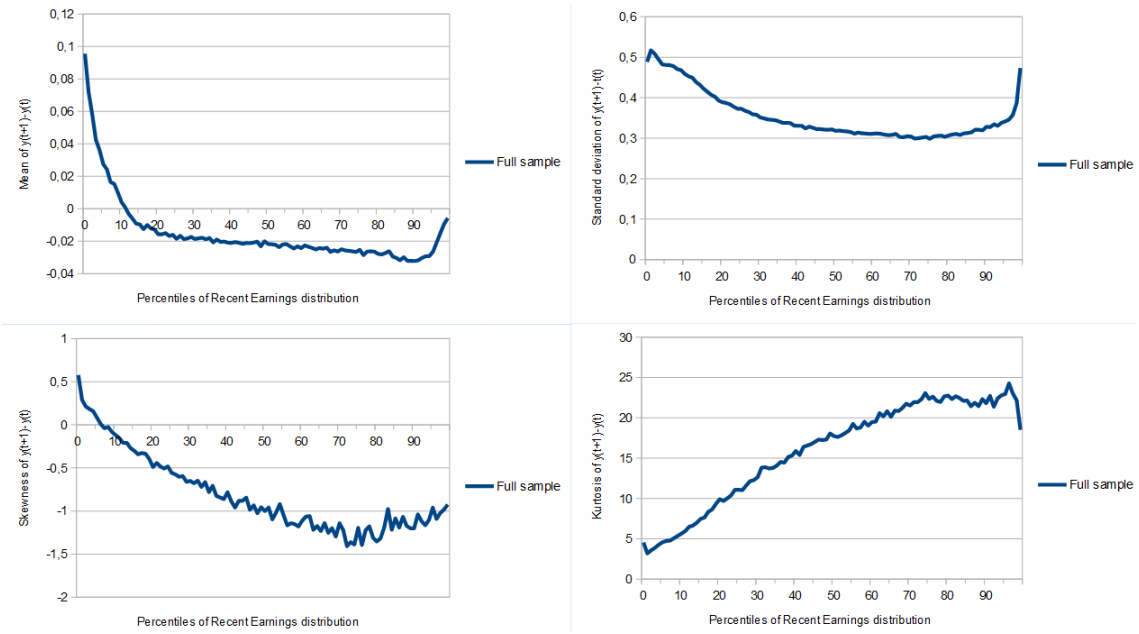

Figure 3: Distribution of annual earnings changes (first moments)

The diagnosis established above remains when we extend the analysis to women and to the public sector (Figure 3). We investigate next how labor earnings risk varies with some dimensions of observable heterogeneity, including age, gender, sector and education.

The findings regarding age-specific exposures to risk described in the American case look also very similar in the French case. The most striking conclusion is that young workers are most exposed to labor earnings uncertainty. On average, earnings changes are most pronounced for them (Figure 4). The dispersion of these shocks is almost decreasing with age everywhere in the distribution. At the bottom of the distribution, shocks tend to be positive on average, and they are

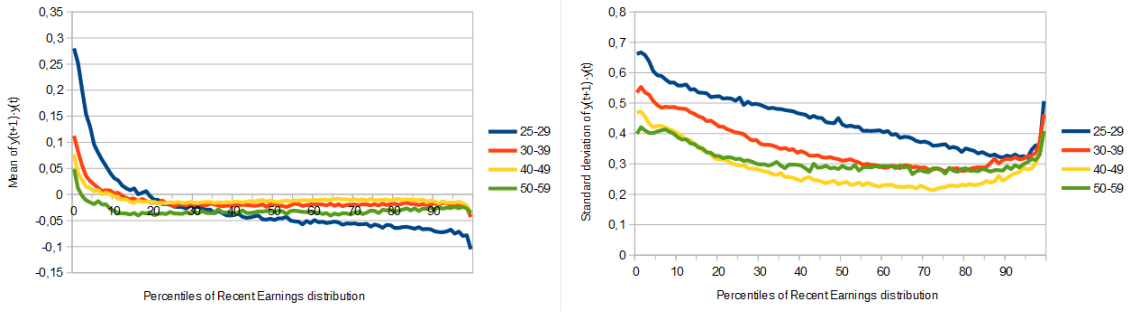

Figure 4: Heterogeneity of exposure to earnings risk (by age)

particularly large for individuals aged 25-29, but the standard deviation is about .7 while it is hardly .4 for older individuals aged between 40 and 59. This result is reminding of the higher volatility of transitory earning shocks documented by both Bonhomme and Robin (2009) and Ceci-Renaud, Charnoz, and Gaini (2014) for young French individuals.

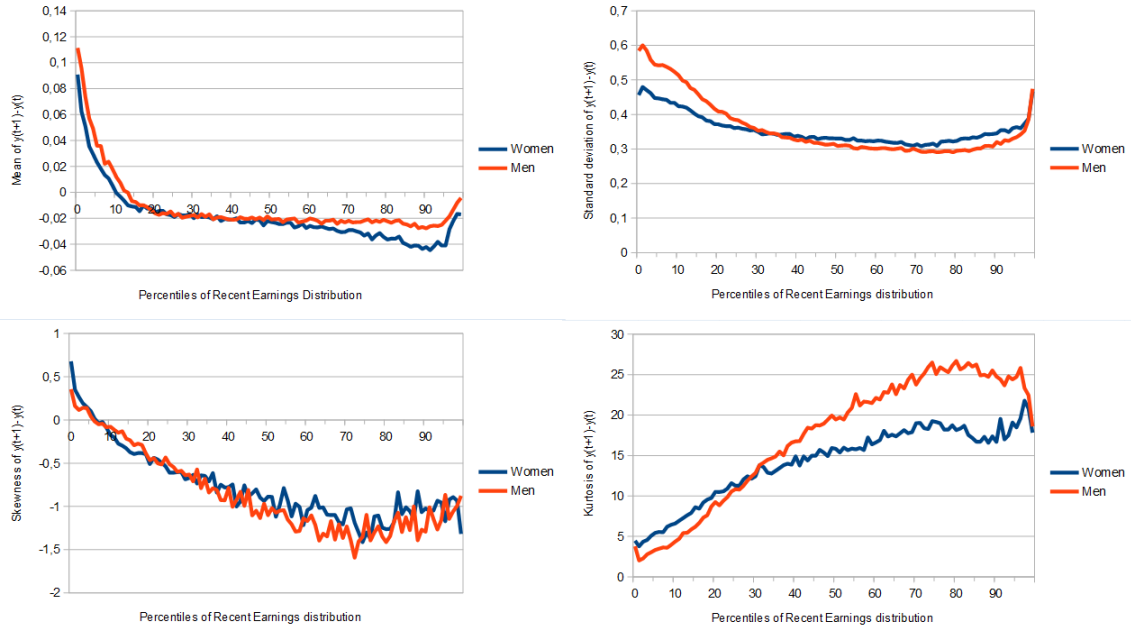

Figure 5: Heterogeneity of exposure to earnings risk (by gender)

There are two main differences as far as average individual earnings shocks affecting either men or women are concerned (Figure 5).

First, at the very bottom but also in the whole first part of the hierarchy, women tend to experience slightly less favorable earnings shocks than men do.

This is consistent with the glass ceiling story, and might help explain the scarcity of women in top positions of the hierarchy. The presence of women is indeed monotonically decreasing as one gets closer to the top of the distribution. In the bottom 1% women are 70% and the only locations where women are more numerous than men are the bottom 30%. From the 40th to the 70th percentile of the distribution, this proportion is roughly stable and amounts about to 43-45% but it falls sharply afterwards: women are only 13% in the top 1%.

Second, women are exposed differently than men to earnings risk. At the bottom of the distribution, the dispersion is much higher for men than for women. A possible explanation for this heterogeneity of exposure might involve different risks of unemployment, which will be investigated somehow in section 5.3. On the contrary, from the 34th percentile to the 98th percentile, women's earnings shocks are slightly more dispersed than men's ones, and no significant difference is observed to this respect in the top 2%. The asymmetry of shocks looks very similar for men and for women, but the kurtosis is much higher for men in the upper part of the distribution.

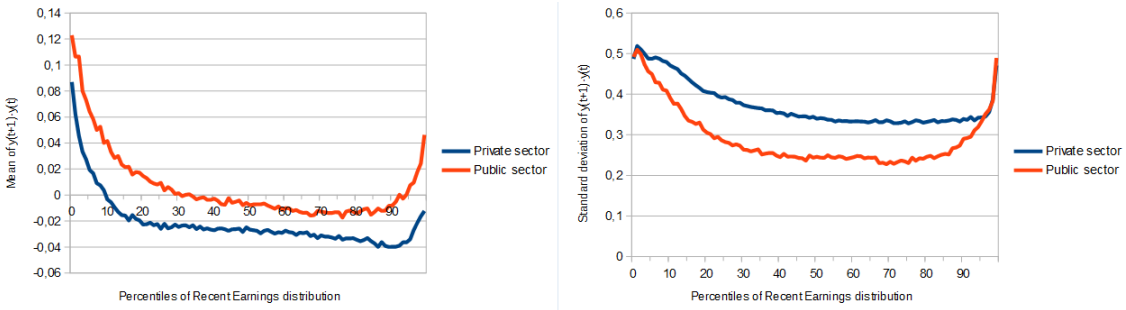

Figure 6: Heterogeneity of exposure to earnings risk (by sector)

Turning now to a comparison between the private and the public sectors, individuals working in the latter are exposed to less favorable earnings shocks. They are also subject to a higher labor earnings risk (Figure 6). However, these rather intuitive findings do not hold at the extremities of the distribution where the risk is almost equal in both sectors. The proportion of individuals working in the public sector is higher in the upper part of the distribution, especially around the 75th percentile, while it is particularly low around the 15th percentile as well as in the

top 1%. The skewness of shocks is almost null at the bottom of the distribution as far as individuals working in the public sector are concerned.

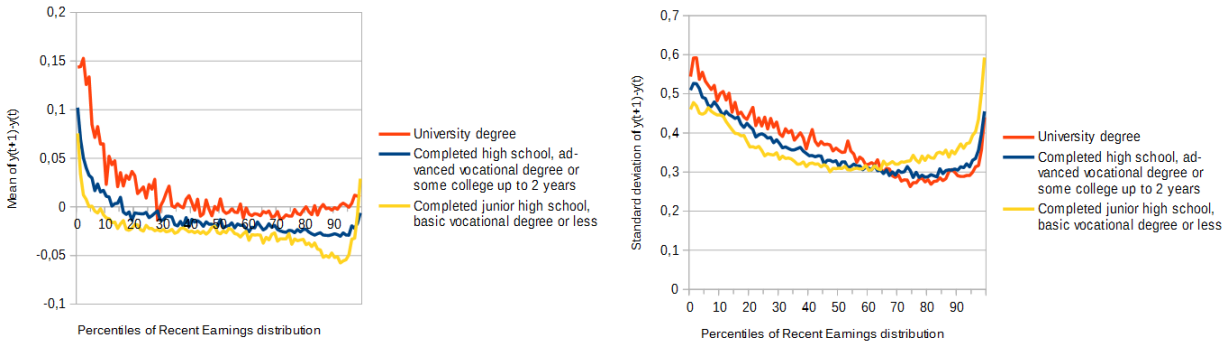

Figure 7: Heterogeneity of exposure to earnings risk (by education)

Does education offer a protection against labor earnings risk? On average, earnings tend to grow more rapidly for graduates (Figure 7). However, the effect of education on the dispersion of annual earnings changes varies much with the position in the earnings hierarchy. The first 70 percentiles experience a higher risk when they have some education while the reverse holds in the last 30 percentiles. Less educated individuals are also more frequently subject to negative shocks, i.e. their skewness is more negative and more pronounced; the same holds yet for most educated people. As another striking feature, in the first 70 (last 30) percentiles, the more educated, the less (more) heavy the tails of the distribution of shocks.

5.3 A decomposition of labor earnings risk

Finally, we would like to go further than the previous analysis by decomposing the labor earnings risk into two risks. We exploit the fact that annual earnings can be written as the product of the hourly wage and of the annual job duration (or working time). The first risk corresponds to an hourly wage risk, whereas the second risk is related to uncertainty about working hours, which includes the risk of unemployment, but also the risk of part-time work, of work hours adjustment, of parental leave, etc. To disentangle both sources of risk, we consider individuals working full-time the whole year only, and we compute the first four moments of the distribution of shocks. The main advantage of this approach is that these

individuals bear no working time risk at all; their sole remaining uncertainty has to do with hourly wage. We face however a constraint driven by data limitations: the information on working time has been available since 1995 and in the private sector only, which leads us to perform the analysis on the corresponding subsample.

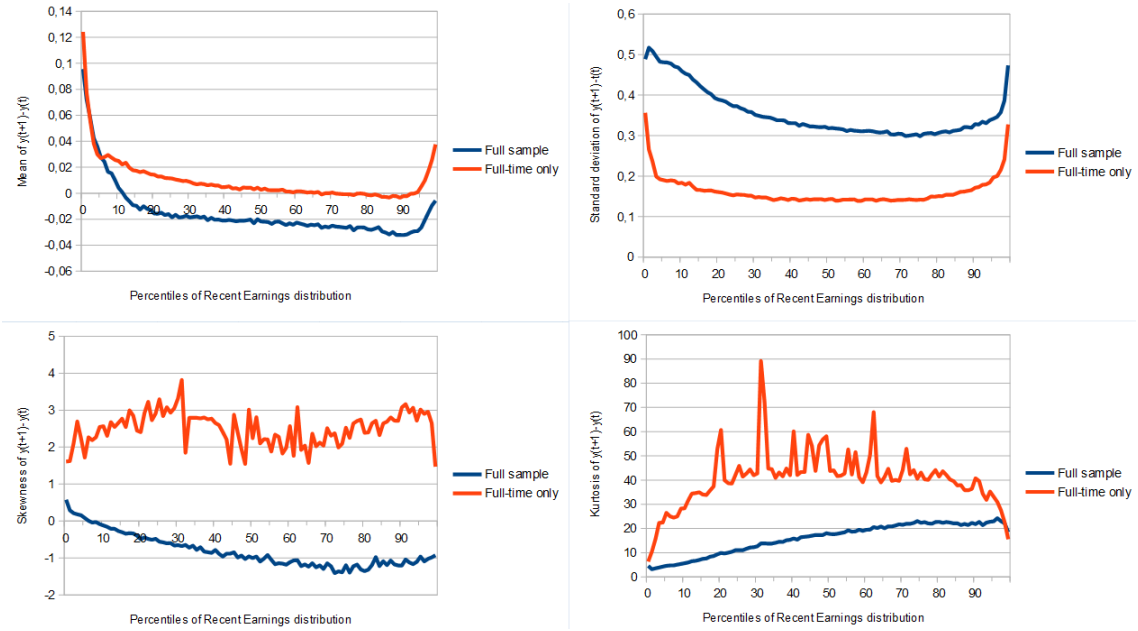

Figure 8: Decomposition of annual earnings risk

The diagnosis relative to the mean of shocks is not much altered (Figure 8): the top 10% is still exposed to shocks that are increasingly more favorable as one gets closer to the top of the hierarchy. However, contrary to what happened previously where only the first 12% could expect positive shocks on average, the hourly wage keeps on increasing on average, regardless of the location in the distribution. The dispersion is much lower but still represents about half of the total risk; it varies much with the location in the hierarchy. Interestingly, in the lower part (the very bottom aside) the risk is much smaller for individuals working full-time during a whole year, which means that locally the main risk corresponds there to uncertainty in working time. More strikingly, and most importantly, the skewness of shocks on hourly wages is everywhere positive, while the skewness of shocks on annual earnings was not: in other words, more favorable changes on hourly wages are more frequent, which is consistent with the rationale. Hence the asymmetry of

annual earnings changes stems entirely from unfavorable shocks on working time. Finally, the kurtosis is even more pronounced as far as shocks on hourly wages are concerned.

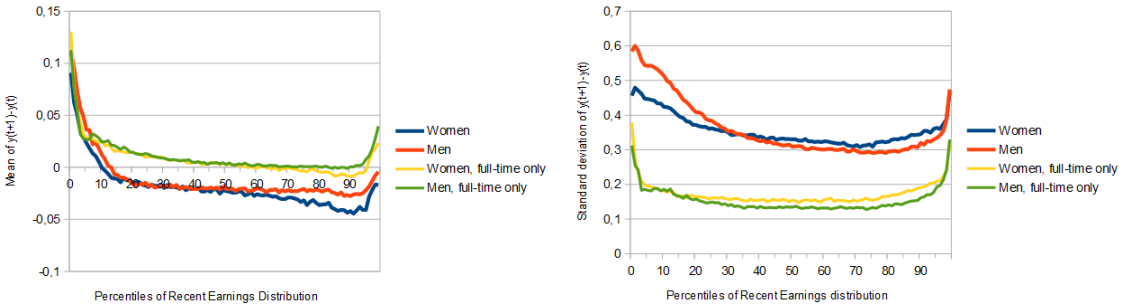

Figure 9: Decomposition of annual earnings risk (by gender)

It is also useful to reconsider the gender analysis above using this decomposition (Figure 9). For instance, there is no difference any longer between men and women at the bottom of the distribution (same mean and same dispersion, which was not the case with annual earnings shocks: women face on average less favorable and more dispersed shocks than men). We can therefore conclude to an heterogeneous, gender-specific exposure to working time risk. Interestingly, after taking into account the fact that women are more likely to experience more unfavorable events to this respect than men do (think of the risk of maternity leave, of part-time work, of unemployment, etc.), there is no significant difference in terms of hourly wage prospects, at least in this part of the distribution. In the top 30% however, one still observes differences in both mean and variance of annual individual earnings changes: these differences reflect here hourly wage differentials, and are consistent with a glass ceiling. Even when they work full-time during the whole year, these women are exposed to less favorable shocks on their hourly wages, which contributes to their eviction from the top locations of the earnings hierarchy.

Similarly, the sectorial analysis above can benefit from the decomposition into the two different kinds of risk (Figure 10). In particular, much of the differences that were pointed out between the private sector and the public sector can be explained by heterogeneous exposures to the risk of unemployment (the latter

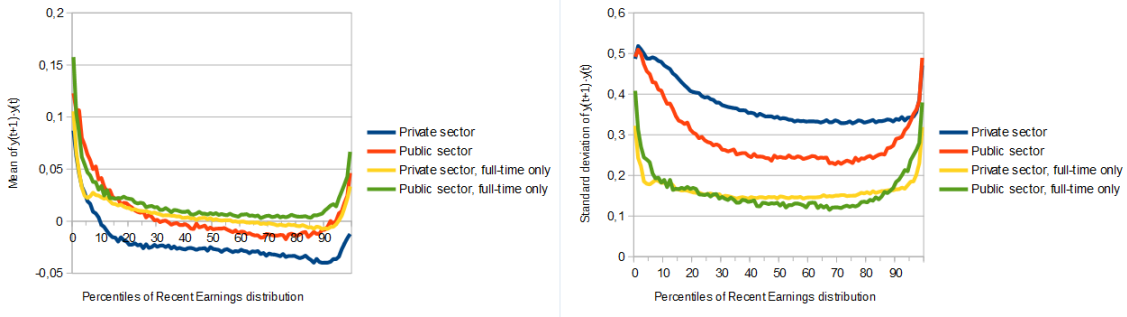

Figure 10: Decomposition of annual earnings risk (by sector)

being included within the working time risk). In terms of hourly wage, the risk faced by workers is roughly the same in both sectors, whereas it is much higher in the private sector in terms of annual earnings. Besides, the more favorable prospects documented previously for public sector workers are somehow attenuated when looking at hourly wages.

6 Robustness

Many robustness checks have been done and will be detailed soon:

- sample/representativeness of the panel: rates 1/12 versus 1/24.
- with respect to the imputation for missing years (1990, 1994).
- we have tried several widths of bins (1/1000, 1/100, 1/20); results are remarkably unsensitive to the chosen partition.
- covariates in equation (1) which might include gender, sector, education, etc.

References

Altonji, J.G., A.A. Smith, and I. Vidangos. 2013. "Modeling earnings dynamics." *Econometrica* 81:1395–1454.

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.

Becker, G. 1964. *Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education*. Chicago: University of Chicago Press.

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.

Ceci-Renaud, N., P. Charnoz, and M. Gaini. 2014. "Évolutions de la volatilité des revenus salariaux du secteur privé en France depuis 1968." Working paper INSEE.

Charnoz, P., E. Coudin, and M. Gaini. 2014. "Changes in the French Wage Distribution 1976-2004: Inequalities within and between Education and Experience Groups." Working paper INSEE.

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.

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.

Low, H., C. Meghir, and L. Pistaferri. 2010. "Wage Risk and Employment Risk over the Life Cycle." *The American Economic Review*, pp. 1432–1467.

Magnac, T., N. Pistolesi, and S. Roux. 2013. "Post schooling human capital investments and the life cycle variance of earnings." IZA Discussion Paper.

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.

Topel, R.H., and M.P. Ward. 1992. "Job Mobility and the Careers of Young Men." *The Quarterly Journal of Economics* 107:439–479.

Verdugo, G. 2014. "The great compression of the French wage structure, 1969–2008." *Labour Economics* 28:131–144.