Journal of Public Economics 169 (2019) 160–171

Contents lists available at ScienceDirect

Journal of Public Economics

journal homepage: www.elsevier.com/locate/jpube

Employment time and the cyclicality of earnings growth☆

Eran B. Hoffmann[a,*], Davide Malacrino[b]

[a] Department of Economics, Hebrew University, Mt. Scopus, Jerusalem 91905, Israel
[b] Internationl Monetary Fund, 700 19th Street NW, Washington DC 20431, United States of America

ARTICLE INFO

Article history:
Received 3 November 2016
Received in revised form 13 September 2018
Accepted 14 September 2018
Available online 11 December 2018

Keywords:
Employment
Labor earnings
Business cycles
Distribution

ABSTRACT

We study how the distribution of earnings growth evolves over the business cycle in Italy. We distinguish between two sources of annual earnings growth: changes in employment time (number of weeks of employment within a year) and changes in weekly earnings. Changes in employment time generate the tails of the earnings growth distribution, and account for its procyclical skewness. In contrast, the distribution of weekly earnings growth is close to symmetric and stable over the cycle. This suggests that studies of earnings risk should carefully model the employment margin to avoid erroneous conclusions on the nature and magnitude of risks underlying individual earnings. We show that the combination of simple employment and wage processes is enough to capture the complex features of the earnings growth distribution.

1. Introduction

A long strand of literature has studied and quantified the individual labor earnings process, and in particular the risk in earnings growth.[1] In a recent paper, Guvenen et al. (2014) use a large administrative dataset from the US Social Security Administration and find that recessions have an asymmetric impact on the distribution of individual earnings growth: large decreases in earnings become more common and large increases less common during recessions, while the distribution of small changes remains stable over the cycle.

A common interpretation of these results is that the downside risk to individual permanent income increases during recessions. This interpretation has spurred a growing number of papers which study its implication on a wide range of economic outcomes, including consumption and wealth dynamics, the effects of public policy, asset pricing, and the transmission of monetary policy.[2]

However, annual earnings growth is the outcome of two distinct components with vastly different statistical properties. Consider the decomposition of individual i's (log) annual earnings growth Δy_{it}, into the change in the (log) employment time Δx_{it}, measured in weeks, and the change in (log) weekly earnings Δw_{it}:

$$\Delta y_{it} = \Delta x_{it} + \Delta w_{it}. \tag{1}$$

Decomposition (1) implies that a change in earnings could be the outcome of either a change in employment time or changes in weekly earnings (or both). The dataset used by Guvenen et al. (2014) does not contain observations of the time spent in employment within a given year, or the weekly earnings when employed, and thus does not allow this decomposition.[3]

☆ We are highly grateful to Monika Piazzesi and Luigi Pistaferri for their valuable guidance. We thank Magne Mogstad (editor) and three anonymous referees for their helpful comments. We also benefited from a discussion with Nick Bloom, John Bluedorn, Raj Chetty, Romain Duval, Mark Duggan, Fatih Guvenen, Caroline Hoxby, Moritz Lenel, Giuseppe Moscarini, Serdar Ozkan, Petra Persson, Itay Saporta-Ecksten, Chris Tonetti, Alonso Villacorta, Eran Yashiv, and the participants at the Society of Economic Dynamics (SED) annual meeting 2016, All California Labor Economics Conference 2016 and the seminar participants at Stanford University. The INPS data were provided by the *Ministero Italiano del lavoro e delle politiche sociali*. Financial support from the Stanford Institute for Economic Policy Research and the Weiland Research Fellowship is gratefully acknowledged.

* Corresponding author.
E-mail addresses: eran.hoffmann@mail.huji.ac.il (E.B. Hoffmann), dmalacrino@imf.org (D. Malacrino).

[1] See for instance MaCurdy (1982), Abowd and Card (1989), Meghir and Pistaferri (2004), Storesletten et al. (2004), Guvenen (2009), Arellano et al. (2017).

[2] DeNardi et al. (2016) and McKay (2017) study consumption and wealth dynamics. McKay and Reis (2016) design an optimal unemployment insurance rule. Constantinides and Ghosh (2014) and Schmidt (2016) study the role of idiosyncratic risk in asset pricing and Berger et al. (2016) study its role in monetary policy.

[3] The Master Earnings File (MEF) data from the US Social Security Administration used by Guvenen et al. (2014) and by others is based on reported income from W-2 tax forms. See for example Song and Manchester (2007), von Wachter et al. (2011) and French and Song (2014). For a discussion of advantages and limitations of these data see Kopczuk et al. (2010).

https://doi.org/10.1016/j.jpubeco.2018.09.009
0047-2727/© 2018 Published by Elsevier B.V.

In this paper, we apply decomposition (1) to measure the contribution of changes in employment time and changes in weekly earnings in shaping the earnings growth distribution. We view the distinction between the employment margin and the wage margin (captured here by weekly earnings) as important for two main reasons. First, the persistence of changes to these two components may be very different. Unemployment spells are typically measured in months, while wage changes have a persistent component which last for many years.[4] Thus, changes in employment time that are not accompanied by changes in weekly earnings are unlikely to have a strong impact on a worker's permanent income. Estimating the earnings process without acknowledging the separate roles of the employment margin and wage margin may therefore be misleading.

Second, the two sources of variation in earnings have different policy implications. For example, unemployment insurance may effectively insure against drops in employment time, but not insure at all against declines in wages. Similarly, wage insurance policies, such as the one suggested by LaLonde (2007), can only reduce the adverse consequences of falls in wages.

We conduct most of our analysis using a large administrative panel dataset from the Italian social security institute (INPS), which covers the period 1985–2012. This dataset includes observations of annual earnings and weeks of employment within every given year for each worker, allowing us to perform decomposition (1) at the worker-year level.

Our analysis of the data is divided into three parts. First, we show that most of the cross-sectional variation in annual earnings growth in Italy is due to changes in employment rather than changes in weekly earnings. In particular, changes in the number of weeks of employment generate the tails of the distribution (see Fig. 1 for a decomposition of a single cross-section), both in recessions and in expansions.

Second, we study how the cross-sectional distribution of annual earnings growth and its components evolve over time. We provide visual and statistical evidence of a strong association between the distribution of changes in employment time and that of annual earnings growth. In particular, the third moments of the distributions, which capture asymmetry, are highly correlated over time and are both procyclical. In contrast, the distribution of changes in weekly earnings around its mean shows little asymmetry and is stable over the cycle.

Third, we propose a model of an earnings process based on the combination of an *employment process* and a *wage process*. The employment process, which is driven by random transitions between labor market states, is enough to generate the tails of the annual earnings distribution, and their cyclical movements. The wage process is the sum of a Gaussian permanent component and a transitory shock, which generate a symmetric wage growth distribution. We demonstrate that this process captures the key features of the earnings growth distribution in Italy with only few parameters. Using labor market flows data from the US, we offer additional suggestive evidence that the proposed employment process captures the cyclical movements of the earnings growth distribution there too.

This paper makes two main contributions to the literature. First, we are able to replicate the main findings of Guvenen et al. (2014) with Italian data, which suggests that they are more universal stylized facts that are also valid outside the context of the US labor market.

Second, and perhaps more importantly, we use the additional information on employment that is available in the Italian data to show that these findings should be carefully interpreted. We show that the distributions of changes in employment time and weekly earnings differ in shape, cyclicality, persistence, and response to policy. Hence, special attention should be given to modeling the two margins as separate processes. Not doing so is likely to lead to erroneous conclusions about the underlying dynamics of annual earnings, and may ultimately lead to misleading inference in richer models of consumption and wealth.

The rest of the paper is organized as follows. Section 2 discusses the related literature. Section 3 describes the theoretical framework behind the decomposition of earnings growth into changes in the employment time and changes in weekly earnings. Section 4 presents the data. Sections 5 reports and discusses the evidence on the role of employment time. Section 6 proposes a model for the earnings process. Section 7 concludes.

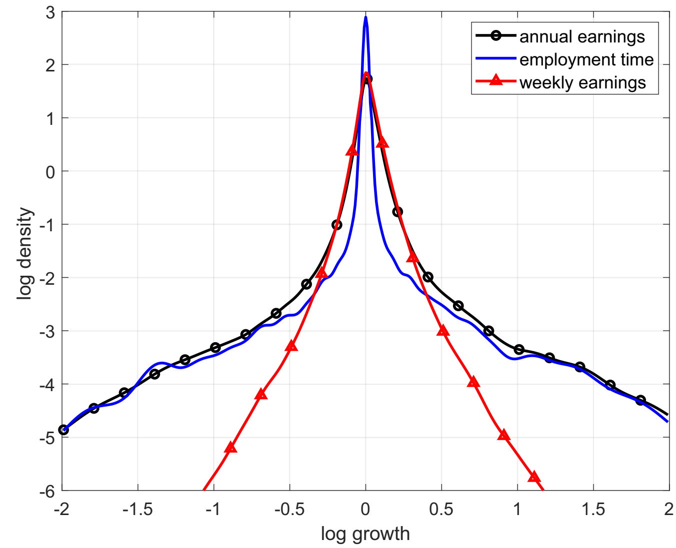

Fig. 1. Decomposition of annual earnings growth distribution.
Notes: Log densities of one-year growth of annual earnings and its components. Based on a representative sample of males 25–60 that includes 300,000 observations (approximately 6.5% of all male workers in the private sector in that age range), for the year 2002. Employment time is the number of weeks of work within a year. Weekly earnings is the annual earnings divided by employment time.

2. Related literature

Many studies rely on time-varying idiosyncratic risk to explain household behavior and, through aggregation, its impact on the macro economy (Constantinides and Duffie, 1996; Storesletten et al., 2007; Krebs, 2007; Low et al., 2010; Constantinides and Ghosh, 2014; Schmidt, 2016). Using data from the PSID, Storesletten et al. (2004) find that the variance of the persistent idiosyncratic shocks to earnings in recessions is twice as high as in expansions. Using US administrative data from the SSA, Guvenen et al. (2014) find that the first and third moments of the cross-sectional distribution of earnings growth are procyclical while the second moment is not, and estimate an earnings process with time-varying non-Gaussian persistent shocks. We provide a new interpretation for the statistical findings of Guvenen et al. (2014) based on data from Italy. While earnings growth in Italy exhibit similar cyclical patterns, a decomposition of individual earnings growth into changes in employment time and changes in weekly earnings reveals a dominant role for employment time in generating the procyclical skewness of earnings growth.

In contrast to our results, Busch et al. (2018) argue that skewness in wage growth is responsible for skewness of annual earnings growth in German data. While some institutional differences

[4] See for example Rothstein (2011) for evidence on unemployment duration and Low et al. (2010) for structural evidence on the persistence of wage shocks.

between the Italian and the German labor market might be responsible for part of the different results, we believe that most of them can be accounted for by two differences in methods. First, we focus on measurement of the third central moment while they base their results on a quantile based measure of skewness. While quantile measures of skewness have the advantage of reducing the influence of outliers, we view the third moment as better fitted for this analysis because it lends itself to a natural decomposition, which quantile based measures lack. Second, we focus on a more direct statistical test of the sources of cyclicality, while they argue by comparing magnitudes of coefficients across two separate regressions. Busch et al. (2018) measure of the cyclicality of quantile differences (L9050,L5010), as presented in their Table III, is substantially smaller for wages than for total earnings. This suggests that, at best, wages can explain only part of the skewness in annual earnings.[5]

The idea that random transitions between labor market states can generate the tails of the earnings growth distribution is also suggested by Hubmer (2018), who proposes a calibrated structural life-cycle model with displacement risk and shows that it captures non-Gaussian features of the cross-sectional distribution of earnings growth. This paper complements his in providing as direct non-parametric evidence that changes in employment time are driving the observed shape of the distribution of earnings growth. The focus of this paper, howver, is focus on the cyclical properties of the distribution, rather than on its life-cycle properties.

3. Decomposing earnings growth

In this section we discuss the decomposition of annual earnings into employment time and weekly earnings and provide some useful definitions for the rest of the analysis.

3.1. The components of earnings

Annual earnings can be separated into three components: employment time in weeks (weeks spent in employment spells, sometimes referred to as the extensive margin), hours worked per week (sometimes referred to as the intensive margin) and the mean wage per hour worked. For individual i at time t, these three components form the following accounting identity:

$$Y_{it} = X_{it} \cdot H_{it} \cdot \tilde{W}_{it} \tag{2}$$

where Y_{it} is annual earnings, X_{it} is employment time in weeks, H_{it} is average hours worked per week and $\tilde{W}_{it}$ is the mean hourly wage.[6]

By taking logs and first differencing, we get the following decomposition:

$$\Delta y_{it} = \Delta x_{it} + \Delta h_{it} + \Delta \tilde{w}_{it} \tag{3}$$

where lowercase denotes logged values, and Δ is the difference between year t and year $t-1$.

In many cases, a direct observation of all three components is absent. The US administrative data from the SSA used by Guvenen et al. (2014), for instance, contain observations of Δy, but not of any of its components. In the INPS data used in this paper, earnings growth, Δy, and changes in employment time, Δx, are observed, but changes in number of hours worked per week and hourly wages are not. Given this limitation of the data, we adopt the decomposition in Eq. (1), in which the mean hours per week and the hourly wage are combined, $\Delta w = \Delta h + \Delta \tilde{w}$, and refer to it as changes in weekly earnings:

$$\Delta y_{it} = \Delta x_{it} + \Delta w_{it}$$

A change in weekly earnings is less straightforward to interpret than a change in mean hourly earnings ($\Delta \tilde{w}$). Weekly earnings confound a potentially endogenous decision margin (hours per week) with the hourly wage, thus Δw cannot be read as a change in "prices" as do hourly wages. Nonetheless, Δw captures changes over a frequency that is typical of employment contracts, often denominated in weekly, monthly or even annual terms rather than hourly pay. In addition, labor economists have found responses in the intensive margin of male workers to be smaller than adjustments in the extensive margin.[7] Thus, we do not see the intensive margin are as crucial for the validity of the results in our paper.

3.2. Decomposing moments of the cross-sectional distribution

In our analysis, we are interested in quantifying the role of Δx in generating both the cross-sectional moments of the distribution of Δy and their variations with the business cycle. More specifically, we would like to measure the contribution of each component to the first three central moments, since these are measures commonly used to describe distributions.

We denote the jth central moment of the cross-sectional distribution as $m_j(\cdot)$, and $m_{k,l}(\cdot,\cdot)$ as the cross-term of order k and l, that is:

$$m_{k,l}(a,b) \equiv \mathrm{E}^i\left[(a - m_1(a))^k (b - m_1(b))^l\right] \tag{4}$$

for any two random variables a and b, and where the superscript i indicates that the expected value is taken with respect to individuals in a single cross section. Here are the decompositions of the first three central moments.

Mean. The first moment of the distribution has an additive decomposition:

$$m_1(\Delta y) = m_1(\Delta x) + m_1(\Delta w) \tag{5}$$

Variance. The second central moment of Δy can be decomposed to a sum of the second moments of its components and an additional cross term.

$$m_2(\Delta y) = m_2(\Delta x) + 2m_{1,1}(\Delta x, \Delta w) + m_2(\Delta w) \tag{6}$$

The cross term for the variance decomposition is two times the covariance. If the covariance is small in magnitude (such as in the case when Δx and Δw are mean-independent) we can measure the approximate contribution of each component as their share of the variance.

Third moment. The third central moment expands to four terms:

$$m_3(\Delta y) = m_3(\Delta x) + 3m_{2,1}(\Delta x, \Delta w) + 3m_{1,2}(\Delta x, \Delta w) + m_3(\Delta w) \tag{7}$$

In the analysis that follows we will refer to the sum $3m_{2,1}(\Delta x, \Delta w) + 3m_{1,2}(\Delta x, \Delta w)$ as the cross term of the third moment. The decomposition of the third central moment of

[5] We replicate their exercise in the Italian data and discuss the results in Online Appendix D.

[6] To give an example, a worker who worked for 39 weeks, 40 h per week (on average) and earned \$20 an hour (on average), would earn \$31.2 K in a given year. i.e. $X_{it} = 39$, $H_{it} = 40$, $\tilde{W}_{it} = \$20$ and $Y_{it} = 39 \cdot 40 \cdot \$20 = \$31,200$.

[7] Most recently, Blundell et al. (2011) find that the elasticity of labor supply on the extensive margin is higher than on the intensive margin across countries. Chetty (2012) argues that even a small adjustment cost can explain the rigidity of labor supply on the intensive margin. Notice that in some cases, particularly during recessions, changes in the extensive margin may be involuntary.

earnings growth is the sum of the third central moment of its components plus this cross term.

4. Data

Our main data source is the Italian social security data (*Istituto Nazionale di Previdenza Sociale* or INPS).[8] INPS collects data on employer and employee relationships in order to compute social contributions and pension benefits. We use a sample dataset covering the period 1985–2012, based on workers who were born in 24 randomly selected birth dates from the universe of all the Italian employees in the non-farm private sector, who are insured at INPS.[9] The data represents a 6.6% sample of this population.

The basic observation in the data is a job relationship with a private employer within a calendar year. For every job relationship we observe the number of weeks of employment and the contributive earnings which include both salary and non-salary components.[10] The earnings from each job relationship are top-coded in accordance with a daily cap of €650 in 2013 (equivalent to individual annual earnings of more than €200K). We find that this cap affects at most 0.5% of all matched observations. In our main analysis we exclude these observations. In Online Appendix A, we show that adding back the observations does not change our analysis.[11]

We obtain an individual level panel including joint observations of annual earnings and weeks of work by performing the following steps. First, we combine information from multiple jobs for the same individual by summing over all records associated with the same worker in a given year. This gives us the annual earnings of that worker. We adjust earnings using CPI and compute the difference in logs to get earnings growth Δy_{it}. Second, we calculate employment time as the number of weeks worked – the minimum between 52 weeks and the sum of all weeks worked across all jobs within a calendar year.[12] We take the difference in logs across every two consecutive years in which the worker is observed to get the changes in employment time Δx_{it}. Lastly, we take the difference between annual earnings growth, and changes in employment time, to recover weekly earnings growth Δw_{it}:

$$\Delta w_{it} \equiv \Delta y_{it} - \Delta x_{it}$$

For comparability to the literature, we restrict the sample to 25–60 years old male workers, who have records both in year t and $t-1$ with earnings above the 2.5 percentile of the income distribution, and who have worked for at least 3 weeks in a given year.[13] This choice also helps mitigate concerns with measurement error. Suppose that work is compensated at a daily rate. Then a worker who worked for two weeks and two days would have a reported three weeks of work and have a measured weekly earnings that is lower than the same worker had he worked for the full three weeks. Fortunately, this type of measurement error is bounded and falls as a percentage of the weekly earnings with the number of weeks of work (1/number of weeks). For example, a worker who works for two weeks has a maximal measurement error of 50%, but a worker that works for 10 weeks has a maximal measurement error of 10%.

We also restrict the sample to employees whose main work contract is full time.[14] Including part-time workers does not change the main results, and dropping these observations reduces the concerns that variation in average weekly earnings could be in fact driven by large changes in hours worked within a week.

We are left with an unbalanced panel of 974,686 workers over 27 years with a total of more than 9 million individual-year observations. In every particular two year period, we have 335,000 observations on average. Workers appear in the sample for 9.7 years on average, and the median worker spends 8 years in the sample. More descriptive statistics about the age and the level of annual earnings, weekly earnings and employment time are reported in Table 1 where all years are pooled together.[15] Mean annual earnings are €27,988 and most workers are employed for the full year (the median worker is employed for 52 weeks). More precisely, 77% of the workers in the sample work for 52 weeks,[16] and to 80% work 50 weeks or more. For comparison, according to the CPS, in the United States 68.9% of prime aged males work 52 weeks, and the figure rises to 82.90% if one only considers individuals working at the time of the interview, those who we think are the relevant sample to compare to ours.[17] Weekly earnings range between €321 at the 10th percentile to €896 at the 90th

Table 1
Summary statistics – INPS panel data.

	Mean	Std.dev.	P10	P50	P90
Age	40.71	9.04	29.00	40.00	54.00
Annual earnings	27,988	18,919	12,181	23,905	45,565
Employment time (weeks)	48.20	9.77	36.00	52.00	52.00
Weekly earnings	572	360	321	474	896
Δ Earnings (Δy)	0.01	0.40	−0.20	0.01	0.22
Δ Employment time (Δx)	−0.01	0.38	−0.10	0.00	0.08
Δ Weekly earnings (Δw)	0.02	0.17	−0.10	0.01	0.15
Observations: 9,293,543					

Notes: The sample is restricted to 25–60 years old males with at least one record at the Italian social security administration between 1985 and 2012. Earnings data are in 2013 euros. Δ terms are in logs. *Source:* INPS data provided by MLPS.

[8] We also use Current Population Survey (CPS) data for some analysis of the United States labor market flows, moments of annual earnings growth provided by Guvenen et al. (2014), the timeseries of GDP growth and CPI from NIPA tables (US), and similar statistics from the Italian National Institute of Statistics (Italy).

[9] We use the LoSai dataset made available by the Italian Ministry of Labor and Social Policy (*Ministero Italiano del lavoro e delle politiche sociali* or MLPS). More information can be found at http://www.cliclavoro.gov.it.

[10] Contributive earnings include the earnings used to compute individual contributions to the social security system (*imponibile previdenziale*) and are different from the taxable earnings (*imponibile fiscale*) as the social security contributions are included in the former but excluded in the latter. While over-time pay is included, some lump sum payments, such as severance payments, are excluded.

[11] Ideally, one would like to include the (unobserved) top-coded earnings in the data and check if the main results are affected. This is not possible, so we perform an exercise where we assign extreme values of the earnings (outside the observed range) to the top-coded observations and replicate some of our results to assess possible consequences of a top coding. The results of our exercise reveal that the effects are likely to be small. See Online Appendix A for a discussion of the results.

[12] If a worker held two jobs at the same time we would not capture it with this data. However, for a subsample of workers we are also able to observe the exact number of months worked, correctly accounting for overlapping jobs. We redo our calculations on this subsample and find that our results are quantitatively very similar and qualitatively unchanged.

[13] The 2.5 percentile is equivalent to an annual income of €800 in 2012 (alternative choices do not affect the results significantly). These sample restrictions are equivalent to the sample restrictions in Guvenen et al. (2014). Removing workers who have worked less than 3 weeks (0.25% of the full sample) corresponds to that definition, and possibly reduces the measurement error in weekly earnings. The removed workers are likely to be individuals receiving payments relative to work done in previous years, accounted in the current year for accounting reasons. 60% of these workers are actually recorded with 0 week of work, leaving only 0.1% of the full sample working 1 or 2 weeks.

[14] Full-time and part-time employment is specified at the record level. We define main work contract to be the one associated with the highest share of overall earnings in a year.

[15] See Online Appendix A for summary statistics for selected years. Year by year statistics are available from the authors upon request. In the appendix Table A.2, we also report more statistics for all the years in the sample.

[16] The share of workers in the sample who work for 52 weeks is 77.2% if we only include people with more than two weeks worked, and 77.3% if we also restrict the sample to include only non-top-coded observations.

[17] We obtained these statistics pooling all available CPS interviews from 1962 to 2016.

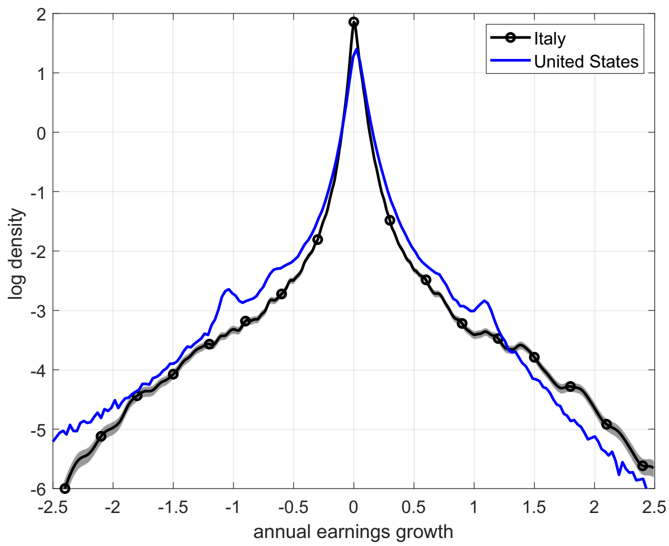

Fig. 2. Log-density of annual earnings growth, Italy and US 1996.
Source: Italy – INPS (as described in the main text). United States – Guvenen et al. (2015). All samples are restricted to males 25–60 years old.

percentile. Annual earnings growth (in log points) is more dispersed than weekly earnings growth: the standard deviation of Δy is more than twice the standard deviation of Δw (0.40 and 0.17 respectively).

One concern is that measurement errors in employment time will carry over to weekly earnings (also known as "division bias"), and will generate a mechanical negative correlation. However, weeks of work are measured accurately in the INPS data, and we excluded observations with particularly low number of weeks of work, which may lead to large measurement errors. Also, if there were substantial measurement errors in employment time, it would have increased the variance of changes in weekly earnings. We measure the standard deviations of weekly earnings at 17%, and it is thus not likely to be a threat to this measurement exercise.

Since our analysis emphasizes the role of changes in employment time in explaining earnings growth, we present some additional descriptive statistics of the employment time dynamics in Online Appendix A. In particular, Table A.6 reports the autocovariance structure of employment time and employment time changes. We also report the decomposition of the variance of changes in employment into within worker and between workers components. The former component accounts for 86% of the variance, and the latter for 14%.

Finally, Fig. 2 compares the log-density of earnings growth, Δy_{it}, in Italy and the United States, for the year 1996.[18] Both distributions display similar heavy left and right tails. The Italian distribution is less dispersed and (in 1996) more symmetric. These similarities imply that the forces driving the particular shape of the distribution are not unique to the United States, and support the external validity of the analysis using the Italian data.

5. Evidence on the role of employment time

In this section we decompose individual level annual earnings growth into changes in employment time and changes in weekly earnings using the INPS data. We divide our analysis in this section into two parts: evidence from the cross-sectional distribution of earnings growth, and evidence on cyclical patterns of earnings growth.

5.1. Evidence from the cross section

Fig. 1 in the introduction displays the log densities of earnings growth and weekly earnings for the year 2002. The graph shows that the distribution of weekly earnings growth has thinner tails than that of annual earnings growth, and that the distribution of changes in employment time matches the tails of the distribution of annual earnings growth. Here we expand this analysis and provide additional non-parametric evidence for the role of employment time.

Our preferred way to explore the role of the employment time in the cross-sectional distribution, is to look at the distribution of earnings growth for workers who have a certain amount of employment time in a given year. In particular, we look at the distribution of earnings growth for three groups, and compare them to the full sample for the reference year ($t = 2002$):

(A) Workers employed for 52 weeks in year $t-1$ (Fig. 3 panel (a))
(B) Workers employed for 52 weeks in year t (Fig. 3 panel (b))
(C) Workers employed for 52 weeks in both year $t-1$ and year t (Fig. 3 panel (c))

Panel (a) in Fig. 3 shows the distribution of earnings growth for group A. Since workers in this group were employed for 52 weeks in year $t-1$, their changes in employment time cannot be positive. That is, positive growth in earnings for this group can only be generated by increases in their weekly earnings. The right tail of the distribution of group A is considerably thinner than that of the full sample, while retaining the shape and magnitude of the left tail. This implies that the right tail of the distribution is generated almost entirely by workers who are in the complimentary group – those that were employed for less than 52 weeks in year $t-1$. Fig. 3 panel (b) is the mirror image of panel (a). Members of group B, who were employed for 52 weeks in year t, are unlikely to experience a drop in earnings larger than the 10th percentile of the full sample distribution (dashed line). The bottom panel of Fig. 3 completes the picture. Workers in group C, who were employed for the full 104 weeks of 2001 and 2002, are unlikely to have experienced large positive *or* negative earnings growth during this period.

To quantify this visual evidence we compare the probability of experiencing large or small earnings growth in each of the groups, and report results in Online Appendix Table A.4. We show that workers in Group A, i.e. those with 52 weeks of employment in year $t-1$ are approximately 19 times less likely to experience large earnings growth than those with some non-employment spell in the same year. Similarly, those *not* employed for 52 weeks in year t (the complimentary group for B) are about 19 times more likely to experience earnings growth lower than the 10th percentile (see Table A.4 column (5)).[19]

Since the analysis so far is completely non-parametric, and is invariant to any monotone transformation of earnings growth (90th percentile remains the 90th percentile), this is perhaps the most convincing evidence that changes in employment time are responsible for the tails of the distribution. However, letting log changes be the metric for the size of changes, as is typically done in the literature, we can also quantify the relative magnitudes by variance decomposition. Variance decomposition for the year 2002 shows a large role for employment time:

$$\underset{0.158}{\mathrm{var}(\Delta y)} = \underset{0.150}{\mathrm{var}(\Delta x)} + \underset{-0.014}{2\mathrm{cov}(\Delta x, \Delta w)} + \underset{0.023}{\mathrm{var}(\Delta w)} \tag{8}$$

where the number below each term is its estimated value.

[18] The density data for the United States is taken from Guvenen et al. (2015).

[19] See Online Appendix Section A.4 for details on these calculations.

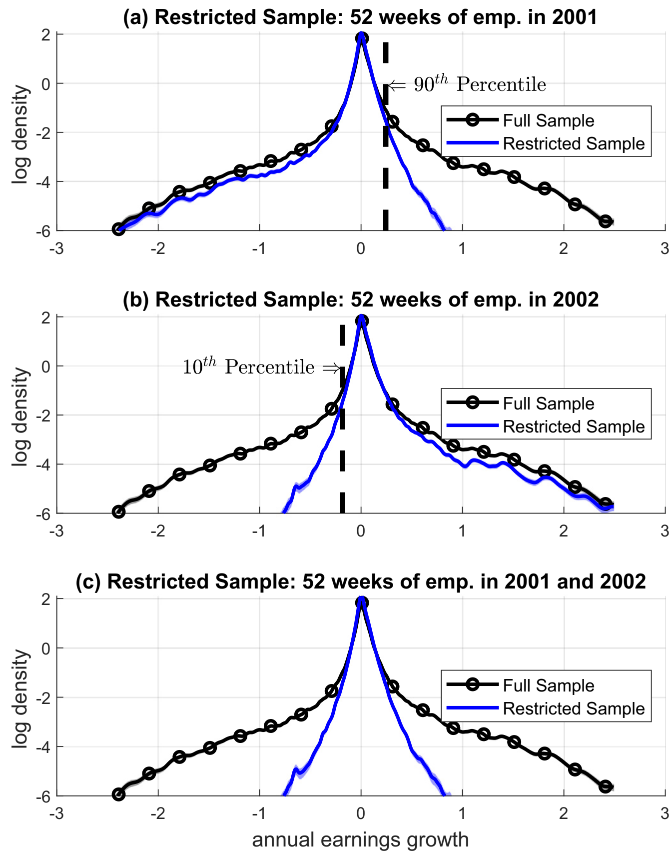

Fig. 3. Log-densities of earnings growth in Italy 2002, Group A, B, C.
Notes: Black line - full-sample distribution, 25–60 years old males in all panels. Panel (a): blue line, sample restricted to workers who have worked for 52 weeks in year 2001 (group A). Panel (b): blue line, sample restricted to workers who have worked for 52 weeks in year 2002 (group B). Panel (c): blue line, sample restricted to workers who have worked for 52 weeks in both years 2001 *and* 2002 (group C).
*Source:*INPS data provided by MLPS.

In 2002, the variance of employment time was 0.158, implying a standard deviation of 40%, compared to a variance of 0.022, or a standard deviation of 15% for weekly earnings.

Table 2 provides values of decomposition (6) for all the years in the sample. Two results for Eq. (8) and Table 2 are notable. First, the variance of changes in employment time is greater than four-fifths of the variance of annual earnings growth in all years in the sample. Second, the cross-term, twice the covariance of the two components, is small and negative throughout most of sample period. Together these two results imply that the variance of log earnings growth is mostly driven by changes in employment time rather than changes in weekly earnings.

A similar decomposition of the third moment (See Table A.3) reveals a similar pattern, with the sign and magnitude of the third moment of employment time growth tracking closely those of earnings growth. Moreover, the magnitude of the third moment of weekly earnings growth is mostly one order of magnitude smaller than that of annual earnings growth and employment time growth, with the only exceptions of year 1998 and 1999, when the third moment of annual earnings growth is notably close to zero and (exceptionally) positive.

5.2. Time series evidence

Since employment time has a dominant role in shaping the tails of the cross-sectional annual earnings growth distribution, we are now interested in measuring its contribution to the cyclicality of annual earnings. We start by confirming that recessions affect the shape of the distribution of annual earnings growth. Fig. 4 shows the cross-sectional distribution of annual earnings growth in Italy for 2002 (expansion year) and 2009 (recession year). In recessions, the right tail of the distribution (large increases) goes down while the left tail (large decreases) goes up. Since the central part of the distribution remains stable, the overall skewness of the distribution becomes more negative. The similarity of the impact of recessions in Italy to that documented by Guvenen et al. (2014) for the US suggests a common mechanism.

We continue by presenting the decompositions of the first three central moments of earnings growth over the sample period in Fig. 5.

Panel (a) shows the decomposition of the mean earnings growth. The mean change in employment time remains negative throughout most of the sample period. It reaches a low of −0.07 in 2009, in the midst of the Great Recession. The mean change in weekly earnings is on average a positive 0.02 during expansions. It drops during recessions, reaching a low of −0.02 in 2012. The dynamic behavior of the means suggests that cyclicality of mean earnings growth reflects cyclical properties of both employment time and weekly earnings.

The variance of earnings growth is decomposed in panel (b), visually repeating the variance decomposition reported in Table 2. The variance of earnings growth and the variance of changes employment time follow a long-term increasing trend. This trend is particularly evident after the beginning of the Great Recession in 2007. There appears to be some amount of countercyclicality in the variance earnings growth and changes in employment time, but both seem to be less pronounced than the long-term trend. The time series of the variance of changes in weekly earnings is relatively flat, with a period of higher variance in the late 1990s early 2000s. The variance of changes in employment time is visually associated with the variance of earnings growth.

Panel (c) presents the decomposition of the third central moment. Controlling for the mean and variance, the third central moment captures the asymmetry of the distribution. The third moment of annual earnings growth and of changes in employment time follow the same path, both in magnitude and in pattern. They are both clearly procyclical and drop to a negative −0.1 in 2009, during the global economic slowdown. The third moment of changes in weekly earnings is relatively flat and close to zero. We interpret these results as employment time being the primary source for the observed cyclical asymmetry of annual earnings growth.

The visual evidence is confirmed with a statistical test using the constructed time series of moments. Table 3 presents the statistical test. We define the cyclicality of a cross-sectional moment as its contemporaneous correlation with GDP growth. The third moment of annual earnings growth in Italy is highly procyclical (see column 1), which is similar to the findings of Guvenen et al. (2014) for the United States. However, when controlling for the third moment of changes in employment time in the regression, the correlation disappears (see column 2).[20] This result offers additional suggestive evidence that employment time is the source of the cyclical skewness of earnings growth. Online Appendix D provides additional statistical tests and robustness exercises that confirm this result.

5.3. Earnings growth at longer horizons

We also evaluate the decomposition of earnings growth at longer horizon. Guvenen et al. (2014) argue that comparing the distribution of five-year earnings growth to the distribution of one-year earnings

[20] This statistical test treats the sample moments of earnings growth and changes in employment time as population moments, and does not adjust for measurement errors. Since the sample includes more than 300 thousand observations for each sample moment this should not have any significant impact on the results.

Table 2
Variance decomposition of earnings growth, Italy 1986–2012.

Year	$m_2(\Delta y)$		$m_2(\Delta x)$		$m_2(\Delta w)$		$2m_{1,1}(\Delta x, \Delta w)$	
	(1)		(2)		(3)		(4)	
1986	0.126	(0.001)	0.113	(0.001)	0.024	(0.000)	−0.012	(0.000)
1987	0.131	(0.001)	0.118	(0.001)	0.024	(0.000)	−0.012	(0.000)
1988	0.131	(0.001)	0.118	(0.001)	0.024	(0.000)	−0.011	(0.000)
1989	0.126	(0.001)	0.116	(0.001)	0.024	(0.000)	−0.014	(0.000)
1990	0.125	(0.001)	0.116	(0.001)	0.022	(0.000)	−0.014	(0.000)
1991	0.137	(0.001)	0.128	(0.001)	0.023	(0.000)	−0.016	(0.000)
1992	0.148	(0.001)	0.143	(0.001)	0.023	(0.000)	−0.019	(0.000)
1993	0.147	(0.001)	0.131	(0.001)	0.025	(0.000)	−0.009	(0.000)
1994	0.149	(0.001)	0.139	(0.001)	0.025	(0.000)	−0.015	(0.000)
1995	0.140	(0.001)	0.131	(0.001)	0.022	(0.000)	−0.013	(0.000)
1996	0.141	(0.001)	0.133	(0.001)	0.021	(0.000)	−0.014	(0.000)
1997	0.141	(0.001)	0.135	(0.001)	0.021	(0.000)	−0.016	(0.000)
1998	0.156	(0.001)	0.145	(0.001)	0.025	(0.000)	−0.014	(0.000)
1999	0.156	(0.001)	0.145	(0.001)	0.024	(0.000)	−0.012	(0.000)
2000	0.166	(0.001)	0.153	(0.001)	0.025	(0.000)	−0.013	(0.000)
2001	0.167	(0.001)	0.156	(0.001)	0.025	(0.000)	−0.014	(0.000)
2002	0.158	(0.001)	0.150	(0.001)	0.023	(0.000)	−0.014	(0.000)
2003	0.164	(0.001)	0.159	(0.001)	0.023	(0.000)	−0.017	(0.000)
2004	0.156	(0.001)	0.146	(0.001)	0.022	(0.000)	−0.012	(0.000)
2005	0.167	(0.001)	0.148	(0.001)	0.024	(0.000)	−0.006	(0.000)
2006	0.165	(0.001)	0.147	(0.001)	0.022	(0.000)	−0.004	(0.000)
2007	0.158	(0.001)	0.140	(0.001)	0.023	(0.000)	−0.004	(0.000)
2008	0.168	(0.001)	0.143	(0.001)	0.023	(0.000)	0.001	(0.000)
2009	0.203	(0.001)	0.168	(0.001)	0.028	(0.000)	0.007	(0.000)
2010	0.193	(0.001)	0.159	(0.001)	0.029	(0.000)	0.005	(0.000)
2011	0.195	(0.001)	0.160	(0.001)	0.027	(0.000)	0.007	(0.000)
2012	0.197	(0.001)	0.167	(0.001)	0.027	(0.000)	0.004	(0.000)

Notes: The variance of earnings growth (1), is decomposed into the variance of changes in employment time (2), the variance of changes in weekly earnings (3) and a cross-term (4). Standard deviations of each component are reported in parentheses. Sample includes all 25–60 years old males who appear in the data for two consecutive years. *Source:* INPS data provided by MLPS.

growth is informative on the persistence of earnings shocks. Therefore, we supplement the one-year earnings growth decomposition with a decomposition of the first three moments of the five-year earnings growth distribution.[21]

Fig. 6 presents the time series of the mean, variance, and third moment in a way that is comparable to Fig. 5. The mean earnings growth at longer horizons, which is presented in panel (a), moves together with the mean changes in weekly earnings. In contrast to the one-year results, the mean of changes in employment time does not contribute much to the five-year earnings growth. The variance decomposition in panel (b) reveals that changes in weekly earnings account for a larger share of the variance of five-year earnings growth. Changes in employment time still generate most of the variance and are responsible for the secular upward trend and mild countercyclicality. But the variance of five-year changes in employment time is roughly the same as the variance of one-year changes in employment time, while the variance of five-year weekly earnings growth (0.05) is roughly double the variance of one-year weekly earnings growth (0.025). Panel (c) shows that the role of employment time in generating the asymmetric response to recession is as important at the longer horizon as it was for the one-year earnings growth. The third moment of five-year changes in employment closely follows the third moment of earnings growth, both in magnitude and in pattern, and is visibly lower after the beginning of the slump in mid-2007. As in the one-year case, the third moment of changes in weekly earnings is flat and close to zero.

These results have implications for the persistence of the changes in earnings growth and its components. Under a simple permanent/transitory framework the increase in the variance of weekly earnings growth with the time horizon suggests that shocks to weekly earnings have a considerable permanent component while the absence of an increase in the variance of changes in employment time suggests that most of the variation in employment time is transitory.[22] Since the distribution of weekly earnings growth appears to be symmetric and acyclical, this suggests that the cyclicality in annual earnings growth is mostly coming from transitory shocks that affect employment, rather than permanent shocks to earnings.

5.4. Job stayers and job switchers

Economists have documented differences in the mean earnings growth between workers who switch jobs and workers that stay with the same employer (Topel and Ward, 1992; Bagger et al., 2014; Low and Pistaferri, 2015). Recently Guvenen et al. (2015) show significant differences in the second, third and fourth moments of earnings growth between job stayers and job switchers. Specifically, job stayers face a standard deviation of earnings growth that is approximately half that of job switchers and their group's earnings distribution has a less negative third central moment. In this subsection we document that most of these differences disappear once

[21] Following Guvenen et al. (2014), we define the n-years earnings growth as $\Delta_n y_t = y_t - y_{t-n}$ where y_t is the logged annual earnings at year t. Similarly, the changes in employment time and weekly earnings are defined as $\Delta_n x_t = x_t - x_{t-n}$ and $\Delta_n w_t = w_t - w_{t-n}$, taking the log differences between the annual values.

[22] Let $z_t = \eta_t + e_t$ be a permanent/transitory process, in which $\eta_t = \eta_{t-1} + u_t$ is the permanent component and e_t and u_t are independent shocks with mean zero and variance σ_e^2 and σ_u^2. The variance of the one-year growth is $\mathrm{Var}[\Delta z_t] = \sigma_u^2 + 2\sigma_e^2$. The variance of the five-year growth is $\mathrm{Var}[\Delta_5 z_t] = 5\sigma_u^2 + 2\sigma_e^2$. Therefore the difference $\mathrm{Var}[\Delta_5 z_t] - \mathrm{Var}[\Delta z_t]$ is four times the variance of the permanent component. Using this result and the variance of one- and five-year weekly earnings growth suggest a permanent component with standard deviations of 7% and transitory component with standard deviations of 10%. Since the variance of one- and five-year changes in employment time is roughly the same, this suggests a negligible permanent component.

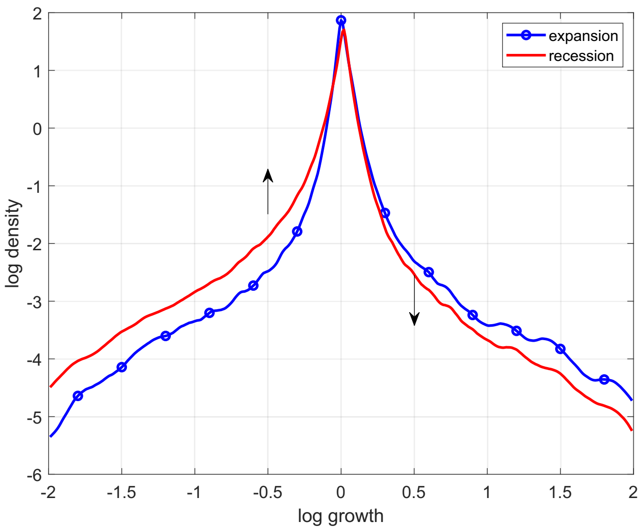

Fig. 4. Annual earnings growth in expansion and recession, Italy 2002 and 2009.
Notes: Annual earnings growth is measured as the difference in logs. The sample includes 25–60 years old males in Italy.
*Source:*INPS data provided by MLPS.

we restrict attention to workers who did not experience large drops or increases in employment time, and therefore the differences are related to spells of non-employment.

Table 3
Statistical test of source of cyclicality.

Dependent variable	Third moment of annual earnings growth (1)	(2)
GDP growth	**0.659**	−0.120
	(0.124)	(0.086)
Third moment of changes in employment time		**1.020**
		(0.091)
R-squared	0.43	0.86
Observations	27	27

Notes: The dependent variable is the third moment of annual earnings growth (columns 1 and 2). All time series are detrended with a linear trend and standardized. Regressors include contemporaneous GDP growth (log difference of real GDP from ISTAT) and the third moment of the distribution of changes in employment time. Data covers the years 1986–2012. Bold text indicates significance at the 0.95 level. Standard errors are computed using Newey-West estimator with two lags.

Guvenen et al. (2015) define a worker to be a "job stayer" if a given employer provides the largest share of his annual earnings (out of all his job relationships in a year) in years $t-1$ through $t+2$, and if the same employer provides at least 80% of his total annual earnings in years t and t + 1. They define as "job switchers" those workers who are not job stayers. We adopt a similar definition, and consider as switchers all workers whose main employer (the employer that provides most of their earnings) changes between year $t-1$ and year t.

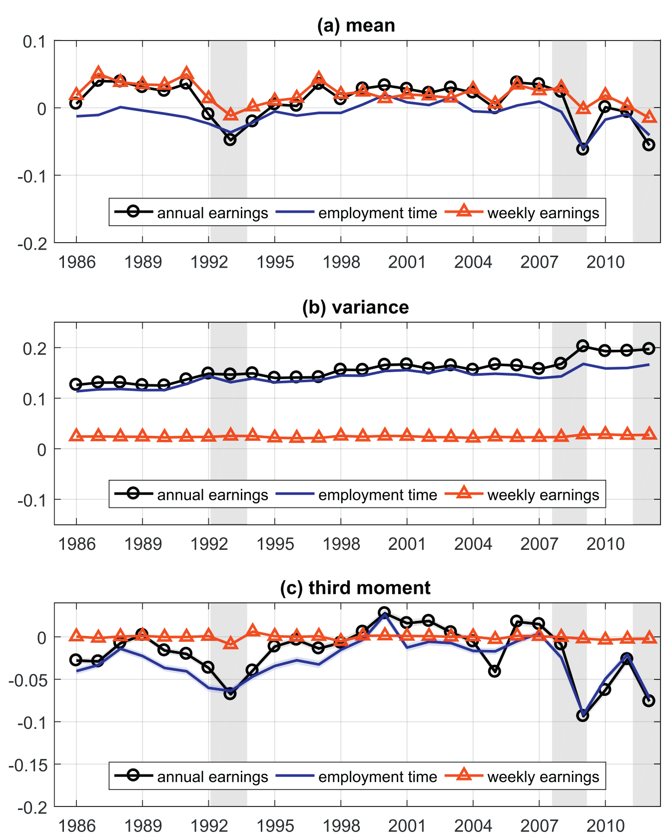

Fig. 5. Moment decomposition of annual earnings growth, Italy 1986–2012.
Notes: Panels (a) through (c) present the time series of moments of the cross-sectional distributions. Panel (a) presents decomposition of the mean, panel (b) presents decomposition of the variance and panel (c) presents decomposition of the third central moment. In each panel, there are three lines: annual earnings (black), employment time (blue) and weekly earnings (red). All variables are measured as difference in logs. The sample includes 25–60 years old males in Italy. Shaded areas represent recessions.
*Source:*INPS data provided by MLPS.

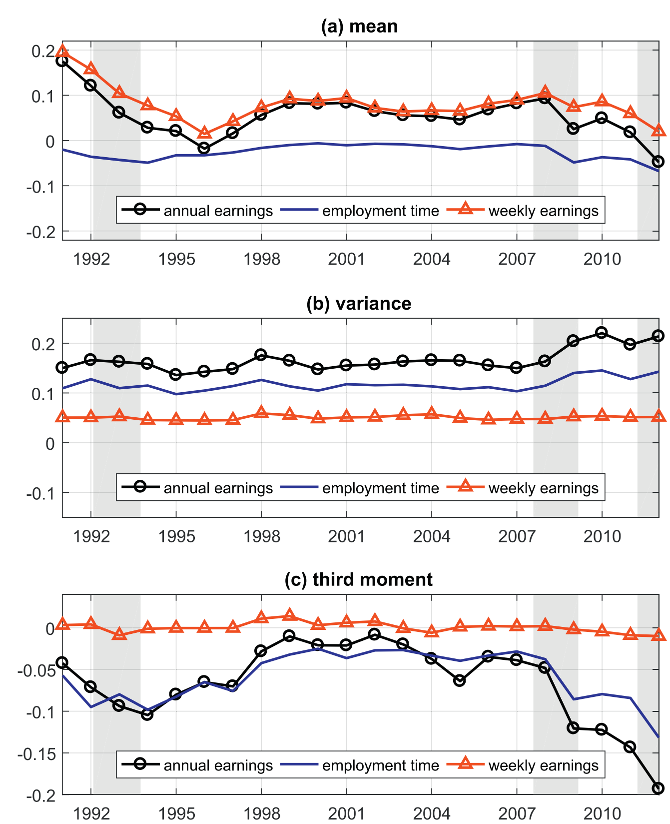

Fig. 6. Moment decomposition of five-year earnings growth.
Notes: Panels (a) through (c) present the time series of moments of the cross-sectional distributions of five-year changes in annual earnings (black, round marker), employment time (blue, smooth), and weekly earnings (red, triangles). Panel (a) presents decomposition of the mean, panel (b) presents decomposition of the variance and panel (c) presents decomposition of the third central moment. All variables are measured as difference in logs. The sample includes 25–60 years old males in Italy. Shaded areas represent recessions.
*Source:*INPS data provided by MLPS.

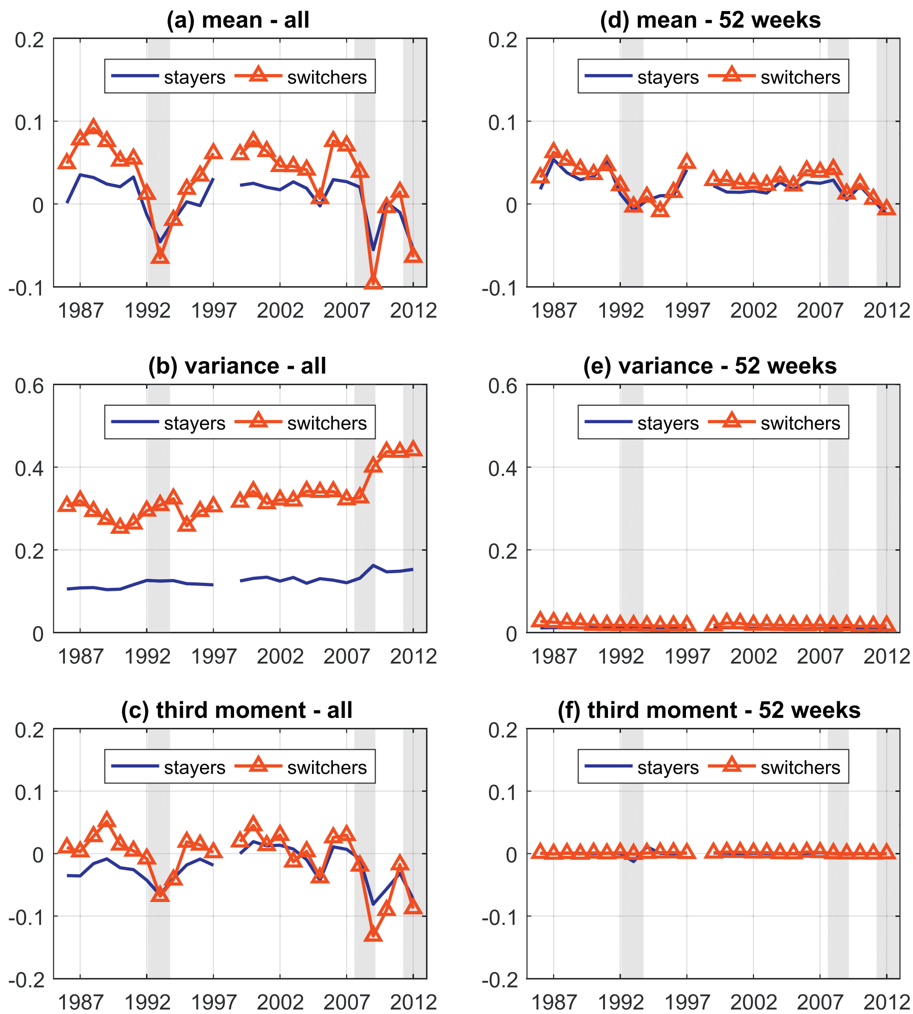

Fig. 7. Job stayers and job switchers, Italy 1986–2012.
Notes: Panels (a), (b), and (c) present the time series of mean, variance and third central moment of the cross-sectional distributions of annual earnings growth for "stayers" and "switchers" over the full sample. Panels (d), (e), and (f) present the time series of the same statistics, only for workers employed 52 weeks both in year $t-1$ and t. The sample includes 25–60 years old males from Italy and covers the period 1986 to 2012. Shaded areas represent recessions. Year 1998 is omitted as in that year the INPS recording system was reformed and it is not possible to precisely define stayers and switchers.
*Source:*INPS data provided by MLPS.

Fig. 7 shows the first, second and third moments of the earnings growth distribution for stayers and switchers. Panels (a), (b), and (c) show the time series of the mean, variance, and third moment of job stayers and job switchers. These results are similar to those in Guvenen et al. (2015): the variance of earnings growth for switchers is twice as large as the that of stayers and the third central moment of switchers is more negative than of stayers. Panels (d), (e), and (f) restrict attention to workers who were employed for 52 weeks in both years t and $t-1$. Within this group the differences narrow significantly, becoming quantitatively negligible. Hence, the difference between the groups observed in the left panels is driven by the composition of workers within the two groups: switchers experience more spells of non-employment and therefore the variance of their earnings growth is higher.[23]

[23] For completeness, Figs. A.8 and A.9 in Online Appendix E also replicate Figs. 1 and 5 including only stayers. These figures are very similar to those in the main text, showing that the relationship between employment time and earnings growth both in the cross-section and in the time series is not driven by workers switching jobs.

Another interesting observation is that during recessions switchers have a lower mean earnings growth, while in expansion they have a higher earnings growth. When looking at continuously employed workers, this difference almost disappears. This difference points to cyclical changes in the composition of job switchers who are laid off, who quit and find a new job, and who are "poached" by other employers while on the job that was highlighted in the literature (see for example Postel-Vinay and Robin, 2002; Holzheu, 2018).

5.5. Additionl analysis: tax and transfers, and the life cycle

5.5.1. Tax and transfers

In a recent paper, Blundell et al. (2015) show that the Norwegian tax system plays a crucial role in reducing the magnitude and persistence of income shocks, especially among low educated and young Norwegian. To reach this conclusion they use a rich collection of administrative data that allows them to compute each Norwegian tax payer's labor income before and after tax and transfers. To assess the impact of tax and transfers on the distribution of earnings growth, in Online Appendix B we perform a similar exercise on the Italian data. The tax system reduces systematically the variance of income growth. Especially following 2007, the combination of unemployment benefit transfers and income taxes systematically reduces the aggregate variance of earnings growth by 25%, from 0.2 to 0.15. Similarly, the third central moment of earnings growth increases following 2007, reflecting an insuring effect of the system during economic downturns. Consistent with our findings that most large negative shocks correspond to large declines in employment time (often mirrored by unemployment spells), a comparison of the earnings growth distribution pre and post tax and benefits reveals that the unemployment benefits are effective in reducing the consequent drop in earnings, making the distribution less left-skewed.

5.5.2. Earnings growth over the life cycle

In our analysis so far we have described the distribution of earnings growth by pooling together workers of different age groups. However, individual earnings follow a predictable age profile – rising when young and flattening when old. This predictable component may inflate the dispersion of the distribution, and contribute to the cyclical patterns if there is some systematic differences in the response to aggregate shocks across age groups. As a robustness check (detailed in Online Appendix C), we remove a time-age fixed effect from annual earnings, employment time, and weekly earnings, at the worker level, and find that the age profile has little to no effect on the results of this paper.

5.6. Summary: evidence from Italy

In this section we present direct evidence indicating that changes in employment time drive the cyclical properties of the distribution of annual earnings growth. First, changes in employment time generate the tails of the earnings growth distribution. Second, the cyclical asymmetry of the earnings growth distribution, measured by its third moment, is generated by cyclical changes in the distribution of changes in employment time. In contrast, we find that changes in weekly earnings have a minor role in driving the tails of annual earnings growth, and that their distribution around the mean is acyclical. We also find that the distribution of changes in weekly earnings displays only small deviations from symmetry. Changes in weekly earnings are also much more persistent than changes in employment time, and therefore may have a larger impact on individual permanent income.

6. Implications: a proposed earnings process

The decomposition of earnings growth into changes in employment time and changes in weekly earnings establishes a set of new stylized facts. Namely, it establishes that changes in employment time are responsible for the tails of the earnings growth distribution and are the source of cyclical asymmetry, while the distribution of changes in weekly earnings around its mean is symmetric and acyclical. This suggests the need to carefully model the employment margin separately from the wage process.

We therefore propose a parsimonious earnings process that is consistent with these stylized facts. Online Appendix F provides a detailed description of the process and its numerical implementation. The key idea is that annual earnings are the product of two independent processes: an employment process and a wage process. The employment process is driven by random transitions between employment and unemployment states at a monthly frequency. The probability of transition into and out of unemployment is affected by the aggregate state of the economy. This type of process is consistent with the Diamond-Mortensen-Pissarides framework of labor search and matching, which is a standard macroeconomic view of labor markets.[24]

The wage process is subject to independent permanent and transitory shocks that generate a symmetric wage growth distribution. This is a common assumption in the literature studying the idiosyncratic earnings process following MaCurdy (1982). Here, we add to this literature by modeling transitory shocks as exponential random variables instead of normal random variables. Non-normal transitory shocks help explain the small deviations from normality in the earnings growth distribution that can be detected in large samples.

The proposed earnings process demonstrates that the combination of two simple processes, with few parameters, captures the cyclical patterns in the data without assuming complex distributions for earnings shocks. Instead of relying on persistent asymmetric shocks, the procyclical skewness of earnings growth is captured through the employment process and its interaction with transitions between expansion and recession aggregate states: In the beginning of recessions, the separation rate (probability of transition from employment to unemployment) increases and the job finding rate (probability of transition from unemployment to employment) declines. The higher separation rate leads more workers to transition into unemployment, and the lower job finding rate generates longer unemployment spells. More employed workers experience a large drop in employment time in the beginning of recessions, and fewer unemployed workers experience an increase in employment. Meanwhile, changes in employment time remain effectively zero for the majority of workers who stay employed. Therefore, the distribution of changes in employment time becomes asymmetric and negatively skewed at the beginning of recessions.

As the share of unemployed workers increases, so does the share of workers who transition out of unemployment and increase their employment time. Thus the distribution of the changes in employment time becomes more symmetric over time. In the long run, the flows into and out of unemployment equalize and the distribution of earnings growth becomes symmetric again. Recoveries have the opposite effect and generate positive skewness in the earnings growth distribution.

Figs. 8 and 9 show how the proposed process captures the new stylized facts. Fig. 8 shows a decomposition of simulated annual earnings growth at an expansion year. The simulated distribution captures the main visual features of the cross section in the data, including heavy tails and many near-zero changes in employment time, and non-Gaussian tails in weekly earnings growth. Fig. 9 compares the distribution of earnings growth in a recession year and in recovery. In the recession year (red) the left tail of the distribution of annual earnings growth is higher and the right tail lower due to high flow of workers into unemployment. Similarly, in the recovery year (blue) the distribution of earnings growth is positively skewed due to high flow of workers from unemployment to employment.

Since this process generates procyclical skewness in earnings growth through the employment process, we can test the mechanism using data on transitions between labor market states. Unfortunately, direct observations of high frequency transition rates between labor market states in Italy are not available. However, this type of data is available for the US through the Current Population Survey.

We therefore evaluate the ability of the employment process to generate the procyclical skewness in using US data. Online Appendix G provides a full description of this exercise. We first construct measures of the share of workers in each one of three

[24] Other recent papers have explored a similar idea (Hubmer, 2018; Holzheu, 2018) aiming at reproducing the cross-sectional properties of the earnings growth distribution. Our model aims to capture both the cross-sectional properties and the cyclical properties of the distribution.

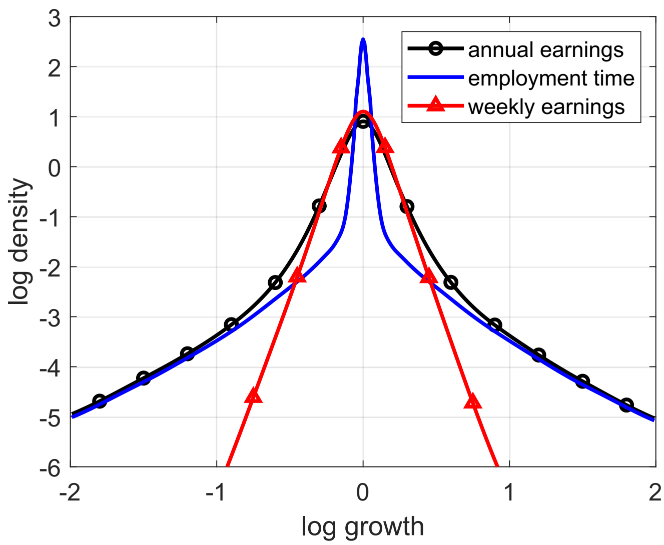

Fig. 8. Numerical example – Decomposition of earnings growth.
Notes: Cross-sectional distribution of one-year log growth of annual earnings, employment time, and weekly earnings based on simulations of the earnings process described in the main text.

labor market states (employed, unemployed, and not-in-the-labor-force) and the transition probabilities between them, at a monthly frequency for the sample period 1976–2015. Then we feed these shares and transition probabilities into an extended employment process and recover the implied distribution of annual changes in employment time.

This exercise suggests that the proposed employment process captures the cyclical asymmetry of the earnings growth distribution. Fig. 10 presents the third moment of changes in employment time implied by the proposed process (blue line), and the actual third moment of earnings growth calculated by Guvenen et al. (2014) (black line, round marker). The implied third moments of changes in employment time match the timing and the magnitude of the third moment of earnings growth. We view this as additional suggestive evidence in support of modeling the employment process separately from the wage process.

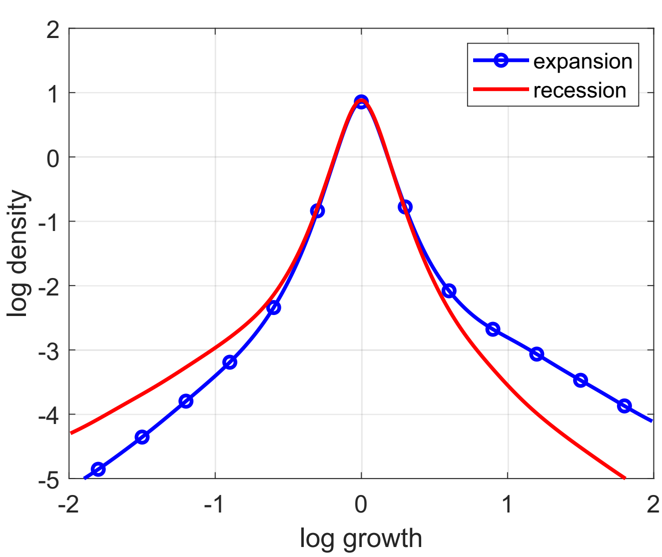

Fig. 9. Simulation – Recession and expansion.
Notes: Cross-sectional distribution of annual earnings growth in recession and expansion. Data simulated from the numerical example.

7. Conclusion

In this paper we study the evolution of the earnings growth distribution over time using administrative data from Italy. We decompose earnings growth into changes in employment time (the number of weeks of work within a year) and weekly earnings. We show that (i) employment changes drive the tails of the earnings growth distribution and (ii) fluctuations in employment explain the co-movement of moments of earnings growth and business cycles. In particular, changes in employment generate the procyclical skewness of earnings growth.

This set of results suggests a new interpretation for the procyclical skewness in the earnings growth distribution found by Guvenen et al. (2014). In particular, it suggests that the aggregate factors that affect the number of workers who lose their jobs and the duration of unemployment spells are also responsible for the observed cyclicality in earnings growth.

In many economic applications the source of earnings growth is important. We find that the distributions of changes in employment time and weekly earnings differ in shape, cyclicality, persistence, and are affected differently by policy. Changes in employment time also explain most of the differences between switchers and stayers. These findings highlight the need to carefully model the employment margin separately from the wage margin, especially in studies of consumption and wealth or public policy.

While the focus of this paper is on decomposing "earnings growth" and not "earnings risk", we view this analysis as a step towards better understanding the latter. A proper study of the risk component of labor earnings requires a structural model in which changes in earnings can be the result of endogenous choices (for example see Low et al., 2010; Lise, 2012). Further research should also analyze the interaction between the employment and the wage process, found to be important in other studies (e.g. Davis and von Wachter, 2011; Saporta-Eksten, 2014).

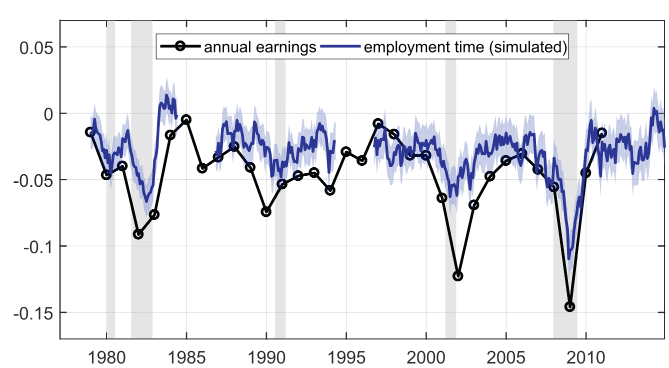

Fig. 10. Third moment of annual earnings growth & changes in employment time (Model).
Notes: The third central moment of annual earnings growth (black) and the distribution of changes in employment time implied by the extended employment process (blue, see Online Appendix G for details on construction). Shaded area shows a 0.95 confidence interval (computed using 500 bootstrap replications for employment time).
*Source:*Guvenen et al. (2014) and IPUMS-CPS (Flood et al., 2015).

Appendices. Supplementary data

Supplementary data to this article can be found online at https://doi.org/10.1016/j.jpubeco.2018.09.009.

References

Abowd, J., Card, D., 1989. On the covariance structure of earnings and hours changes. Econometrica 57 (2), 411–445.

Arellano, M., Blundell, R., Bonhomme, S., 2017. Earnings and consumption dynamics: a nonlinear panel data framework. Econometrica 85 (3), 693–734.

Bagger, J., Fontaine, F., Postel-Vinay, F., Robin, J.-M., 2014. Tenure, experience, human capital, and wages: a tractable equilibrium search model of wage dynamics. Am. Econ. Rev. 104 (6), 1551–1596.

Berger, D., Dew-Becker, I., Schmidt, L.D., Takahashi, Y., 2016. Layoff Risk, the Welfare Cost of Business Cycles, and Monetary Policy. Available at SSRN 2659941.

Blundell, R., Bozio, A., Laroque, G., 2011. Labor supply and the extensive margin. Am. Econ. Rev. 101 (3), 482–486.

Blundell, R., Graber, M., Mogstad, M., 2015. Labor income dynamics and the insurance from taxes, transfers, and the family. J. Public Econ. 127, 58–73. The Nordic Model.

Busch, C., Domeij, D., Guvenen, F., Madera, R., 2018. Asymmetric business-cycle risk and social insurance. Working Paper 24569. National Bureau of Economic Research, May.

Chetty, R., 2012. Bounds on elasticities with optimization frictions: a synthesis of micro and macro evidence on labor supply. Econometrica 80 (3), 969–1018.

Constantinides, G.M., Duffie, D., 1996. Asset pricing with heterogeneous consumers. J. Polit. Econ. 219–240.

Constantinides, G.M., Ghosh, A., 2014. Asset pricing with countercyclical household consumption risk. Working Paper 20110. National Bureau of Economic Research, May.

Davis, S.J., von Wachter, T.M., 2011. Recessions and the cost of job loss. Technical Report. National Bureau of Economic Research.

DeNardi, M., Fella, G., Pardo, G.P., 2016. The implications of richer earnings dynamics for consumption, wealth, and welfare. Working Paper 21917. National Bureau of Economic Research, January.

Flood, S., King, M., Ruggles, S., Warren, J.R., 2015. Current Population Survey: Version 4.0. [Machine-readable database].

French, E., Song, J., 2014. The effect of disability insurance receipt on labor supply. Am. Econ. J. Econ. Pol. 6 (2), 291–337.

Guvenen, F., 2009. An empirical investigation of labor income processes. Rev. Econ. Dyn. 12 (1), 58–79.

Guvenen, F., Karahan, F., Ozkan, S., Song, J., 2015. What do data on millions of U.S. workers reveal about life-cycle earnings risk? NBER Working Papers 20913. National Bureau of Economic Research, Inc., January.

Guvenen, F., Ozkan, S., Song, J., 2014. The nature of countercyclical income risk. J. Polit. Econ. 122 (3), 621–660.

Holzheu, K., 2018. Separations on the Job Ladder. Job Market Paper.

Hubmer, J., 2018. The job ladder and its implications for earnings risk. Rev. Econ. Dyn. 29, 172–194.

Kopczuk, W., Saez, E., Song, J., 2010. Earnings inequality and mobility in the United States: evidence from social security data since 1937. Q. J. Econ. 125 (1), 91–128.

Krebs, T., 2007. Job displacement risk and the cost of business cycles. Am. Econ. Rev. 97 (3), 664–686. June.

LaLonde, R.J., 2007. The Case for Wage Insurance.Number 30. Council on Foreign Relations Press.

Lise, J., 2012. On-the-job search and precautionary savings. Rev. Econ. Stud. 80 (3), 1086–1113.

Low, H., Meghir, C., Pistaferri, L., 2010. Wage risk and employment risk over the life cycle. Am. Econ. Rev. 100 (4), 1432–1467.

Low, H., Pistaferri, L., 2015. Disability insurance and the dynamics of the incentive insurance trade-off. Am. Econ. Rev. 105 (10), 2986–3029.

MaCurdy, T.E., 1982. The use of time series processes to model the error structure of earnings in a longitudinal data analysis. J. Econ. 18 (1), 83–114.

McKay, A., 2017. Time-varying idiosyncratic risk and aggregate consumption dynamics. J. Monet. Econ. 88, 1–14.

McKay, A., Reis, R., 2016. Optimal automatic stabilizers. Working Paper 22359. National Bureau of Economic Research, June.

Meghir, C., Pistaferri, L., 2004. Income variance dynamics and heterogeneity. Econometrica 72 (1), 1–32.

Postel-Vinay, F., Robin, J.-M., 2002. Equilibrium wage dispersion with worker and employer heterogeneity. Econometrica 70 (6), 2295–2350.

Rothstein, J., 2011. Unemployment insurance and job search in the Great Recession. Brook. Pap. Econ. Act. 43 (2 (Fall)), 143–213.

Saporta-Eksten, I., 2014. Job loss, consumption and unemployment insurance. Unpublished manuscript

Schmidt, L., 2016. Climbing and Falling Off the Ladder: Asset Pricing Implications of Labor Market Event Risk.

Song, J.G., Manchester, J., 2007. New evidence on earnings and benefit claims following changes in the retirement earnings test in 2000. J. Public Econ. 91 (3), 669–700.

Storesletten, K., Telmer, C.I., Yaron, A., 2004. Cyclical dynamics in idiosyncratic labor market risk. J. Polit. Econ. 112 (3), 695–717.

Storesletten, K., Telmer, C.I., Yaron, A., 2007. Asset pricing with idiosyncratic risk and overlapping generations. Rev. Econ. Dyn. 10 (4), 519–548.

Topel, R.H., Ward, M.P., 1992. Job mobility and the careers of young men. Q. J. Econ. 107 (2), 439–479.

von Wachter, T., Song, J.G., Manchester, J., 2011. Trends in employment and earnings of allowed and rejected applicants to the social security disability insurance program. Am. Econ. Rev. 101 (7), 3308–3329.