NBER WORKING PAPER SERIES

FAMILY AND GOVERNMENT INSURANCE:
WAGE, EARNINGS, AND INCOME RISKS IN THE NETHERLANDS AND THE U.S.

Mariacristina De Nardi

Giulio Fella

Marika G. Knoef

Gonzalo Paz-Pardo

Raun Van Ooijen

Working Paper 25832

<http://www.nber.org/papers/w25832>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue

Cambridge, MA 02138

May 2019

The authors thank NETSPAR for supporting this research project. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research, any agency of the federal government, or the Federal Reserve Bank of Minneapolis.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2019 by Mariacristina De Nardi, Giulio Fella, Marika G. Knoef, Gonzalo Paz-Pardo, and Raun Van Ooijen. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Family and Government Insurance: Wage, Earnings, and Income Risks in the Netherlands and the U.S.

Mariacristina De Nardi, Giulio Fella, Marike G. Knoef, Gonzalo Paz-Pardo, and Raun Van Ooijen

NBER Working Paper No. 25832

May 2019

JEL No. H31

ABSTRACT

We document new facts on the distributions of male wages, male earnings, and household earnings and income (before and after taxes) in the Netherlands and the United States. We find that, in both countries, wages display rich dynamics, including substantial asymmetries and nonlinearities by age and previous earnings levels. Individual-level male wage and earnings risk is relatively high for younger and older people, and for those in the lower and upper parts of the income distribution. In the Netherlands, the behavior of hours and family labor supply have noticeable effects on earnings persistence and on the skewness and kurtosis of wage changes, but government transfers are a major source of insurance. Instead, the role of family insurance is much larger in the U.S. and also affects the standard deviation of wage changes, in addition to its skewness and kurtosis, and wage persistence. Family and government insurance reduce, but do not eliminate these non-linearities in household disposable income by age and previous earnings in both countries.

Mariacristina De Nardi
Federal Reserve Bank of Minneapolis
90 Hennepin Ave
Minneapolis, MN 55401
and University College London
and also NBER
denardim@nber.org

Gonzalo Paz-Pardo
University College London
Drayton House
30 Gordon St
London WC1H 0AX
United Kingdom
gonzalo.pardo.13@ucl.ac.uk

Giulio Fella
Queen Mary University of London
Mile End Road
London E1 4NS
United Kingdom
and Centre for Macroeconomics - CFM
and also Institute for Fiscal Studies - IFS
g.fella@qmul.ac.uk

Raun Van Ooijen
University of Groningen
P.O. Box 800
9700 AV Groningen
The Netherlands
r.van.ooijen@rug.nl

Marike G. Knoef
Leiden University
Steenschuur 25
2311 ES, Leiden
The Netherlands
and Netspar
m.g.knoef@law.leidenuniv.nl

1 Introduction

Wage risk affects key economic decisions, including consumption, saving, and labor supply, and is an important determinant of household's welfare. Households can self-insure against these shocks. That is, single people can adjust their own labor supply and savings, while couples can adjust the labor supply of both partners, in addition to savings. Furthermore, governments can supplement or partly replace the need for self-insurance through progressive taxes and transfers.

This paper studies the size and distribution of wage shocks and the role of insurance mechanisms against these shocks in the Netherlands and the U.S. We start by documenting the distribution of wage shocks at the individual level by analyzing distributional measures of wage changes, including the standard deviation, skewness, kurtosis, and persistence, by age and previous earnings. To understand the role of individual-level labor supply, we compare the distribution of individual wage shocks with that of individual-level earnings. To analyze the role of family insurance through the labor supply of both partners, we compare the distributions of individual-level and household-level earnings. To examine the role of government insurance, we compare the distribution of household income, pre- and post-taxes, and transfers, by age group and previous earnings.

Our results show clear evidence of non-linearity and age dependence in both countries, with high wage and earnings risk for the people with lowest and highest wages and earnings. In the Netherlands income risk is lower for the broader middle of the income distribution. Wage and earnings persistence vary a lot by age in both countries. Comparing family and government insurance, we find that in the U.S. the role of the family in reducing risk is much more important, while in the Netherlands the government plays a much larger role in reducing risks, particularly through transfers (rather than taxes). This could indicate transfers and perhaps taxes might crowd out the insurance provided by the family.

Our high-quality administrative data on income, taxes, and government transfers on individuals and households for the Netherlands (IPO) enables us to get precise estimates of the dynamics of wage shocks and the role of private and public insurance mechanism to mitigate these shocks. Estimates from a household survey from the Netherlands (DNB Household Survey) show similar results. We compare the results with estimates for the U.S. Panel Study of Income Dynamics (PSID).

This study builds upon papers by Guvenen, Karahan, Ozkan and Song (2016), Arelano, Blundell and Bonhomme (2017), Halvorsen, Holter, Ozkan and Storesletten (n.d.), De Nardi, Fella and Paz-Pardo (2019), and Busch, Domeij, Guvenen and Madera (2018) which show that the distribution of earnings shocks display rich dynamics and, particularly, depend on age and previous earnings.

Our contribution to the literature is fourfold. First, whereas previous studies mainly investigated shocks in individual earnings, we distinguish between shocks in wages and changes in hours worked. Second, we investigate the degree of insurance provided by spousal labor supply (by comparing individual earnings and total earnings at the household level) and insurance provided by the tax and transfer system (by comparing pre- and after-taxes household income). Third, we compare two countries: the Netherlands and the U.S. This is an interesting comparison because these two countries differ substantially in the size of their welfare state and the progressiveness of their tax system.¹ Fourth, our analysis provides data that rich models of risks and insurance should match to be consistent with the key features of the micro-data that we document.

The remainder of the paper proceeds as follows. Section 2 describes our approach. Section 3 presents the data, after which sections 4 and 5 present the results. Section 6 concludes.

2 Approach

As standard in the literature, we purge age and time effects from log wages by running the following regression and identifying its estimated residuals as wage shocks

$$\log w_{it} = \beta_1 \text{age}_{it} + \beta_2 \text{age}_{it}^2 + \alpha_t + u_{it}. \quad (1)$$

The subscript i refers to an individual, t is year, α_t represents year fixed effects, and the error term u_{it} captures the stochastic component of wages.

Because the widespread modeling of wage shocks as an AR(1) process implies strong

¹Although eligibility has become more restrictive over the past two decades, the Dutch welfare system is one of the most comprehensive in Europe (see Kalwij, Kapteyn and de Vos, 2018 for a detailed and up-to-date description of social security reforms in the Netherlands). The OECD Social Expenditure Database 2016 shows that public social expenditure on family support, disability, unemployment and active labor market policies as a percentage of GDP is twice as high in the Netherlands compared to the U.S.

restrictions on wage changes that previous work has found to be violated in the data, rather than making this functional form assumption, we compute key moments of wage shock changes ($\hat{u}_{it} - \hat{u}_{it-1}$), including their standard deviation, skewness, and kurtosis.

These moments derive from interesting and important economic mechanisms. For instance, negative skewness can come from a job ladder model (see, for instance, Graber and Lise (2015)) in which people staying on the job experience small wage raises most of the time, but people losing their job often experience a large wage and earnings drop. This kind of model can also explain some kurtosis: most wage changes are small, but then there is a small fraction of people experiencing large wage changes, due to job loss, or job and career switches, for instance. In addition, the persistence of these wage changes might depend on one's age and current earnings. A young worker is more likely to switch jobs and careers to figure out what he or she is best at, which tends to lower the persistence of their wage changes. An old worker might switch to a part-time or less demanding job, thus also having lower wage shocks persistence. Finally, earnings persistence might depend also on previous earnings. For instance, high earners might be experiencing more wage uncertainty.

To measure skewness, we compute the conventional measure of skewness (Pearson's or standardized third moment). Because this measure is very sensitive to outliers (deviations from the mean are cubed), we also compute the robust Kelley's coefficient of skewness, which is given by

$$S_K = \frac{P_{90} + P_{10} - 2P_{50}}{P_{90} - P_{10}}, \quad (2)$$

where a zero implies a symmetric distribution, positive values represent right skewness, and negative values represents left skewness.

To measure kurtosis, we start with the conventional measure given by the fourth standardized moment, but we also compute the robust Crow-Siddiqui kurtosis which is given by

$$S_{CS} = \frac{P_{97.5} - P_{2.5}}{P_{75} - P_{25}}. \quad (3)$$

The term S_{CS} is large if $P_{97.5} - P_{2.5}$ is large relative to the probability mass that is concentrated between P_{75} and P_{25} , corresponding to heavy tails.

Finally, we analyze persistence by age, by regressing $\hat{u}_{it+1}$ on $\hat{u}_{it}$ for different ages.

To investigate insurance mechanisms, after studying wages, we repeat the analysis for individual-level earnings, household earnings, household pre-tax income (earnings and income from savings) and household after-tax (disposable) income. The comparison of wages and earnings is informative about self-insurance through labor supply. The comparison of individual-level and household-level earnings is informative about family insurance through the labor supply of the spouse. The comparison between household pre-tax income and household disposable income helps shed light on the role of insurance by the government through transfers and progressive taxation.

3 Data and sample selection

This section describes the data, the derived income measures, and our sample selection criteria. We use administrative tax records from the Dutch Income Panel Study (IPO) for the Netherlands and household survey data from the Panel Study of Income Dynamics (PSID) for the U.S.

The IPO data set contains a representative sample of about 95,000 individuals, randomly selected by Statistics Netherlands based on their national security number and followed over time, together with their household members. Detailed information is available on income, wealth (as from 2005), gender, age, marital status, children, ethnicity, home ownership, and labor market status. Data on income is available in the IPO for the years 1989 to 2014, but the data on hours worked is available from 1999. Because of a major tax reform that took place in 2001, which affected income definitions and therefore the measures that we use, we use data from 2001.

The IPO data set presents several important advantages. First, the data is often collected from or checked with a third party. For instance, the wealth measures are derived from tax records from the Dutch National Tax Administration complemented with information provided by banks and other financial institutions. The data on hours are obtained by linking the IPO data to social insurance records on hours worked, reported as the proportion of a yearly full-time job. In addition, Statistics Netherlands performs several checks on the data to guarantee its quality. This drastically reduces or even eliminates measurement error and errors due to non-reporting. Second, individuals are followed for as long as they are residing in the Netherlands (as of December 31 of the

sample year). We thus have little to no endogenous panel attrition. Panel attrition only occurs as a result of migration or death. New panel members enter the panel for the first time in the year of their birth, and immigrants to the Netherlands in the year of their arrival. Third, and very importantly, the IPO data set contains a detailed decomposition of labor and asset income, taxes and social insurance premia paid, and government transfers received for all household members. It also contains a detailed transfers breakdown, including unemployment insurance, disability insurance and social assistance. These features of the data allow us to measure the value of both family and government insurance. Thus, to summarize, the main advantages of IPO are its large sample size, long panel dimension, data accuracy, representativeness, and richness.

Adding up the detailed income components in all data sets we derive the following income definitions: individual gross wages, individual gross earnings, household gross earnings, household pre-tax (primary) income, and household after-tax (disposable) income. Individual gross earnings is defined as paid work on a contractual basis. We exclude self-employed work. Gross earnings include employee's contribution for health insurance and premia for unemployment and disability benefits and pensions. Household gross earnings equals total earnings of all household members. Household pre-tax income equals the sum of earnings and income from savings. Household after-tax income equals household pre-tax income minus taxes on income and social security premia plus allowances (healthcare, rent, child and childcare, study costs, and alimony) and transfers. Transfers are the sum of employment benefits (including benefits for civil servants and supplementary benefits for unemployed elderly), disability benefits (including benefits for persons who were already disabled at the age of 17), social assistance and pension benefits.

Our household survey for the Netherlands, the DHS, was launched in 1993 and is a representative Internet-based panel of over 2000 households administered by CentERdata at Tilburg University and sponsored by The Dutch Central Bank. We use the DHS to confirm that the patterns that we document for the Netherlands in our administrative data set also hold in survey data for the Netherlands over the period 2001 to 2014. Given that our data for the U.S. comes from a household survey, this reassures us that the differences that we document across countries are not due to the nature of the data set but rather to interesting economic differences across countries.

Turning to the U.S. data, the PSID began in 1968 with a representative sample of 18,000 individuals living in 5,000 families. We use it for the period 1968 to 1992. We exclude the years 1993-1997, because of a major redesign of the survey and those after 1997 because the PSID became bi-yearly after that date. To confirm that the results are not driven by the different time periods for the U.S and the Netherlands (see Heathcote, Perri and Violante, 2010 for a discussion of changes in the distribution of wages and earnings in the U.S. across this period of time), we also study our statistics of interest for the period after 1997 for two-years income changes in both countries. This robustness check shows that the cross-country differences that we document come from different cross-country features and not from comparing different sample periods.

For each data set, we select a sample of male earners age 25 to 60 to abstract from education and retirement decisions. We exclude self-employed workers (that is those for whom income out of self-employment is their main income source) and impose a lower limit for labor earnings of 2720 dollar a year (2200 euro) (in 2014 prices) to include people with a minimum of attachment to the labor market. We equalize household earnings, and household pre- and after-tax income.

4 Results: Netherlands

In this section, we first discuss the properties of male wage changes. Then, we compare them with those of male earnings, household earnings, pre-tax household income and disposable household income. We also discuss what they imply in terms of family and government insurance.

4.1 Male wages

The top left hand side of Figure 1 displays wage persistence by age and shows large age variation in male wage persistence (unlike typically assumed by a standard AR(1) process). Wage persistence starts from a low value of 0.6 at age 25, consistently with younger people switching jobs and careers to figure out what job is the best fit for them. Many of them also have temporary contracts. It then increases fast, reaching 0.85 at age 35, and gradually keeps increasing to 0.9 until age 45. Only after that age, it remains flat. Thus, bad (and good) wage shocks at younger ages are not as long-lived as they

Figure 1: Dutch male wages. Wage persistence (top left) and following moments of wage changes: standard deviation (top right), skewness (middle left), Kelley's skewness (middle right), kurtosis (bottom left), and Crow-Siddiqui kurtosis (bottom right), by age group and previous earnings percentile.

would be if their persistence were much higher, as assumed by many models evaluating policy intervention and taxation.

The top right hand side of Figure 1 displays the standard deviation of wage changes by age group and previous earnings.² Several features of the data are worth noticing. First, the variability of wage changes is over three times larger at the lowest percentiles of previous earnings (0.87) than for workers in the middle percentiles of previous earnings (.23). Second, this variability increases by a factor of about two (from 0.23 to almost 0.5) from previous median earnings to percentiles of previous earnings above the 80th percentile. Third, workers with previous earnings below the median and in the youngest age group (25-34) experience the largest wage change volatility, which goes down conditional on these previous earnings levels until age 54. Fourth, volatility by previous earnings starts rising again after age 55 and, among workers above the 30th percentile of previous earnings, it is highest for the oldest workers.

The patterns for wage volatility for young workers at lower previous earnings be may due, at least partly, to flexible contracts among young workers. In contrast, workers at late stages of their career might have a higher prevalence of absences due to longer-lived health problems. In the Netherlands, during the first two years of sick leave, the employer is required to continue paying at least 70% of their wage. After that, one may be eligible for disability benefits.

The middle left hand side of Figure 1 plots the skewness of wage changes and shows that wage skewness starts around zero at low levels of previous earnings but becomes more and more negative as previous earnings increase to the 40th percentile, to then flatten out between -3 and -5 after that. This means that workers with higher previous earnings are much more likely to experience a negative wage change, or a wage drop, rather than a positive one, or wage increase. Also, young workers with previous earnings below the 50th percentile, have less negative skewness than their older counterparts, indicating that more of them might be climbing up the career ladder and thus being less likely to experience negative wage changes.

The conventional measure of skewness can be difficult to interpret because it is very sensitive to the tails of the distribution. Kelley's skewness (middle right hand side panel of this figure) is a measure that is robust to outliers. It turns out to be rather flat at

²In all figures we use the same horizontal axis, which is the distribution of previous earnings. Using previous wages on the horizontal axis of Figure 1 would not change the results.

zero for most age groups and most percentiles of the earning's distribution; indicating a symmetric distribution of wage shocks outside of the tails of wage changes. Only the oldest age group, in fact, displays negative skewness, indicating that negative wage shocks are more prevalent than positive ones in this group, for example because they face demotions or decide to step back their work efforts, for instance due to a long-term sickness affecting their wages. The Kelley measure for the oldest age group is rather flat for previous earnings above the 10th percentile; a number of -0.2 means that the lower tail of the distribution (negative shocks) accounts for 60 percent of the overall dispersion (between the 90th and 10th percentiles) and the upper tail (positive shocks) account for the remaining 40 percent.

The bottom left hand side of Figure 1 shows the kurtosis of wage changes. It is increasing by previous earnings percentiles, indicating that wage changes become less prevalent but larger at higher percentiles of previous earnings. Workers in their prime working lives (aged 35-54) have the highest kurtosis and thus face the distribution of wage changes with the fattest tails. Because the kurtosis is also sensitive to the tails, we also study a robust version of it.

The Crow-Siddiqui kurtosis (bottom right panel) displays a slightly higher kurtosis for workers at the bottom of previous earnings (up to the 10th percentile), but then the kurtosis drops and becomes flat by previous earnings, indicating that, while workers at or close to minimum wage experience small but more frequent wage shocks, when people move beyond the minimum wage, these shocks become larger but much less prevalent. This pattern is particularly strong for workers in the older age group. This may be due to less flexible contracts and higher employment protection for older age groups which might make wages of older and higher income workers more rigid.

Interestingly, comparing our standard and robust measures of kurtosis and skewness reveals that, outside of the tails of wage changes, most of the skewness is experienced by older workers and that the pattern that the kurtosis of wage changes is increasing by age is more evident outside the tail below the 25th percentile of previous earnings, but still present for all of the distribution for most age groups. The increasing kurtosis by age could be explained by positive wage changes associated with career switches up until people's primary working years, whereas older workers (with long job tenures) are more likely to experience relatively large wage cuts when they find a new job after displacement.

Thus, Figure 1 shows strong evidence in favor of age-variation, non-linearity and non-normality of wage changes: wage persistence is the lowest for the youngest; wages are more variable for the highest and lowest earnings; there are hardly any shocks for most and large shocks for some; large negative shocks are more likely than large positive shocks.

4.2 Male wages, male earnings, and household earnings

The literature has focused on the features of earnings shocks. A natural question is whether these come from wages or hours. Figure 2 compares our statistics for wages (left panel) with those for male earnings (middle panel) and household earnings (right panel).

The top panel compares persistence. There are several noticeable differences. First, the persistence of wages before age 35 is much lower than the persistence of male earnings at the same age, which in turn is somewhat less persistent than household income. Second, this persistence increases much more rapidly by age for wages than for earnings, and in turn, even less rapidly for household earnings than for male earnings. Thus, male labor supply and, to some extent, the labor supply of the secondary earner, generate much more persistent earnings than wages and less of an increase in persistence over the life cycle. This is an important feature of the data that should be matched in a model aiming at understanding and affecting wage and earnings dynamics and their inequality.

The middle panel compares male wages with earnings. The figures are very similar; only for the oldest age group the variability in earnings changes is slightly higher than the variability in wage changes, also the skewness is slightly more negative. This pattern could be explained by a higher prevalence of job-loss among the oldest age group.

The bottom panel of the figure shows that the labor supply of the secondary earner plays an important role in reducing the kurtosis of household earnings compared to that of male earnings and wages. For older workers we find that households' labor supply makes Kelley's skewness somewhat less negative for household earnings than for male earnings. These features of the data suggest an increase in the labor supply of women when their husbands experience a negative wage shock (added worker effect). For young workers we find more negative skewness for household earnings compared to male earnings. This could reflect female spouses who reduce working hours after the birth of children. Thus, even in the Netherlands, we find that household labor supply does affect the risks that households face.

Figure 2: NL, male wages (left), male earnings (middle), and household earnings (right). Persistence (top row), standard deviation (second row), Kelley's skewness (third row), and Crow-Siddiqui kurtosis (bottom row).

4.3 Household pre-tax income and household disposable income

To investigate the role of government insurance, Figure 3 compares household pre-tax income (left panel) with disposable income (right panel) for the Netherlands.³

While earnings persistence is much more affected by the labor supply of the primary and secondary earner, taxes and transfers make a huge difference for the other statistics that we study, especially at the lower end of the income distribution and for older households.

For disposable income, the standard deviation and the kurtosis of the income shocks decline and skewness becomes almost zero. For instance, the standard deviation of household income changes at the lowest percentiles of previous earnings declines from about 0.75 before taxes and transfers to a little over 0.37. It also becomes much more compressed towards lower levels at the highest percentiles of previous earnings, and especially for the older workers. The reduction in Kelley's skewness is especially apparent for the workers in the oldest age group. The Crow-Siddiqui kurtosis further drops from about 8 at the household level (it peaked at about 17 for wages and male earnings) to well below 7 after taxes and transfers.

This picture makes clear that the government provides a lot of insurance in the Netherlands, above and beyond that already provided by the family. Progressive taxation reduces earnings variability and the benefit system (unemployment and disability insurance and welfare) reduces earnings variability and effectively eliminates large negative shocks which eliminates negative skewness and reduces the kurtosis.

Given that government insurance is especially prevalent in the Netherlands and especially so at older ages, Figure 4 further breaks down the role of various government programs for our 55 to 59 age group, by sequentially adding specific transfer programs or taxes. The graphs show that disability insurance greatly reduces the standard deviation of household earnings changes below the 20th percentile of previous earnings, while unemployment insurance generates a significant reduction even at higher levels of previous earnings. It also shows that, for this age group, pension transfers play a much larger role in reducing variation in household income than taxes.

³Household pre-tax income contains earnings and income from savings. In Appendix A we show that allowing for capital income makes little difference for household income dynamics.

Figure 3: NL, household pre-tax (left), and disposable income (right). Persistence (top row), standard deviation (second row), Kelley's skewness (third row), and Crow-Siddiqui kurtosis (bottom row).

Figure 4: NL, age 55-59, Relative contribution of transfers and taxes to the standard deviation of household income. Red line, household gross income, gold line: including disability insurance, green line: also including unemployment insurance, dotted green line: also including social assistance, dotted blue line also including pensions, dotted red line: also net of taxes.

5 Results: Netherlands versus U.S.

Figure 5 compares our summary statistics for the Netherlands and the U.S., with the figures on the left summarizing the results for the Netherlands for all age groups together. For the Netherlands, these graphs confirm and clarify our results conditional on age-groups. That is, the various moments show that, in the Netherlands, labor supply does not have a large effect on the standard deviation of wage changes. In contrast, skewness and kurtosis are both affected by the labor supply of the primary and secondary earner in the household. Notably, the choice of hours of the primary earner increases the negative skewness and the kurtosis in earnings changes compared to wage changes (over a wide range or percentiles of previous earnings), while the labor supply of the secondary earner reduces them. In the Netherlands taxes and transfers have large effects on all of the summary statistics that we consider and reduce both risk and inequality in wages and earnings. The patterns that we observe are consistent with the primary earner taking on more earnings risk in terms of additional negative skewness and larger kurtosis while being able to count both on spousal insurance through labor supply and insurance through the government via taxes and transfers to help insure against the risks related to these choices.⁴

⁴For comparison, we examined income dynamics using survey data from the DHS for the Netherlands. The results are quantitatively very similar. The main difference is that the DHS data only shows a

The standard deviations of all income measures are higher in the U.S than in the Netherlands, except at very low previous earnings. Turning to the various moments, we find that the standard deviation of male earnings is higher than that of male wages at higher and lower levels of previous earnings, indicating volatility in hours that amplifies the dispersion in wage changes in the U.S. We also find much more of a role in spousal labor supply in the U.S. in reducing the standard deviation of male earnings over all levels of previous earnings. In terms of skewness, in contrast to the Netherlands, we find that the labor supply behavior of the primary earner increases the negative skewness of male wages (and thus increases the chance of a negative wage change) over a large range of previous earnings, while that of the secondary earner reduces it. Also, the labor supply of the primary earner increases the kurtosis of earnings changes. Our results indicate that while the labor supply of the primary earner reduces the variance of male earnings changes compared to that of male wages, it increases tail risks, both in terms of negative skewness and a larger kurtosis. The labor supply of the secondary earners, in contrast, tend to compress both the volatility and the tails of the household earnings distribution in the U.S. Whereas the government accounts for a very large fraction of the insurance of disposable earnings in the Netherlands, the family contributes much more as a source of insurance in the U.S.

To confirm that results are driven by cross-country differences and not by period of observation, we also examined income dynamics for the PSID in the post 1997 period covering the same time frame as the IPO data. In Appendix B we show that the results are very similar, if anything, the role of family insurance has slightly reduced over time in the U.S, indicating that the results are driven by cross-country differences and not by the nature of the data set.

negligible difference between gross-income and net-income. The small difference between gross-income and net-income might be explained by the different definition of gross-income in the DHS which does not include social security contributions from employers.

Figure 5: NL (Left), U.S. (right)

Figure 6: NL (Left), U.S. (right) Robust measures of Kurtosis and Skewness

6 Conclusions

The results show clear evidence of non-linearity and age dependence of wages and earnings in both countries, with higher risk for the lowest and highest earners. In the Netherlands income risk is relatively low for the broader middle of the income distribution. Wage and earnings persistence vary a lot by age in both countries. Comparing family and government insurance we find that the government plays a much larger role in reducing wage risk in the Netherlands, while in the U.S. the role that the family plays is much more important. The results suggest that taxes and transfers may crowd out insurance that could be made available within the family.

References

- Arellano, Manuel, Blundell, Richard and Bonhomme, Stéphane (2017), ‘Earnings and consumption dynamics: A non-linear panel data framework’, *Econometrica* **85**(3), 693–734.
- Busch, Christopher, Domeij, David, Guvenen, Fatih and Madera, Rocio (2018), Asymmetric business-cycle risk and social insurance, Working Paper 24569, National Bureau of Economic Research.
- De Nardi, Mariacristina, Fella, Giulio and Paz-Pardo, Gonzalo (2019), ‘Nonlinear household earnings dynamics, self-insurance, and welfare’, *Journal of the European Economic Association* . Forthcoming.
- Graber, Michael and Lise, Jeremy (2015), ‘Labor market frictions, human capital accumulation and consumption inequality’, mimeo, University College London.
- Guvenen, Fatih, Karahan, Fatih, Ozkan, Serdar and Song, Jae (2016), What do data on millions of U.S. workers reveal about life-cycle earnings risk? Working paper, University of Minnesota.
- Halvorsen, Elin, Holter, Hans, Ozkan, Srdar and Storesletten, Kjetil (n.d.).
- Heathcote, Jonathan, Perri, Fabrizio and Violante, Giovanni L (2010), ‘Unequal we stand: An empirical analysis of economic inequality in the united states, 1967–2006’, *Review of Economic dynamics* **13**(1), 15–51.
- Kalwij, Adriaan, Kapteyn, Arie and de Vos, Klaas (2018), *Why Are People Working Longer in the Netherlands?*, University of Chicago Press.

A Household earnings and pre-tax (primary) income

Figure 7: NL, household earnings (left) and household pre-tax (primary) income (right)

B Two-year changes in the Netherlands and the U.S.

Figure 8: Netherlands, IPO after 2000 period (left), U.S., new PSID after 1997 period (right)

C Non-robust measures of skewness and kurtosis

Figure 9: The Dutch data: Non-robust measures of Skewness (left) and Kurtosis (right): male wages (first row) male earnings (second row) household earnings (third row) disposable income (fourth row).