

DISCUSSION PAPER SERIES

IZA DP No. 11402

**Workers, Firms and Life-Cycle Wage
Dynamics**

Paul Bingley
Lorenzo Cappellari

MARCH 2018

DISCUSSION PAPER SERIES

IZA DP No. 11402

Workers, Firms and Life-Cycle Wage Dynamics

Paul Bingley

VIVE Copenhagen

Lorenzo Cappellari

Università Cattolica Milano and IZA

MARCH 2018

Any opinions expressed in this paper are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but IZA takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The IZA Institute of Labor Economics is an independent economic research institute that conducts research in labor economics and offers evidence-based policy advice on labor market issues. Supported by the Deutsche Post Foundation, IZA runs the world's largest network of economists, whose research aims to provide answers to the global labor market challenges of our time. Our key objective is to build bridges between academic research, policymakers and society.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

Workers, Firms and Life-Cycle Wage Dynamics*

Studies of individual wage dynamics typically ignore firm heterogeneity, whereas decompositions of earnings into worker and firm effects abstract from life-cycle considerations. We study firm effects in individual wage dynamics using administrative data on the population of Italian employers and employees. We propose a novel identification strategy for firm-related wage components exploiting the informative content of the wage covariance structure of coworkers. Wage inequality increases three-fold over the working life; firm effects are predominant while young, but sorting of workers into firms becomes increasingly important, explaining the largest share of lifetime inequality. Static models that do not allow for life-cycle dynamics underestimate the importance of sorting and overstate match and firm effects.

JEL Classification: J24, J31

Keywords: wage inequality, wage dynamics, co-workers' covariance

Corresponding author:

Lorenzo Cappellari
Department of Economics and Finance
Università Cattolica Milano
Largo Gemelli 1
20123 Milano
Italy
E-mail: lorenzo.cappellari@unicatt.it

* We thank Massimo Antichi, Tito Boeri, David Card, Mariella Cozzolino, Francesco Devicienti, Edoardo Di Porto, Pietro Garibaldi, Luigi Guiso, Paolo Naticchioni, Kostas Tatsiramos, Joseph Zweimüller and audiences at the Italian Social Security Institute (INPS) and the EC-JRC in Ispra for comments and discussions. Funding and data access were provided through the VisitINPS Scholars A programme, whose support we gratefully acknowledge. The usual disclaimers apply.

1. Introduction

Understanding the dynamics of individual wages is the subject of a large literature. Labor economists have studied the earnings process to distinguish the long-term determinants of wage inequality from wage instability, relating the former to heterogeneity in human capital investments and returns (Baker and Solon, 2003; Moffitt and Gottschalk, 2012; Blundell, Graber and Mogstad, 2015). Household economists and macroeconomists have characterised the statistical properties of labor incomes to understand patterns of consumption and the role of alternative policies for insuring income risk (Meghir and Pistaferri, 2004, 2011; Huggett, Ventura and Yaron, 2011; Blundell, 2014). The evolution of income inequality over the life cycle represents a policy concern also because it may lead to unequal opportunities for future generations, for example through household decisions about investment in child human capital (Corak, 2013).

Studies of individual earnings often acknowledge the relevance of firm heterogeneity. There are several mechanisms through which the interaction between workers and firms may have an impact on life-cycle wages. On-the-job, or within worker-firm matches, wages are affected by employer learning, informal insurance provision and firm-specific human capital accumulation. Between jobs, worker mobility following employer and worker search for better matches is reflected in wage changes. While these theoretical mechanisms have found varying degrees of empirical support, a direct empirical assessment of the role played by firms in affecting life-cycle wages is still missing in the literature.

We are among the first to study how firm heterogeneity explains individual wage inequality over the working life. We use recently released data on the population of Italian employers and employees to estimate the individual wage process as the result of individual-specific, match-specific and firm-specific shocks occurring over the life cycle.¹ Our

¹ Papers allowing for match-specific effects without accounting for firm-specific heterogeneity include Low, Meghir and Pistaferri (2010), Altonji, Smith and Vidangos (2013) and Cappellari and Leonardi (2016).

identification strategy exploits the empirical auto-covariance structure of wages for both individuals and co-workers. While the covariance structure of individual wages has been widely studied (see among others Baker and Solon, 2003, and Moffitt and Gottschalk, 2012) ours is the first paper to estimate the intertemporal covariance of co-worker wages, which we show can identify firm and worker components of the wage process without needing to further specify worker- and firm- specific wage effects.

The literature on firm effects in individual wage dynamics is in its infancy. In a study that is closely related to our own, Friedrich et al. (2016) use Swedish register data to estimate the relationship between life-cycle wage shocks and shocks to firm productivity, allowing for endogenous employment and mobility between firms. We only model the wage process, which allows for weaker assumptions - about second moments, rather than about the functional form of the shocks distribution as would be required with more processes. Furthermore, our model allows for the sorting of workers into firms and distinguishes life-cycle shocks from the effect of the business cycle.

Our paper is also closely related to the literature decomposing wage inequality into worker- and firm- specific effects abstracting from life-cycle considerations. In a perfectly competitive labor market, firm-specific wage premia would be wiped out and the existence of firm-related wage dispersion is usually interpreted as violation of the ‘law of one price’ (Card et al., 2018). A related question concerns the degree of sorting – the extent to which high/low wage workers work for high/low wage firms. Abowd, Kramarz and Margolis (1999, AKM henceforth) were the first to successfully tackle the estimation challenges entailed by such decompositions, showing that separate identification of workers and firms effects is feasible with two-way fixed effects models when data are available on sets of firms that are linked by workers mobility. Their results for France show that individual heterogeneity is more important

than firm heterogeneity in explaining wage inequality, and that indeed high wage workers sort into high wage firms.

Card, Heining and Kline (2013) spurred renewed interest in these research questions, adapting the analytical framework to the analysis of changes in wage inequality.² Estimating the AKM model on various sub-periods over the past 30 years, they showed that in Germany the increase in wage inequality has been driven by widening distributions of both worker- and firm-specific wage premia.³ In contrast to these papers, we do not rely on mobility to identify firm-related variance components; we provide a framework for modelling changes of the wage distribution over time allowing life-cycle effects in wage mean and variance, and without needing to split the data into sub-panels.⁴

Our analysis is set in the context of the Italian labor market. The Italian wage setting system features national contracts bargained at the industry level (with a coverage rate of about 80 percent) and complimentary firm-level bargaining meant to adjust wages to local economic conditions. There is no legal minimum wage and wage floors are established in national contracts. Italy has been undergoing institutional changes in the last 20 years, mostly aimed at increasing employment flexibility through the diffusion of temporary work contracts and, lately, reductions in firing restrictions for permanent contracts. Earlier reforms were focussed on wage flexibility. For example, during the 1970s an egalitarian system of wage indexation against inflation known as *Scala Mobile* caused a great compression of wage differentials between skill groups and it was only in the following decade the system was reformed and eventually abolished (Erickson and Ichino, 1995; Manacorda, 2004). These changes formed the

² See also Card, Cardoso and Klein (2016) and Song et al. (2016).

³ Card et al. (2018) discuss the parametric restrictions needed in two way fixed effects models when one wants to control for age effects in mean wages. Because fixed effects include year of birth, indicators for age and calendar time can be included among regressors to control for time-varying heterogeneity only through parametric restrictions, and results of variance decompositions turn out to be sensitive to the specific restrictions adopted.

⁴ Card, Heining and Klein (2013) discuss tests for exogeneity of workers' mobility. Macis and Schivardi (2016) perform those tests on a sample originating from the same administrative data source as our own, finding evidence in support of the exogeneity assumption.

background for a re-opening of wage differentials throughout the 1990s and early 2000s (Cappellari and Leonardi, 2016), but, as we document, differentials have not increased in recent years.

We focus on men aged 25 to 55 to limit issues of endogenous labour force participation and we keep with the literature on earnings dynamics in estimating wage differentials by year of birth to separate life-cycle shocks from calendar time variation. In our model, individual wages evolve over the life cycle through the arrival of shocks. Shocks can be long-lasting or purely transitory, and can be individual- or firm-specific. Individual-specific shocks accumulate both during the overall life cycle and within worker-firm matches. The life-cycle process may result from the accumulation of general human capital or mobility between firms, while match specific accumulation may reflect mechanisms such as firm-specific human capital, employer learning or employer-provided insurance.⁵ Individual wages also evolve because of firm-specific shocks. Differently from match-specific shocks that are purely idiosyncratic, firm-specific shocks are common to all co-workers. As in papers using the AKM approach, these effects are meant to capture firm-level factors affecting the whole workforce of the firm, and are allowed to be correlated with individual-specific shocks to capture the possibility of sorting of certain types of workers in certain firms. Differently from those papers, in ours, besides modelling life-cycle wages, we also allow for dynamics in firms effects exploiting information about the age of the firm. Finally, the model includes purely transitory wage shocks both at the individual and the firm level. All types of shocks are allowed to impact on the wage process through period-specific loading factors that capture aggregate changes in the wage distribution.

Results show that on average over the life cycle and across time, factors related to idiosyncratic workers' ability account for about 39 percent of overall inequality, of which 28 percent is individual-specific and 11 percent match-specific. The sorting component alone

⁵ See Lange (2007) and Guiso, Pistaferri and Schivardi (2005) for discussions of employer learning and employer-provided insurance.

accounts for 38 percent of wage inequality, while firm-specific factors only explain 15 percent of wage inequality, and the remaining 8 percent is due to transitory shocks. The estimated prevalence of sorting is in line with recent findings by Borovičková and Shimer (2017), while we attribute a lower share of inequality to individual-specific heterogeneity and a larger share to sorting compared to Card, Heining and Klein (2013). Wage inequality grows considerably over the life cycle, consistent with findings in the literature (see e.g. Huggett, Ventura and Yaron, 2011). Wage differentials at age 25 are almost entirely due to firm-related factors, induced by either idiosyncratic firm effects or workers' sorting into firms. Individual-specific factors account for much of the growth of life-cycle wage inequality. Idiosyncratic match effects are characterised by larger wage dispersion among young workers, consistent with the possibility that early employers face larger uncertainty on match quality and there is more to be learnt about young workers than about older ones.

Our model nests static specifications akin to the AKM model. Compared with the dynamic model, estimates from these restrictive versions attribute larger shares of wage inequality to individual- and firm-related heterogeneity and lower shares to sorting, resembling the findings of Card, Heining and Kline (2013) for Germany. Comparing these findings to the baseline estimates suggests that life-cycle dynamics are an important feature of workers sorting into firms. In the last part of the paper we look at differences in the relative importance of firm effects across occupational groups. We find that for blue collar workers wage inequality is largely reflecting heterogeneity between firms, while among white collars also individual ability matters in shaping wage dynamics.

2. Data

We draw data from the archives of the Italian Social Security Institute (*Istituto Nazionale di Previdenza Sociale*, INPS) covering the population of employers and employees in the private

non-agricultural sector of the Italian economy.⁶ The main source of information is the form that employers have to fill in order to pay state pension contributions for their employees, which is available in digitalized form since 1983. In this form employers report the gross take home pay, which is the full net pay including all forms of monetary compensation, grossed up with labor income taxes and pension contributions levied on the employee. Besides the amount of gross pay, for each employment spell the data report total working days, initial and end dates (initial dates are recorded since February 1974) and broad occupational category (apprentices, blue collar, white collar or manager). This is supplemented with employer-level information about industry, location, date of establishment and date of closure. Information is available at both the firm- and plant-level and we maintain firms as the unit of analysis on the employer side.

Throughout the period, average firm size is about 9 employees, while the average firm age goes from 8 years in 1983 to 13 years in 2016. Demographic information on workers include gender and year of birth, but, as typical with many administrative sources, do not include education. We focus on men. For each men in each year, we determine the *prevalent employer* as the one with which he has spent the largest number of working weeks, and we exclude matches with prevalent employers lasting less than 8 working weeks. In order to derive a precise measure of workplace tenure, we keep fresh spells in the 1983-2016 and exclude men whose first observed job in that window started prior to February 1974. The resulting dataset matches all employers and male employees in Italy between 1983 and 2016, including 3,493,326 firms and 14,021,258 workers, totalling 192,112,742 observations over the period.

We use data since 1985 because of incompleteness of digital records for 1983 and 1984. In keeping with much of the literature on individual wage dynamics we consider working careers between age 25 and 55 to limit endogenous participation of men due to educational choices or early retirement. We exclude apprentices and managers, which represent 0.48 and

⁶ Population data are accessible through the VisitINPS programme.

1.52 percent of observations in that age range. We derive gross daily wages as the ratio between the gross annual pay with the prevalent employer and the number of working days, rescaled to full-time equivalents in the case of part-time jobs.

Because we are chiefly interested in life-cycle dynamics, we select individuals on the basis of year of birth, such as we can observe a certain portion of the life cycle for each cohort. To ease identification of life-cycle profiles, we require that each cohort is observed for at least 10 years. In this way, cohorts participating in the analysis range from 1939 (aged 46 in 1985 and observed ten times until it reaches age 55 in 1994) to 1982 (aged 34 in 2016 and observed ten times since 2007 when it is aged 25). The cohort structure of the data is represented in Figure 1.A. Two birth cohorts (1959 and 1960) are observed throughout the 25-55 age range, while the number of observations on other cohorts progressively decreases moving backward or forward from these two central cohorts. Each cohort born before 1954 or after 1979 accounts for between 1.5 and 2 percent individuals in the sample, whereas each remaining cohort account for between 2 and 3 percent of individuals in the sample. Besides requiring 10 potential observation on individual wages based on birth year, we also require that individuals are actually observed for at least 5 times consecutively; we further winsorize the resulting wage distribution at the first and last half-percentile by year.

Applying the above selection rules gives a sub-population of 12,339,989 men and 3,067,753 firms between 1985 and 2016 , corresponding to 152,470,973 observations. Figure 1.B contrasts average age in the sub-population under analysis with average age for the population of men aged 16-65. While there is some difference between the two, the former being slightly younger before the early 2000s and slightly older afterwards due to the revolving-by-cohort design of the data, especially after 2007, differences are not major, suggesting that wage-age patterns estimated from the sub-population are a good proxy for the overall population of men. We next compare hourly wages for the sub-population used in estimation

with their counterparts in the full population of men aged 16-65.⁷ Using information on the starting dates of job spells we derive job tenure, whose average in the selected sub-population fluctuates between 4 and 6 during the years 1985 to 2007, and increases to 8 years between 2008 and 2016, paralleling the increase in average age in those years. Figure 1.C shows no substantial difference in average hourly wages between the two groups. Figure 1.D shows that similar trends between the selected sub-population and the full population emerge also considering the standard deviation of logs. There is a difference of about a fortieth of a standard deviation between the two plots (larger figures corresponding to the full population), which is roughly constant throughout the period.

3. Econometric model

Wages evolve over the life cycle through the arrival of individual-specific shocks and of firm-specific shocks that are common to all the workers employed in a given firm. Wage shocks may be long-lasting or purely transitory. In the first case, they reflect persistent or slowly changing wage determinants. Transitory wage shocks capture the impact of economic volatility.

Persistent individual shocks are modelled as unit root processes, that can reflect the accumulation of human capital or other sources of persistent wage change, like employer learning. Wage dynamics may occur both between or within firm, the latter corresponding to idiosyncratic worker-firm match effects. Between firms, wages may change thanks to the acquisition of general human capital or workers mobility. Within matches, workers accumulate firm-specific human capital, while employer learn about workers ability. We assume that match effects are orthogonal to the rest of the model and we estimate the accumulation of wage differentials within matches. We allow the weight of individual shocks to change flexibly over

⁷ We reflate nominal values to 2015 using the CPI.

calendar time; for example, these time trends may reflect aggregate changes in human capital returns.

Firm-specific effects reflect the wage premium or penalty of working for a given employer. These may originate from rent extraction, efficiency wages, or other mechanisms generating heterogeneity among employers. Because these effects are common to all employees, they do not depend on workers' age. Monitoring technology or unions' ability to extract rents may not be time-invariant, however, and we allow for this possibility by allowing firm-specific effects to change over time according to the age of the firm, capturing the idea that a firm's ability to impact the wages of its employees may be different at different stages of its life. Moreover, we also allow for aggregate changes in the weight of firm effects in wage determination through period-specific loading factors. Finally, we allow firm effects to be correlated with life-cycle permanent shocks to account for the possibility of sorting of workers into firms, and its variation through the life cycle.

Transitory shocks can be individual-specific or firm-specific white noise processes, orthogonal with the rest of the model; we allow for aggregate changes of their distribution over time.

3.1 Model Specification

Let w_{ijt} denote the residualized log of daily wages for person i in year t and let $j = J(i, t)$ denote the firm in which this person is employed in that year. We residualize raw log-wages on a set of age dummies by birth cohort, such as residuals are centered on time-by-cohort means. Our model allows for multiple sources of time-varying heterogeneity (not only time-invariant one), therefore we do not need to control for time-varying regressors as is done for example in the AKM framework. Note that although wages are indexed by calendar time, the couple (i, t)

unambiguously identifies the age of person i in year t . We summarise the mechanisms described above in the following model of life-cycle wages

$$w_{ijt} = \alpha_t(\lambda_{it} + \mu_{ijt}) + \delta_t \phi_{jt} + \gamma_t(\varepsilon_{it} + \xi_{jt}) \quad (1)$$

where α , δ and γ are time shifters on wage components allowing for aggregate changes in the wage distribution.

The first term in brackets represents the idiosyncratic persistent component resulting from life-cycle (λ) and match-specific (μ) wage shock. Life cycle shocks permanently affect the age-wage profile irrespective of the firm at which person i is employed in a given year, while match specific shocks capture the deviation from the life cycle profile that depends on the specific firm that is employing person i in year t . Both processes are modelled as unit root (Random Walk, RW) because these shocks have long lasting effects.

The life-cycle RW in age captures shock accumulation since age 25:

$$\lambda_{it} = \lambda_{i(t-1)} + u_{it} = \lambda_{i(c+25)} + \sum_{k=c+26}^t u_{ik}; \quad \lambda_{i(c+25)} \sim (0, \sigma_\lambda^2); \quad u_{it} \sim (0, \sigma_{u(t-c)}^2) \quad (2)$$

where $c = c(i)$ is the year of birth of person i , such as $c + 25$ is the year in which the income trajectory conventionally starts, and $(t - c)$ is i 's age in year t . The initial variance σ_λ^2 measures idiosyncratic wage dispersion at the conventional starting point and will in effect also reflect the accumulation of shocks occurring prior to that age. Permanent shocks u_{it} are age specific-shocks and are drawn from age-specific distributions with variance $\sigma_{u(t-c)}^2$ to allow for the possibility that their importance in shaping the wage distribution changes with age. For example differences in human capital may exert a greater impact on wage inequality for young workers than for middle age ones.

Match effects are represented by a match specific RW that accumulates only within matches

$$\mu_{ijt} = S_{i(t,t-1)}\mu_{ij(t-1)} + v_{ijt} = \sum_{k=t-\tau}^t v_{ijk}; \quad S_{i(t,t-1)} = 1[J(i,t) = J(i,t-1)] \quad (3)$$

where $1[\]$ is an indicator function, S is a dummy for job stayers and τ is tenure at the firm. We allow match-specific shocks to be drawn from age-specific distributions in order to see if matches formed earlier on in the life cycle are more or less important than later matches in determine wage hierarchies in the long run:

$$v_{ijt} \sim (0, \sigma_{v(t-c)}^2).$$

One possible mechanism is employer learning: larger uncertainty on match quality for young workers implies that there is much to be learnt about their ability during the life of the match, resulting in more heterogeneous wage progressions within the firm compared to older workers whose underlying ability can be more precisely measured by employers already at the start of the worker-firm relationship. Other mechanisms such as faster accumulation of firm-specific human capital among young workers can account for such patterns.

Firm effects captures firm-specific wage policies common to all co-workers. These effects determine the average position in the wage distribution for the employees of a given firm. Clearly, the ability of the firm to affect the wage position of its employees within the overall wage distribution may change over time in idiosyncratic ways, examples are changes in firm's market power or in its industrial relation practices. We capture these possibilities by allowing the distribution of firm effects to shift with the age of the firm, and we characterize its intertemporal persistence:

$$\phi_{jt} \sim (0, \sigma_{\phi(t-d)}^2) \qquad E(\phi_{jt} \phi_{jt'}) = \sigma_{\phi(t-d)(t'-d)}$$

where $d = d(j)$ is the year in which the firm is established, such as $(t - d)$ represents the age of the firm in year t .

Transitory shock are specified as White Noise (WN) processes with innovations drawn from age-specific distributions:

$$\varepsilon_{it} \sim (0, \sigma_{\varepsilon(t-c)}^2); \qquad \xi_{jt} \sim (0, \sigma_{\xi(t-d)}^2).$$

While it is customary in the earnings dynamics literature to include some form of autoregression in “transitory” wage shocks to accommodate slow reversion to the mean, our model already allows for multiple sources of persistence in the permanent part of the process (all three parts of the permanent wage are highly persistent, and additional persistence will be allowed through the correlation between individual and firm effects, see below), in ways that are unparalleled by models using ARMA specification of transitory shocks, leaving little room for additional persistence in wage shocks.

We assume that transitory shocks and matching effects are uncorrelated among themselves and with anything else:

$$\begin{aligned} E(\mu_{ijt}, \lambda_{it}) &= E(\mu_{ijt}, \phi_{jt}) = E(\mu_{ijt}, \varepsilon_{it}) = E(\mu_{ijt}, \xi_{jt}) = 0 \\ E(\varepsilon_{it}, \lambda_{it}) &= E(\varepsilon_{it}, \phi_{jt}) = E(\xi_{jt}, \lambda_{it}) = E(\xi_{jt}, \phi_{jt}) = 0 \end{aligned} \tag{4}$$

We assume that firm and individual effects are correlated to allow for the possibility of sorting of workers into firms, with different degrees of intensity depending on workers age:

$$E(\lambda_{i(c+25)}, \phi_{jt}) = \rho_{\phi 25}; E(u_{it}, \phi_{jt}) = \rho_{\phi(t-c)}.^8 \quad (5)$$

3.2 Moment restrictions and identification

The model is estimated by GMM (Minimum Distance) matching empirical second moments of the wage distribution across cohorts and time periods to their counterparts implied by the model.⁹ The individual covariance structure between year t and $t' \geq t$ reflects all sources of wage dispersion:

$$\begin{aligned} E(w_{ijt}w_{ij't'}) &= \alpha_t \alpha_{t'} \left(\sigma_\lambda^2 + 1[t > c + 25] \sum_{k=26}^{(t-c)} \sigma_{uk}^2 + S_{i(t,t')} \sum_{k=t-\tau}^t \sigma_{v(k-c)}^2 \right) + \quad (6) \\ &S_{i(t,t')} \delta_t \delta_{t'} \left(\sigma_{\phi(t-d)}^2 1[t = t'] + \sigma_{\phi\phi(t-d)(t'-d)} 1[t \neq t'] \right) + \\ &(\alpha_t \delta_{t'} + \alpha_{t'} \delta_t) \rho_{25\phi} + 1[t > c + 25] \alpha_t \delta_{t'} \sum_{k=26}^{(t-c)} \rho_{\phi k} + \\ &1[t' > c + 25] \alpha_{t'} \delta_t \sum_{k=26}^{(t'-c)} \rho_{\phi k} + \\ &1[t = t'] \gamma_t^2 (\sigma_{\varepsilon(t-c)}^2 + \sigma_{\xi(t-d)}^2) \end{aligned}$$

Permanent and transitory shocks are identified by the contrast between variances ($t = t'$) and covariances ($t \neq t'$). Life-cycle shocks are identified separately from match-specific shocks because they accumulate over different running variables, age and tenure. Empirically, we will exploit variation of tenure across birth cohorts and time periods of wages residualized on cohort and period fixed effects, which is less of a concern in terms of endogenous mobility than if we were using tenure variation within a cohort.

⁸ While feasible in principle, we refrain from allowing sorting effects to depend also on the age of the firm to preserve model tractability.

⁹ We use Equally Weighted Minimum Distance (EWMD) and a robust variance estimator $\text{Var}(\hat{\beta}) = (G'G)^{-1} G'VG(G'G)^{-1}$, where $\hat{\beta}$ is the vector collecting all model parameters, V is the fourth moments matrix and G is the gradient matrix evaluated at the solution of the minimisation problem (see Chamberlain, 1984 and Haider, 2001).

It is clear from Equation (6) that information on individual wages is not enough to separate worker-specific effects from firm-related ones, and additional information is needed. This can be generated by considering the wage covariance structure of co-workers, that is persons i and h working for the same firm at some point of their lives:

$$\begin{aligned}
E(w_{ijt}w_{hjt'}) &= \delta_t\delta_{t'} \left(\sigma_{\phi(t-d)}^2 1[t = t'] + \sigma_{\phi\phi(t-d)(t'-d)} 1[t \neq t'] \right) + & (7) \\
& (\alpha_t\delta_{t'} + \alpha_{t'}\delta_t)\rho_{25\phi} + 1[t > c + 25]\alpha_t\delta_{t'} \sum_{k=26}^{(t-c)} \rho_{\phi k} + \\
& 1[t' > c + 25]\alpha_{t'}\delta_t \sum_{k=26}^{(t'-c)} \rho_{\phi k} + 1[t = t']\gamma_t^2\sigma_{\xi(t-d)}^2
\end{aligned}$$

This expression does not contain individual-specific variance components because idiosyncratic components are independently distributed across persons, but only firm-related ones, that is the firm effect and the sorting parameters. Therefore, using the information provided by the individual covariance structure and the one generated by the co-workers covariance structure in conjunction identifies individual specific elements separately from firm-related ones. Noteworthy, such separation is identified without the need of workers mobility between firms, because co-workers moments are defined irrespective of workers moving.

Equation (7) shows that while life-cycle sorting effects ($\rho_{\phi(t-c)}$) are identified by age variation in co-worker wage covariances, sorting at age 25 ($\rho_{25\phi}$) is not identified separately from the variance of firm-specific wage effects ($\sigma_{\phi(t-d)}^2$ and $\sigma_{\phi\phi(t-d)(t'-d)}$), both terms being constant in the co-workers covariance structure. Put another way, individual and co-workers wage covariances are enough to identify individual, firm and life-cycle sorting components, but not initial sorting, which needs extra information (moment restrictions) for identification. One possibility for identifying sorting at age 25 is to generate additional moment restrictions using

information on wage covariances before individuals actually join a given firm. Consider worker i who is employed by firm x in year t and joins firm j in year $t+k$, and worker h who is an employee of firm j in year t . Any covariance between their wages in year t cannot reflect a firm effect, because in that year they are employed by two different firms. Such covariance instead emerges from their sorting into similar firms, which can be thought of as a form of assortative mating. To operationalize these ideas, first we define the *future employer* as the one that a person will join in two years' time, conditional on not having been employed in that same firm in the prior two years; next for each individual with a future employer, we compute the covariance between his current wage and the current wage of his *future co-workers*. Clearly, this strategy requires workers mobility, because the future employer is not defined for stayers, which may raise issues of endogenous mobility. However, because the parameter identified by future co-workers refers to sorting at age 25, empirically we only use information on wage covariances of movers aged 25 joining the future employer at age 27, an age range when workers mobility is more common, mitigating concerns about the endogeneity of workers mobility. Also, we will supplement results with those from a model that assumes no sorting at age 25 and that does not use the covariance of future co-workers, which will provide a sense of the bias induced by young workers mobility in the estimates of the fully identified model. Using future co-workers generates the following additional moment restriction:

$$E(w_{ixt}, w_{hjt'}) = (\alpha_t \delta_{t'} + \alpha_{t'} \delta_t) \rho_{25\phi}, \quad (8)$$

$$\forall i: J(i, t+2) = j, J(i, t+s) \neq j, (t-c) = 25; s = -2, \dots, 1$$

which identifies sorting at age 25.

4. Empirical covariance structures

We estimate cohort-specific wage covariance structures that we match to the set of moments discussed in the previous section to estimate the parameters of the model. There are three sets of moments of interest. Individual moments (denoted with I) are estimated by averaging the cross products of residualized log-wages across individuals:

$$m_{tt'}^I = \frac{\sum_i \omega_{ijt} \omega_{ijt'}}{\sum_i p_{ijt} p_{ijt'}} \quad (9)$$

where ω is the empirical counterpart of w , while p indicates whether person i is observed in period t , because individuals may leave and/or join the panel provided they are observed at least five consecutive times.¹⁰

The co-workers covariance structure (indexed by C) is estimated by adapting the algorithm of Page and Solon (2003) for the estimation of neighbourhood covariances in outcomes. First, the firm-specific covariance is estimated by averaging cross-products of log-wage residuals for all pairwise matches that can be formed across co-workers; next, firm-specific covariances are averaged across firms using the square root of the number of pairwise matches as weight. The weighting procedure attributes more weight to larger firms and makes inference person-representative. For a given cohort, the co-workers covariance structure is given by:

$$m_{tt'}^C = \sum_j \theta_j \frac{\sum_i \sum_{h>i} \omega_{ijt} \omega_{hjt'}}{\sum_i \sum_{h>i} p_{ijt} p_{hjt'}} \quad (10)$$

¹⁰ A discussion of GMM estimation of earnings dynamics model with unbalanced panels is provided by Haider (2001).

where $\theta_j = \sqrt{\sum_i \sum_{h>i} p_{ijt} p_{hjt'}} / \sum_j \sqrt{\sum_i \sum_{h>i} p_{ijt} p_{hjt'}}$ is the firm-specific weighting factor. For cohorts of co-workers up to 200 individuals, all co-workers are used in the estimation of (10), while for larger cohorts, a random sample of 200 co-workers stratified by occupation is used in estimation.

The third set of moments (indexed by F) is the one between 25 years old employees and their future co-workers, i.e. the current employees of the firm they will join in two years since the year of observation and for which they have not been working for in the two years prior to the period of observation:

$$m_{tt'}^F = \sum_i \frac{\frac{\sum_{h \neq i} \omega_{ixt} \omega_{hjt'}}{\sum_{h \neq i} p_{hjt'}}}{p_{ixt}}, \quad (11)$$

$$\forall i: J(i, t+2) = j, J(i, t+s) \neq j, (t-c) = 25, s = -2, \dots, 1$$

There are 21,302 empirical moments in total, 10,582 each for individuals (estimated with equation 9) and co-workers (estimated with equation 10), and 138 for future co-workers (equation 11, which is estimable only for 25 years old individuals and for cohorts older than 1959). We report estimated empirical moments in Figure 2. Panel A portrays the historical evolution of the variance. The plot labelled “Individuals” is the overall wage dispersion estimated using deviations of individual wages from cohort-specific means, averaging cohort-specific variances across cohorts. This shows an increasing trend between the mid-1980s and the early 2000s, followed by a slight decrease until 2008 and a slight increase in the more recent years. The same panel reports the plot labelled “Co-workers” which is obtained using equation (10) and which essentially provides a measure of how much co-workers *jointly* deviate from the overall (cohort-specific) mean. In other words, this plot documents the evolution of wage dispersion between firms, due to either idiosyncratic firm effects or the similarities of wage

generating characteristics among co-workers, emerging from the sorting of similar workers in the same firm. It does not reflect wage differentials due to idiosyncratic individual (or match) characteristics. Remarkably, dispersion between clusters of co-workers follows a time pattern that parallels the one of overall dispersion, and at a level that oscillates approximately between one half and two-thirds of overall wage inequality implying that firm-related wage differentials go a long way in accounting for the general trends in the wage distribution. The difference between the “Individuals” and “Co-workers” plots provides an account of how much of wage inequality is due to factors that are not common among co-workers, but are instead individual specific and/or match specific. Empirical moments among clusters of future co-workers (not reported) are constant at a lower level of about 0.01. Lastly, we also report empirical moment for “Placebo” co-workers, which we obtain by using Equation (10) to match employees across firms randomly extracted from the economy. There is no wage association whatsoever among placebo co-workers, which reassures that the patterns recovered for co-workers are effectively a reflection of being exposed to a common firm effect or of the sorting process underlying firm-worker matches, and not simply a statistical coincidence linking individuals born in the same year.

Noteworthy, the contrast between “Individuals” and “Co-workers” in Figure 2 enables distinguishing individual idiosyncratic from firm-related components of wage variance in a fully non parametric way, without specifying a model of workers and firms heterogeneity. Specifying a model as we did in the last Section is needed, instead, to separate individual from match effects within the purely idiosyncratic component, and firm effects from sorting within the co-workers covariance structure.

In Panel B of Figure 2 we consider the evolution of wage variances between ages 25 and 55, which represents the central object of interest for our analysis. We report the overall variance of wages using individuals as the unit of analysis, and the variance obtained between

clusters of co-workers. There is a remarkable increase in overall wage inequality over the life cycle, which is a symptom of underlying heterogeneous wage dynamics across workers. The growth appears to be faster in the 30-s then in the early 40-s; also, there is an acceleration after age 45. Looking at co-workers, we can still see an increase over the life cycle, but not so steep as the one of overall wage dispersion. At age 25, the difference between overall inequality and inequality between clusters of co-workers is rather limited, suggesting that to a large extent wage differentials at the beginning of the working life stem from factors that are common among co-workers. As workers age, the relevance of individual idiosyncratic influences increases, and by age 55 almost half of total dispersion is due to individual-specific factors.

Panel C of Figure 2 reports empirical autocovariances for individuals and for co-workers clusters. The autocovariance function of individual wages shows a sharp drop when moving from variances to first order covariances, consistent with the existence of shortly lived wage shocks. The decline of covariances over lags is essentially linear, a pattern that can be generated by a RW process for individual earnings dynamics. Wage autocovariances for co-workers display a similarly declining linear pattern, although characterised by slower decline compared with individual covariances, suggesting that the factors determining purely idiosyncratic individual wage persistence have a faster decay than factors that are shared among co-workers, which is closer to a simpler age-constant (Random Effect) process than to a RW. After 30 years, because of the differential speed of decay of covariances for individuals and co-workers there is virtually no difference in persistence for individual or co-workers wages.

5. Results

Before discussing the parameter estimates of life-cycle earnings, we provide an overview of results by illustrating in Figure 3 the variance decomposition predicted by the model between 1985 and 2016. These predictions are obtained by using parameter estimates of both wage

shocks and of time shifters to estimate the wage covariance structure using equation (6), and then averaging prediction across birth cohorts by year. Estimates of time shifters are reported in Table A1, while “core” parameter estimates of stochastic wage processes are discussed in detail later in this Section. There is a close overlap between raw and predicted variances. The most relevant sources of wage inequality over the period are the sorting component and the life-cycle component, which points to the relevance of personal characteristics in shaping the overall wage distribution. Personal attributes matter both because they generate idiosyncratic wage differentials between individuals and because they drive the sorting of co-workers into firms. Match-specific effects also contribute to explain the growth of wage inequality in the 1990s, but their level is much lower compared to the life-cycle and sorting components, as is the one of firm-specific effects and transitory shocks. On average, the total predicted variance of log-wages is 0.11, the sorting component accounts for 38 percent of it, the life-cycle component for 28 percent, the match component for 11 percent and the firm component for 15 percent, the remaining 8 percent being due to transitory shocks. Interpreting the share of inequality explained by sorting as the correlation between workers and firms types, our finding lies close to the lower bound of the estimates of Borovičková and Shimer (2017), who report a sorting correlation between 0.4 and 0.6. Instead, our findings are different from the ones reported by Card et al. (2013) for Germany in that we attribute lower share of wage inequality to the purely idiosyncratic individual component.

In Table 1 we report parameter estimates for the permanent component of the wage model. The individual life-cycle component has an initial condition at age 25 (absorbing wage heterogeneity at labour market entry or the accumulation of shocks up to age 25 for earlier entrants) and age-specific shocks. We balance model flexibility and tractability by allowing the distribution of age-specific shocks to change by 10-year age groups up to age 45, and we further split the last ten years of the observed life cycle in two (46-50 and 51-55) to allow for the

acceleration in wage inequality growth after age 45 observed in Figure 2.B. We maintain a similar age groupings also for the tenure-related shocks of the match-specific component. The distribution of firm-specific shocks changes with the age of the firm, and we use three groups (younger than 10, between 10 and 20, older than 20) that correspond to terciles of the distribution of firm age across workers. Firm-specific effects are assumed to be constant within those terciles and to be correlated between terciles as firms age. Finally, sorting effects are estimated as the covariance between life-cycle shock and the firm effect, and we restrict the covariance to be constant in the 46-55 age range (rather than changing after 50) to solve some empirical identification issue.

Heterogeneity at labor market entry is substantial, see Panel A of Table 1. The dispersion of life-cycle shocks diminishes between 26 and 45 years of age, reflecting the concave evolution of wage differentials before age 45 shown in Figure 2.B. Both diminishing heterogeneous returns to general human capital or a reduction in the rate of workers mobility and of the associated heterogeneous wage changes could explain this pattern. The evidence is different after age 50, when we estimate a substantial increase in the dispersion of shocks that is again in line with the evidence of Figure 2.B. Heterogenous labor supply behaviour may explain this finding, as some workers may start reducing daily hours approaching retirement. In the limit, some may also start leaving the labor force, and the observed pattern may be the outcome of selective participation, if surviving workers are polarised at the tails of the wage distribution. Having model parameters that are specific to this age range ensures that any bias induced by selective participation is isolated and not transmitted to the rest of the model.

The next set of estimates (Table 1.B) refers to the match component, and illustrates the change of wage inequality associated with one year of tenure and its variation over the life cycle. Differently from the life-cycle component there is a monotone decline of tenure-related shocks dispersion as workers age, without any re-opening of the shocks distribution among

older workers. These patterns may mean that older workers have lower ability to accumulate firm-specific human capital, but are also consistent with an exhaustion of employers' uncertainty about the quality of the worker. Earlier employer may extract information on workers quality that is then passed on to later employers (for example because wages paid in each match are common knowledge), who therefore face smaller uncertainty, resulting in lower employer learning when workers are old.

Panel C in Table 1 reports estimates of the dispersion of firm-specific wage effects. These are larger for younger firms, possibly reflecting greater heterogeneity in monitoring technologies or market power. Also, the finding is consistent with differential rates of firms survival across the distribution of firm-specific wage effects. Firm effects are also characterised by substantial persistence. Panel D of Table 1 reports parameter estimates of the sorting coefficients, that is the covariance between age-specific life-cycle shocks and the average firm-specific effect. These estimates reflect the declining life-cycle pattern observed for the dispersion of life-cycle shocks. It is worth noting that parameter estimates for sorting are all positive, indicating that indeed high wage workers work for high wage firm (and viceversa); this is different from studies applying the AKM framework which sometimes report negative estimates of sorting parameters.¹¹

We summarise the inequality implications of the life-cycle wage model in Figure 4. The plot labelled "Firm-related" shows the life-cycle variance prediction obtained using firm and sorting parameters. We collapse these two components of the model because there is no variation with workers' age in firm-specific effects, whose variance is *de facto* an intercept shift for the life-cycle plot. In particular, we use the variance estimated for "middle-aged" firms to generate the prediction shown in the graph. The resulting plot describes the life-cycle evolution

¹¹ Table 2 also reports Newey's (1985) χ^2 statistic for a test of the model against the alternative hypothesis of an unspecified covariance structure. As noted in Baker and Solon (2003, footnote 25), with many empirical moments and relatively few parameters, the test is bound to reject the maintained specification.

of the variance if the only wage differentials were the ones between clusters of co-workers. Adding parameter estimates for the RW in age we obtain the plot labelled “Mover” which describes wage inequality over the life cycle in the absence of match-specific effects. This is the life-cycle variance profile that would emerge if all workers changed employer each year. Finally, the plot labelled “Stayer” considers the hypothetical case of life-cycle inequality if all workers were spending the entire career (between ages 25 and 55) with the same employer, such as one additional year of age corresponds to one year of tenure and implies an increase in the variance as the one predicted by the corresponding estimate of match-specific shocks.

Figure 4 shows that when workers are young, wage inequality is almost entirely stemming from firm-related factors, especially from the dispersion of firms specific effects; individual-specific or match specific effects account only for a minor share of wage inequality at age 25. As workers age, the firm-related component still grows because firm effects are linked to life-cycle effects via the sorting parameters. However, the growth of the individual-specific and match-specific components is faster because it also depends on the age-specific and tenure-specific unit root shocks. The graph suggests that it takes time to individual ability to display its inequality enhancing effects, both between and within firms. Overall through the life-cycle, firm-related wage components account for only 60 percent of overall wage inequality if we ignore match effects. If instead we consider the hypothetical case of workers continuously working for the same firm, firm-related components account for 40 percent of total inequality.

In Table 2 we report parameter estimates for the transitory components. We model the life cycle profile of the individual-specific WN shocks using the same age groupings that we used for the sorting effects; moreover, we group firm-level shocks into two categories of firms’ age, young and middle-old. Individual-specific shocks evolve according to a declining pattern over the life cycle; on the other hand, firm-specific volatility is larger among the workers of older firms. Considering that older workers tend to work for older firms (the correlation

between the two variables is 0.71), our overall evidence on wage instability reproduces the findings of previous research of a u-shaped life-cycle pattern of shocks volatility (e.g. Baker and Solon, 2003). However, thanks to our specification that distinguishes between individual- and firm-level transitory shocks, we can characterise the different sources of volatility of young and older workers.

6. Sensitivity checks and heterogeneity by occupation

6.1 Sensitivity checks

As discussed in Section 3, identification of sorting is achieved because firm effects vary with the life cycle of the firm but are common among co-workers at different stages of their working careers. This leaves sorting at age 25 (the initial condition) under-identified, and we supplement moment restrictions with information on young future co-workers, who are similar because they will sort into the same firm at some point in the future, but are not exposed to the same firm effect at the time wages are observed. Clearly, this requires reliance on workers mobility and job stayers do not contribute to identification because their future and current co-workers coincide. As a way to check the sensitivity of results to the use of movers for identification of sorting at age 25, we re-estimate the model without future co-workers and without separating initial sorting from firm effects. While we expect estimates of parameters related to the initial condition of the process to be sensitive to the exclusion of future co-workers, parameters for earnings dynamics and for life-cycle sorting should not be sensitive to the use of the under-identified specification. Results from this under-identified specification are reported for the permanent wage component in Table 3, Column 1, and indeed show that while both the variance of initial conditions of life-cycle shocks and the variance of firm effects are larger compared with the estimates for the baseline model in Table 1 because now they also reflect sorting, estimated variances of life-cycle shocks are unaffected. Notably, there is a reduction of

matching effects compared with the baseline, but the pattern of declining match-specific dispersion with age is unaffected.

We further address the implications of life-cycle considerations for the implied variance decompositions by estimating a nested version of the model that assumes away any life cycle variation, by setting $\lambda_{it} = \lambda_i$, $\mu_{ijt} = \mu_{ij}$, and $\phi_{jt} = \phi_j$. The resulting specification is akin to the one of AKM and Card, Heining and Kline (2013), differences being that our model allows for secular shifts in variance components and explicitly deals with match effects.¹² Results for the permanent components of this AKM-type model are reported in Column 2 of Table 3. Perhaps the best way to compare these estimates with the ones from the baseline model is to look at the implied average inequality decomposition for the overall period 1985 - 2016. To ease comparisons, Table 4 reports these decompositions for the baseline model and for the AKM version of the model without life-cycle variation. The overall variance imputable to individual heterogeneity (life-cycle and match components) is larger in the AKM model (50 vs 40 percent), which also attributes more than half of this variance to the match effect, while the baseline model weights more the life-cycle component. What is striking is the different balance between firm and sorting effects: while the baseline gives more weight to the latter, the AKM specification attributes more weight to the ‘pure’ firm effect. It is worth noting that the sorting effect has a life-cycle aspect in the baseline and that empirically wage inequality has a strong life-cycle variation, that cannot be detected by the life-cycle constant AKM specification. The estimated share of inequality imputable to the sorting component more than halves compared with the baseline and is now in lines with figures reported by Card, Heining and Kline (2013, about 10 percent on average)

¹² Another similarity with the AKM specification is that in this model we do not distinguish transitory shocks between person- and firm-specific and set $\varepsilon_{it} + \xi_{jt} = \psi_{ij}$.

6.2 *Heterogeneity by occupation*

We now turn to an analysis of heterogeneity in life-cycle wage dynamics across groups of workers. Specifically we split the data according to broad occupational groups (blue collar and white collar workers) assigning individuals to the group to which they most frequently belongs over the observed portion of their careers. In this way, we assign 8,768,457 men to the blue collar group (corresponding to 63% of men in the main analysis) and 5,252,801 to the white collar group. These groups are characterised by different levels of human capital (and particularly different levels of education, which is missing in the data) and this might have an impact in shaping their wage dynamics.

Figure 5 reports empirical moments for these two occupational groups, showing sharp differences. For blue collar workers in Panel I there is hardly any secular trend of inequality, and also the life-cycle evolution is flatter than what we observed for the whole labor market; the autocovariance plot is also lower than in the general case, and also declines at slower paces. The other striking feature for blue collars in Panel I is that there is little difference between individuals and co-workers in each of the three plots, providing a first indication that for low-skilled workers the firm one works for is the predominant determinant of wage differentials. The evidence for white collars is remarkably different. First of all, Panel II.A shows a dramatic increase of overall wage inequality between 1985 and 2016, of more than 100 percent, reproducing the trends that have been observed in other countries, like Germany. Second, the life-cycle profile of wage inequality clearly shows two distinct phases: a great expansion before age 40 (the level of dispersion almost doubles between 25 and 35) and a substantial flattening thereafter, with only a slight sign of rewidening differentials among older workers. Third, and this is perhaps the most striking fact about white collars, there is a large difference between wage covariances for individuals and co-workers, showing that in this group individual specific effects matter more than for other workers, consistently with the fact that these are highly skilled

individuals. Also, it is interesting to note that the co-workers autocovariance function is nearly constant.

We estimate the life-cycle wage model by skill group and report parameter estimates of permanent wages in Table 3 (Columns 4 and 5) while Figure 6 works out the predicted profiles of life-cycle inequality. On average over the period the level and components of inequality strikingly differ by occupation. The overall level of dispersion almost doubles between blue and white collar workers, 0.06 vs 0.10, see Table 4, Columns 3 and 4. For blue collar workers, the individual and match components explain 18 and 13 percent of overall inequality, while firm effects absorb 40 percent of overall inequality, the rest being due to sorting (17 percent) and transitory shocks (12 percent). Among white collars, individual-specific components alone account for 32 percent of overall inequality, and match -specific effects explain another 20 percent, such as more than half of wage inequality can be ascribed to permanent factors that are not shared among co-workers, pointing to the importance of individual ability in shaping the wage careers of high skilled workers. Heterogeneous returns to human capital investments (both within the firm and between firms) can be one mechanism behind these patterns for white collar workers that are characterised by higher levels of human capital compared to blue collar workers. Conversely, the “pure” firm effect explains 20 percent of inequality, sorting explains another 21 percent and the remaining 7 percent is volatility. It is interesting to note that the share due to sorting is much lower within each skill group than it was in the labor market as a whole, suggesting that much of the overall sorting occurs because different occupational groups sort into different firms. It is also interesting to see that for blue collar workers that are characterised by negligible life-cycle variation, the baseline model produces variance decompositions that are more similar to the AKM specification.

The evidence in Figure 6 confirms that wage differentials have a different nature within occupational groups. For blue collars, at age 25 there is virtually no role for individual-specific

heterogeneity and all of wage dispersion comes from firm-related differences. Instead for white collars there is a substantive share of initial inequality that is not coming from factors shared among co-workers. Over the life cycle, wage differentials induced by individual-specific effects play little role for blue collar workers, as shown by the narrow gap existing between the “Firm-related” and “Mover” plots. For white collar workers, there is a sharp increase of wage dispersion due to individual specific-components in the early years of the career, while after age 35 the evolution of inequality flattens, consistently with the patterns of empirical moments (Figure 5). Matching effects are also more relevant for skilled workers, as shown by the “Stayer” plot obtained in the hypothetical scenario of workers never changing employer.

We provide an additional characterization of differences in wage inequality across occupational groups by estimating the AKM-type model with constant wages by occupation; we report the predicted variance decomposition in Table 4, while we provide the underlying parameter estimates of permanent wage components in Table A2. Moving from the specification with life-cycle dynamics to the one with constant permanent wages reproduces the pattern observed for the overall sample also within occupations. For both groups life-cycle related components of the wage process (individual life-cycle component and sorting component) lose relevance in explaining overall wage inequality, while the firm and match component become predominant. In the case of blue collar workers, the constant model predicts that permanent individual shocks are even less important than transitory shocks in explaining wage dispersion.

Taken together, results of the heterogeneity analysis reveal a sharp contrast between occupational groups. For blue collar workers wage inequality is mostly a matter of where you work, while for white collars it also matters who you are, because to a great extent one’s position in the wage ladder depends on individual-specific effects.

7. Conclusion

Exploiting newly accessible data on the population of Italian employers and employees, we model the life-cycle wage dynamics of men allowing wages to depend on firm-specific, individual-specific and match-specific shocks. We develop a novel identification strategy based on the wage covariance structure of co-workers for isolating firm-related inequality.

We show that firm-related heterogeneity is a dominant factor for explaining wage differences among young workers, and that it takes time to idiosyncratic individual ability to display its inequality enhancing effects, both within and between firms. Within firms, shocks dispersion to individual wages declines with age, consistently with the exhaustion of employer learning. The sorting of workers into firms is characterised by strong life-cycle variation. We show that ignoring life-cycle wage dynamics leads to underestimate the importance of sorting in accounting for overall wage inequality.

Our results present some distinctive difference from other papers in the literature investigating the relevance of firm-related wage inequality without allowing for life-cycle dynamics (AKM models), namely our identification strategy attributes a greater share of wage inequality to workers-firms sorting and a lower share to individual idiosyncratic factors. However, when we use our strategy on models without life-cycle dynamics we reconcile the discrepancy between approaches, suggesting that life-cycle dynamics are an important feature of workers sorting into firms.

There is a sharp heterogeneity in the patterns of life-cycle wage inequality between occupational groups. For workers in low-skilled occupations, life-cycle wage differentials are rather flat and mostly traceable to firm-related factors. Conversely, in highly skilled occupations individual ability matters more, both at the start of the working career and throughout the life-cycle. Sorting is less prevalent within skill groups than in general, highlighting the importance of firms' occupational structure as driver of wage inequality.

References

- Abowd, John, Francis Kramarz, and David Margolis. 1999. "High Wage Workers and High Wahge Firms." *Econometrica*, 57(2): 411-45.
- Altonji Joseph G., Anthony A. Smith Jr., and Ivan Vidangos. 2013. "Modeling Earnings Dynamics." *Econometrica*, 81(4): 1395–454.
- Baker, Michael, and Gary Solon. 2003. "Earnings Dynamics and Inequality among Canadian Men, 1976-1992: Evidence from Longitudinal Income Tax Records." *Journal of Labor Economics*, 21 (2): 267-88.
- Blundell, Richard. 2014. "Income Dynamics and Life-Cycle Inequality: Mechanisms and Controversies." *Economic Journal*, 124: 289-318.
- Blundell, Richard, Michael Graber, and Magne Mogstad. 2015. "Labor Income Dynamics and the Insurance from Taxes, Transfers, and the Family." *Journal of Public Economics*, 127: 58-73.
- Borovičková, Katarina and Robert Shimer. 2017. "High Wage Workers Work for High Wage Firms", mimeo, November
- Cappellari, Lorenzo and Marco Leonardi. 2016. "Earnings Instability and Tenure." *Scandinavian Journal of Economics*, 118(2), 202–34.
- Card, David, Joerg Heining, and Patrick Kline. 2013. "Workplace Heterogeneity and the Rise of West German Wage Inequality", *Quarterly Journal of Economics*, 128: 967-1015.
- Card, David, Ana Rute Cardoso, and Patrick Kline. 2016. "Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women." *Quarterly Journal of Economics*, 131: 633-86.
- Card, David, Ana Rute Cardoso, Joerg Heining, and Patrick Kline. 2018. "Firms and Labor Market Inequality: Evidence and Some Theory", *Journal of Labor Economics*, 36, S13-70.

- Chamberlain, Gary. 1984. "Panel Data." in *Handbook of Econometrics*, Vol. 2, edited by Zvi Griliches and Michael D. Intriligator, 1248-1318. Amsterdam: North Holland.
- Corak, Miles. 2013. "Income Inequality, Equality of Opportunity, and Intergenerational Mobility" *Journal of Economic Perspectives*, 37(3): 79-102.
- Erickson, Christopher and Andrea Ichino. 1995. "Wage Differentials in Italy: Market Forces, Institutions, and Inflation." in *Differences and Changes in Wage Structures*, edited by Richard B. Freeman and Lawrence F. Katz, Chicago: University of Chicago Press.
- Friedrich, Benjamin, Lisa Laun, Costas Meghir, and Luigi Pistaferri. 2016. "Earnings Dynamics and Firm-Level Shocks." mimeo, July.
- Guiso, Luigi, Luigi Pistaferri, and Fabiano Schivardi. 2005. "Insurance within the Firm." *Journal of Political Economy*, 113(5): 1054-87.
- Haider, Steven J. 2001. "Earnings Instability and Earnings Inequality of Males in the United States: 1967-1991." *Journal of Labor Economics*, 19 (4): 799-836.
- Huggett, Mark, Gustavo Ventura and Amir Yaron. 2011. "Sources of Lifetime Inequality", *American Economic Review* 101: 2923–54.
- Lange, Fabian. 2007. "The Speed of Employer Learning," *Journal of Labor Economics*, 25(1): 1–35.
- Low, Hamish, Costas Meghir and Luigi Pistaferri. 2010. "Wage Risk and Employment Risk over the Life Cycle." *American Economic Review*, 100(4): 1432-67.
- Macis, Mario and Fabiano Schivardi. 2016. "Export and Wages: Rent Sharing, Workforce Composition or Returns to Skills?", *Journal of Labor Economics*, Vol. 34(4): 945–978.
- Manacorda, Marco. 2004. "Can the Scala Mobile Explain the Fall and Rise of Earnings Inequality in Italy? A Semiparametric Analysis, 1997-1993." *Journal of Labor Economics*, 22(3): 585-613.

- Meghir, Costas, and Luigi Pistaferri. 2004. "Income Variance Dynamics and Heterogeneity." *Econometrica* 72 (1): 1-32.
- Meghir, Costas, and Luigi Pistaferri. 2011. "Earnings, Consumption, and Life-cycle Choices." in *Handbook of Labor Economics*, Vol. 4, edited by Orley Ashenfelter and David Card, 773-854. Amsterdam: North Holland.
- Moffitt, Robert, and Peter Gottschalk. 2012. "Trends in the Transitory Variance of Male Earnings: Methods and Evidence." *Journal of Human Resources*, 47 (1): 204-36.
- Newey, Whitney. 1985. "Generalized Method of Moments Specification Testing" *Journal of Econometrics* 29 (3): 229-256.
- Page, Marianne, and Gary Solon. 2003. "Correlations between Brothers and Neighboring Boys in their Adult Earnings: The Importance of Being Urban." *Journal of Labor Economics*, 21 (4): 831-56.
- Song, Jae, David J. Price, Fatih Guvenen, Nicholas Bloom, and Till von Wachter. 2016. "Firming Up Inequality." mimeo, July.

Figure 1: Descriptive statistics

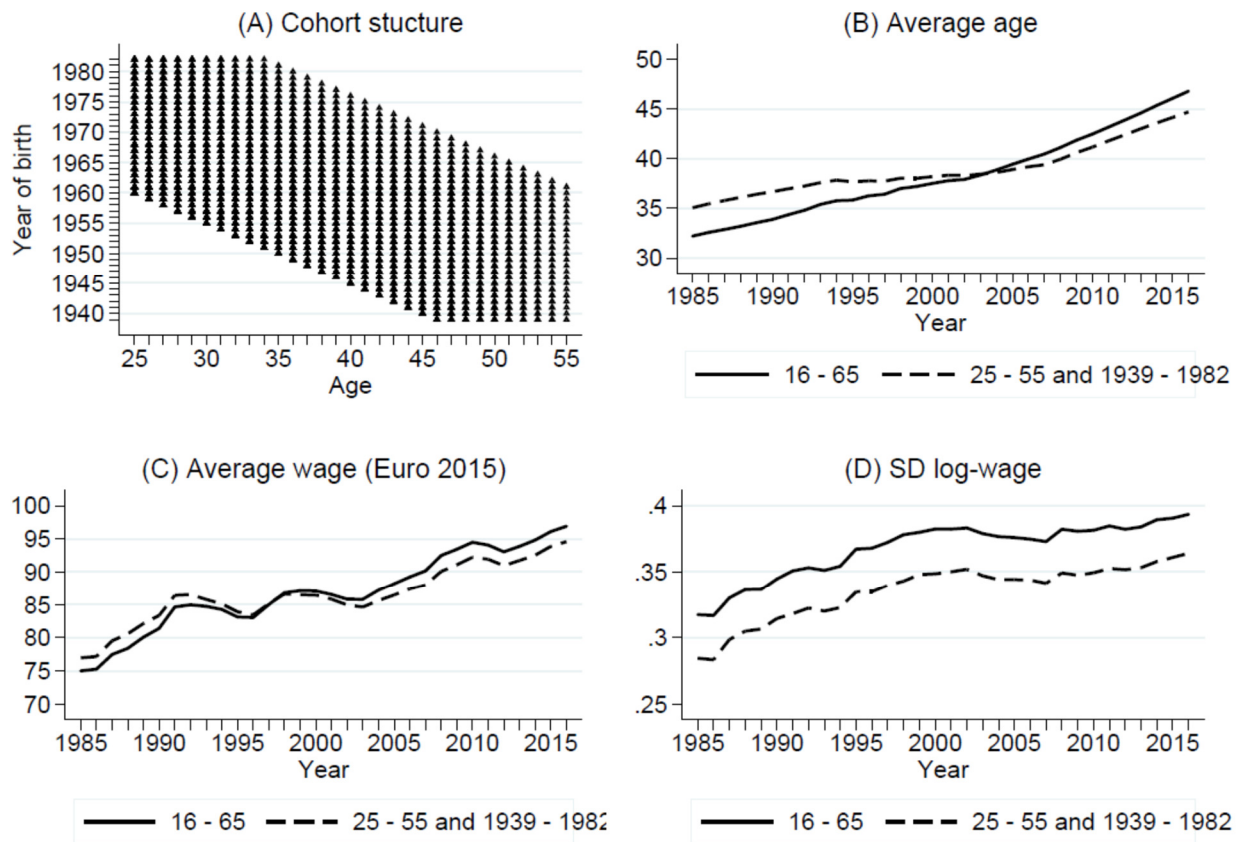


Figure 2: Empirical moments

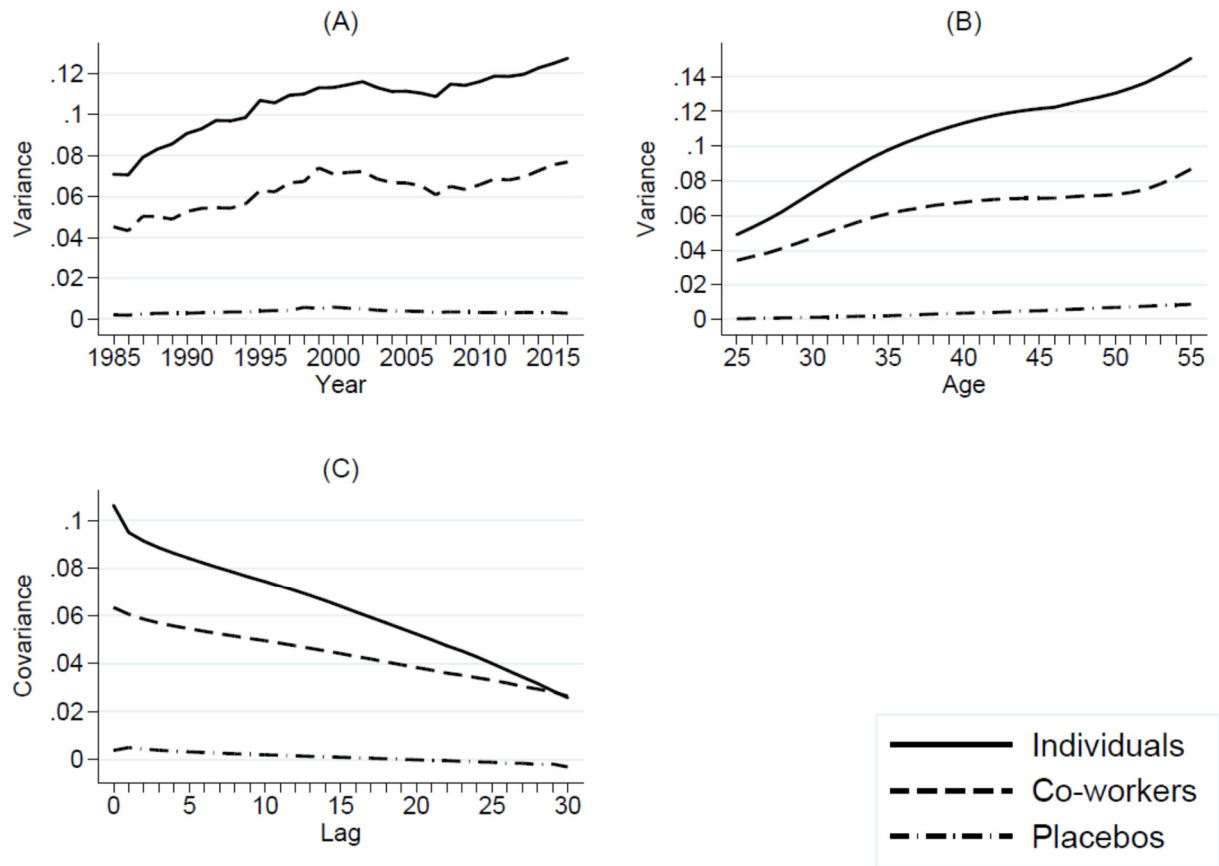


Figure 3: Variance decomposition, 1985 - 2016

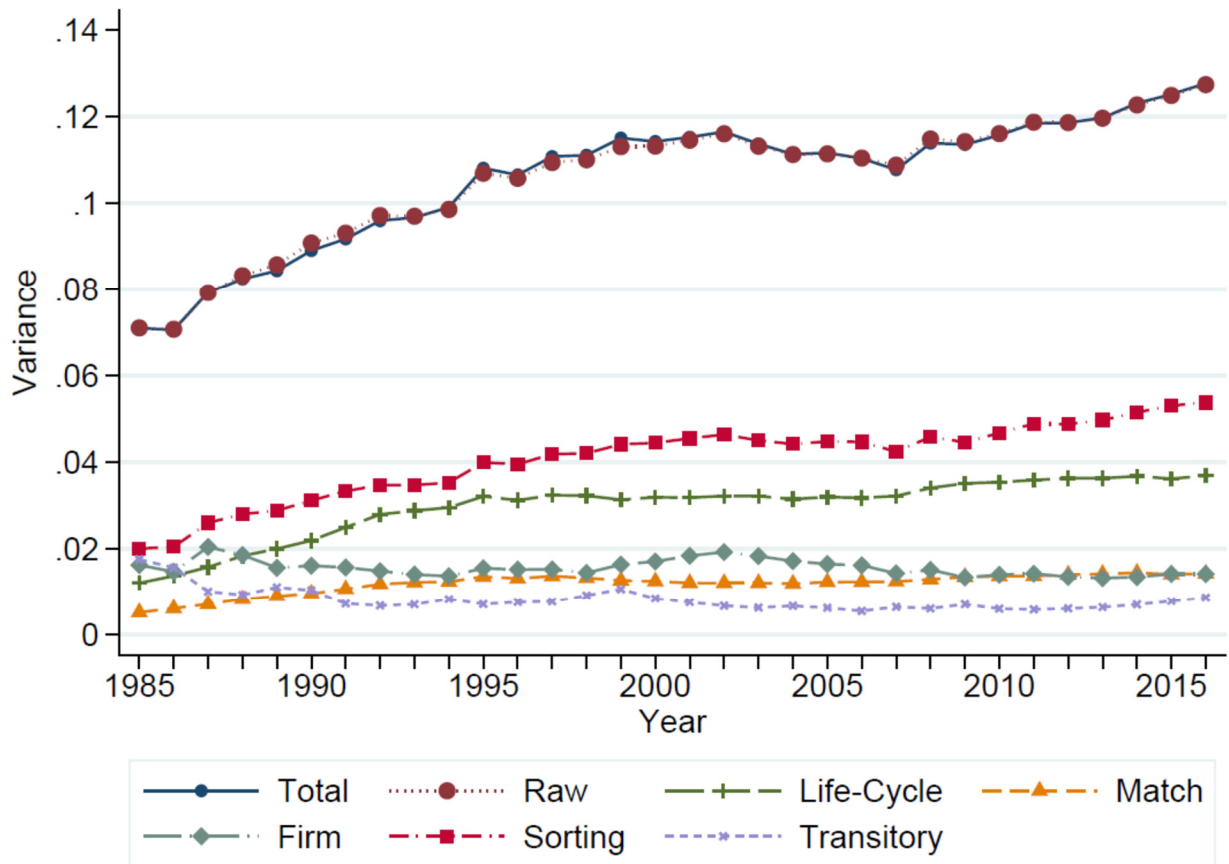


Figure 4: Life-cycle wage inequality

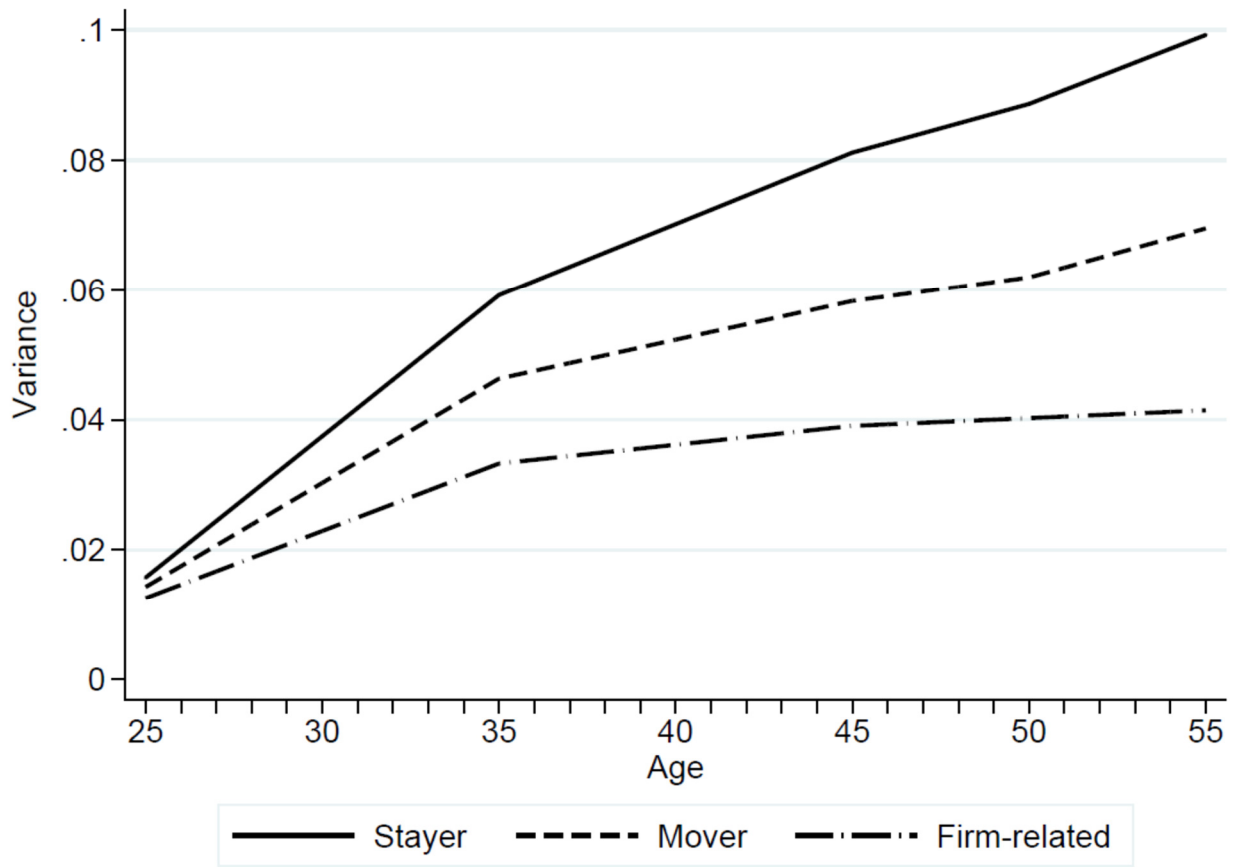
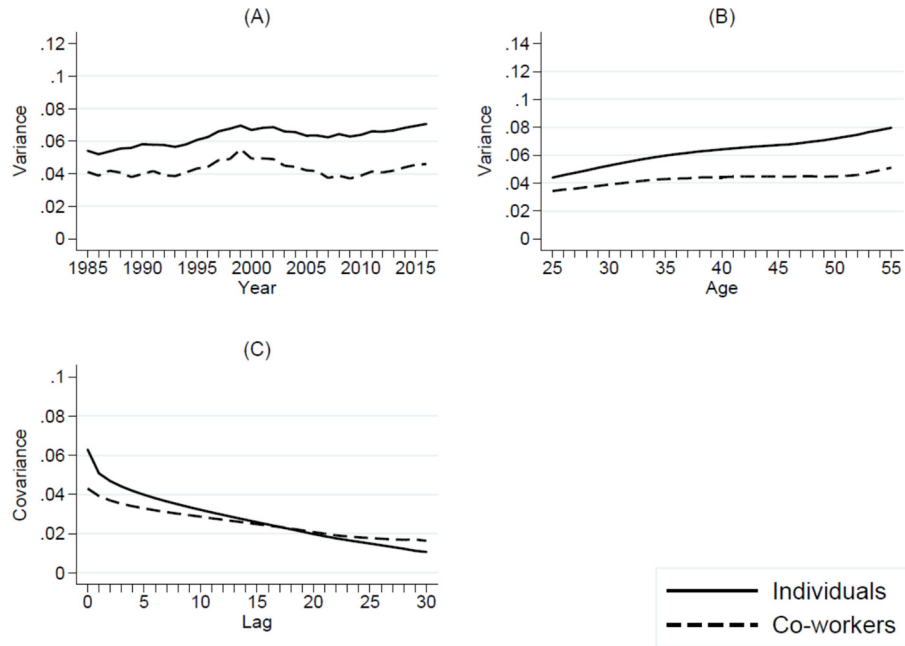


Figure 5: Empirical moments by occupation

I) Blue collar workers



II) White collar workers

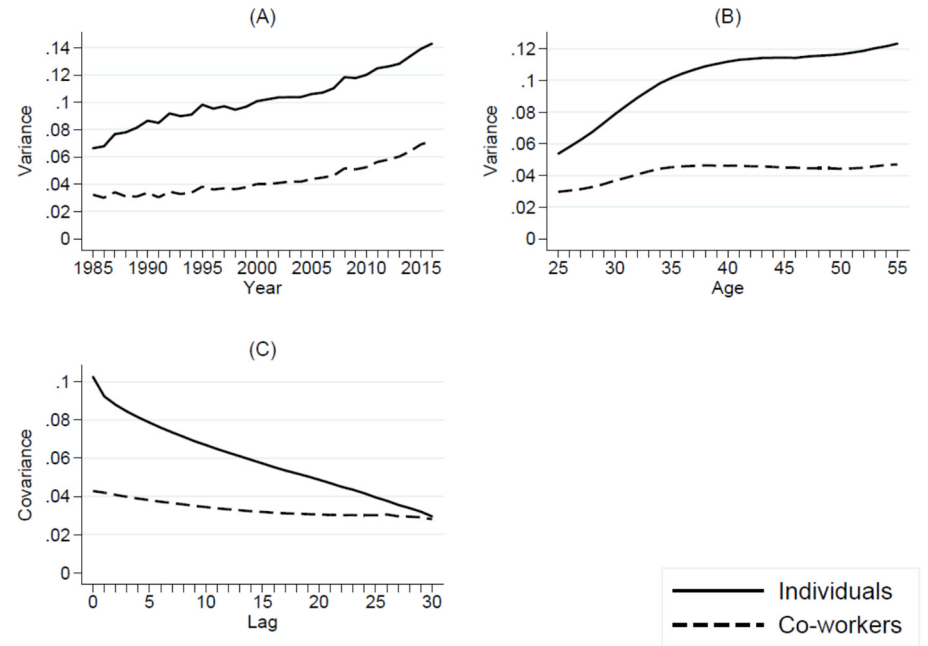
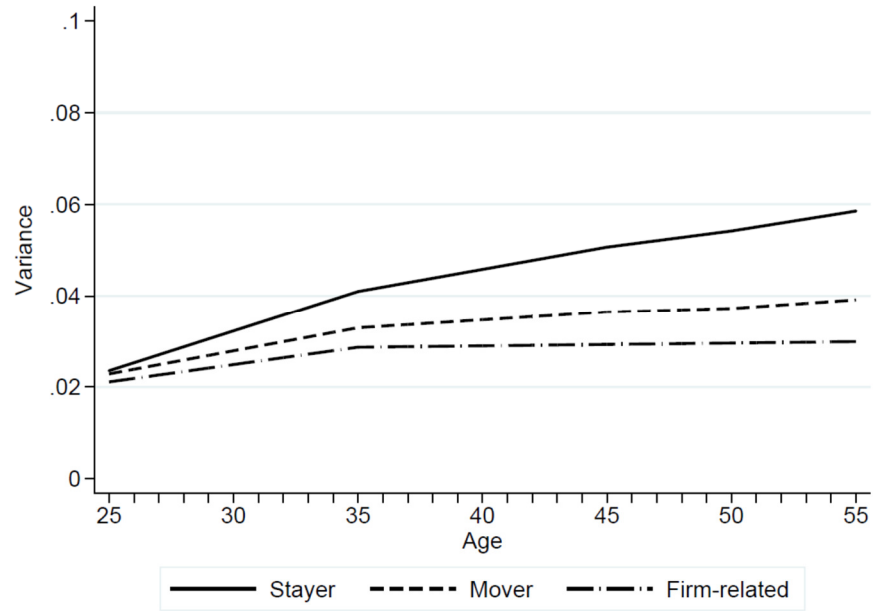


Figure 6: Life-cycle wage inequality by occupation

I) Blue collar workers



II) White collar workers

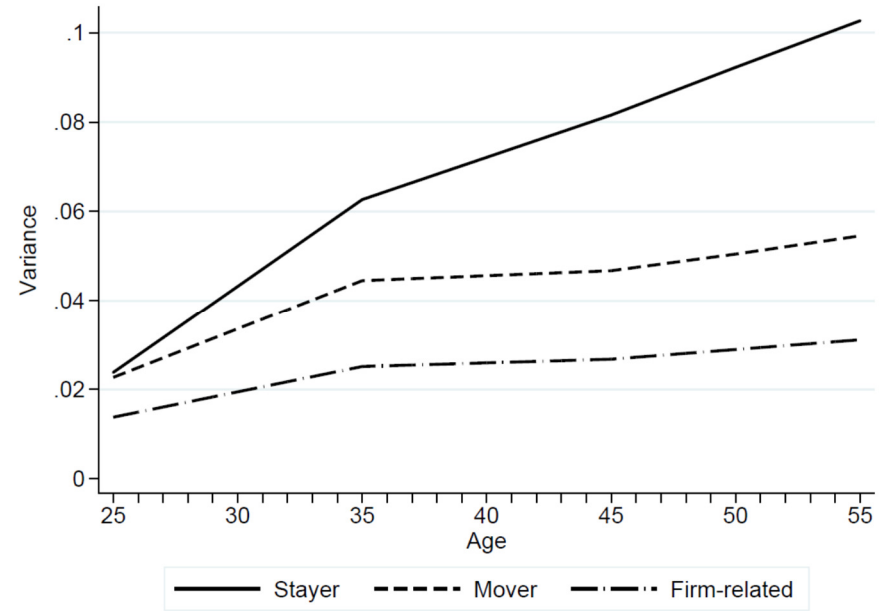


Table 1: Parameter estimates of permanent wage components

	Coeff.	S.E.x10
A) Life-cycle		
σ_{λ}^2	0.0017	0.0006
σ_{u26-35}^2	0.0011	0.0001
σ_{u36-45}^2	0.0006	0.0001
σ_{u46-50}^2	0.0005	0.0002
σ_{u51-55}^2	0.0013	0.0002
B) Match		
σ_{v25}^2	0.0015	0.0002
σ_{v26-35}^2	0.0011	0.0002
σ_{v36-45}^2	0.0010	0.0002
σ_{v46-50}^2	0.0008	0.0002
σ_{v51-55}^2	0.0006	0.0002
C) Firm		
$\sigma_{\phi young}^2$	0.0162	0.0011
$\sigma_{\phi middle}^2$	0.0098	0.0010
$\sigma_{\phi old}^2$	0.0074	0.0010
$\sigma_{\phi\phi young-middle}$	0.0045	0.0008
$\sigma_{\phi\phi young-old}$	0.0105	0.0009
$\sigma_{\phi\phi middle-old}$	0.0060	0.0008
D) Sorting		
$\rho_{\phi 25}$	0.0014	0.0003
$\rho_{\phi 26-35}$	0.0010	0.00003
$\rho_{\phi 36-45}$	0.0003	0.00003
$\rho_{\phi 46-55}$	0.0001	0.00004
Equally Weighted Minimum Distance (EWMD) estimates for the parameters of the permanent wage model of Section 3. Number of observations 152,470,973; number of individuals 12,339,989; number of firms 3,067,753; number of empirical moments 21,302; overall number of model parameters 119; $\chi^2(21183)=956692.64$.		

Table 2: Parameter estimates of transitory wage components

	Coeff.	S.E. x 10
A) Individual		
$\sigma_{\varepsilon 25}^2$	0.0223	0.0022
$\sigma_{\varepsilon 26-35}^2$	0.0125	0.0016
$\sigma_{\varepsilon 36-45}^2$	0.0045	0.0017
$\sigma_{\varepsilon 46-55}^2$	0.0012	0.0025
B) Firm		
$\sigma_{\xi young}^2$	0.0087	0.0012
$\sigma_{\xi middle-old}^2$	0.0271	0.0036
Equally Weighted Minimum Distance (EWMD) estimates for the parameters of the transitory wage model of Section 3. Number of observations 152,470,973; number of individuals 12,339,989; number of firms 3,067,753; number of empirical moments 21,302; overall number of model parameters 119; $\chi^2(21183)=956692.64$.		

Table 3: Sensitivity checks and heterogeneity by occupation: parameter estimates of permanent wage components

	No Sorting at age 25		Constant wages		Blue collar workers		White collar workers	
	Coeff.	S.E. x 10	Coeff.	S.E. x 10	Coeff.	S.E. x 10	Coeff.	S.E. x 10
A) Life-cycle								
σ_{λ}^2	0.0036	0.0004			0.0018	0.0005	0.0089	0.0011
σ_{u26-35}^2	0.0011	0.0001			0.0003	0.0001	0.0010	0.0001
σ_{u36-45}^2	0.0006	0.0001	0.0148	0.00071	0.0003	0.0001	0.0001	0.0001
σ_{u46-50}^2	0.0005	0.0002			0.0001	0.0001	0.0003	0.0002
σ_{u51-55}^2	0.0013	0.0002			0.0003	0.0002	0.0004	0.0003
B) Match								
σ_{v25}^2	0.0010	0.0002			0.0007	0.0002	0.0011	0.0004
σ_{v26-35}^2	0.0007	0.0001			0.0007	0.0002	0.0017	0.0003
σ_{v36-45}^2	0.0007	0.0001	0.0190	0.0011	0.0006	0.0001	0.0017	0.0003
σ_{v46-50}^2	0.0004	0.0001			0.0005	0.0001	0.0014	0.0003
σ_{v51-55}^2	0.0003	0.0002			0.0005	0.0001	0.0013	0.0003
C) Firm								
$\sigma_{\phi young}^2$	0.0190	0.0010			0.0231	0.0015	0.0165	0.0018
$\sigma_{\phi middle}^2$	0.0126	0.0009			0.0207	0.0017	0.0121	0.0019
$\sigma_{\phi old}^2$	0.0105	0.0009	0.0216	0.0011	0.0211	0.0018	0.0112	0.0018
$\sigma_{\phi\phi young-middle}$	0.0074	0.0006			0.0131	0.0012	0.0085	0.0015
$\sigma_{\phi\phi young-old}$	0.0133	0.0008			0.0194	0.0015	0.0121	0.0016
$\sigma_{\phi\phi middle-old}$	0.0089	0.0007			0.0167	0.0014	0.0097	0.0016
D) Sorting								
$\rho_{\phi 25}$					0.0002	0.0003	0.0008	0.0005
$\rho_{\phi 26-35}$	0.0010	0.00003	0.0037	0.0001	0.0004	0.00003	0.0006	0.0001
$\rho_{\phi 36-45}$	0.0003	0.00003			0.00003	0.00002	0.0001	0.0000
$\rho_{\phi 46-55}$	0.0001	0.00004			0.00003	0.0000	0.0002	0.0001

Table 4: Variance decomposition, 1985 - 2016

	(1) Life-cycle wages		(2) Constant wages	
	Variance component	Share of total	Variance component	Share of total
A) Full sample				
Life-Cycle	0.030	28.04	0.025	22.35
Match	0.012	11.23	0.032	28.58
Firm	0.016	14.72	0.038	33.97
Sorting	0.041	38.36	0.013	11.37
Transitory	0.008	7.64	0.004	3.73
Total	0.106	100.00	0.113	100.00
B) Blue collar workers				
Life-Cycle	0.011	18.31	0.003	4.45
Match	0.008	13.11	0.024	36.08
Firm	0.024	39.12	0.022	33.34
Sorting	0.011	17.41	0.008	12.76
Transitory	0.008	12.06	0.009	13.37
Total	0.062	100.00	0.066	100.00
C) White collar workers				
Life-Cycle	0.034	32.23	0.021	19.81
Match	0.021	19.76	0.045	42.42
Firm	0.021	19.99	0.027	25.63
Sorting	0.023	21.46	0.010	9.54
Transitory	0.007	6.59	0.003	2.60
Total	0.105	100.00	0.106	100.00

Table A1: Estimates of time shifters

Year (1985=1)	Individual and Match (α)		Firm (δ)		Transitory shock (γ)	
	Coeff.	S.E.	Coeff.	S.E.	Coeff.	S.E.
1986	1.05	0.002	0.95	0.003	0.95	0.003
1987	1.11	0.003	1.12	0.003	0.77	0.003
1988	1.18	0.003	1.11	0.003	0.68	0.003
1989	1.21	0.004	1.09	0.004	0.66	0.003
1990	1.26	0.004	1.12	0.004	0.63	0.003
1991	1.32	0.004	1.12	0.004	0.52	0.004
1992	1.38	0.004	1.11	0.004	0.49	0.004
1993	1.38	0.004	1.09	0.004	0.50	0.003
1994	1.38	0.004	1.10	0.004	0.53	0.003
1995	1.44	0.005	1.19	0.004	0.48	0.003
1996	1.42	0.005	1.19	0.004	0.49	0.003
1997	1.45	0.005	1.24	0.004	0.49	0.003
1998	1.44	0.005	1.25	0.004	0.53	0.003
1999	1.42	0.005	1.33	0.005	0.57	0.003
2000	1.44	0.005	1.33	0.004	0.51	0.003
2001	1.43	0.005	1.36	0.005	0.48	0.003
2002	1.44	0.005	1.38	0.005	0.46	0.003
2003	1.44	0.005	1.34	0.005	0.44	0.003
2004	1.43	0.005	1.33	0.005	0.45	0.003
2005	1.44	0.005	1.34	0.005	0.43	0.003
2006	1.43	0.005	1.34	0.005	0.40	0.003
2007	1.44	0.005	1.26	0.005	0.43	0.003
2008	1.46	0.005	1.31	0.005	0.43	0.003
2009	1.46	0.005	1.24	0.005	0.46	0.003
2010	1.45	0.005	1.28	0.005	0.43	0.003
2011	1.44	0.005	1.31	0.005	0.42	0.004
2012	1.43	0.005	1.29	0.005	0.43	0.004
2013	1.42	0.005	1.30	0.005	0.44	0.004
2014	1.41	0.005	1.33	0.005	0.47	0.004
2015	1.38	0.005	1.37	0.006	0.50	0.004
2016	1.38	0.005	1.37	0.006	0.53	0.005

Equally Weighted Minimum Distance (EWMD) estimates for the time shifters of the model of Section 3. Number of observations 152,470,973; number of individuals 12,339,989; number of firms 3,067,753; number of empirical moments 21,302; overall number of model parameters 119; $\chi^2(21183)=956692.64$.

Table A2: Estimates of wage components from wage constant model by occupation

	Blue Collar Workers		White Collar Workers	
	Coeff.	S.E.x10	Coeff.	S.E.x10
Life-cycle (σ_λ^2)	0.0019	0.0005	0.0117	0.0010
Match (σ_μ^2)	0.0154	0.0015	0.0251	0.0018
Firm (σ_ϕ^2)	0.0156	0.0013	0.0141	0.0015
Sorting (ρ)	0.0029	0.0002	0.0027	0.0003
Transitory (σ_ψ^2)	0.0176	0.0012	0.0113	0.0015