Boom, Bust, Rebound: Two Decades of Earnings Inequality, Risk and Mobility in **Ireland**

Brian E. Higgins
Institute for International Economic Studies

Tara McIndoe-Calder
ESRI and Trinity College Dublin

Barra Roantree
Trinity College Dublin

Kun Wu
University College Dublin

September 8, 2026

Abstract

This paper studies income inequality in Ireland using new administrative social security records spanning 2002–2022. We provide the first comprehensive set of harmonized micro statistics on earnings inequality, dynamics, and mobility for Ireland, contributing to the Global Repository of Income Dynamics (GRID). We document how inequality evolved through Ireland's boom, bust, and recovery, highlighting that the boom in the early 2000s lifted all parts of the distribution, the following recession disproportionately reduced earnings at the bottom, and the subsequent recovery brought strong catch-up at lower percentiles but left persistent gaps at the top. We uncover a sharp rise in inequality at top earnings percentiles, and find that Ireland has moderate levels of income risk and economic mobility compared with other advanced economies. We analyze how the distribution of earnings growth comoves with aggregate output and inflation, and how this varies across worker types. We further show that migration can distort standard inequality measures, and that weekly earnings provide a useful complementary perspective on living standards in a labour market with significant entry and exit.

JEL: E24, J24, J31.

KEYWORDS: Ireland, income inequality, dynamics, mobility.

This paper studies income inequality in Ireland using new administrative social security records. We document for the first time a wide range of micro statistics on earnings inequality, dynamics, and mobility in Ireland across two decades.[1] The data forms the Irish contribution to the Global Repository of Income Dynamics (GRID) (Guvenen, Pistaferri and Violante, 2022a,b), a new open-access, cross-country database that contains harmonized statistics for 13 other countries, with further countries expected to join alongside Ireland. All the underlying statistics will be made available on the GRID database and our own repository.[2] GRID is built on micro panel data drawn from administrative records in each country, is designed with the goals of harmonization and cross-country comparability, and includes granular descriptions of income inequality and income dynamics for finely defined subpopulations such as age and gender.

Ireland provides us with a novel laboratory for studying income inequality during periods of both economic convergence and deep contractions. Ireland experienced rapid economic growth during the 1990s: GNP expanded almost 70 percent in the decade from 1987, and was associated with convergence of GNP per head from 59 percent of the EU15 average to 88 percent (Barry, 1999). Ireland also experienced one of the sharpest declines during the Great Recession: house prices fell 54 percent, unemployment reached a peak of 15.8 percent in 2012, and the resulting banking and fiscal crisis led the government to enter into a lending program with the European Commission, European Central Bank and International Monetary Fund. Indeed, Donovan, Lu, Pedtke and Schoellman (2024) rank Ireland's experience in the Great Recession as the 9th most severe labour market downturn out of their sample of 311 recessions among 83 countries.

We make four contributions. First, we provide the first comprehensive set of micro statistics on Irish earnings using harmonized GRID methods, documenting the evolution of inequality across the entire distribution—including the very top percentiles—and across men and women over the boom–bust–recovery cycle. Second, we establish new evidence on earnings dynamics and mobility by characterizing the full distribution of earnings changes over short and long horizons, by gender, and across the permanent-income distribution, thereby quantifying the risks faced by different workers. Third, we situate Ireland in a comparative perspective, benchmarking levels and changes in inequality, risk, and mobility against the first wave of GRID countries. Fourth, we highlight

[1]We are especially grateful to Fatih Guvenen, Luigi Pistaferri, and Gianluca Violante for including us in the GRID project, to Serdar Ozkan and Sergio Salgado for providing the master code to produce the main GRID statistics, and to Ciaran Judge for facilitating access to the data. For helpful comments, we thank Laura Boyd and Kieran Larkin, and participants at several seminars and conferences including the Econometric Society World Congress, ESCoE Conference on Economic Measurement, GRID Conference 2024, Irish Economic Association, Central Bank of Ireland, and UCD PhD Workshop. Part of this paper was completed while McIndoe-Calder was employed with the Central Bank of Ireland. The views expressed here do not necessarily reflect the views of the Central Bank of Ireland, the Eurosystem of Central Banks, nor the Department of Social Protection. Some results are based on analysis of strictly controlled Research Microdata Files provided by the Central Statistics Office (CSO). The CSO does not take any responsibility for the views expressed or the output generated from this research.

[2]Accessible at `www.grid-database.org/` and `github.com/higginsbrian/GRID-Ireland`.

the challenges in measuring inequality when large changes in selection are present, and show that weekly earnings can be particularly informative in such settings.

We use a new administrative panel dataset based on social security records covering two decades of earnings from 2002 to 2022. The data is a 10 percent sample of all employees. To be consistent with the GRID database, our primary measure of income is annual individual labour earnings before tax from all employments (excluding self-employment).[3] We show that the data is representative of the Irish labour market as it captures a large and stable share of the employment and labour market earnings reported in the national accounts.

Income inequality is multifaceted, and its significance may depend on whether it reflects differences in income at a given point in time, in exposure to income risk over time, or in the persistence of income differences over longer periods. To reflect these conceptual differences, we measure inequality across three dimensions. First, we study cross-sectional income *inequality* as measured by the distribution of earnings. Second, we study income *dynamics* as measured by (residual) changes in earnings. We interpret this as a measure of the risk that individuals face, with the caveat that part of these changes may be foreseen by the individual or may result from their labour supply choices. Third, we study income *mobility* as measured by how persistent permanent income is over time. An advantage of the large administrative dataset is that for each of these measures we can investigate heterogeneity by gender, age and permanent income.

Our data reveal four novel facts about the evolution of Ireland's earnings distribution over this period. First, gains during the boom of the early 2000s were broadly shared across the distribution, but inequality has increased markedly since then. The nature of this increase has changed over time. During the Great Recession, the initial rise in inequality was driven by larger declines at the bottom of the distribution than at the top. Since then, however, the bottom has caught back up with the median. This degree of catch-up is not observed in many other countries.[4] By contrast, the increase in inequality between the median and the top of the distribution persisted throughout the recovery from the Great Recession.

Second, we uncover a striking rise in top earnings inequality—a new finding that is enabled by our large, administrative data. Relative to the median, earnings of the top 1 and 0.1 percent grew by more than 30 and 50 log points respectively. This is striking because it is larger than the rise in any of the first wave of GRID countries over the same time period—yet, it still leaves Ireland with much less top income inequality than high-inequality countries like the United States.

Third, inflation plays an important role in determining whether increases in *aggregate* output

[3] Given this definition of income we use the terms earnings and income interchangeably, though in principle individual income may include labour earnings as well as self-employment, business or investment income.

[4] Unlike in Ireland, P10 had not caught up with P50 by the end of the respective sample periods in Norway (2017), Spain (2018), or the UK (2020).

translate into increases in *real labour* earnings. During the boom, inflation drove a wedge between nominal and real earnings, with strong nominal gains in the mid-2000s masking stagnant or declining real earnings. Conversely, during the downturn, deflation was sufficiently large that many percentiles experienced *increases* in real labour earnings even as aggregate output fell and unemployment rose. Overall, these patterns suggest that much of the apparent prosperity of the boom years reflected nominal price increases rather than improvements in purchasing power, while the post-2015 recovery marked a period of genuine real earnings growth.

Fourth, looking at subpopulations reveals large gender differences and persistent scarring for recession cohorts. The cohort entering the labour force at the *onset* of the Great Recession experienced weak growth that is consistent with hysteresis and scarring (von Wachter, 2020). Despite lower initial earnings, the cohort entering at the *end* of the recession enjoyed faster growth during the recovery, nearly overtaking the earlier recession cohort by 2022, despite having less potential labour market experience. Women were much less affected by the Great Recession: men at most points of the income distribution fell below their starting level in 2002 and did not recover until the late 2010s, while the entire female distribution remained above its 2002 level—in fact, throughout the recession women's *real* median earnings never fell below their level at the start of the recession.

Turning to income risk, we find that Ireland exhibits low-to-moderate amounts of earnings risk. For example, the dispersion of income shocks lies in the middle of the range of countries in the first wave of GRID. As in all of those countries, the distribution of income risk is non-normally distributed exhibiting fatter tails ("excess kurtosis") and pro-cyclical skewness. This means that good times feature a higher incidence of positive shocks, while bad times are associated with more negative shocks and fewer positive shocks.

We measure economic mobility by estimating how persistent permanent income is over an individual's working life. The slope on a rank-rank regression captures this persistence, with $\beta = 0$ reflecting perfect mobility and $\beta = 1$ perfect immobility. Consistent with other countries, there is substantial mobility early in the life cycle and greater persistence at older ages. The average slope of the rank–rank regression, around $\beta = 0.75$, places Ireland near the middle of the GRID distribution: less mobile than Scandinavia but more mobile than countries such as Brazil and the UK, and broadly similar to Canada, Germany, and the United States. Relative to its level of inequality, Ireland's mobility is comparatively high—particularly when contrasted with Spain and Italy, which exhibit similar inequality but lower mobility—underscoring that high dispersion in earnings coexists with meaningful movement across the distribution over longer horizons.

In the second part of the paper, we go beyond the GRID statistics to study issues of general interest to the literature on inequality. First, we study the responsiveness of earnings risk—as measured by the distribution of earnings growth—to aggregate output and inflation. We estimate the average responsiveness "betas", and use our data on weeks worked to provide a novel decomposition into the responsiveness of *weekly earnings* and of *weeks worked*. Second, we use quantile regression to estimate the sensitivity of the distribution of earnings risk—as measured by volatility and skewness—to these aggregate variables.

In the case of aggregate output, our average results are consistent with the prior literature (Guvenen, Schulhofer-Wohl, Song and Yogo, 2017) showing that the earnings growth of men and low-income workers is especially sensitive to output. Our novel decomposition shows that while most of the risk comes from weekly earnings, a sizable share comes from weeks worked: 30 percent in the case of men. The relationship between total responsiveness and the responsiveness of weeks worked is non-monotonic. For example, men are more responsive in total and have a higher share attributed to weeks worked than women, while some of the least responsive groups have the highest shares accounted for by weeks worked. For example, weeks worked contribute more than half of the responsiveness of the least responsive permanent income group (P90). Since wage changes affect weekly earnings, and unemployment spells affect the number of weeks worked, these differences suggest that there is substantial heterogeneity in the types of risks faced by different workers. In terms of the distribution, we find that the skewness of earnings is pro-cyclical for most worker types, which is consistent with Busch et al. (2022). This is true for both weekly earnings and weeks worked, though the skewness of weekly earnings is substantially more cyclical.

In the case of aggregate inflation, we find that *real* earnings are mostly unaffected by inflation, implying complete pass-through of inflation to *nominal* earnings. However, this average result masks a new and striking pattern across worker types. While most workers experience (almost) complete pass-through, workers with the lowest permanent incomes experience *more* than complete pass-through. Indeed, a one percentage point increase in inflation is associated with a 0.8 percentage point increase in *real* earnings for those at the P10. This pattern is consistent across sub-groups, including men and women as well as all age groups. The decomposition shows that it is driven by increases in weekly earnings for these groups.

Lastly, we examine the challenge of measuring living standards, even with administrative data, in economies with a lot of labour force selection arising from migration. Ireland experienced large inward migration following the expansion of the EU in 2004, causing changes in the composition of the population in our data. By comparing our results with balanced samples that do not include selection, we show that this selection masked some of the growth in earnings and exaggerated changes in inequality. We also show that weekly earnings were less affected by this selection than annual earnings, and thus conclude that weekly earnings may be a complementary measure of living standards in a setting with substantial selection.

Contribution to literature. We provide the first estimates of income inequality, risk and mobility using Irish administrative microdata. The prior literature has relied on survey data

(Roantree and Barrett, 2023, 2024; Roantree, Maître and Russell, 2024) and aggregate tax data (Nolan, 2007; Callan, Doorley and McTague, 2020).[5] Given the lack of administrative data, the prior literature measuring inequality has focused on estimates of the Gini coefficient and P90-P10, and therefore our estimates of income growth in the upper percentiles and comparisons across sub-groups (cohorts and gender) are entirely novel. Strikingly, these new estimates also uncover an increase in inequality at the top which is larger than any other country in the first wave of GRID. Likewise, the panel nature of our data allows us to provide the first estimates of income risk and mobility in Ireland.

We contribute to a literature measuring and comparing inequality across countries (Krueger, Perri, Pistaferri and Violante, 2010; Guvenen, Pistaferri and Violante, 2022b). Our data is a notable addition to the international evidence since we provide detailed evidence on the Irish labour market experience in the Great Recession, which Donovan, Lu, Pedtke and Schoellman (2024) report as the 9th worst labour market contraction in recent history. The closest work that allows for international comparisons is Nolan (2007, 2018)'s contributions to the World Inequality Database (2025). They use aggregated tax data to construct measures of top income shares, which we also provide comparable measures of. To the international literature, we highlight challenges to constructing comparable metrics in countries with a higher level of selection due to migration and unemployment. A rare feature of the Irish data is the availability of weeks worked (Hoffmann and Malacrino, 2019; Broer, Kramer and Mitman, 2023), and we contribute new evidence on how changes in the number of weeks worked contribute to the higher-order moments of the earnings risk distribution (Guvenen, Ozkan and Song, 2014; Guvenen, Karahan, Ozkan and Song, 2021). We also contribute to a literature on the pass-through of inflation to earnings across the distribution (Blanco, Drenik and Zaratiegui, 2025; Afrouzi, Blanco, Drenik and Hurst, 2026).

1 Data and Institutional Background

1.1 Data Description

Dataset. Our data comes from the Welfare and Work Longitudinal Database (WWLD), an administrative panel dataset held by Ireland's Department of Social Protection (DSP). Earnings information in the WWLD is based on records relating to Pay Related Social Insurance (PRSI) payments, which determine workers' eligibility for the state pension and other social insurance programs.[6] The raw WWLD data are held at the job level: one record per employment. Earnings are reported by employers at the pay-period level and aggregated by DSP into an annual measure

[5] Notable exceptions using Irish administrative data are Doris, O'Neill and Sweetman (2015), who study nominal earnings flexibility in the Great Recession, and Hargaden (2020) who studies the responsiveness of earnings to taxes.

[6] Further information regarding PRSI and Irish social insurance can be found at `https://www.citizensinformation.ie/en/social-welfare/irish-social-welfare-system/social-insurance-prsi/`.

for each employment.

Earnings Definition. Our measure of earnings is annual individual labour earnings before tax. This includes taxable bonuses and benefits-in-kind, but excludes pensions and severance payments.

Timeframe. The sample we use covers the years 2002–2022.

Sampling Frame. We use a 10 percent random sample drawn from the WWLD, selected on the basis of the last digit of individuals' (anonymised) Personal Public Service Number (PPSN). The WWLD contains information on all earnings subject to PRSI, in addition to the 'class' of PRSI that the earnings are subject to.

We restrict attention to earnings subject to PRSI Classes A through H, which includes both public- and private-sector employments but excludes self-employment income (which is subject to Class S PRSI). We observe (uncensored) annual earnings from each of these employments net of *employer* PRSI contributions but before any income tax, Universal Social Charge or employee PRSI contributions. We aggregate across all of an individual's employments with earnings subject to PRSI Classes A through H to construct a single annual worker-year earnings record per individual. This gives rise to a sample that ranges in size from approximately 233,000 to 342,000 individuals per year, before the GRID sample restrictions described below are applied.

GRID Covariates and Other Observables. We use five-year age groups and gender as covariates to calculate residual earnings and for sub-group analysis. Unlike some of the other GRID countries, we do not observe education.

Supplemental Data. We also provide GRID statistics using a second administrative dataset EAADS, which we describe in Appendix F. The coverage of both datasets is very similar. EAADS is top-censored and has a shorter time series, so we focus our discussion in the paper on the results using WWLD.

1.2 GRID Sample Restrictions

The previous description applies to all data available in the WWLD dataset. Next we describe our variable definitions and additional sample restrictions, where we do our best to be consistent with the other GRID countries, as described by Guvenen, Pistaferri and Violante (2022b). Table 1 summarizes the data.

We impose the same three sample restrictions imposed by other GRID countries. First, we focus on workers between 25 and 55 years old, a range within which most education choices are usually completed and after which workers tend to leave the labour force for retirement. Second, for most of the analysis we drop observations with earnings (defined next) below a threshold (call it $\underline{y}$) to avoid using records from workers without a meaningful attachment to the labour force or with

Table 1: Summary Statistics

(a) Sample Size

	Original	*CS* sample		*LX* sample		*H* sample	
Year	N	N	% Women	N	% Women	N	% Women
2002	233,289	125,161	48.0	113,175	48.2	—	—
2007	284,174	159,210	46.9	142,037	47.8	138,147	47.2
2012	250,971	138,811	50.7	125,148	51.1	126,551	51.2
2017	294,119	160,734	48.7	147,993	48.9	143,677	49.0
2022	342,242	182,907	48.8	—	—	160,529	48.8

(b) Descriptive Statistics

					Mean earnings		Age shares (%)		
Year	Obs. (Thou.)	Mean Earnings	Std. Dev. Earnings	Women Share(%)	Men	Women	[25,34]	[35,44]	[45,55]
2002	125	32,174	25,792	48.0	37,828	26,062	44.6	30.6	24.8
2007	159	34,252	31,423	46.9	38,841	29,064	45.8	30.3	23.9
2012	139	35,140	38,281	50.7	39,551	30,854	38.2	34.9	26.9
2017	161	37,377	42,873	48.7	41,841	32,682	33.8	37.0	29.2
2022	183	41,125	43,779	48.8	45,924	36,090	33.0	34.9	32.1

Notes. Table summarizes the dataset and the three GRID samples (CS, LX and H).

very low earnings, which could skew log-based statistics. Specifically, we discard observations with earnings below what workers would earn if they worked 260 hours at the Irish national minimum wage (which is equivalent to one and a half months full time). Third, we exclude self-employment income, using the aforementioned restriction to employment income subject to PRSI Classes A–H. After applying the GRID sample restrictions above, the final sample ranges from approximately 125,000 to 183,000 individuals per year.

We construct three separate samples to be used for different parts of the analysis:

1. **The cross-sectional (CS) sample**: This sample is used to compute cross-sectional inequality statistics. All individuals who satisfy the criteria above are included at date t. This is the most comprehensive sample, utilizing the longest possible time series available.

2. **The longitudinal (LX) sample**: This sample is used to study the distribution of earnings changes. It includes all individuals in the CS sample who also have 1-year and 5-year forward

earnings changes.[7]

3. **The heterogeneity (H) sample**: This sample examines variation across demographic groups defined by observable characteristics such as age, gender, and permanent income. It includes individuals in the LX sample who also have a permanent earnings measure (definition provided below). We pool observations across years for this analysis.

1.3 Variable Definitions

We construct the following measures of earnings for worker i in year t:

1. Raw real earnings in levels, y_{it}, and logs, $\log(y_{it})$. Real earnings are computed from nominal earnings and deflated to 2018 euros using Irish CPI provided by the CSO.

2. Residualized log earnings, ϵ_{it}. This measure is the residual from a regression of log real earnings on a full set of age dummies, separately for each year and gender. It is intended to control for predictable changes in individual earnings (life cycle and business cycle effects).

3. Permanent earnings, P_{it-1}. Defined as average earnings over the previous 3 years, $P_{it} = \sum_{s=t-2}^{t} y_{is}/3$, where y_{is} can include earnings below $\underline{y}$ for at most 1 year. The measure is intended to average over transitory income changes and proxy for skill levels.

4. 1-year change in residualized log earnings, g_{it}^1. This is the 1-year forward change in ϵ_{it}, defined as $g_{it}^1 = \Delta\epsilon_{it} = \epsilon_{it+1} - \epsilon_{it}$, where earnings must be above $\underline{y}$ for both years.[8]

5. 5-year change in residualized log earnings, g_{it}^5. This is the 5-year forward change in ϵ_{it}, defined as $g_{it}^5 = \Delta\epsilon_{it} = \epsilon_{it+5} - \epsilon_{it}$, where earnings must be above $\underline{y}$ for both years.

The above measures of earnings are universal to the GRID project. Since our data includes the number of weeks worked in each employment, we also construct a measure of weekly real earnings following Hoffmann and Malacrino (2019):

6 Real weekly earnings in levels, y_{it}^W, and logs, $\log(y_{it}^W)$. Weekly earnings are defined as annual earnings y_{it} divided by the number of weeks worked, $y_{it}^W = 52 \times y_{it}/w_{it}$, where weeks worked w_{it} are defined as the sum over weeks worked at each employment j and capped at 52 weeks

[7] When reporting results that only depend on 1-year changes, we do not impose that individuals have 5-year forward changes. Thus, technically we could refer to these as separate LX^1 and LX^5 samples.

[8] As noted by Guvenen, Pistaferri and Violante (2022b), using "leads" avoids mechanical mean reversion when conditioning on permanent earnings. In other words, for a given year t, the statistics we are interested in are calculated using t-years forward, whereas the permanent earnings measure which the statistics are conditioned on is computed for years $t-1$ and earlier, avoiding any overlapping years.

$w_{it} = \min \left\{ \sum_j w_{itj}, 52 \right\}$. We multiply by 52 to return the measure to an annual scale and thus ensure that the minimum income thresholds are applied in a consistent manner to weekly and annual earnings.

1.4 Representativeness

The WWLD dataset is broadly representative of total employment and earnings. Figure 1 compares the time series for total earnings and total employment in national accounts with our estimate from aggregating WWLD. Relative to national accounts, the WWLD dataset captures over two thirds of total earnings and includes more employments. For employments, the difference is likely because the national accounts measure is based on a survey. For earnings, we identify some differences in the sampling frame and definitions of earnings that may explain these discrepancies. The national accounts uses a much broader definition of earnings. In addition to the income included in WWLD, the national accounts include employers' social insurance contributions (PRSI), defined contribution pension contributions actually made by employers, as well as imputed contributions for defined benefit schemes (including public sector schemes). Given the substantial differences in these definitions of earnings, we consider the coverage rate of WWLD to be reasonably high.

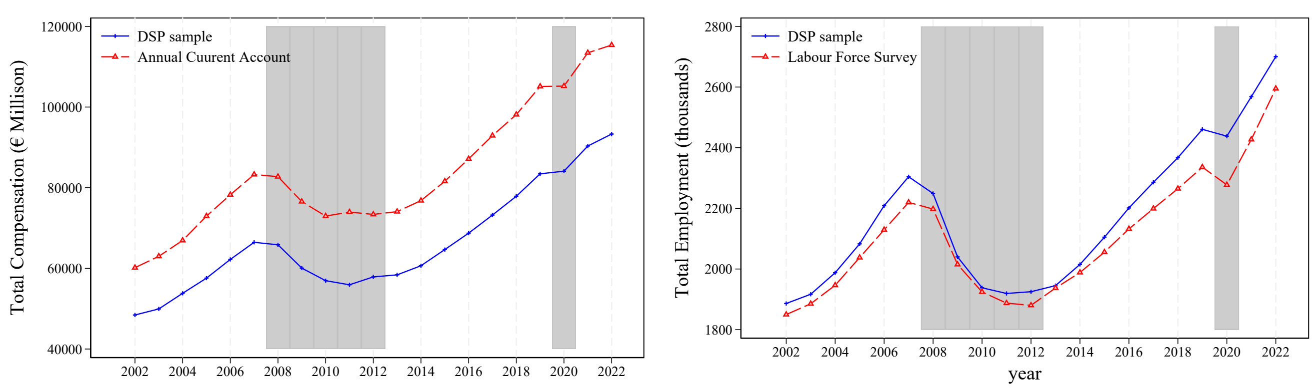

(a) Total Earnings

(b) Total Employment

Figure 1: Comparison of WWLD with Aggregate Statistics

Notes. Figure compares the coverage of the aggregated WWLD dataset with aggregate statistics. The left plot compares against total earnings in the national accounts. The right plot compares against total employment in the Labour Force Survey, which is the official measure of employment and unemployment in the state. See the text for a discussion of differences between the definitions in each data source.

1.5 Irish Economy

Over our sample period 2002-2022, the Irish economy can be categorised by four periods: (i) a boom until 2007, (ii) a bust between 2008 and 2012, (iii) a robust recovery since 2013, and (iv) a brief interruption to that recovery due to the Covid-19 pandemic. Figure 2 plots the corresponding trends in panels (a) output, (b) unemployment, (c) the minimum wage, and (d) inflation. We use Modified Gross National Income (GNI*) as our preferred measure of output. Due to the large presence of multinational firms in Ireland, the Central Statistics Office produces GNI* and Modified Domestic Demand (MDD) statistics, which both capture the domestic component of output.[9]

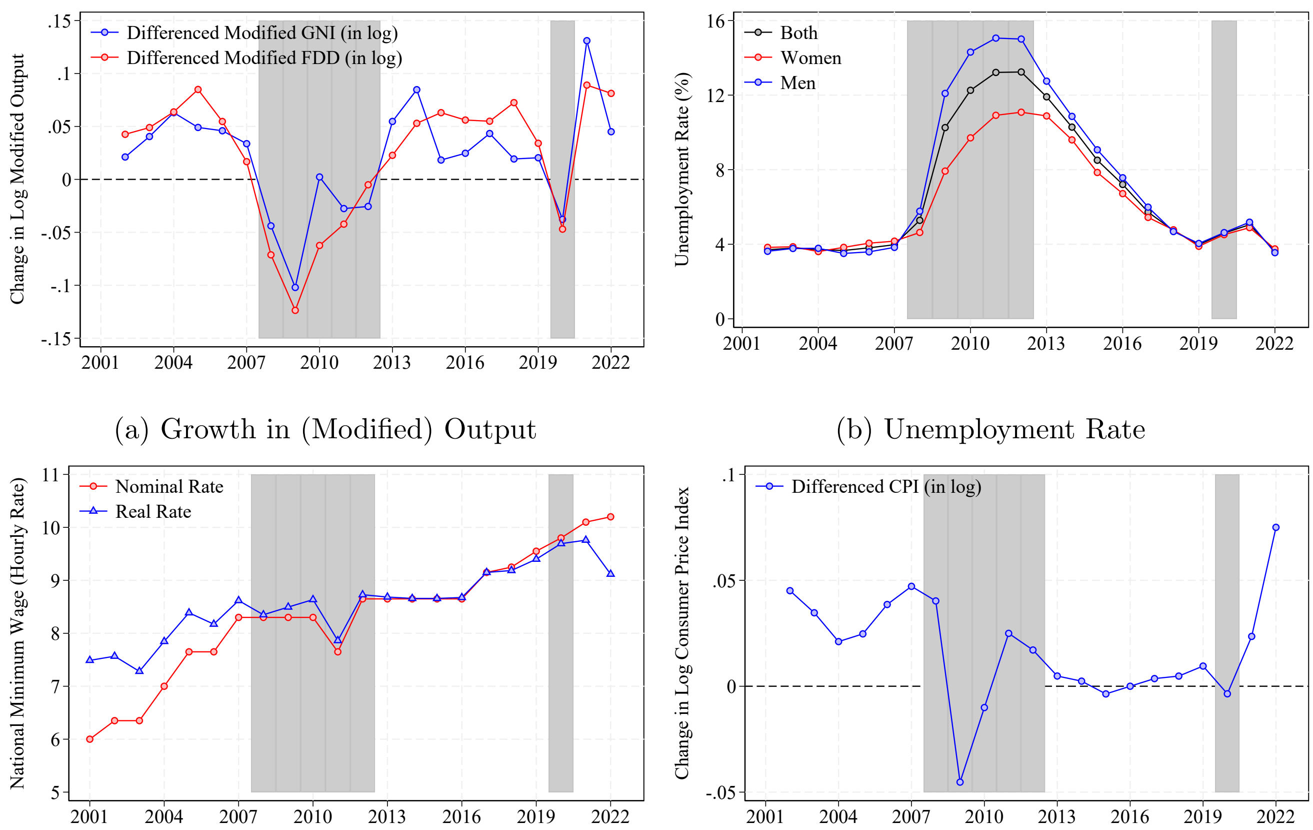

(a) Growth in (Modified) Output

(b) Unemployment Rate

(c) Minimum Wage

(d) Inflation Rate

Figure 2: Trends in the Irish economy

Notes. Figure shows aggregate trends in the Irish economy and labour market since 2000. Panel (a) shows Modified GNI and Modified Domestic Demand, preferred measures of national income that remove the influence of multinational companies that distort Irish GDP and GNI. Panel (b) shows the unemployment rate. Panel (c) shows the minimum wage. Panel (d) shows the inflation rate. All data are the official series provided by the Central Statistics Office (CSO).

From the 1990s Ireland experienced rapid catch-up growth in output (Barry, 1999; Honohan and Walsh, 2002; Ó Gráda and O'Rourke, 2022), a period often referred to as the *Celtic Tiger*. This can be seen by the growth in output averaging more than 5 percent from 2000 to 2007.

[9] See Lane (2017) and Honohan (2021) for discussion of the issues, and https://www.cso.ie/en/interactivezone/statisticsexplained/nationalaccountsexplained/modifiedgni/ for further details on modified national accounts. GNI* approaches GDP using the income method and MDD approaches GDP using the expenditure method. We define recession years (shaded in plots) based on declines in MDD; the same years coincide with declines in GNI* apart from 2010, which exhibited marginally positive growth in GNI*.

This period was associated with a decline in unemployment from above 8 percent in the 1990s to between 4 and 5 percent, and inflation that was high relative to the Eurozone average.

Ireland experienced one of the sharpest declines during the Great Recession: house prices fell 54 percent, unemployment reached a peak of 15.8 percent in 2012, and the resulting banking and fiscal crisis led the government to enter into a lending program with the European Commission, European Central Bank and International Monetary Fund (Lane, 2012; Honohan, 2019). As shown in Figure 2, the peak of the boom in 2007 was followed by four consecutive years of declines in output, including a decline of more than 10 percent in 2009. Unemployment for men rose almost 8 percentage points in a single year (2009), which was larger than the rise for women because men were more exposed to the particularly sharp decline in employment in the construction sector that coincided with the fall in property prices (Bergin and Kelly, 2018).[10]

Though the recession was pronounced, the recovery has been robust. Output started to recover in 2013, and apart from the Covid recession there has been growth near or above five percent per year since 2014. Unemployment has been slower to recover staying above ten percent until 2016 and not returning to pre-recession levels until 2019.

As in many countries, the impact of the Covid recession differed to past recessions due to the differential effect on sectors that could adapt to working from home (Barrero, Bloom and Davis, 2023). Policy also differed in Ireland as summarized by Alamir, McGinnity and Russell (2024). A special Pandemic Unemployment Payment was introduced between March 2020 and March 2022 with a higher rate of benefit and wider eligibility than the existing unemployment benefit. There were also two Wage Subsidy Schemes (TWSS and EWSS), which subsidized up to 70 percent of the net wages of eligible employees between March 2020 and May 2022.[11]

Panel (c) reports the minimum wage over our time period. The national minimum wage was introduced in April 2000, which replaced minimum wages set by industry-specific Joint Labour Committees (Redmond, 2020).[12] In principle, increases in a binding minimum wage could determine growth in earnings at the bottom of the distribution. We do not find strong evidence that the changes in the earnings distribution are primarily driven by the changes in the minimum wage.[13]

[10] Bergin and Kelly (2018) note that more than half of the fall in employment between 2007 and 2012 was in the construction sector.

[11] For details, see `https://www.citizensinformation.ie/en/employment/unemployment-and-redundancy/employment-wage-subsidy-scheme/` and `https://www.citizensinformation.ie/en/social-welfare/unemployed-people/covid19-pandemic-unemployment-payment/`

[12] Redmond (2020) summarizes minimum wage policy over our time period.

[13] This can be seen by comparing the minimum wage series with the lower percentiles of the income distribution in Figure 3: despite increases in the minimum wage throughout the boom period, the P10 falls. Table A1 reports the euro values of the P1 and P10, and the ratio of these percentiles changes throughout the period and differs for men and women. Taken together, these suggest that changes at the bottom of the distribution are not mechanically driven by changes in the minimum wage.

Panel (d) reports the inflation rate over our time period. Inflation was high during the boom and also after the Covid recession. It is also worth noting that there was a year of almost five percent *deflation* in 2009. As we will see later, this mitigated the decline in nominal wages identified by Doris et al. (2015) in that year resulting in, rather surprisingly, increases in *real* earnings at some percentiles during the worst years of the crisis.

2 Trends in Inequality, Risk and Mobility from GRID

2.1 Earnings in Levels

Figure 3 presents changes in the log earnings distribution since 2002 for men (left) and women (right). The top panels report changes in inequality, as measured by the P90-P10 ratio and standard deviation; the middle panels report changes in specific percentiles between the 10th and 90th percentiles; and the bottom panels focus on top inequality above the 90th percentile. We also report the euro amounts of relevant percentiles in Table A1 in Appendix A.

There are five key patterns. First, the business cycle is evident in earnings across the distribution. Median earnings rise in the early 2000s, fall to a trough in 2014, and rise until the spike in inflation at the end of the sample. That said, the timing of the rise and fall of real earnings across the distribution does not obviously track the growth in output or rise in unemployment. For example, earnings at the bottom of the distribution (P10) fall throughout the boom, which contemporary accounts do not appear to notice. In Section 4 we focus the attention of our country-specific analysis on these peculiar patterns. We find that they are driven by selection due to new migrants. Once accounted for, the patterns suggest that the boom was associated with rising earnings across the distribution and a bust that affected the bottom of the distribution more than the top.

Second, inequality has risen throughout the sample period, but the source of the inequality has varied over time. The largest increase in the P90-P10 for men came during the bust and was associated with a fanning of both the top and the bottom of the income distribution. This was mostly driven by larger declines in earnings at the bottom during the recession, while incomes at the top were mostly flat. In the subsequent years, the 10th percentile has largely caught up with the median, while the gap between the 90th percentile and the median has remained constant.

The rise in inequality at the very top is even more stark, with the bottom panels illustrating a clear fanning out of the distribution above the 90th percentile. Over the sample period, the top one percent for men has grown by more than four times the increase at the median, and the top 0.1 percent has grown more than 6 times as much. The majority of this rise has occurred after the Great Recession. These facts on top percentiles are particularly novel for Ireland given that past

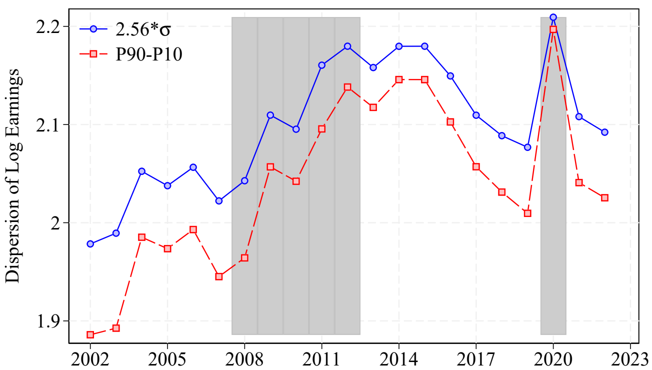

(a) Men

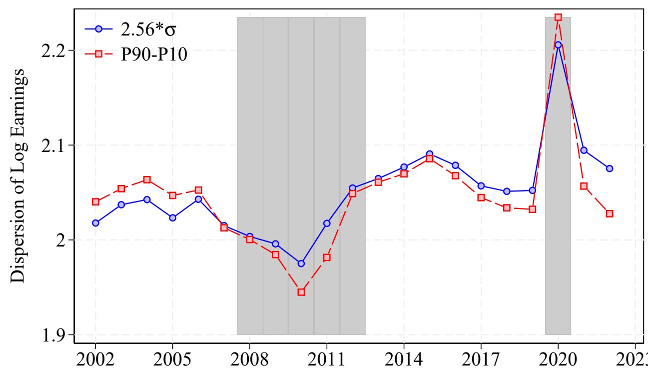

(b) Women

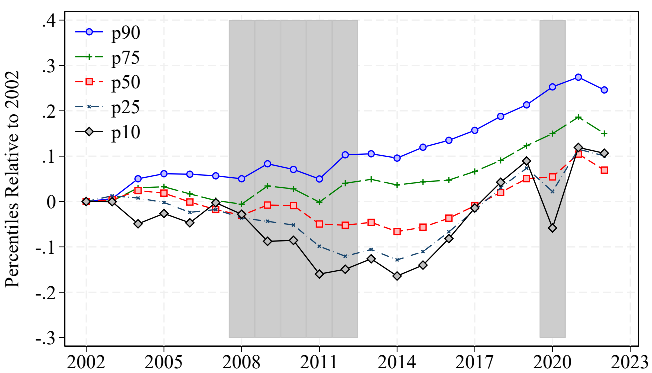

(c) Men

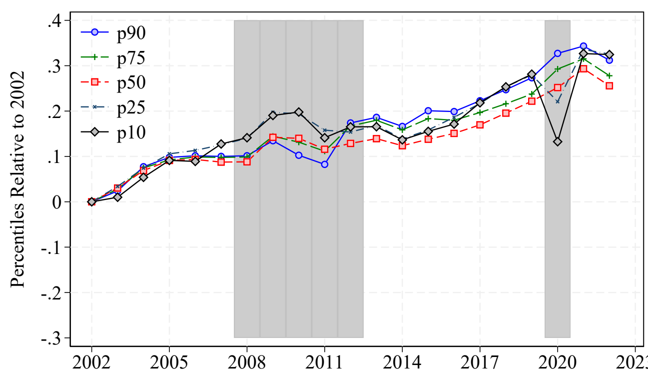

(d) Women

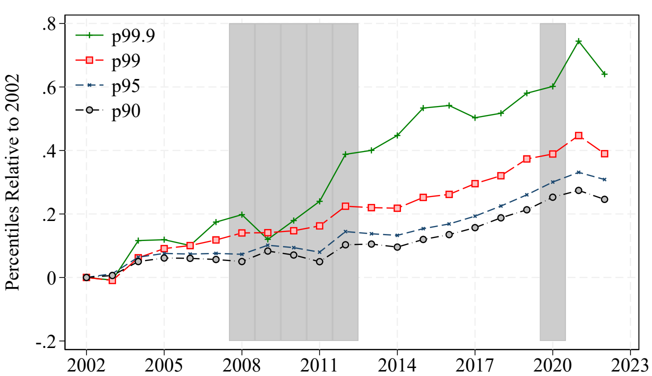

(e) Men

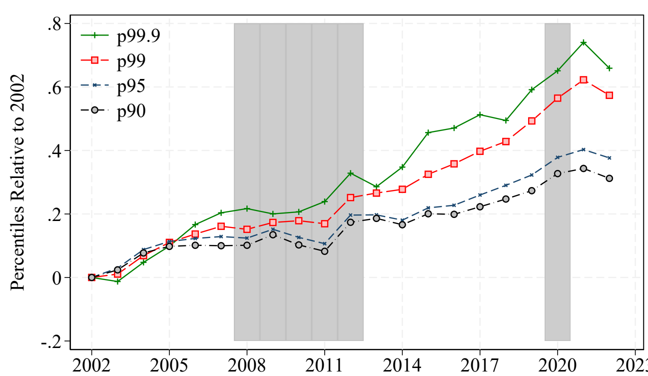

(f) Women

Figure 3: Changes in the percentiles of the log real earnings distributions

Note: Log income is normalised to 0 in the base year 2002. Top panels (A) and (B) show the P90-P10 ratio and standard deviation of log income; middle panels (C) and (D) show P10, P25, P50, P75, P90; bottom panels (E) and (F) show P90, P95, P99, P99.9. Shaded areas indicate recession years, based on declines in Modified Domestic Demand.

research has mostly used surveys (Roantree, Maître and Russell, 2024) and aggregated tax records (Callan, Doorley and McTague, 2020; Nolan, 2007).

Third, inflation has been an important determinant of whether increases in real aggregate *output* passed through to real labour *earnings*—which can be seen by comparing Figure 3 with its equivalent for nominal earnings in Figure A12. The rise in real earnings in the early 2000s is modest relative to the large increase in output, and many percentiles show declines in real earnings from 2005 despite real aggregate output growing by almost ten percent in that year. However, Figure A12 shows the boom in output was associated with a large increase in nominal earnings: median nominal earnings rose almost 20 percent and did not fall until after the recession in 2008. This implies that much of the nominal earnings growth did *not* translate into real earnings, especially in the latter years of the boom. Indeed, deflation also cushioned real earnings declines in the recession: the increase in real earnings at higher percentiles in 2009 is entirely driven by deflation because nominal earnings fell in that year—the increase in real earnings being especially notable because it occurred during the worst year of the recession when real output fell by 15 percent and unemployment rose to 10 percent. Fourth, there are striking differences between men and women. Women were much less affected than men by the recession. At no point during the crash does any female percentile fall below the level it started at in 2002, whereas it takes until 2017 before men below the median returned to incomes higher than they had started with in 2002. There is also much less of an increase in inequality in the central part of the distribution, with almost no increase in the P90-P10. The pattern that is most similar for men and women is the rising inequality at the very top. Indeed, this rise is even more striking among women with all percentiles between the 90th and 99th growing more than their male equivalents.

Fifth, the Covid recession mostly affected earnings in the bottom half of the distribution. Median earnings for men were flat, and earnings fell most for those at the tenth percentile, while earnings at the top grew. This is consistent with higher-income workers being able to work from home while lower income service workers saw declines in their earnings (Dingel and Neiman, 2020).

2.2 Cohorts

Figure 4 examines whether changes in inequality and income growth are driven by new cohorts arriving or by time trends that affect all cohorts uniformly. The patterns are consistent with signs of scarring for cohorts entering in the Great Recession (von Wachter, 2020).

Figure 4 reports how the trends at the 50th and 90th percentiles differ by cohort over time. The coloured lines follow a given cohort over their life (for example the 2002 cohort in blue) while the gray lines track equivalently aged individuals across cohorts (for example the bottom gray line follows age 25 individuals, which is the age at entry for each cohort). Figure A1 in Appendix A

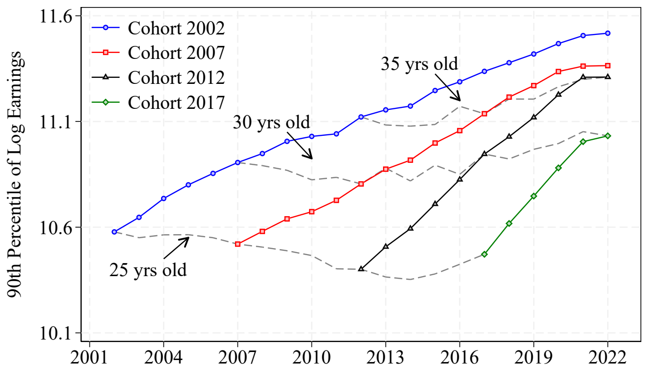

(a) Men P90

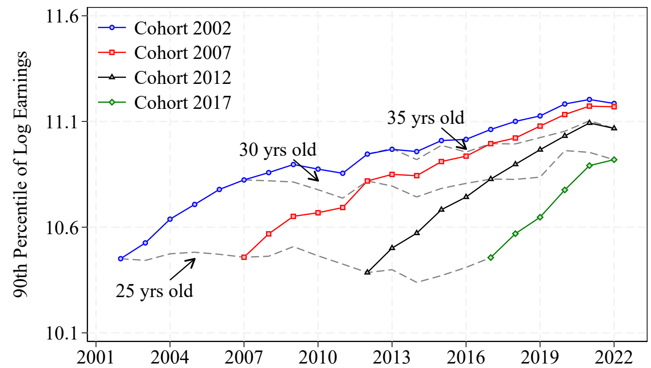

(b) Women P90

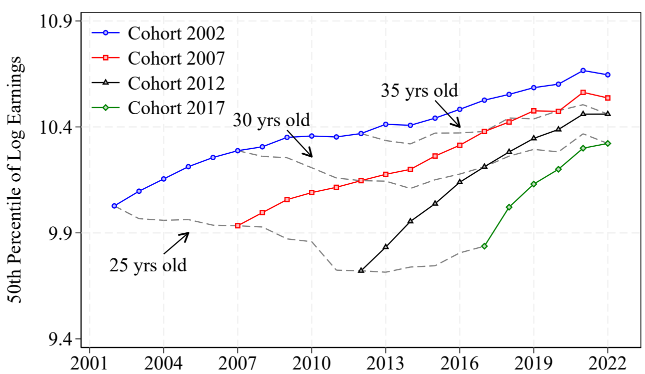

(c) Men P50

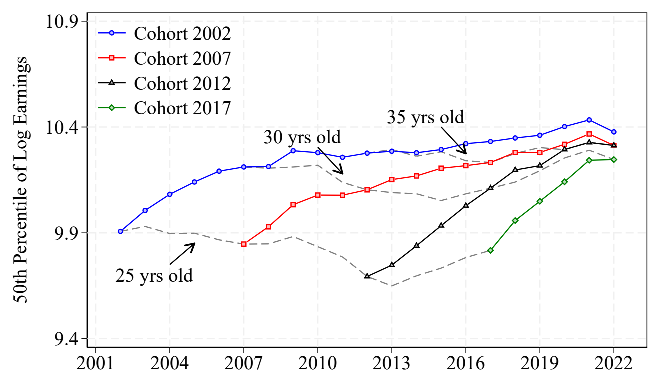

(d) Women P50

Figure 4: Life-cycle inequality over cohorts

Note: Panels show changes in log income at the 90th and 50th percentiles. Dashed lines correspond to earnings of 25-, 30-, and 35-year-olds as indicated by arrows. Each solid line corresponds to an individual cohort where "cohort X" represents the cohort aged 25 in year X.

reports equivalent trends at the bottom (P10) and in inequality (P90-P10) over time and cohorts. Consistent with the flat profile of P90-P10 inequality in Figure 3, there is little change in inequality across either cohorts or time.

There are clear signs of scarring for cohorts entering around the Great Recession. We find that cohorts entering at the beginning of the recession (2007, red) experience both lower entry earnings and less growth. Despite entering at lower earnings, the cohort entering at the end of the recession (2012, black) experienced faster growth in the rebound than the earlier recession cohort. Indeed, by 2022, this cohort in black has almost overtaken the cohort in red despite having five years less labour market experience. These patterns are similar for men and women, and for individuals at each of the 10th, 50th and 90th percentiles.

2.3 Earnings Dynamics and Risk

The second set of statistics measures earnings risk. We residualize changes in log earnings to remove predictable changes in income arising from gender, life-cycle, or (average) business cycle effects. The remaining variation can be interpreted as shocks that represent risk to the individual.[14]

Figure 5 plots trends in the (a) volatility, (b) skewness and (c) kurtosis of these shocks. As documented in the first set of GRID countries (Guvenen, Pistaferri and Violante, 2022b) the distribution of shocks does not resemble a normal distribution because it exhibits more skewness and fatter tails (kurtosis) than a normal distribution.[15]

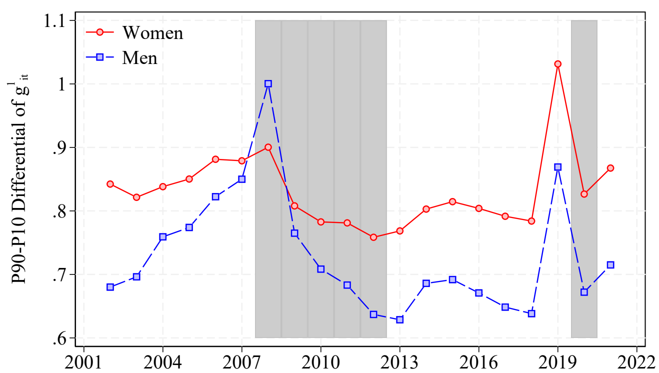

(a) Volatility

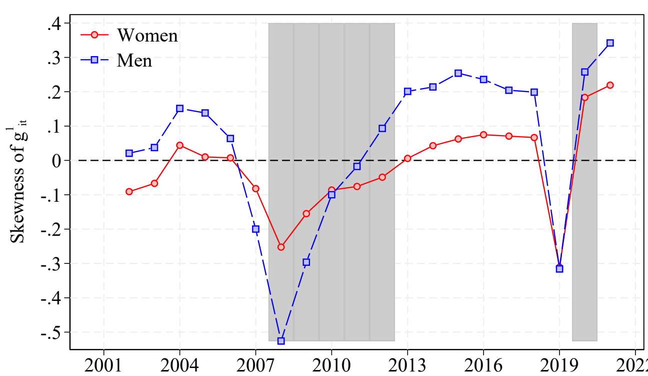

(b) Skewness

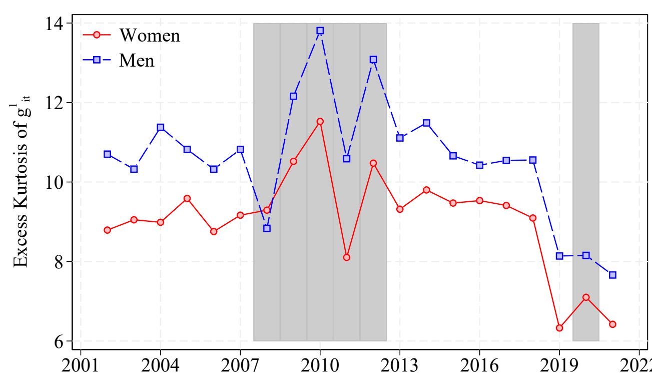

(c) Kurtosis

Figure 5: Trends in Volatility, Skewness and Kurtosis of Earnings Growth

Note: Volatility is the P90-P10 of earnings growth, skewness is Kelley skewness of earnings growth, and kurtosis is the 4th moment of the earnings growth distribution. Shaded areas indicate recession years, based on declines in Modified Domestic Demand.

[14]This interpretation comes with the caveat that some of these earnings changes may still reflect changes that individual workers have more information about (like the likelihood of being promoted) or changes arising from labour supply choices on behalf of the individual.

[15]The distributions in 2005 are reported in Appendix Figure A9 alongside a normal distribution with the same standard deviation.

Overall, Ireland can be characterized as exhibiting low to modest levels of earnings risk. Panel (a) of Figure 5 documents a modest rise in the volatility of shocks in the lead-up to 2008. Even with this rise, Ireland has low to modest levels of volatility, ranking 6th relative to the first wave of GRID and closest to Sweden (Mexico and Brazil are highest and Germany is lowest). The higher volatility among women is common in the countries in the first wave of GRID. This likely reflects that women are more likely to work part time and thus (voluntary and involuntary) changes in hours will pass through to changes in earnings.

Panels (b) and (c) of Figure 5 show that skewness is pro-cyclical, as is commonly found across countries (Guvenen et al., 2022b).[16] In good times, the positive skewness means that good shocks are more common than bad shocks. But in bad times, the distribution becomes strongly negatively skewed and the likelihood of bad shocks is more pronounced. Men are more likely to be exposed to these cyclical shocks as illustrated by the bigger changes in skewness and kurtosis.

The extent to which the tails of the risk distribution exhibit fat tails or excess kurtosis (relative to a normal distribution) tells us how likely *very* large changes in earnings are to occur. One measure of fatness is the slope of the log-density (as reported in Figure A7) where a smaller absolute value implies a fatter tail. Ireland's tails might be described as intermediate, as they lie between the extremes of the United States and Sweden (McKinney, Abowd and Janicki, 2022; Friedrich, Laun and Meghir, 2022). The slope for the right tail for Irish men is -2.5, while the US is -2.1 and Sweden is -3.37.[17] In all three countries, the left tail is fatter than the right, implying that extremely large negative earnings shocks are more likely than extremely large positive shocks.

2.4 Earnings Risk and Permanent Income

Figure 6 documents how the same measures of volatility, skewness and kurtosis vary over the permanent earnings distribution (where permanent income is estimated using the past three years of earnings). The patterns in all three plots are broadly consistent with the first wave of GRID.

The top panel of Figure 6 illustrates that volatility is higher for younger age groups and those with lower permanent incomes. As in other countries, the exception to falling volatility for higher permanent incomes is in the highest permanent income percentile where volatility is higher than the 99th percentile. The increase at lower permanent income percentiles in Ireland is less convex than some other countries.

The middle panel shows that negative skewness is larger for older and higher permanent earnings

[16] The figures for skewness and kurtosis are noisy. We believe that this is due to our short time horizon and reasonably small sample. For example, the estimates are noisier when using our secondary dataset (EADDS), which has a shorter time horizon. Similar levels of noise are present in Argentina (Blanco et al., 2022), which has the same time horizon as us (20 years) and a slightly smaller sample size (in the earliest year, Argentina and Ireland have approximately 100,000 and 120,000 observations respectively).

[17] The Pareto tail coefficient for the left tail is 1.54, while the US is 1.38 and Sweden is 1.87.

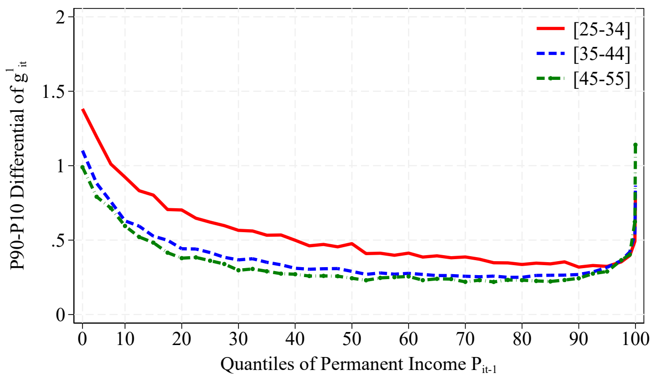

(a) Men

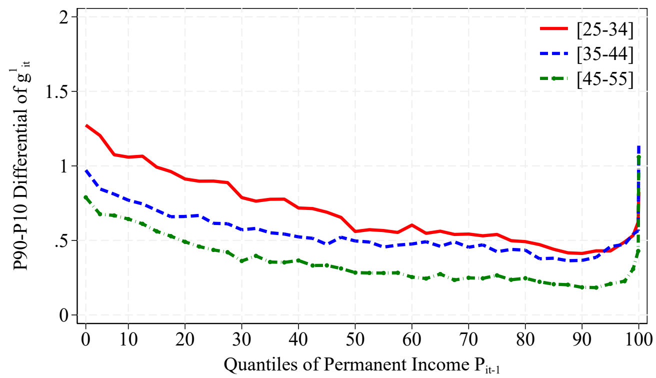

(b) Women

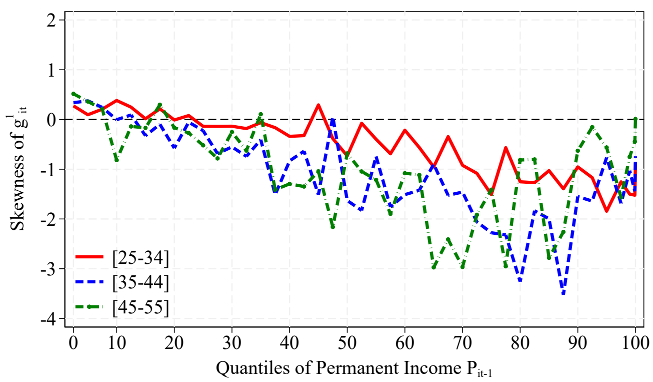

(c) Men

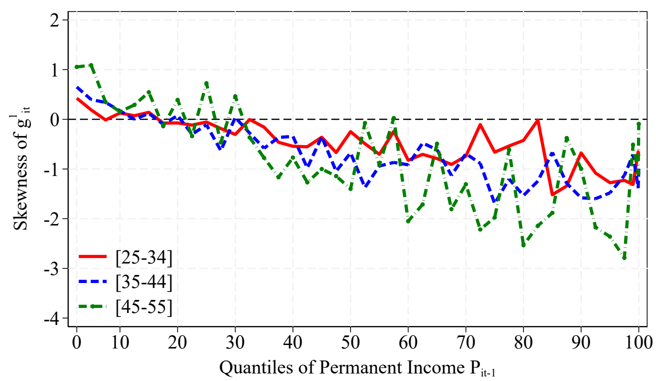

(d) Women

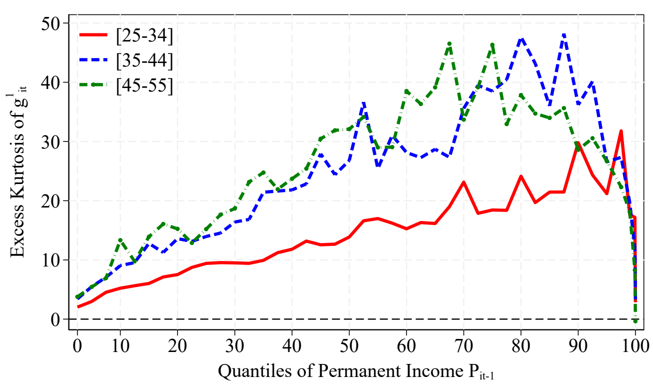

(e) Men

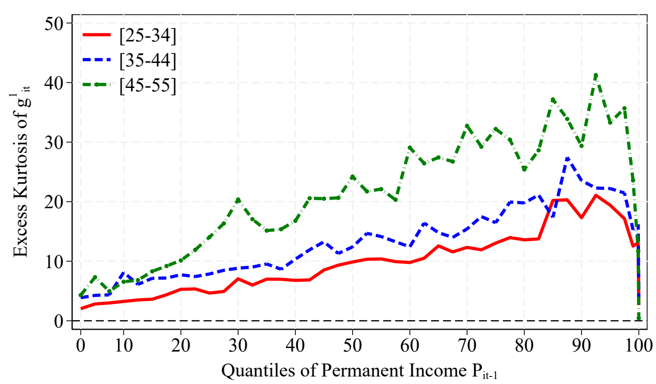

(f) Women

Figure 6: Dispersion of earnings growth by permanent earnings and age

Note: Every panel shows the relationship between permanent income and moments of the 1-year earnings growth distribution. Top panels show the volatility of earnings growth; middle panels show the skewness of earnings growth; and bottom panels show the excess kurtosis of earnings growth.

individuals. This may reflect that large negative shocks become more likely as individuals progress in their career. The bottom panel shows that large shocks—as measured by excess kurtosis— are also more likely for higher-income and older individuals.

Interestingly, the patterns tend to change for the very top percentile, which exhibits higher volatility and lower skewness and kurtosis than the percentiles just below it. This is a robust fact across countries.

2.5 Earnings Mobility

Figure 7 documents the persistence of earnings over time which is a measure of economic mobility. The figure is a rank-rank plot showing the mean percentile rank of individual earnings against their rank five years earlier. The slope of the line measures the degree of persistence (or immobility) across time. A slope of 1 (the dashed line) means no mobility; individuals stay at the same relative position over time. A slope of 0 means perfect mobility; one's past rank does not predict their current rank.

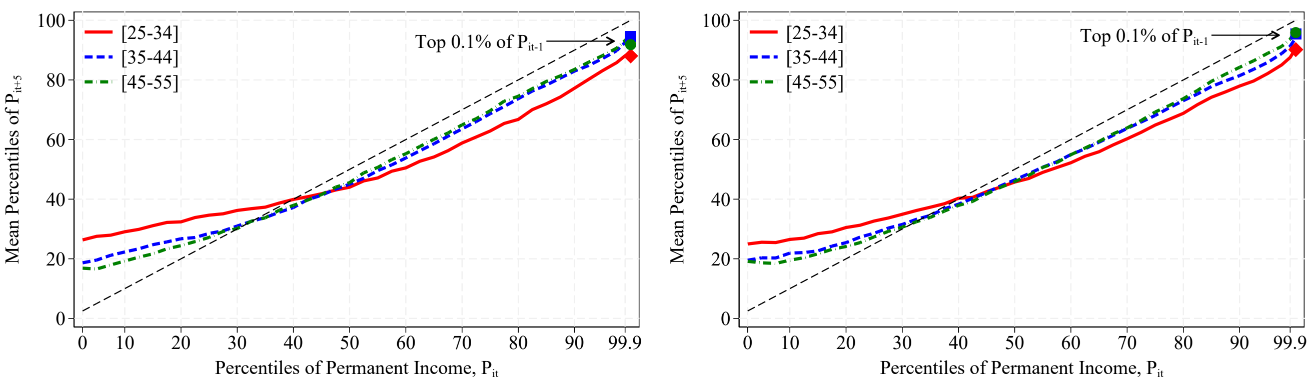

(a) Men

(b) Women

Figure 7: Evolution of 5-year mobility over the life cycle

Note: The figure shows average rank-rank for men and women and for different age groups. Sample includes only those observed at both t and t+5. Permanent income is based on 3 years of income, t-2, t-1, t. The black dashed line is the 45-degree line inserted to indicate what would be expected if there is no mobility.

The patterns are broadly consistent with the first wave of GRID. There is more mobility in the first decade of working life and less so at older ages. Individuals who begin in the very lowest permanent income percentiles are likely to rise to at or above the 20th percentile. While individuals at higher percentiles are likely to fall to lower percentiles, this mean reversion becomes less likely at the very top.

We compare across countries using the linear slope of the rank-rank line.[18] Ireland's slope

[18]We estimate the slope β from the following regression, $rank_{t+5,p} = \alpha + \beta rank_{t,p} + \epsilon_p$, where p is a percentile

$\beta = 0.75$ is in the middle of the pack. This is similar to the Canada, the United States, and Germany. Less mobile countries like Brazil and the UK have a slope coefficient of 0.87 while more mobile countries like Sweden have a slope coefficient of 0.67.

2.6 Comparing Ireland with Other Countries

A key contribution of the GRID database (Guvenen, Pistaferri and Violante, 2022a) is that statistics are directly comparable across countries. In this section, we compare the Irish statistics with the countries in the first wave of GRID.[19]

Inequality over Time Figure 8 compares Ireland in green with the time series of inequality across countries using three measures: earnings inequality (P90-P10 of log earnings), top earnings inequality (P99.9-P50 of log earnings), and earnings risk (P90-P10 of residualized earnings growth). The left panels show the trends in levels (for the entire time series available for each country) and the right panels show the change relative to the start of our sample (2002).

The comparisons reveal three patterns. First, the level of inequality is relatively high and stable in Ireland. The P90-P10 difference in log incomes is consistently around 2. Of the advanced economies, only the US and Germany are consistently higher at around 2.4, while Denmark and Sweden are substantially lower at or below 1.5.[20] It is worth noting that our data is for *individual pre-tax* earnings and Roantree and Barrett (2023) show that inequality in *household disposable* income is much lower after accounting for taxes, transfers and household composition.

Second, while the level of top inequality was moderately low at the beginning of the sample, it has grown faster than any of the countries in the first wave of GRID over the same period. In 2002, the P99.9-P50 difference in log incomes was close to 2.1, with only the Scandinavian countries (Denmark, Norway, Sweden) exhibiting lower levels. Since then, it has grown 50 log points which is more than any other country in the sample over the same period. Its final level, at just above 2.5, is still considerably below the United States which is above 3.

Third, while Ireland's overall level of earnings risk generally falls in the middle range relative to other countries, it saw the most pronounced rise in earnings risk in the lead-up to the financial

rank.

[19]To avoid repeated citations as we compare countries we list the relevant GRID country studies here: Argentina (Blanco, Diaz de Astarloa, Drenik, Moser and Trupkin, 2022), Brazil (Engbom, Gonzaga, Moser and Olivieri, 2022), Canada (Bowlus, Gouin-Bonenfant, Liu, Lochner and Park, 2022), Denmark (Leth-Petersen and Sæverud, 2022), France (Kramarz, Nimier-David and Delemotte, 2022), Germany (Drechsel-Grau, Peichl, Schmid, Schmieder, Walz and Wolter, 2022), Italy (Hoffmann, Malacrino and Pistaferri, 2022), Mexico (Puggioni, Calderón, Zurita, Bujanda, González and Jaume, 2022), Norway (Halvorsen, Ozkan and Salgado, 2022), Spain (Arellano, Bonhomme, De Vera, Hospido and Wei, 2022), Sweden (Friedrich, Laun and Meghir, 2022), the UK (Bell, Bloom and Blundell, 2022), and the US (McKinney, Abowd and Janicki, 2022).

[20]The emerging economies Argentina, Brazil and Mexico exhibit higher inequality on most measures in Figure 8.

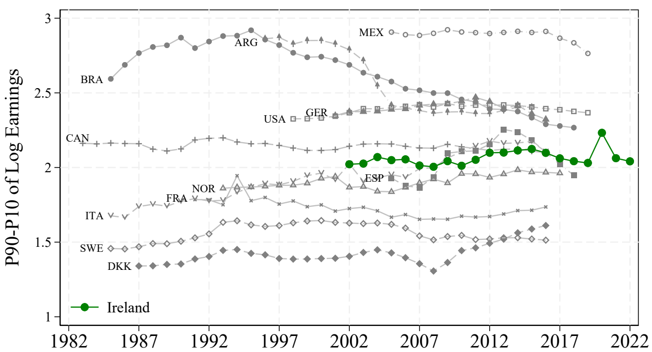

(a) Earnings Inequality (Level)

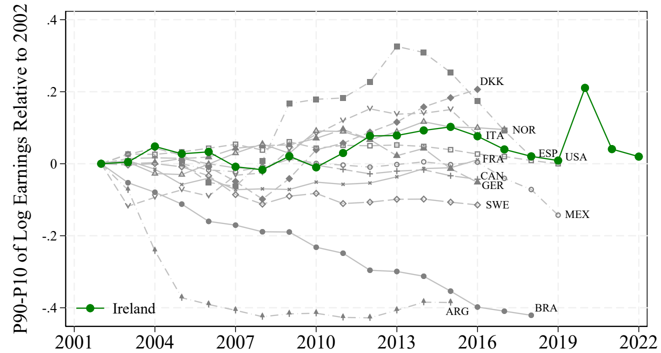

(b) Earnings Inequality (Change since 2002)

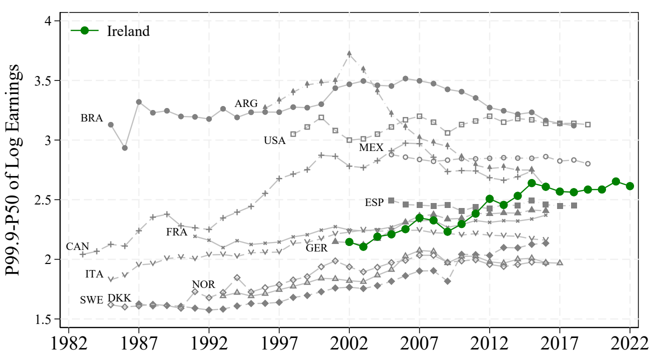

(c) Top Earnings Inequality (Level)

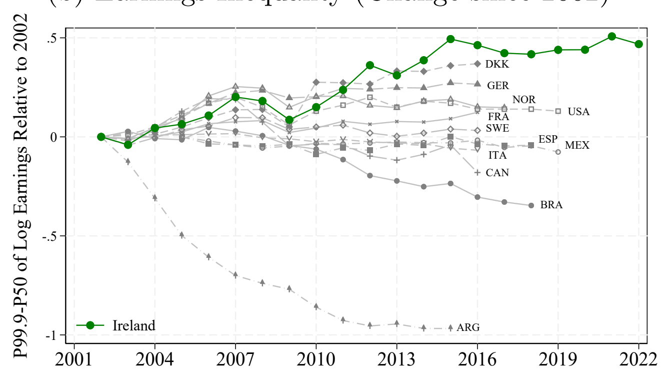

(d) Top Earnings Inequality (Change since 2002)

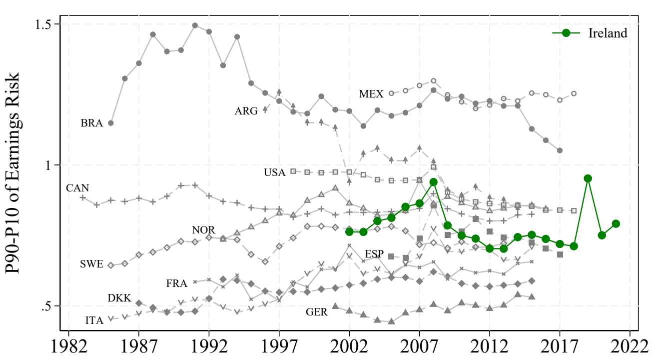

(e) Earnings Risk (Level)

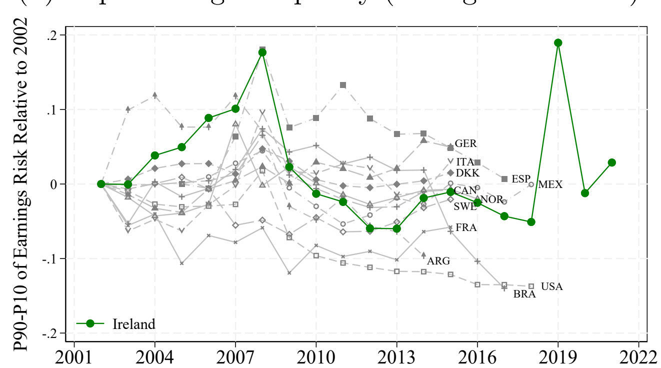

(f) Earnings Risk (Change since 2002)

Figure 8: Inequality and Risk across Countries and Time

Note: Figure compares inequality across countries in GRID and Ireland (in green). Top panel compares earnings inequality as measured by the P90-P10 of log earnings from Section 2.1. Middle panel compares top earnings inequality as measured by P99.9-P50 of log earnings. Bottom panel compares the earnings risk, as measured by the P90-P10 of (residual) earnings growth from Section 2.3. Each statistic is for all workers. Data comes from GRID database and Irish analysis.

crisis. In 2002, the P90-P10 of (residual) log earnings growth was around 0.75, which was in between the United States (close to 1) and Germany (close to 0.5). In the subsequent years risk increased almost 20 log points, with only Spain experiencing a similarly large increase in risk, before falling back to around or below its starting level. Notably, much of the increase in earnings risk actually occurred in the latter half of the boom from 2003 onward— which is unlike others who mostly saw increases in 2007 and/or 2008. Our results in Section 4 suggest this increase may be related to inward migration during these years.

Relationship between Inequality, Risk and Mobility Figure 9 plots the relationship between earnings inequality (P90-P10 log earnings), risk (P90-P10 of residual earnings growth), and mobility (the slope of a rank-rank regression on earnings discussed in Section 2.5). A higher rank-rank beta implies higher persistence of earnings and less mobility. There is a positive relationship between the rank-rank slope and both inequality and (to a lesser extent) risk—implying countries with higher inequality are also less mobile.

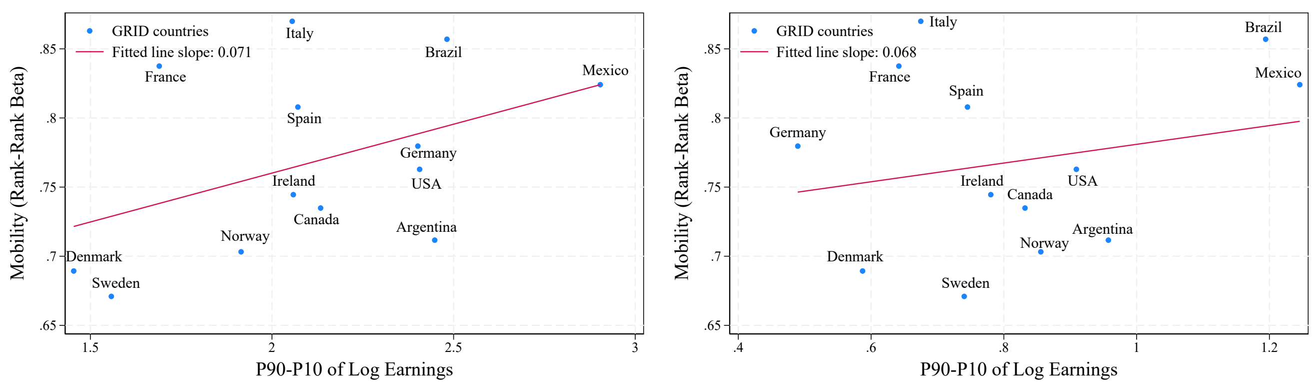

(a) Inequality and Mobility

(b) Risk and Mobility

Figure 9: Earnings Inequality, Risk and Mobility

Note: Figure compares cross-country patterns in earnings mobility against earnings inequality (left) and earnings risk (right). Earnings mobility is measured using the coefficient from a rank-rank regression discussed in Section 2.5, earnings inequality is the P90-P10 of log earnings from Section 2.1 and risk is measured by the P90-P10 of residual earnings growth from Section 2.3. Data comes from GRID database and Irish analysis.

The Scandinavian countries (Denmark, Sweden, and Norway) are in the bottom left quadrant of panel (a) illustrating that they are both mobile and relatively equal. Ireland and Canada are next closest, indicating that they also perform relatively well on these metrics. For instance, Ireland has higher than expected mobility given its level of inequality, especially compared with Spain and Italy who have similar inequality. It also has much less inequality than the United States who have a similar level of mobility.

3 Responsiveness of Earnings to Output and Inflation

In this section, we study how earnings growth comoves with aggregate output growth and inflation, both on average and across the distribution of earnings growth. Our data enable two extensions. First, because we observe weeks worked, we can decompose the responsiveness of annual earnings into the responsiveness of weekly earnings and weeks worked. Second, we extend the traditional output-beta approach of Guvenen, Schulhofer-Wohl, Song and Yogo (2017) by including inflation as an additional aggregate regressor, allowing us to estimate the relationship between inflation and the distribution of earnings growth.

3.1 Estimation Framework

Consider worker i at time t from group g, where groups are the combination of gender, age, and permanent income. We estimate the sensitivity of income growth Δy_{it} to *aggregate* output Y_t

$$\Delta y_{it} = \alpha^g + \beta^g \Delta Y_t + \varepsilon_{it}, \tag{1}$$

where Δy_{it} is the change in log earnings, or its equivalent for log weekly earnings Δy_{it}^W, or log weeks worked Δw_{it}.[21] We estimate this equation separately for sub-groups and therefore α^g and β^g are allowed to depend flexibly across groups g, though we suppress the g superscript unless necessary for exposition.

Average. Our first approach estimates the average responsiveness by estimating (1) using ordinary least squares (OLS) (Guvenen, Schulhofer-Wohl, Song and Yogo, 2017). The coefficient α corresponds to the average growth rate of earnings when aggregate output growth is zero, and β corresponds to the elasticity of average earnings growth to output growth.

A contribution of our paper is to decompose the responsiveness of earnings to aggregate variables into the responsiveness of *weekly earnings* and the responsiveness of *weeks worked*. Changes in (log) real earnings Δy_{it} can be decomposed into changes in (log) weekly real earnings Δy_{it}^W and (log) weeks worked Δw_{it}, $\Delta y_{it} = \Delta y_{it}^W + \Delta w_{it}$. Since OLS is a linear operator, this means that the responsiveness of total earnings β_y can also be decomposed into these components

$$\beta_y = \beta_{y^W} + \beta_w. \tag{2}$$

We call the ratio of the beta on weeks and annual earnings β_w/β_y, the *weeks share* of exposure.

[21] Earnings, weeks worked, and weekly earnings are as defined in Section 1, whereas the non-residual Δy_{it} differs from the changes in residual log income g_{it} which was defined in the same section. We continue to use modified domestic demand as our preferred measure of aggregate output. To be consistent with the GRID income variables, we use one-year forward changes.

In addition to being inherently informative about nature of income changes, the comparison will also be informative about the patterns investigated in the next Section (4). In particular, entering new migrants (who work only part of the year) are more likely to alter the distribution of annual earnings than that of weekly earnings.[22]

Distribution. Our second approach estimates the heterogeneous effects of aggregate variables on the distribution of earnings growth, by estimating (1) with quantile regression (Floccari and Karpati, 2026).[23] The coefficient β_q corresponds to the elasticity of earnings growth at the q-th quantile to changes in output. As Floccari and Karpati (2026) show, a convenient advantage of the quantile regression approach is that we can use the estimated coefficients to construct any quantile-based higher-order moments of the distribution, such as volatility and skewness.[24] Since quantile regression is not a linear operator, we cannot *decompose* the estimates into share contributed by weeks worked and weekly earnings; instead we present these estimates separately.

3.2 Results: Output

Average. Table 2 reports the estimated coefficients and standard errors. The first column reports OLS betas for annual earnings, and the second column reports the weeks share. Since we estimate β coefficients separately for each combination of gender, age group, and permanent income percentile, the table reports averages of these coefficients within the groups shown in each panel, and we report the full set of coefficients in Figure C16.[25]

First, we find that the exposure to output is significantly higher for workers with lower permanent income and modestly higher for men and for younger workers. The responsiveness of the 10th percentile is eight times larger than the responsiveness of the 90th percentile; the responsiveness of men is twice as large as for women; and the responsiveness of the youngest group (25-34) is almost twice as large as the oldest group (45-55). The results are consistent with Guvenen, Schulhofer-Wohl, Song and Yogo (2017) who highlighted the U-shaped pattern for exposure across the permanent income: exposure is generally falling at higher permanent income percentiles apart from the very top where it rises.

Second, weeks worked contribute a substantial share of the exposure for most workers, with the

[22]It is not immediately obvious that selection that affects the level of percentiles of the *annual* cross-sectional distribution will substantially affect the growth rate of earnings based on percentiles of the *permanent* income distribution.

[23]An alternative to quantile regression is to construct quantiles and regress them on output directly—an approach taken by Busch et al. (2022). We find similar patterns with the quantile approach—in particular the cyclicality of skewness is important— suggesting these patterns are robust to the specific methodology used.

[24]See Appendix C and Floccari and Karpati (2026) for details.

[25]Since these are OLS regressions for heterogeneous sub-groups, the weeks share decomposition can be applied, unlike in the case of quantile regressions.

Table 2: Decomposition of Average Responsiveness to Output

Group	**Mean** β_y	**Weeks share,** β_w/β_y
Panel A: Gender		
Men	0.827	0.293
	(0.056)	(0.024)
Women	0.405	0.114
	(0.112)	(0.071)
Panel B: Percentiles		
p10	1.250	0.168
	(0.033)	(0.019)
p50	0.494	0.441
	(0.015)	(0.018)
p90	0.152	0.612
	(0.011)	(0.044)
Panel C: Age Groups		
Age 25–34	0.867	0.156
	(0.181)	(0.050)
Age 35–44	0.513	0.245
	(0.042)	(0.034)
Age 45–55	0.467	0.369
	(0.029)	(0.038)

Notes. Table reports OLS output beta estimates within each group.The β coefficients are estimated separately for each gender $\times$ age $\times$ permanent income group, in a regression including output. The weeks share $= \beta_w/\beta_y$ is the weeks-worked beta as a fraction of the annual earnings beta. Standard errors are based on 100-replicate bootstrap.

weeks share ranging between 5 percent and 61 percent. The relationship between total exposure and the contribution of weeks worked is non-monotonic. Men are more exposed to cyclical earnings risk, and weeks worked account for a higher fraction of the exposure than for women: 29 percent compared to 11 percent for women. On the other hand, even though older workers are less exposed to cyclical earnings risk, weeks worked account for a higher share of that exposure: 37 percent for the older compared to 16 percent for the youngest. Since wage changes affect weekly earnings, and unemployment spells affect the number of weeks worked, these differences suggest that there is substantial heterogeneity in the types of risks faced by different workers.

Distribution. Table 3 reports the results for the responsiveness of the volatility (P90-P10) and skewness of earnings growth to output. Figures C17 and C18 report the coefficients for each combination of gender, age group, and permanent income percentile, and the table reports averages of these coefficients within the groups shown in each panel. Since the decomposition doesn't apply to quantiles, we separately report the responsiveness of the *weekly earnings* growth distribution, and of the *weeks worked* growth distribution.

Figure C17 shows that volatility of earnings growth is counter-cyclical, while the skewness of earnings growth is pro-cyclical, and both patterns are more pronounced for men than women. The pro-cyclicality of skewness is consistent with evidence from the United States, Germany and Sweden (Busch et al., 2022) and the Netherlands (Floccari and Karpati, 2026). We find that the counter-cyclicality of volatility in Ireland is larger and more consistently present than Busch et al. (2022), who argue that it is small and flat. For example, they find different signs depending on the sample country and gender, whereas we find the sign to be consistently negative.

Comparing weekly earnings and weeks worked in Table 3 highlights a qualitative difference between these two measures. The counter-cyclicality of volatility is associated with counter-cyclical volatility of *weeks worked*, while the volatility of *weekly earnings* is strongly pro-cyclical at the bottom of the permanent income distribution. The different signs for volatility of weekly earnings and weeks worked may help explain the different signs on earnings growth in other countries (Busch et al., 2022).

Table 3: Responsiveness of Higher Moments to Output

	$\beta_{P90-P10}$		β_{Skewness}	
Group	**Weekly earn.**	**Weeks**	**Weekly earn.**	**Weeks**
Panel A: Gender				
Men	-0.088	-0.687	1.940	0.798
	(0.055)	(0.045)	(0.067)	(0.094)
Women	0.011	-0.406	1.319	0.386
	(0.039)	(0.043)	(0.068)	(0.091)
Panel B: Percentiles				
P10	1.242	-0.722	2.594	1.272
	(0.049)	(0.103)	(0.067)	(0.123)
P50	-0.219	-0.716	1.665	0.830
	(0.024)	(0.046)	(0.068)	(0.292)
P90	-0.206	-0.126	1.026	0.262
	(0.016)	(0.024)	(0.063)	(2.361)
Panel C: Age Groups				
Age 25–34	0.130	-0.872	1.671	0.655
	(0.075)	(0.107)	(0.085)	(0.048)
Age 35–44	-0.053	-0.480	1.584	0.589
	(0.042)	(0.048)	(0.075)	(0.105)
Age 45–55	-0.165	-0.242	1.715	0.477
	(0.045)	(0.029)	(0.086)	(0.262)

Notes. Table reports higher-order moment betas computed from quantile intercepts and slopes via the formulas in the text. Each column represents a separate regression for weekly earnings and weeks worked. We estimate coefficients separately for each combination of gender, age group, and permanent income percentile, and this table reports averages within each group.

Skewness, on the other hand, is consistently pro-cyclical throughout the distribution—consistent with the evidence in Guvenen et al. (2014) and Busch et al. (2022). The cyclical skewness is stronger for men than women, and for those at the bottom of the permanent income distribution. Both weeks worked and weekly earnings exhibit pro-cyclical patterns, though the covariance is stronger for weekly earnings.

3.3 Results: Inflation

We now extend the analysis to include both aggregate output growth and inflation, denoted by Π_t.[26] We estimate

$$\Delta y_{it} = \alpha^g + \tilde{\beta}^g \Delta Y_t + \beta^g \Pi_t + \varepsilon_{it}. \quad (3)$$

[26] Inflation is defined as the log change in the CPI: $\Pi_t = \log(\text{CPI}_t) - \log(\text{CPI}_{t-1})$.

Consistent with the beta interpretation used above, the coefficients summarize the conditional comovement of real earnings growth with aggregate output growth and inflation. In particular, β is the inflation beta conditional on contemporaneous output growth. Because the dependent variable is real earnings, $\beta = 0$ is a natural benchmark: it implies that real earnings are unchanged as inflation varies or, equivalently, that nominal earnings keep pace one-for-one with prices. A negative coefficient indicates that nominal earnings do not fully keep pace with inflation, resulting in lower real earnings.[27]

Table 4: Decomposition of Average Responsiveness to Inflation

Group	**Mean** β_y^π	**Weeks share,** β_w^π/β_y^π
Panel A: Gender		
Men	0.003	-17.138
	(0.118)	(3.397)
Women	0.148	-0.703
	(0.268)	(740.919)
Panel B: Percentiles		
p10	0.808	-0.547
	(0.066)	(0.104)
p50	-0.105	-0.311
	(0.032)	(0.454)
p90	-0.224	0.001
	(0.024)	(0.082)
Panel C: Age Groups		
Age 25–34	0.240	-0.943
	(0.426)	(3177.797)
Age 35–44	-0.038	0.947
	(0.087)	(10.775)
Age 45–55	0.024	1.423
	(0.060)	(7.500)

Notes. Table reports OLS inflation beta estimates within each group. The β coefficients are estimated separately for each gender × age × permanent income group, in a regression including both output and inflation. The weeks share = β_w/β_y is the weeks-worked beta as a fraction of the annual earnings beta. Standard errors are based on 100-replicate bootstrap.

[27] Inflation and output growth are positively correlated in our sample. Consequently, an inflation-only specification partly reflects the comovement common to inflation and output growth and produces estimates similar to those from the output-only analysis above. Including both variables allows us to ask whether inflation has an additional association with earnings after accounting for the relationship with output explored in the previous section.

Average. Table 4 reports the estimated average inflation-beta, while Figure 10 presents the full set of coefficient estimates.[28] We focus our discussion on the relationship with inflation—Table C2 shows that the estimated output coefficient, $\tilde{\beta}$, is very similar to the specification without inflation.

On average, we find that nominal earnings move one-for-one with inflation. The coefficient for men is zero, and the coefficient for women is positive but statistically indistinguishable from zero, implying that real earnings are unaffected by changes in inflation.[29]

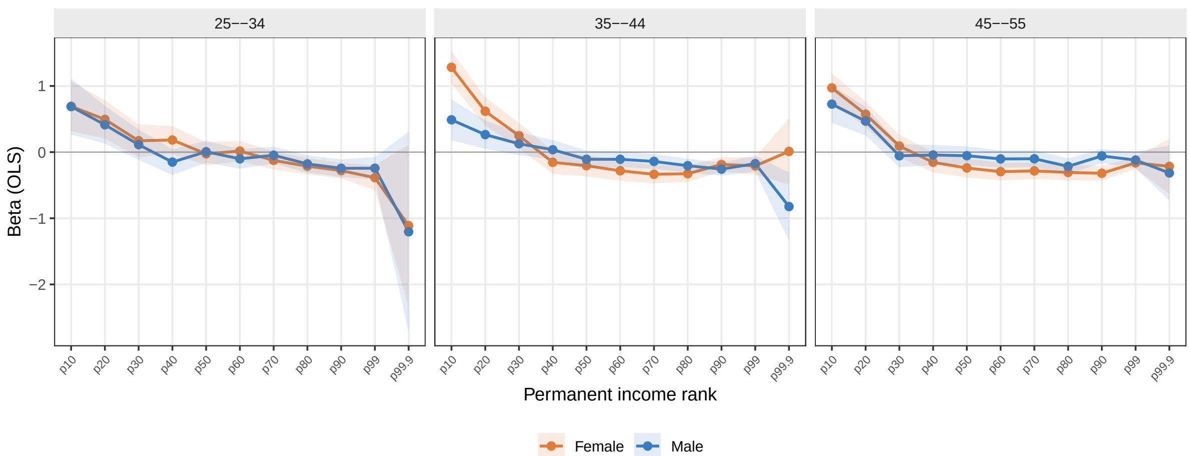

(a) Annual Earnings

Figure 10: Average Responsiveness to Inflation (OLS Beta)

Note: Figure plots average (OLS) betas for annual earnings, for each age group, gender and permanent income percentile.

Figure 10 shows that there is striking heterogeneity in the average effect of inflation across the permanent income distribution. For most permanent income percentiles, the impact of inflation is close to zero, or modestly negative, implying almost complete pass-through. However, there is more than complete pass-through at the bottom of the permanent income distribution: for those at the P10, a one percentage point increase in inflation is associated with an increase in *real* incomes of 0.8 percentage points. The pattern is consistent across men and women, and for each age group. In Figure C19, the decomposition between weeks worked and weekly earnings shows that it is driven by weekly earnings. At the P10 the weeks share is -0.54, meaning that the weeks worked in fact fall in response to inflation, and thus that weekly earnings rise by even more than

[28] Since the weeks share is a ratio, the estimates and standard errors are imprecise when the denominator β_y^π is close to zero.

[29] There may appear to be a tension between the result of complete pass-through in this section, and the results in Appendix A which suggested that inflation eroded nominal gains. However, the regressions estimate the average contemporaneous covariance of inflation over the entire sample period, and as such do not take into account nonlinearities, changing relationships, or lagged effects. Indeed, to estimate the *causal* pass-through of unexpected inflation would require identifying shocks to inflation.

the coefficient on annual earnings.[30]

Section 4 indicates that migration-related selection affects the measurement of lower percentiles, so it is worth considering whether the selection highlighted in the next section explains the patterns. Section 4 also indicates that such migration-related selection is most pronounced in the male annual-earnings distribution, whereas the patterns for inflation-betas at the bottom appear across both sexes (and across age groups). This makes it less likely that migration-related composition alone accounts for the results, although it does not eliminate selection as a contributing mechanism.

Table 5: Responsiveness of Higher Moments to Inflation

	$\beta_{P90-P10}$		β_{Skewness}	
Group	**Weekly Earn.**	**Weeks**	**Weekly Earn.**	**Weeks**
Panel A: Gender				
Men	-0.249	0.007	1.636	-0.427
	(0.082)	(0.060)	(0.114)	(0.165)
Women	0.028	-0.021	0.940	-0.509
	(0.082)	(0.069)	(0.104)	(0.214)
Panel B: Percentiles				
P10	0.800	0.660	1.097	-0.977
	(0.199)	(0.204)	(0.145)	(0.220)
P50	-0.688	-0.360	1.883	-0.406
	(0.105)	(0.102)	(0.156)	(0.400)
P90	0.125	0.072	0.829	-0.880
	(0.061)	(0.045)	(0.169)	(2.849)
Panel C: Age Groups				
Age 25–34	-0.473	-0.016	0.574	-0.408
	(0.142)	(0.150)	(0.133)	(0.102)
Age 35–44	-0.057	-0.018	1.371	-0.481
	(0.096)	(0.084)	(0.136)	(0.184)
Age 45–55	0.052	0.016	2.000	-0.026
	(0.074)	(0.035)	(0.142)	(0.548)

Notes. Table reports higher-order moment betas computed from quantile intercepts and slopes via the formulas in the text. Each column represents a separate regression for weekly earnings and weeks worked. We estimate coefficients separately for each combination of gender, age group, and permanent income percentile, and this table reports averages within each group.

[30] To the best of our knowledge, we are the first to document such a relationship between inflation pass-through and permanent income. A fruitful open question is to explore whether these patterns hold systematically across countries or whether they reflect the particulars of Ireland or this time period, and whether similar patterns hold for causally identified *shocks* to inflation.

Distribution. Table 5 reports the results for the relationship between inflation and the volatility and skewness of earnings growth. The table reports averages of these coefficients within the groups shown in each panel, and we report the full set of coefficients in Figures C20 and C21.

In terms of the distribution, inflation is consistently associated with an increase in the skewness of weekly earnings growth: every estimate in the third column of Table 5 is positive and statistically significant. The contribution from weeks worked is generally negative, although these estimates are considerably less precise and are not statistically distinguishable from zero for several groups. This pattern is consistent with staggered wage adjustment: if many workers receive no nominal adjustment in a given period while a smaller group receives relatively large increases, the distribution of earnings growth becomes more right-skewed.

In terms of volatility, there is no uniform pattern. However, Figure C20 shows that volatility increases most at the bottom percentiles of the permanent income distribution. At most percentiles, inflation is associated with a fall in volatility or little change. Thus, inflation appears to increase volatility primarily for low permanent-income workers, the same part of the distribution that exhibits the strongest average response.

4 Measurement with Selection

As we noted in Section 2, the patterns in Figure 2 were puzzling because the boom-bust pattern that is clearly evident in output was not especially striking for labour earnings. For example, incomes at the bottom actually fell during the boom. To explore the role of selection, we reconstruct the statistics from Section 2 using alternative measures of earnings, each of which illuminates different mechanisms at work.

4.1 Annual Earnings, Weekly Earnings, and Weeks Worked

Our focus on selection is motivated by Figure 11, which plots the distribution and P90-P10 ratio for men using our original measure of annual earnings and weekly earnings for men, where weekly earnings are defined as annual earnings divided by the number of weeks worked.[31] Weekly earnings remove changes in annual earnings that arise from changes in the number of weeks worked (which may include voluntary changes and involuntary changes like unemployment spells). Weekly earnings are closer to the concept of an hourly wage than annual earnings. However, since we do not observe hours, weekly earnings still include changes in the number of hours worked per week.

The bottom panels illustrate a sharper relationship between inequality and the business cycle when earnings are measured in weekly (right) rather than annual earnings (left). Based on weekly earnings, the boom in the early 2000s was broad-based with flat inequality, while inequality rose

[31] See Section 1 for the precise definition relating to multiple job holders.

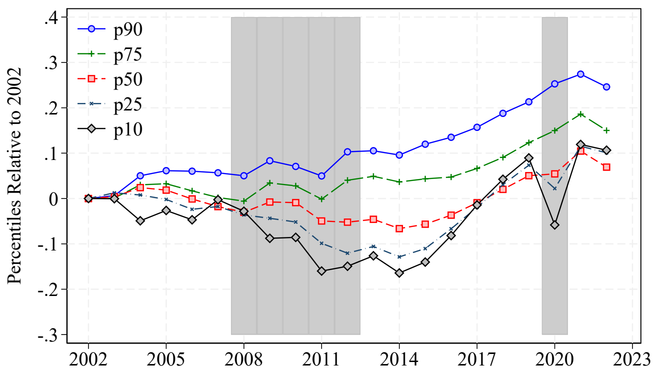

(a) Distribution - annual

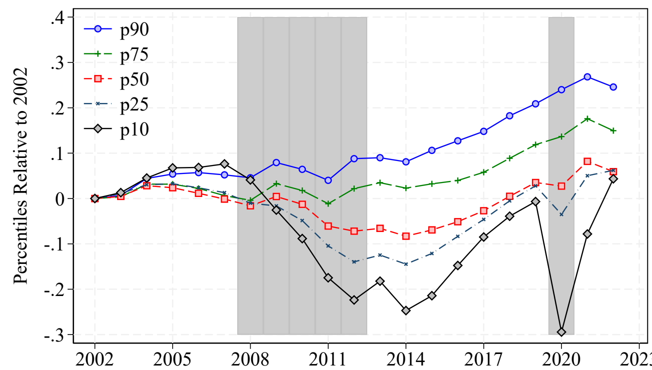

(b) Distribution - weekly

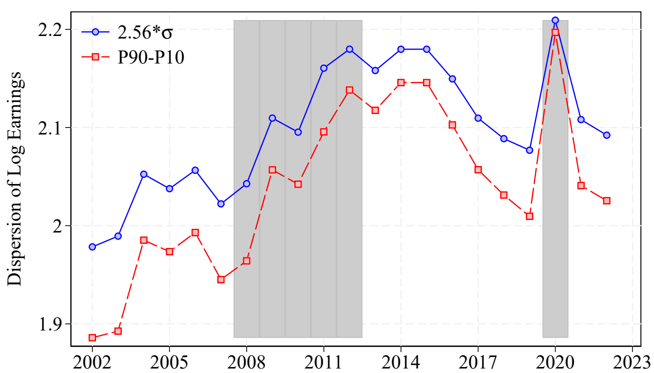

(c) P90-P10 - annual

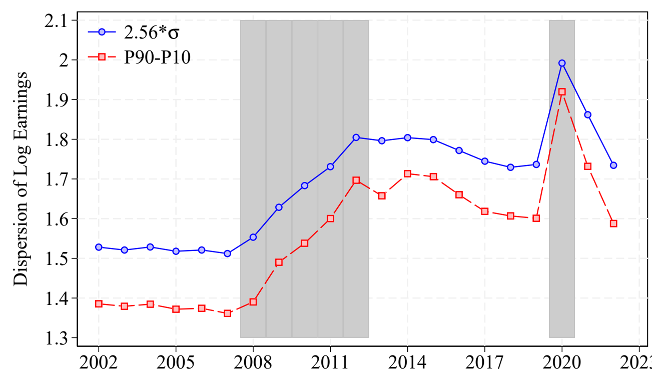

(d) P90-P10 - weekly

Figure 11: Annual vs Weekly Earnings, Men, percentiles over time

Note: Figure ranks workers using the relevant metric of earnings so the same individuals may not be at the same percentiles in each plot.

significantly during the Great Recession; the P90-P10 ratio rose 30 log points. In contrast, annual earnings suggest that about a third of the increase in inequality occurred during the boom years. The top panels of Figure 11 show that the differences during the boom are largest at the P10, which falls in annual earnings but rises in weekly earnings. The patterns during the rebound are more consistent across weekly and annual earnings, with the bottom of the distribution catching up with the median, whilst the top of the distribution remains above the median.

We report all the equivalent figures for women in Appendix D. Figure D22 shows that the differences in the patterns of annual and weekly earnings are smaller for women. While the 10th percentile also rises more when measured in weekly earnings, the broad patterns of inequality are unchanged. This suggests that the patterns of selection are less important for women and, for this reason, we focus on men for the remainder of this section.

Figure 12 shows that the difference between weekly and annual earnings is driven by changes in weeks worked at the bottom of the distribution. By definition, the difference between weekly and annual earnings is in the number of weeks worked, but it could also have been driven by differences in the ranking of workers between annual and weekly earnings. Average weeks worked clearly fall during the boom at the 10th and 25th percentiles for men. Weeks do not fall for equivalent women, and can therefore explain why there was such a big difference in annual earnings between men and women at the bottom. The changes in weeks are also large: at the tenth percentile, average weeks worked fell by more than 10 percent during the boom which is large enough to explain why annual earnings fell despite rising weekly earnings.

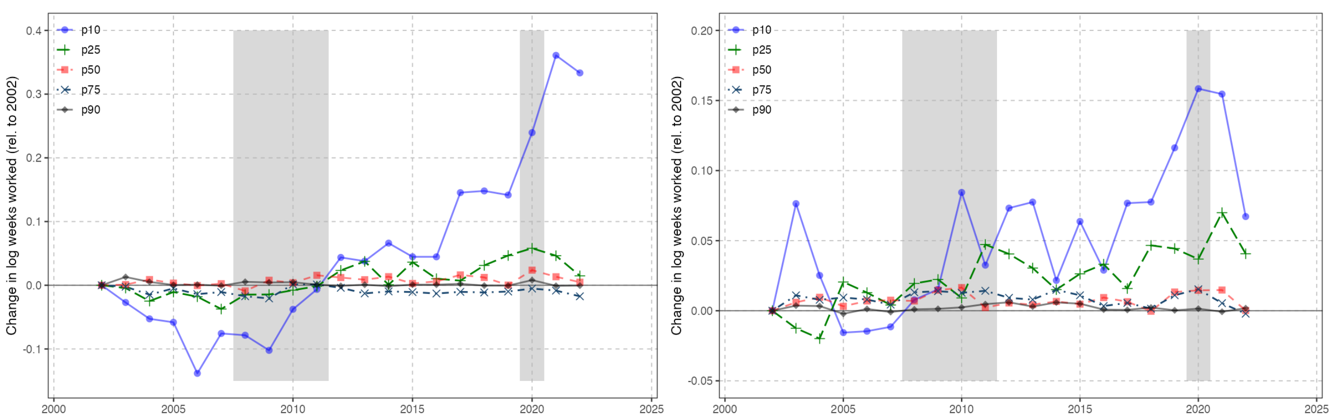

(a) Men

(b) Women

Figure 12: Average Number of Weeks Worked by Percentile

Note: Shows the average number of weeks worked (capped at 52 and summed across employments) by percentile of the annual earnings distribution over time, P10, P25, P50, P75, P90. Percentiles are formed on annual earnings each year.

4.2 Selection and Migration

We have established that changes in the number of weeks worked can account for the patterns in earnings and inequality during the boom and bust. The natural question is what economic force caused the decline in weeks worked during these periods. Figure 13 shows changes in weeks worked are associated with migration. Since new migrants may arrive at any point in the year, an increase in the number of new migrants could lead average weeks worked to fall. And this would reconcile why weeks worked fall as earnings per week increase.[32]

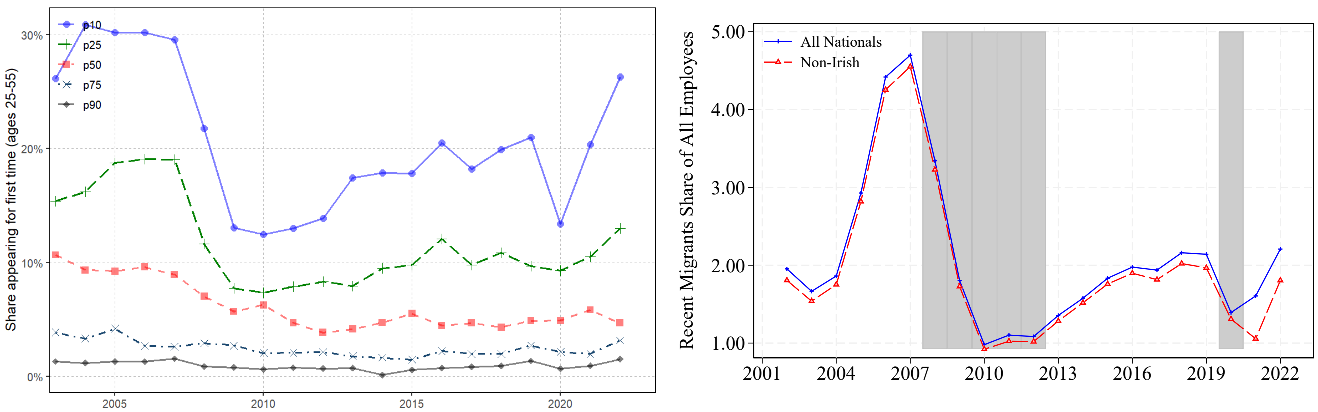

(a) New Entrants to WWLD Data

(b) Aggregate Migration Patterns in LFS Data

Figure 13: Migrants and New Entrants to Data

Note: Left panel (A) shows the share of "new entrants" (first appearance in our administrative data) by percentile and year; right panel (B) shows the share of recent migrants (resident for less than 12 months prior to the survey) in total employment from the Labour Force Survey; this is split into all migrants and non-Irish migrants, the gap representing returning Irish nationals.

In our data, we do not directly observe migrants or other indicators like nationality. However, we can observe *new entrants* to the data, which will include migrants and others (like young people) who enter the sample for the first time.[33] Panel (a) of Figure 13 plots the share of new entrants by income percentile over time. There is a sharp increase in new joiners between 2004 and 2007, and the biggest increases are in the bottom percentiles of the earnings distribution. Figure D24 shows that this pattern is robust to restricting the sample to 35-54 year-olds, which is consistent with it being driven by new migrants and not new young entrants.

To provide further evidence that the spike of new entrants is driven by migrants, Panel (b) of Figure 13 shows the share of migrants out of total employment using data from the Labour Force Survey (LFS). It shows a large increase from one percent in 2000 to over four percent in 2007.

[32] Alternatively, one would need a large income effect to explain why weeks worked and annual earnings would fall in response to a rise in weekly earnings.

[33] As noted in the data section, we begin our sample in 2002 because earnings are consistently measured since then. However, we can construct our measure of new joiners using data from before 2002.

The increase is driven by the arrival of EU migrants following the accession of ten countries in May 2004 and followed by the entry of Romania and Bulgaria in January 2007. The timing of the largest increases in 2005, 2006, and 2007 corresponds directly with the increase in new entrants in our main sample in Panel (a).

To further understand the role of selection, Figure 14 plots earnings for two restricted samples. Panel (a) is restricted to individuals working *full years* (i.e. 52 weeks) for whom changes in weeks are not relevant. Panel (b) shows a balanced panel of *survivors* who worked continuously from 2007 to 2012 which completely removes the role of new entrants to the sample. This second group includes individuals working less than 52 weeks per year.

Both restricted samples show a pattern of earnings that are more consistent with the weekly earnings measure than the annual earnings measure. There is a clear boom in the early 2000s and a corresponding bust. Most of the increase in inequality started during the fall in earnings and continues to persist. The magnitudes are also notably different compared to Figure 3: the rise in real earnings is almost twice as large as rise in the baseline samples. The fall in earnings was also substantial, even for those who were continuously employed throughout the recession.

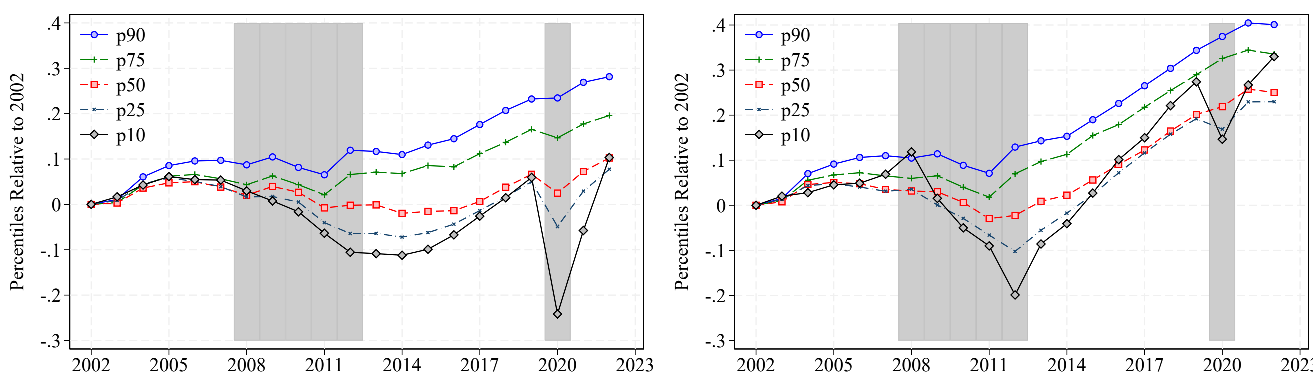

(a) Men — full-year workers

(b) Men — survivors employed 2007-12

Figure 14: Changes in the percentiles of the log real earnings distributions
Note: Log income is normalised to 0 in the base year 2002. Top panels (A) and (B) show P10, P25, P50, P75, P90. Shaded areas indicate recession years, based on declines in Modified Domestic Demand.

Another source of selection that may contribute to the differences is unemployment. Given the sharp rise in unemployment—7 percentage points in 2009—there may be significant reshuffling of individuals across percentiles as many fall down the distribution and some leave the data due to earnings below our threshold. Panel (b) of Figure 14 removes this type of selection, by examining a balanced panel of continuously employed men. A notable feature is that the increase in real earnings in 2009 (e.g. the 90th percentile in blue) is smaller than in the previous plots. This suggests that selection from reshuffling of higher-income workers to lower percentiles accounts for part of this peculiar increase at the start of the recession; the remainder is explained by deflation.

Another striking pattern is that the recovery has been much more pronounced for this group of continuously employed workers, evident in the relative steepness across all percentiles. By the end of the time horizon, all percentiles of this sample had experienced more than 20 percent growth, which is only achieved by those above the 90th percentile in the baseline sample.

Overall, the evidence in this section paints a clearer picture of earnings growth and inequality over the past two decades in Ireland. The boom was associated with an increase in earnings that appears to have been evenly spread across the distribution. Much of the large *nominal* gains were eroded by inflation resulting in modest increases in *real* earnings. The bust was associated with a fall in earnings across the distribution, with the biggest falls at the bottom of the distribution resulting in increased inequality. Since the trough of the bust, there has been strong growth in real earnings across the wage distribution, which has partly offset the inequality that arose during the bust. While the bottom of the distribution, in relative terms, has caught back up with the median earners, inequality has grown between the median and the upper percentiles of the earnings distribution.

5 Conclusion

This paper provides comprehensive, harmonized evidence on the evolution of income inequality, dynamics, and mobility in Ireland over the past two decades, contributing a new country dataset to the Global Repository of Income Dynamics (GRID). The findings reveal a distinctive pattern of rising inequality driven by gains at the top, a recovery of the lower half of the distribution following the Great Recession, and a coexistence of relatively high inequality with moderate economic mobility. Ireland's experience—shaped by both extraordinary growth and one of the most severe labour market crises in the developed world—offers a valuable case for understanding how macroeconomic changes interact with micro-level earnings dynamics.

This paper lays the groundwork for future research on the determinants of inequality and mobility in Ireland. Our analysis illustrates the power of administrative datasets to uncover important dimensions of heterogeneity that are difficult to identify in aggregate or survey data. Fruitful next steps include extending the analysis to account for the role of taxes and transfers in shaping disposable income inequality and investigating the underlying economic forces behind the documented facts. Many of the underlying forces will require the development of new administrative linkages between employees and employers, datasets that are not yet available to researchers in Ireland but that would substantially deepen our understanding of how labour markets generate and transmit inequality.

References

Afrouzi, Hassan, Andres Blanco, Andres Drenik, and Erik Hurst, "A Theory of How Workers Keep Up with Inflation," *The Quarterly Journal of Economics*, 2026, *141* (2), 945–1004.

Alamir, Anousheh, Frances McGinnity, and Helen Russell, "The Equality Impact of the Covid-19 Pandemic on the Irish Labour Market," Technical Report, The Economic and Social Research Institute (ESRI) and the Irish Human Rights and Equality Commission (IHREC), Dublin 2024.

Arellano, Manuel, Stéphane Bonhomme, Micole De Vera, Laura Hospido, and Siqi Wei, "Income risk inequality: Evidence from Spanish administrative records," *Quantitative Economics*, November 2022, *13* (4), 1747–1801.

Barrero, José María, Nicholas Bloom, and Steven J. Davis, "The Evolution of Work from Home," *Journal of Economic Perspectives*, Fall 2023, *37* (4), 23–50.

Barry, Frank, "Introduction," in Frank Barry, ed., *Understanding Ireland's Economic Growth*, Houndmills, Basingstoke, Hampshire and London: Macmillan Press, 1999, pp. 1–11.

Bell, Brian, Nicholas Bloom, and Jack Blundell, "Income dynamics in the United Kingdom and the impact of the Covid-19 recession," *Quantitative Economics*, November 2022, *13* (4), 1849–1878.

Bergin, Adele and Elish Kelly, "The labor market in Ireland, 2000–2016," *IZA World of Labor*, 2018.

Blanco, Andres, Andres Drenik, and Emilio Zaratiegui, "Nominal Devaluations, Inflation, and Inequality," *American Economic Journal: Macroeconomics*, 2025, *17* (3), 270–310.

_ , Bernardo Diaz de Astarloa, Andres Drenik, Christian Moser, and Danilo R Trupkin, "The evolution of the earnings distribution in a volatile economy: Evidence from Argentina," *Quantitative Economics*, November 2022, *13* (4), 1361–1403.

Bowlus, Audra, Émilien Gouin-Bonenfant, Huju Liu, Lance Lochner, and Youngmin Park, "Four decades of Canadian earnings inequality and dynamics across workers and firms," *Quantitative Economics*, 2022, *13* (4), 1447–1491.

Broer, Tobias, John Kramer, and Kurt Mitman, "The Curious Incidence of Shocks Across the Income Distribution," 2023.

Busch, Christopher, David Domeij, Fatih Guvenen, and Rocio Madera, "Skewed Idiosyncratic Income Risk over the Business Cycle: Sources and Insurance," *American Economic Journal: Macroeconomics*, April 2022, *14* (2), 207–242.

Callan, T., K. Doorley, and A. McTague, "Top Incomes in Ireland: Reconciling Evidence from Tax Records and Household Survey Data," *Journal of the Statistical and Social Inquiry Society of Ireland*, 2020, *L.* Read before the Society, 22nd October 2020.

Dingel, Jonathan I. and Brent Neiman, "How Many Jobs Can be Done at Home?," *Journal of Public Economics*, 2020, *189*, 104235.

Donovan, Kevin, Will Jianyu Lu, Joseph H. Pedtke, and Todd Schoellman, "Labor Market Anatomy of a Macroeconomic Crisis," October 2024. Manuscript.

Doris, Aedín, Donal O'Neill, and Olive Sweetman, "Wage flexibility and the Great Recession: The response of the Irish labour market," *IZA Journal of European Labor Studies*, 2015, *4* (1), 1–29.

Drechsel-Grau, Moritz, Andreas Peichl, Kai D Schmid, Johannes F Schmieder, Hannes Walz, and Stefanie Wolter, "Inequality and income dynamics in Germany," *Quantitative Economics*, 2022, *13* (4), 1593–1635.

Engbom, Niklas, Gustavo Gonzaga, Christian Moser, and Roberta Olivieri, "Earnings inequality and dynamics in the presence of informality: The case of Brazil," *Quantitative Economics*, November 2022, *13* (4), 1405–1446.

Floccari, Giuseppe and Daniel Karpati, "Higher Order Income Risk and the Business Cycle: Evidence from the Netherlands," *Working paper*, 2026.

Friedrich, Benjamin, Lisa Laun, and Costas Meghir, "Earnings dynamics of immigrants and natives in Sweden 1985–2016," *Quantitative Economics*, 2022, *13* (4), 1803–1847.

Gráda, Cormac Ó and Kevin Hjortshøj O'Rourke, "The Irish economy during the century after partition," *The Economic History Review*, 2022, *75* (2), 336–370.

Guvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song, "What do data on millions of U.S. workers reveal about life-cycle earnings dynamics?," *Econometrica*, sep 2021, *89* (5), 2303–2339.

_ , Luigi Pistaferri, and Giovanni L Violante, "The Global Repository of Income Dynamics," *https://www.grid-database.org*, 2022.

_ , _ , and _ , "Global trends in income inequality and income dynamics: New insights from GRID," *Quantitative Economics*, 2022, *13* (4), 1321–1360.

_ , Sam Schulhofer-Wohl, Jae Song, and Motohiro Yogo, "Worker Betas: Five Facts about Systematic Earnings Risk," *American Economic Review*, 2017, *107* (5), 398–403.

_ , Serdar Ozkan, and Jae Song, "The nature of countercyclical income risk," *Journal of Political Economy*, jun 2014, *122* (3), 621–660.

Halvorsen, Elin, Serdar Ozkan, and Sergio Salgado, "Earnings dynamics and its intergenerational transmission: Evidence from Norway," *Quantitative Economics*, 2022, *13* (4), 1707–1746.

Hargaden, Enda Patrick, "Taxpayer responses in good times and bad," *Journal of Economic Behavior Organization*, aug 2020, *176*, 653–690.

Hoffmann, Emanuel and Davide Malacrino, "Employment time and the cyclicality of earnings growth," *Journal of Public Economics*, 2019, *169*, 160–171.

Hoffmann, Eran B, Davide Malacrino, and Luigi Pistaferri, "Earnings dynamics and labor market reforms: The Italian case," *Quantitative Economics*, November 2022, *13* (4), 1637–1667.

Honohan, Patrick, *Currency, Credit and Crisis: Central Banking in Ireland and Europe* Studies in Macroeconomic History, Cambridge University Press, 2019.

_ , "Is Ireland Really the Most Prosperous Country in Europe?," Economic Letter Vol 2021, No. 1, Central Bank of Ireland 2021.

_ and Brendan Walsh, "Catching Up with the Leaders: The Irish Hare," *Brookings Papers on Economic Activity*, 2002, *2002* (1), 1–57.

_ and Philip R. Lane, "Divergent inflation rates in EMU," *Economic Policy*, 2003, *18* (37), 357–394.

Kramarz, Francis, Elio Nimier-David, and Thomas Delemotte, "Inequality and earnings dynamics in France: National policies and local consequences," *Quantitative Economics*, November 2022, *13* (4), 1527–1591.

Krueger, Dirk, Fabrizio Perri, Luigi Pistaferri, and Giovanni L. Violante, "Cross-sectional facts for macroeconomists," *Review of Economic Dynamics*, jan 2010, *13* (1), 1–14.

Lane, Philip R., "The European Sovereign Debt Crisis," *Journal of Economic Perspectives*, 2012, *26* (3), 49–68.

__ , "The Treatment of Global Firms in National Accounts," *Economic Letters*, 2017, *2017* (No. 1). Issue 01/EL/17; available at: https://centralbank.ie/docs/default-source/publications/economic-letters/economic-letter-vol-2017-no-1.pdf?sfvrsn=6.

Leth-Petersen, Søren and Johan Sæverud, "Inequality and dynamics of earnings and disposable income in Denmark 1987–2016," *Quantitative Economics*, November 2022, *13* (4), 1493–1526.

McKinney, Kevin L., John M. Abowd, and Hubert P. Janicki, "U.S. Long-Term Earnings Outcomes by Sex, Race, Ethnicity, and Place of Birth," *Quantitative Economics*, November 2022, *13* (4), 1879–1945.

Nolan, Brian, "Long-Term Trends in Top Income Shares in Ireland," in Anthony B. Atkinson and Thomas Piketty, eds., *Top Incomes over the Twentieth Century: A Contrast Between Continental European and English-Speaking Countries*, Oxford University Press, 2007, chapter 12, pp. 501–530. Section 12.1 Introduction.

__ , "Updating Top Incomes for Ireland to 2015," Technical Report 2018/02, World Inequality Lab / WID.world 2018.

Puggioni, Daniela, Mariana Calderón, Alfonso Cebreros Zurita, León Fernández Bujanda, José Antonio Inguanzo González, and David Jaume, "Income dynamics and inequality: The case of Mexico," *Quantitative Economics*, 2022, *13* (4), 1669–1705.

Redmond, Paul, "Minimum wage policy in Ireland," 2020, (2021/2).

Roantree, Barra and Michelle Barrett, "Deaton Review Country Study: Ireland," Technical Report, Institute for Fiscal Studies 2023.

__ **and** __ , "Income inequality in Ireland, 1987–2019," *Fiscal Studies*, jun 2024, *45* (2), 143–153.

__ **, Bertrand Maître, and Helen Russell**, "Poverty, Income Inequality and Living Standards in Ireland: Fourth Annual Report," September 2024.

von Wachter, Till, "The Persistent Effects of Initial Labor Market Conditions for Young Adults and Their Sources," *Journal of Economic Perspectives*, November 2020, *34* (4), 168–94.

World Inequality Database, "WID.world," https://wid.world 2025. World Inequality Database, World Inequality Lab, Paris School of Economics.

ONLINE APPENDIX

A Appendix to Section 2

Table A1: Percentiles of real annual earnings distribution in Ireland

Year	P1	P5	P10	P25	P50	P75	P90	P95	P99	P99.9
(a) All genders										
2002	2,607	4,912	7,736	15,984	27,830	42,053	58,418	71,681	120,404	237,857
2007	2,843	5,227	8,361	17,033	28,782	43,582	62,611	78,301	137,063	300,865
2012	2,794	5,077	8,028	16,184	28,475	45,374	65,488	82,164	147,639	349,420
2017	3,018	5,527	8,802	17,698	30,013	46,871	69,166	88,015	161,160	391,488
2022	3,104	6,172	9,868	19,835	32,480	50,868	76,003	98,559	183,666	443,411
(b) Women										
2002	2,480	4,236	6,362	12,417	22,864	35,172	48,937	57,202	84,579	168,190
2007	2,730	4,669	7,227	14,081	24,955	38,784	54,088	65,073	99,377	206,185
2012	2,794	4,884	7,505	14,481	26,005	41,574	58,229	69,628	108,795	233,585
2017	2,935	5,155	7,914	15,420	27,091	42,824	61,148	74,166	125,856	280,854
2022	3,003	5,553	8,802	17,198	29,527	46,440	66,866	83,358	150,131	325,119
(c) Men										
2002	2,813	6,105	10,106	20,680	32,858	47,965	66,621	82,341	141,149	295,702
2007	2,998	5,938	10,080	20,318	32,281	48,072	70,507	88,830	158,851	352,059
2012	2,795	5,323	8,704	18,331	31,190	49,936	73,847	95,151	176,668	435,848
2017	3,108	6,063	9,964	20,280	32,553	51,253	77,943	99,886	189,699	489,089
2022	3,217	7,012	11,242	22,873	35,215	55,729	85,205	112,106	208,458	560,964

Note: Table presents percentiles of the earnings distribution in euros for selected percentiles and years.

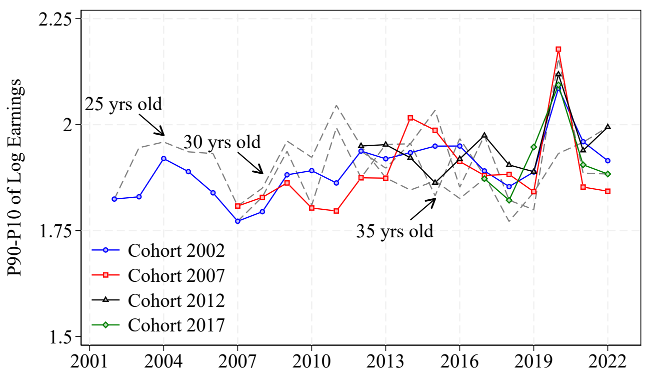

(a) Men P90-P10

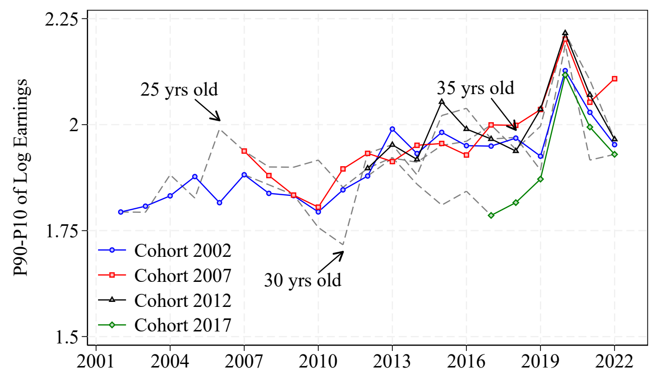

(b) Women P90-P10

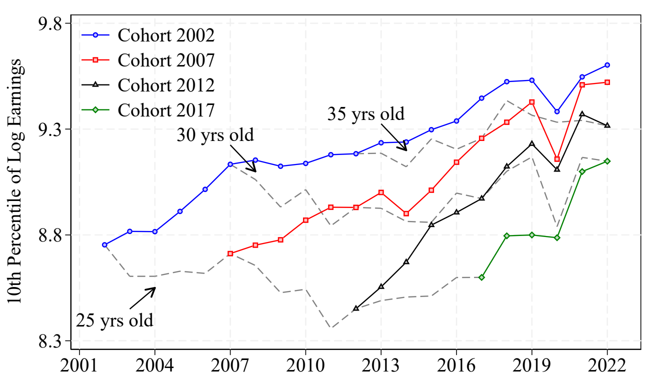

(c) Men p10

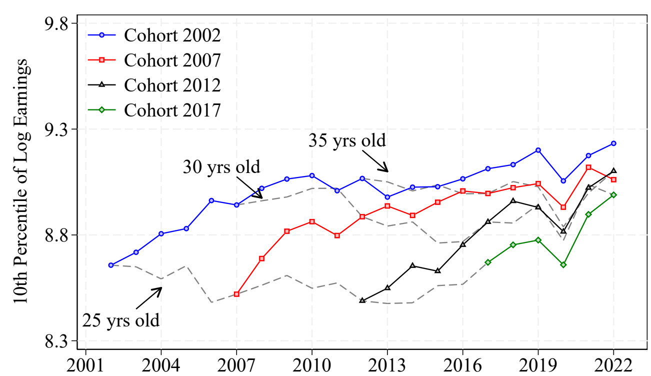

(d) Women p10

Figure A1: Life-cycle inequality over cohorts

Note: Figure shows changes in the P90-P10. Dashed lines correspond to earnings of 25-, 30-, and 35-year-olds as indicated by arrows. Each solid line corresponds to an individual cohort where "cohort X" represents the cohort aged 25 in year X.

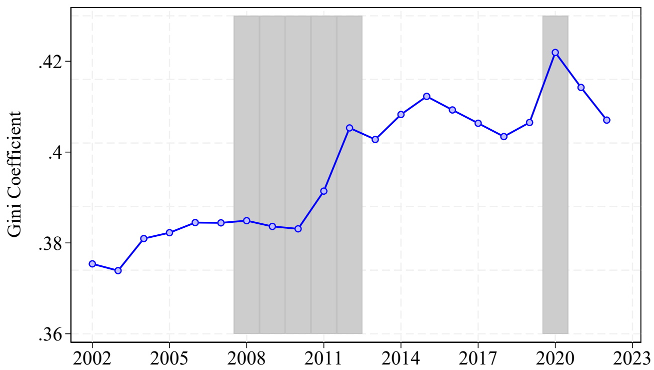

(a) All

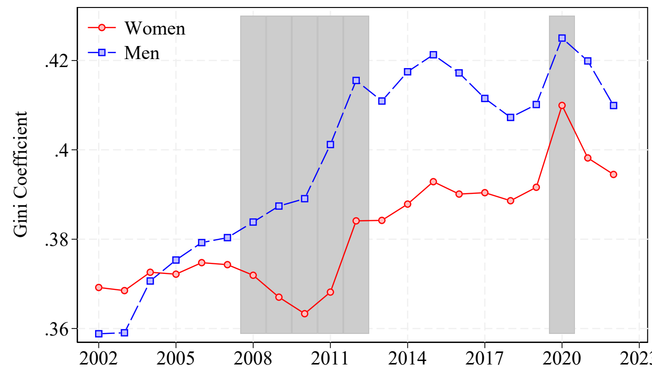

(b) Men and Women

Figure A2: Gini Coefficient

Note: Figure presents the Gini coefficient over time for men and women. Shaded areas indicate recession years, based on declines in Modified Domestic Demand.

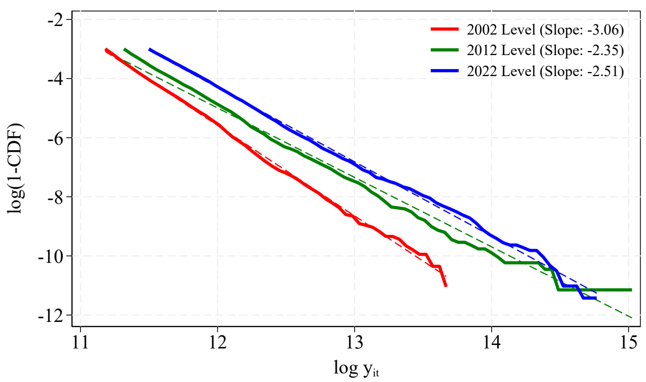

(a) Men

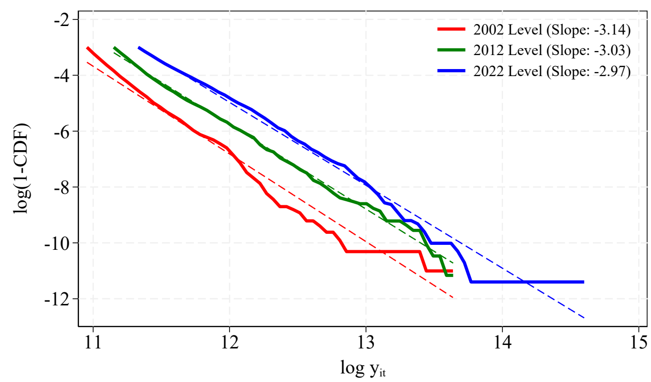

(b) Women

Figure A3: Pareto Tail, top 5 percent

Note: Figure shows the Pareto tail of the earnings distribution as measured by the log CDF for the top 5 percent of earners.

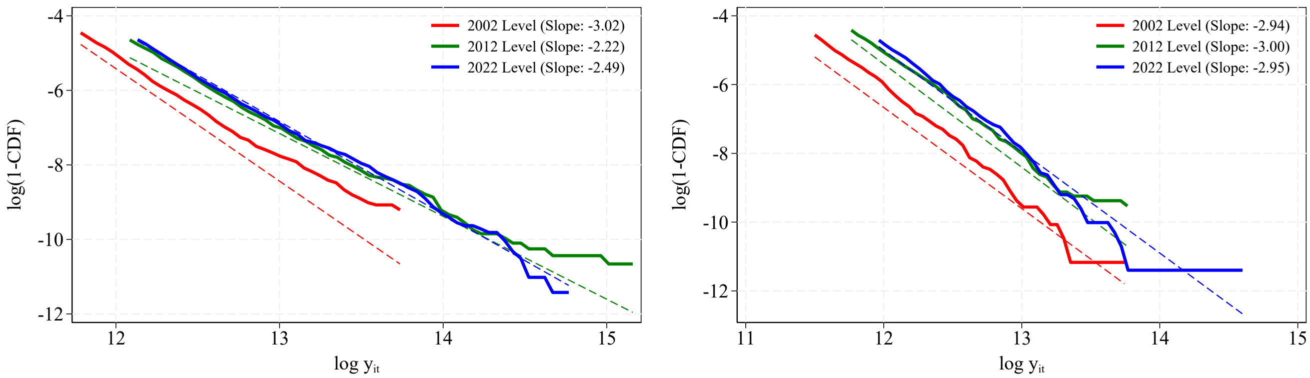

(a) Men

(b) Women

Figure A4: Pareto Tail, top 1 percent

Note: Figure shows the Pareto tail of the earnings distribution as measured by the log CDF for the top 1 percent of earners.

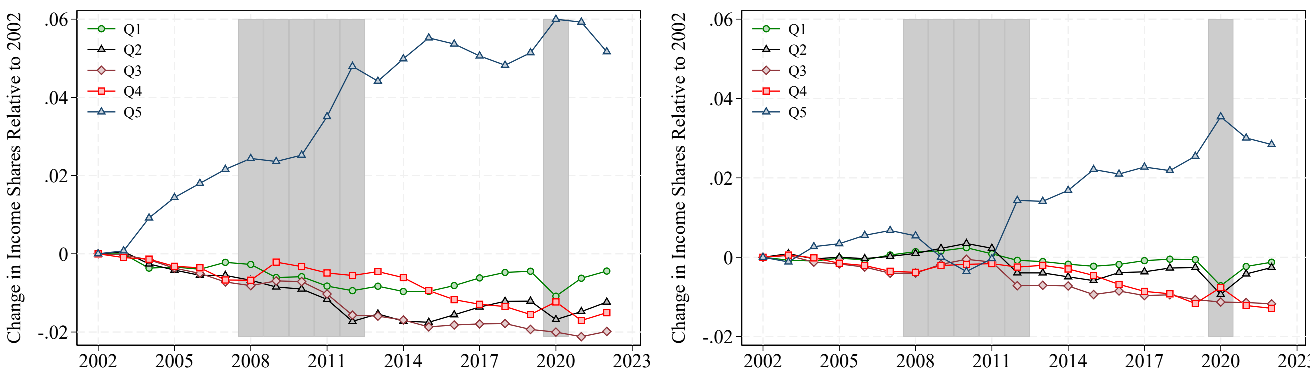

(a) Men

(b) Women

Figure A5: Income Shares by Quintile

Note: Figure shows the income shares by quintile for men and women. Shaded areas indicate recession years, based on declines in Modified Domestic Demand.

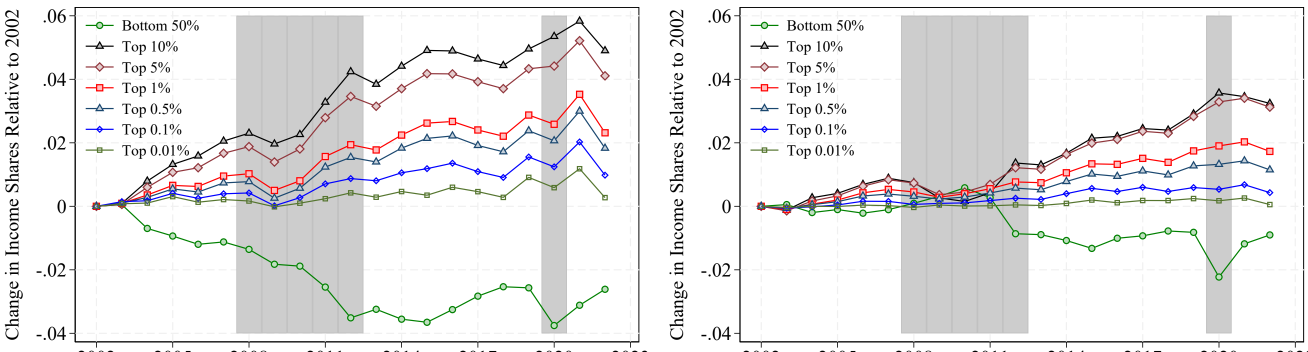

(a) Men (b) Women

Figure A6: Income Shares by Percentile

Note: Figure shows income shares for selected percentiles. Shaded areas indicate recession years, based on declines in Modified Domestic Demand.

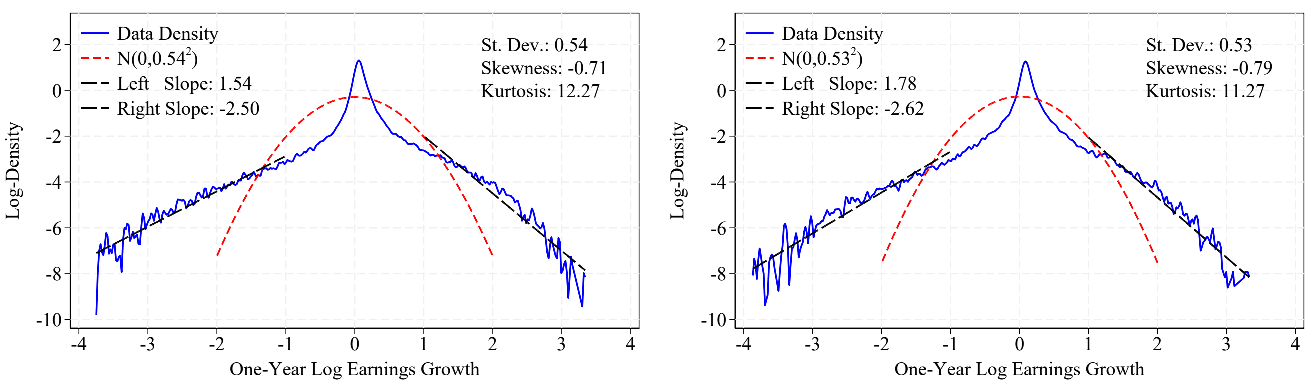

Figure A7: Empirical Log-Densities of One-Year Earnings Growth

Note: Figure shows the log-density of one-year earnings growth in 2005. LX sample.

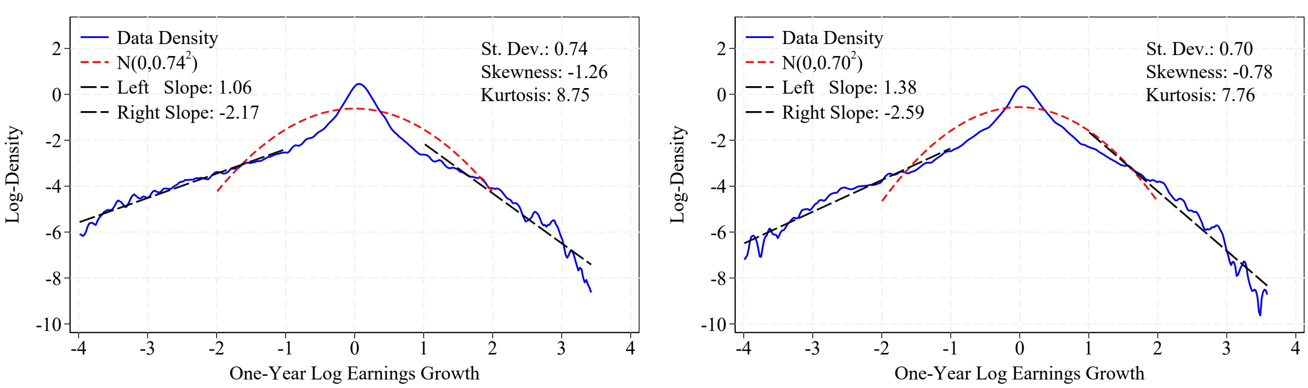

Figure A8: Empirical Log-Densities of Five-Year Earnings Growth

Note: Figure shows the log-density of five-year earnings growth in 2005. LX sample.

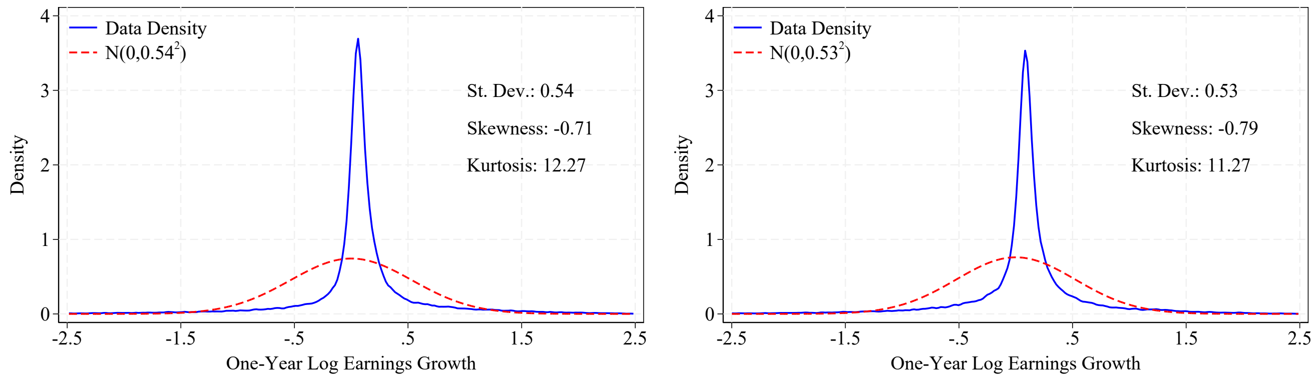

Figure A9: Empirical Densities of One-Year Earnings Growth

Note: Figure shows the density of one-year earnings growth in 2005. LX sample.

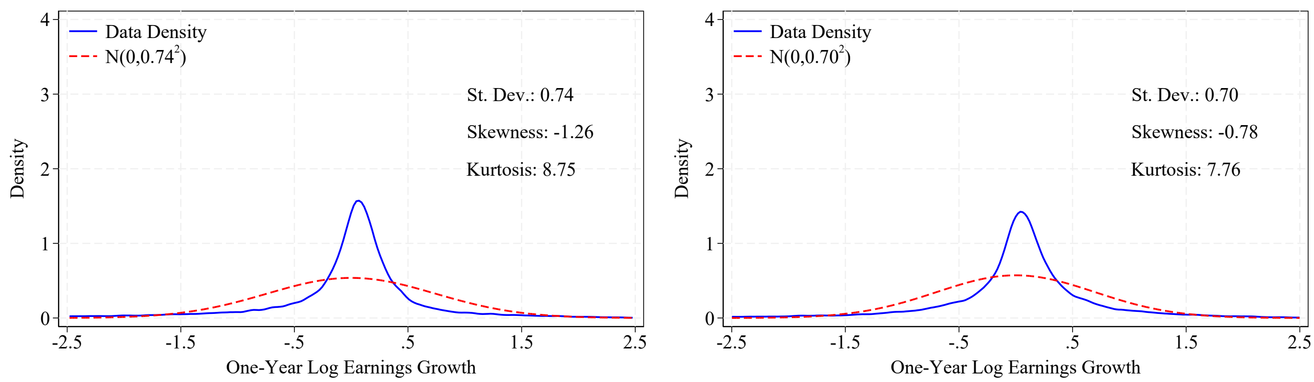

Figure A10: Empirical Densities of Five-Year Earnings Growth

Note: Figure shows the density of five-year earnings growth in 2005. LX sample.

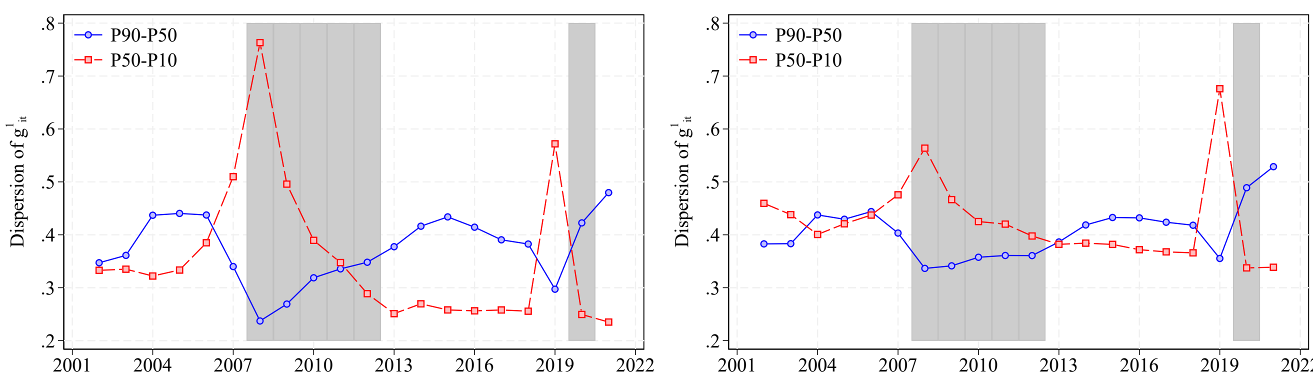

(a) Men

(b) Women

Figure A11: Pro-cyclical Skewness and Good and Bad Shocks

Note: Figure decomposes skewness into positive shocks and negative shocks. Shaded areas indicate recession years, based on declines in Modified Domestic Demand.

A.1 Nominal vs Real Earnings Growth

In this appendix we document that the high and occasionally volatile inflation rates in Ireland (as shown in Figure 2) affect the interpretation of patterns of log earnings growth.[34] Figure A12 compares the real growth (panel a) and nominal growth (panel b). The boom and bust are clearly apparent in nominal earnings: earnings rise at every percentile from 2002 to 2008 and fall at every percentile through at least 2011.

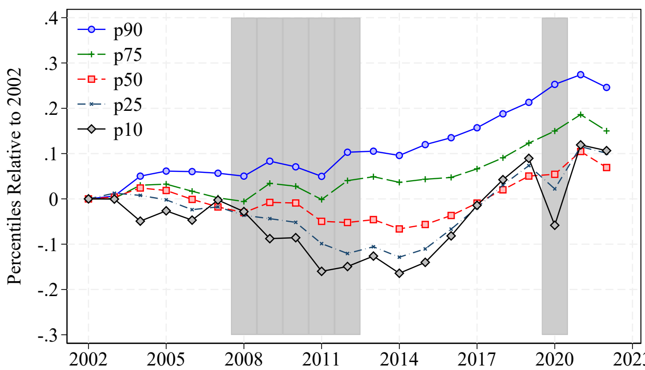

(a) Men - real

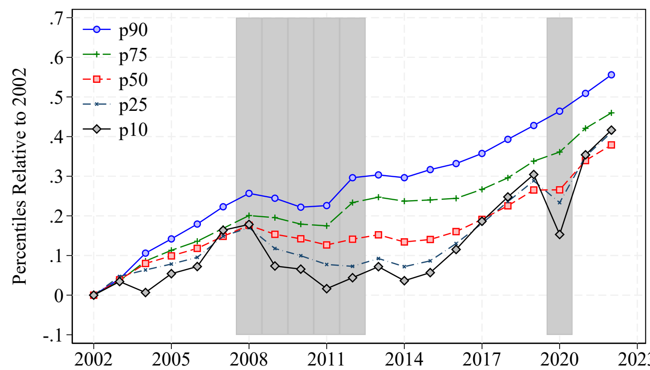

(b) Men - nominal

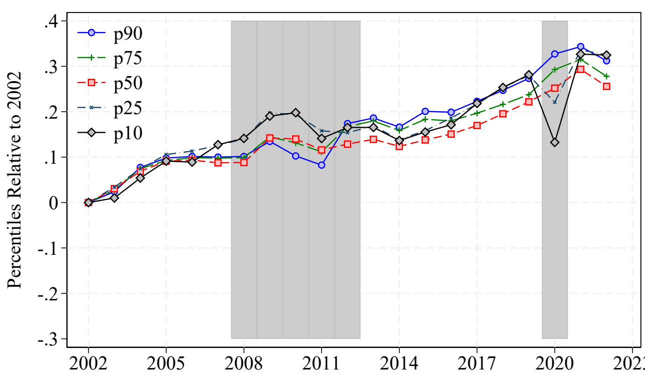

(c) Women - real

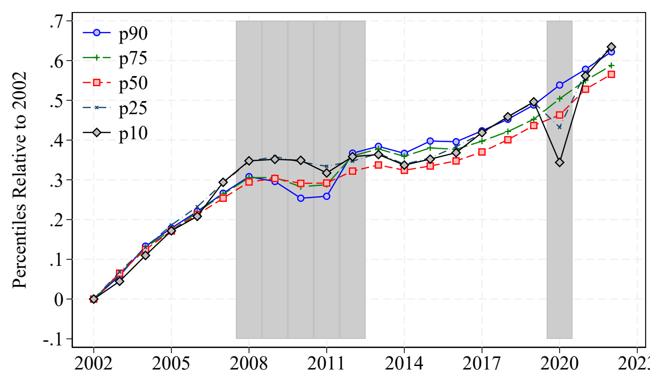

(d) Women - nominal

Figure A12: Comparison of nominal versus real log earnings distributions

Note: Log income is normalised to 0 in the base year 2002 and is shown at P10, P25, P50, P75, P90. Income is real in the left panel and nominal in the right panel. Shaded areas indicate recession years, based on declines in Modified Domestic Demand.

The comparison of nominal and real earnings highlights three important patterns. First, the last three years of the boom (2005-08) resulted in stagnant or even declining *real* earnings despite large nominal earnings growth. For example, the median for men grew 18 percent to its peak in

[34] Honohan and Lane (2003) discuss the forces leading to high inflation after Ireland entered the European Monetary Union (EMU). They argue that inflation increased because the real exchange declined upon entry to the EMU, and inflation may have remained high due to a mixture of falling interest rates, house price increases and fiscal expansion.

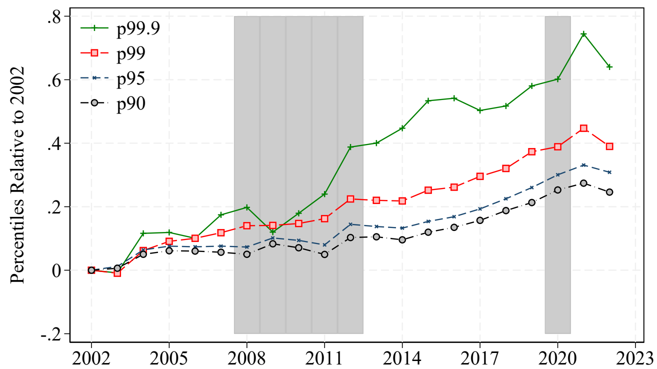

(a) Men - real

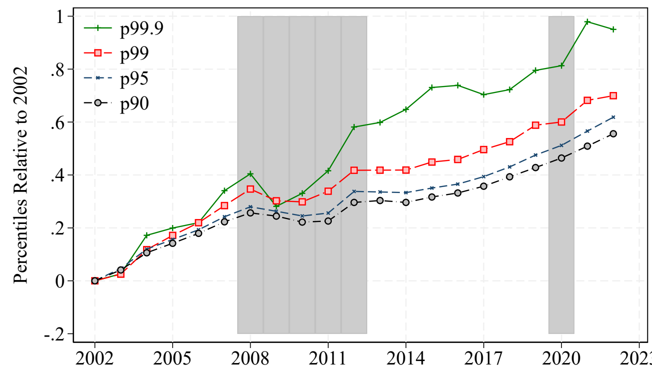

(b) Men - nominal

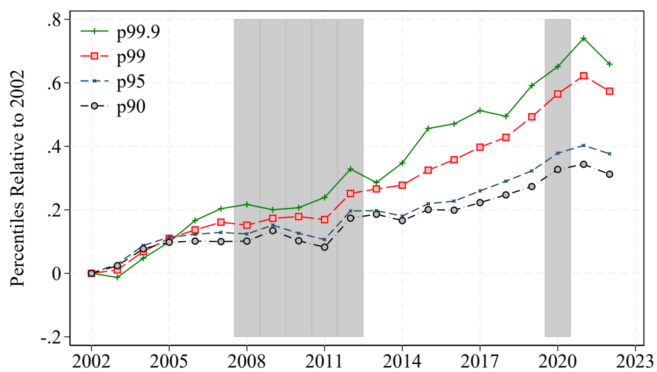

(c) Women - real

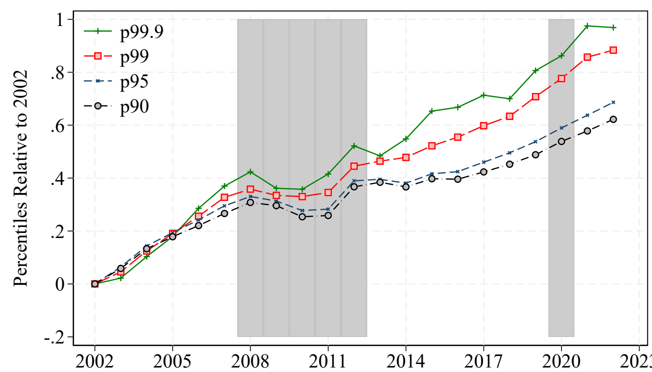

(d) Women - nominal

Figure A13: Comparison of nominal versus real log earnings at the *top* of the distribution

Note: Log income is normalised to 0 in the base year 2002 and is shown at P90, P95, P99, and P99.9. Income is real in the left panel and nominal in the right panel. Shaded areas indicate recession years, based on declines in Modified Domestic Demand.

2008, while real earnings fell 2 percent. This highlights that much of the benefits from the growth in real output over this time period did not translate into real earnings growth because of the high inflation rates.[35]

Second, deflation mitigated nominal income declines during the bust. For example, at the start of the bust in 2008, real median earnings actually *increased* by 2 percent despite a 2 percent decline in nominal. The peak-to-trough decline of 4.2 percent in real terms is one third less than the 6.7 percent nominal decline. That said, it is worth noting that nominal changes in earnings may be more relevant for individuals with nominal debt obligations (such as mortgages). Thus the distinction between real and nominal changes in earnings is particularly relevant to Ireland given the high levels of mortgage default during this period.

Third, the real gains that have occurred since the beginning of the recovery in 2015 are much higher than any real gains in the latter part of the boom. At the median, real earnings have grown 14 percent since 2015 leaving them 7.5 percent higher than at the beginning of the sample in 2002.

[35] Honohan and Lane (2003) note that real wage growth was also modest in the lead-up to our sample with real wages growing only 15 percent from 1996 to 2002.

B Appendix to: Section 2.6

B.1 Relationship between Boom and Bust

Figure B14 plots the relationship between the fall in median earnings in the Great Recession and the subsequent rebound to 2016.[36] There is a positive slope, implying that countries who did well in the recession also did well afterwards, although this appears to be mostly driven by the emerging economies Brazil and Argentina.

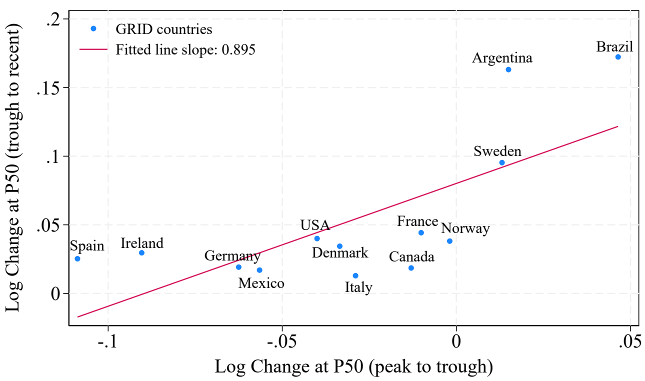

(a) P50 men

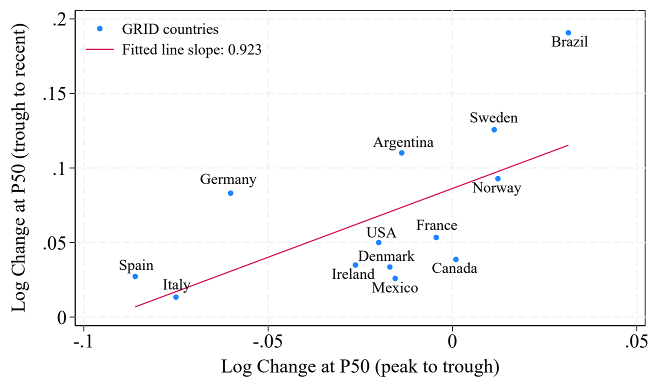

(b) P50 women

Figure B14: Earnings Growth in the Great Recession at P50. Peak to Trough vs Trough to Rebound

Note: For each country, *Peak* is defined using the highest values between 2002 and 2009, and *Trough* is defined using the lowest values after 2009. For comparability and data availability, *Recent* is defined using the values in 2016. Data comes from GRID and Irish analysis.

The figure shows that Ireland's rebound has been better than expected, as illustrated by the dot for men being above the fitted red line. For men, Ireland experienced one of the worst falls in median earnings—coming second only to Spain. On the other hand, its recovery was in line with most others with only Argentina, Brazil and Sweden growing considerably more. For women, the rebound is below the fitted line, but the fall itself was also more in line with other countries.

It is worth noting these comments are limited by many other countries in GRID ending their time series before 2016 when the Irish recovery was only partially complete. Ireland's recovery is more remarkable when compared to the United Kingdom, our closest neighbour geographically and, arguably, economically. While the median Irish man had recovered his lost earnings by 2017 and continued to grow, the entire distribution (P10 to P99.9) of male earnings in the UK had still not recovered by 2020.[37] Indeed, unlike in Ireland, the 10th percentile also did not catch up with

[36] We end in 2016 to maximize coverage across countries.

[37] We do not include the United Kingdom in the plots because their data are not yet available to download from GRID but the results are reported in their paper (Bell, Bloom and Blundell, 2022).

the median by the end of the samples in Norway (2017), Spain (2018) nor the UK (2020).

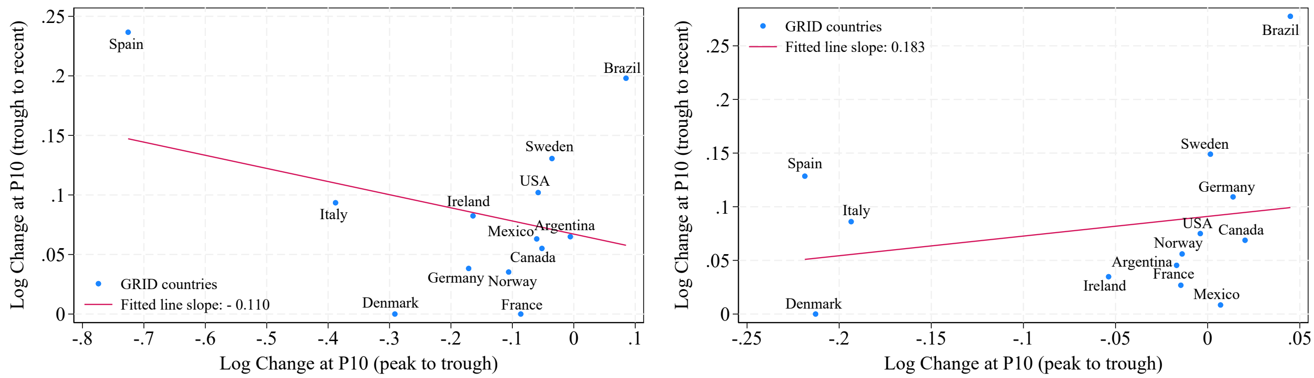

(a) P10 men

(b) P10 women

Figure B15: Earnings Growth in the Great Recession at P10. Peak to Trough vs Trough to Rebound

C Appendix to: Section 3

C.1 Estimation Framework

We present more details on the estimation framework here. Consider worker i at time t from group g, where groups are the combination of gender, age, and permanent income. We estimate the sensitivity of income growth Δy_{it} to *aggregate* output Y_t

$$\Delta y_{it} = \alpha^g + \beta^g \Delta Y_t + \varepsilon_{it}, \tag{4}$$

where Δy_{it} is the change in log earnings, or its equivalent for log weekly earnings Δy_{it}^W, or log weeks worked Δw_{it}.[38] We estimate this equation separately for sub-groups and therefore α^g and β^g are allowed to depend flexibly across groups g, though we suppress the g superscript unless necessary for exposition.

Average. Our first approach estimates the average responsiveness by estimating (4) using ordinary least squares (OLS) (Guvenen, Schulhofer-Wohl, Song and Yogo, 2017). The coefficient α corresponds to the average growth rate of earnings when aggregate output growth is zero, and β corresponds to the elasticity of average earnings growth to output growth

$$\alpha = \mathbb{E}[\Delta y_{it} \| \Delta Y_t = 0], \tag{5}$$

$$\beta = \frac{\partial \mathbb{E} \Delta y_{it}}{\partial \Delta Y_t}. \tag{6}$$

A contribution of our paper is to decompose the responsiveness of earnings to aggregate variables into the responsiveness of *weekly earnings* and the responsiveness of *weeks worked*. Changes in (log) real earnings Δy_{it} can be decomposed into changes in (log) weekly real earnings Δy_{it}^W and (log) weeks worked Δw_{it}, $\Delta y_{it} = \Delta y_{it}^W + \Delta w_{it}$. Since OLS is a linear operator, this means that the responsiveness of total earnings β_y can also be decomposed into these components

$$\beta_y = \beta_{y^W} + \beta_w. \tag{7}$$

We call the ratio of the beta on weeks and annual earnings β_w/β_y, the *weeks share* of exposure. In addition to being inherently informative about the nature of income changes, the comparison will also be informative about the patterns investigated in the next Section (**??**). In particular,

[38] Earnings, weeks worked, and weekly earnings are as defined in Section 1, whereas the non-residual Δy_{it} and differs from the changes in residual log income g_{it} which was defined in the same section. We continue to use modified domestic demand as our preferred measure of aggregate output. To be consistent with the GRID income variables, we use one-year forward changes.

entering new migrants (who work only part of the year) are more likely to alter the distribution of annual earnings than that of weekly earnings.[39]

Distribution. Our second approach estimates the heterogeneous effects of aggregate variables on the distribution of earnings growth, by estimating (4) with quantile regression (Floccari and Karpati, 2026).[40] The coefficient α_q corresponds to the growth of earnings at the q-th quantile Q_q when output growth is zero. The coefficient β_q corresponds to the elasticity of earnings growth at the q-th quantile to changes in output.

$$\alpha_q = Q_q(\Delta y_{it} \| \Delta Y_t = 0), \tag{8}$$

$$\beta_q = \frac{\partial Q_q(\Delta y_{it})}{\partial \Delta Y_t}. \tag{9}$$

Since quantile regression is not a linear operator, we cannot *decompose* the estimates into shares contributed by weeks worked and weekly earnings; instead we present these estimates separately.

As Floccari and Karpati (2026) show, a convenient advantage of the quantile regression approach is that we can use the estimated coefficients to construct any quintile-based higher-order moments of the distribution, such as volatility and skewness. With estimates of α_q and β_q for every percentile q, the higher-order moments and their responsiveness to aggregate variables follow. The responsiveness of the P90-P10 ratio of earnings growth is given by

$$\beta_{9010} = \beta_{90} - \beta_{10}, \tag{10}$$

and the response of Kelley skewness is

$$\beta_{KS} = \frac{2\left(\alpha_{50} - \alpha_{10}\right)\left(\beta_{90} - \beta_{50}\right)}{\left(\alpha_{90} - \alpha_{10}\right)^2} - \frac{2\left(\alpha_{90} - \alpha_{50}\right)\left(\beta_{50} - \beta_{10}\right)}{\left(\alpha_{90} - \alpha_{10}\right)^2}. \tag{11}$$

[39] It is not immediately obvious that selection that affects the level of percentiles of the *annual* cross-sectional distribution will substantially affect the growth rate of earnings based on percentiles of the *permanent* income distribution.

[40] An alternative to quantile regression is to construct quantiles and regress them on output directly—an approach taken by Busch et al. (2022). We find similar patterns with the quantile approach—in particular the cyclicality of skewness is important— suggesting these patterns are robust to the specific methodology used.

C.2 Responsiveness to Output

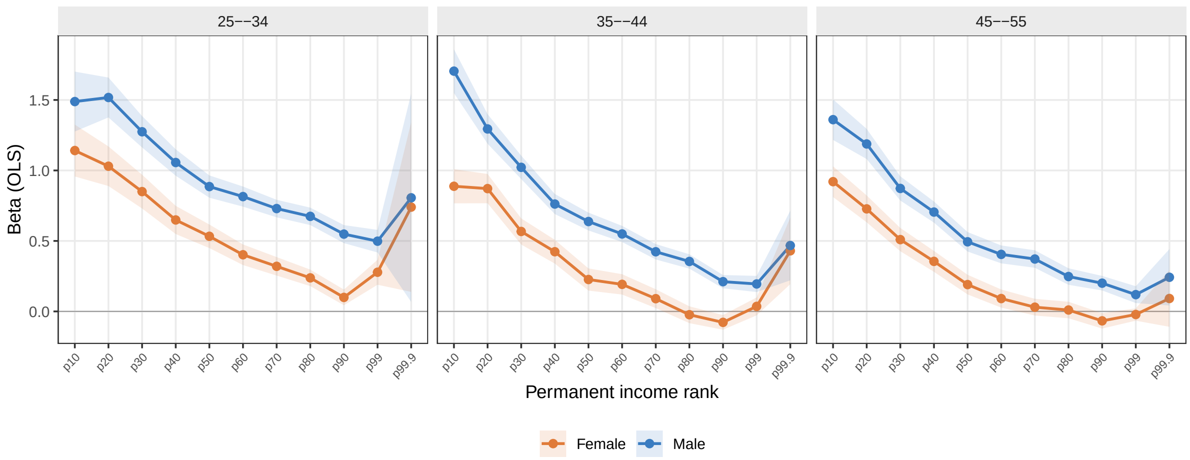

(a) Annual Earnings

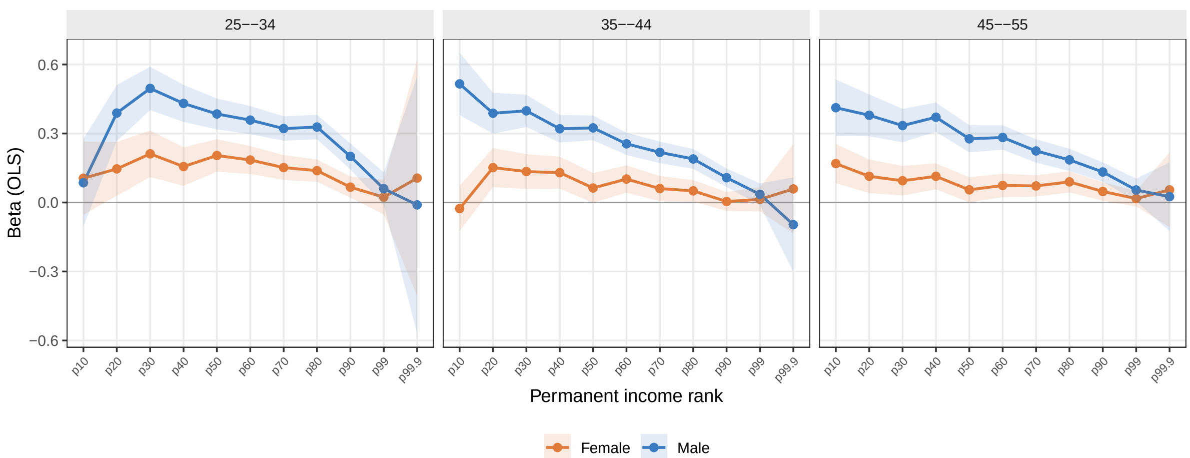

(b) Weeks Worked

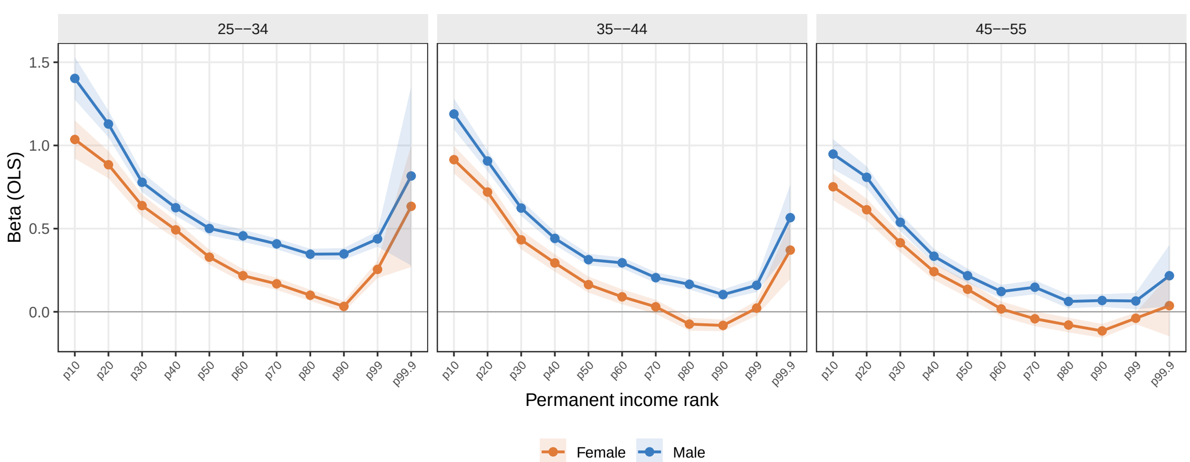

(c) Weekly Earnings

Figure C16: Average Responsiveness to Output (OLS Beta)

Note: Figure plots average (OLS) betas for annual earnings, weeks worked and weekly earnings.

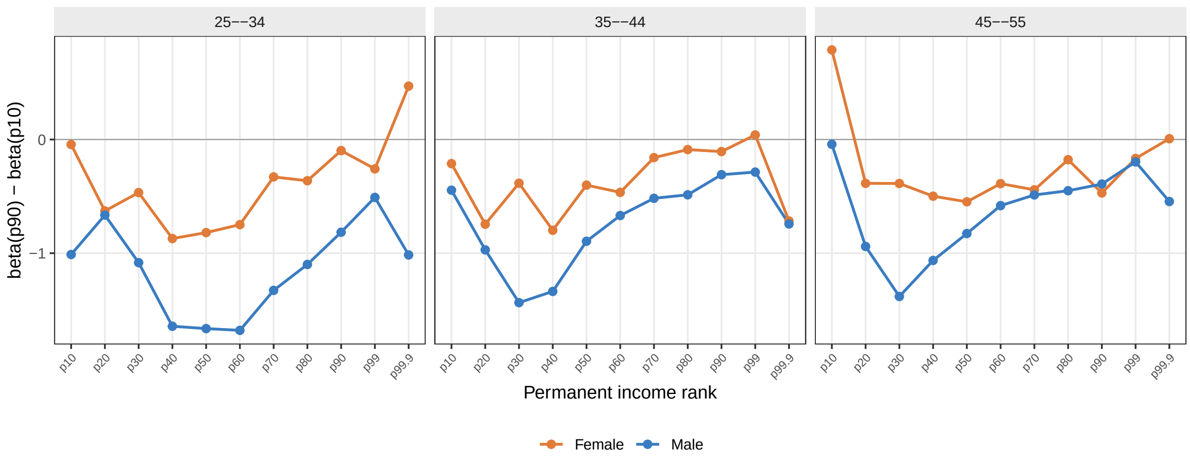

(a) Annual Earnings

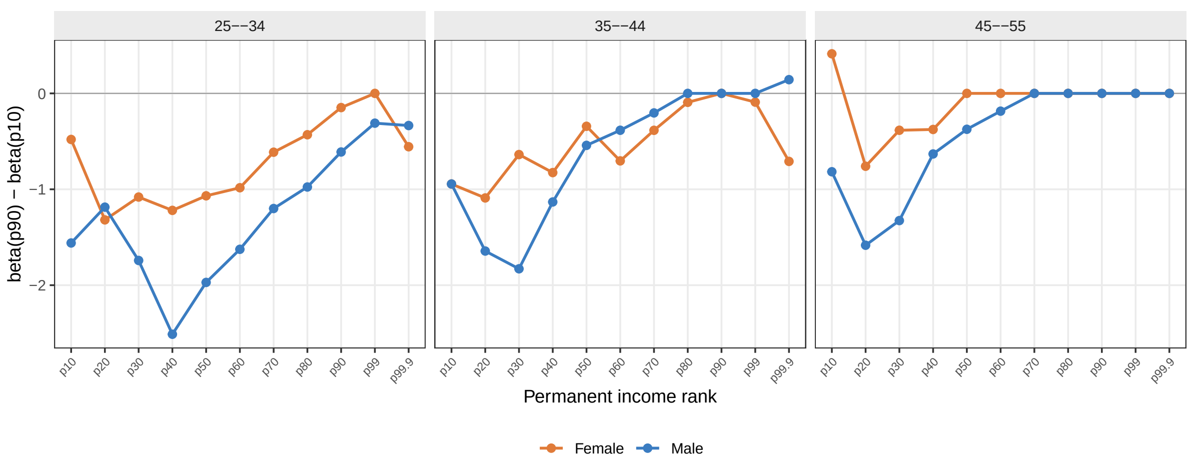

(b) Weeks Worked

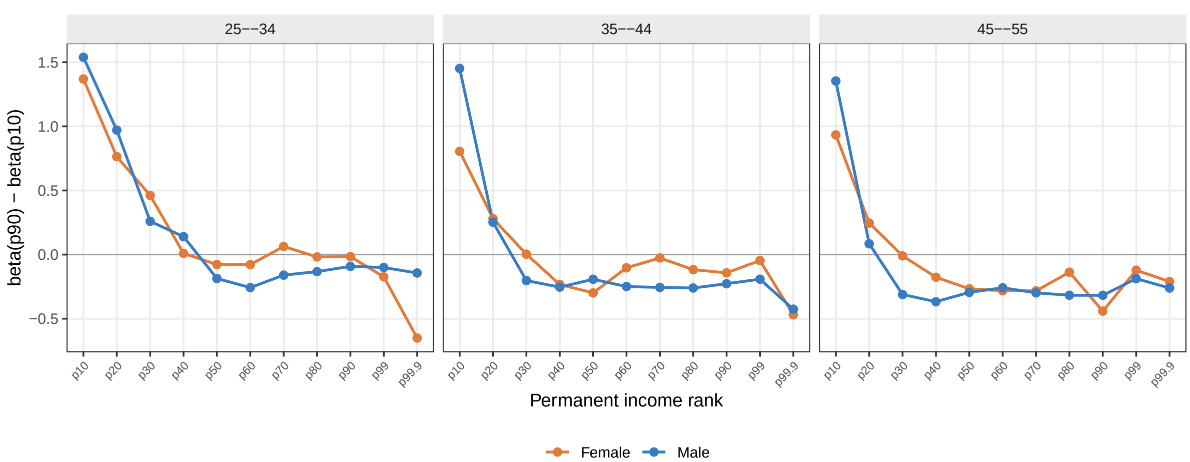

(c) Weekly Earnings

Figure C17: Responsiveness of Volatility (P90-P10) to Output

Note: Figure plots quantile betas for annual earnings, weeks worked and weekly earnings. We report $\beta_{\text{P90-P10}}$ as defined in Section (10).

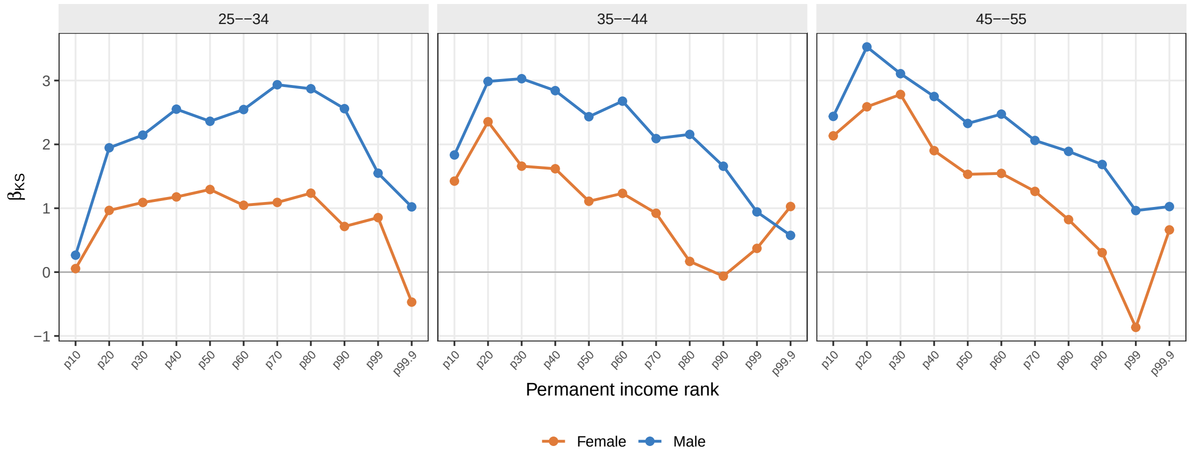

(a) Annual Earnings

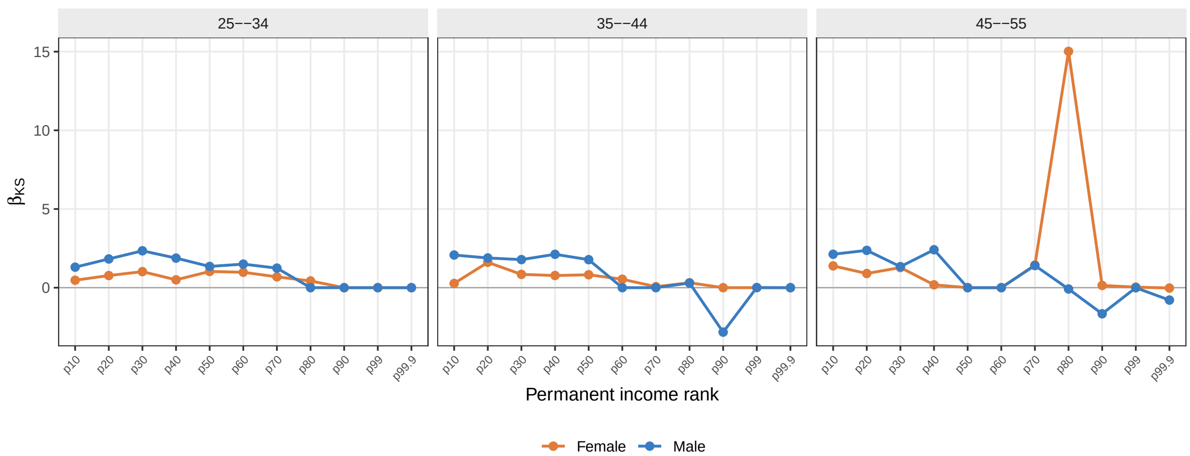

(b) Weeks Worked

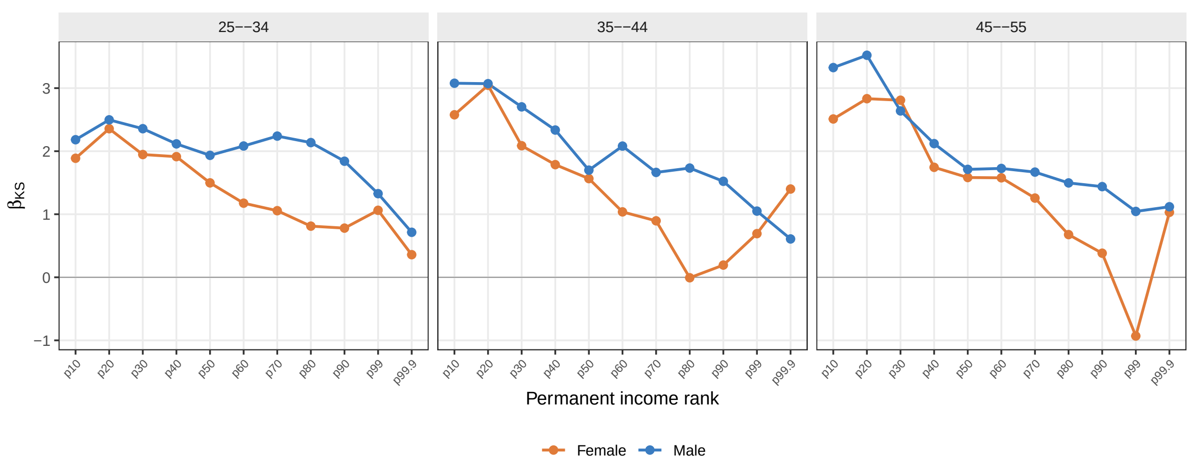

(c) Weekly Earnings

Figure C18: Responsiveness of Skewness to Output

Note: Figure plots quantile betas for annual earnings, weeks worked and weekly earnings. We report β_{KS} as defined in Section (11).

C.3 Responsiveness to Inflation

Table C2: Output Beta: Baseline vs. Inflation-Controlled Spec

Group	**Baseline** β_y	**Infl.-controlled** β_y
Panel A: Gender		
Men	0.827	0.830
	(0.056)	(0.063)
Women	0.405	0.384
	(0.112)	(0.121)
Panel B: Percentiles		
p10	1.250	1.116
	(0.033)	(0.035)
p50	0.494	0.512
	(0.015)	(0.016)
p90	0.152	0.193
	(0.011)	(0.012)
Panel C: Age Groups		
Age 25–34	0.867	0.833
	(0.181)	(0.197)
Age 35–44	0.513	0.522
	(0.042)	(0.046)
Age 45–55	0.467	0.465
	(0.029)	(0.031)

Notes. Table compared coefficients on output in regressions with and without inflation. Regressions are run by sub-group and entries are average beta estimates for each sub-group. Bootstrapped standard errors are in parenthesis.

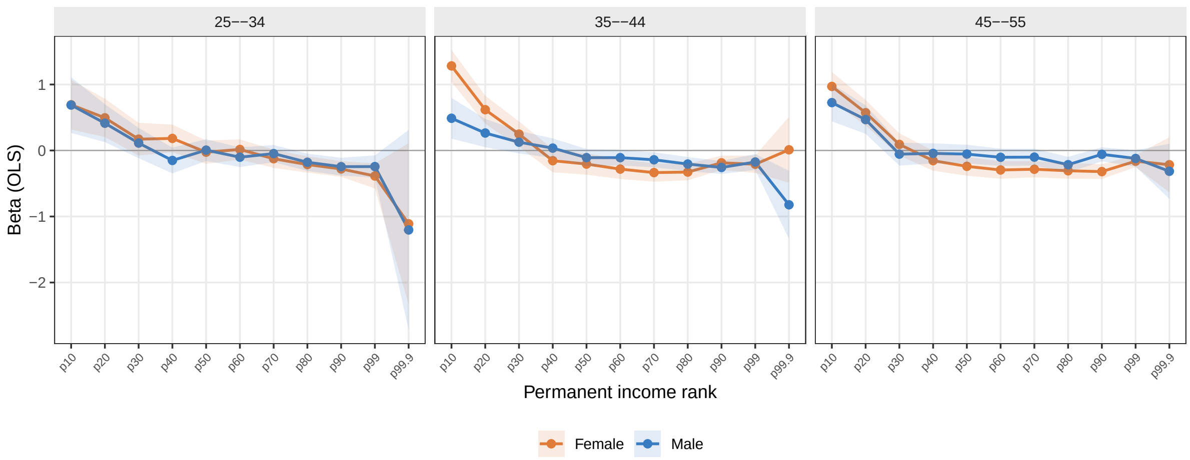

(a) Annual Earnings

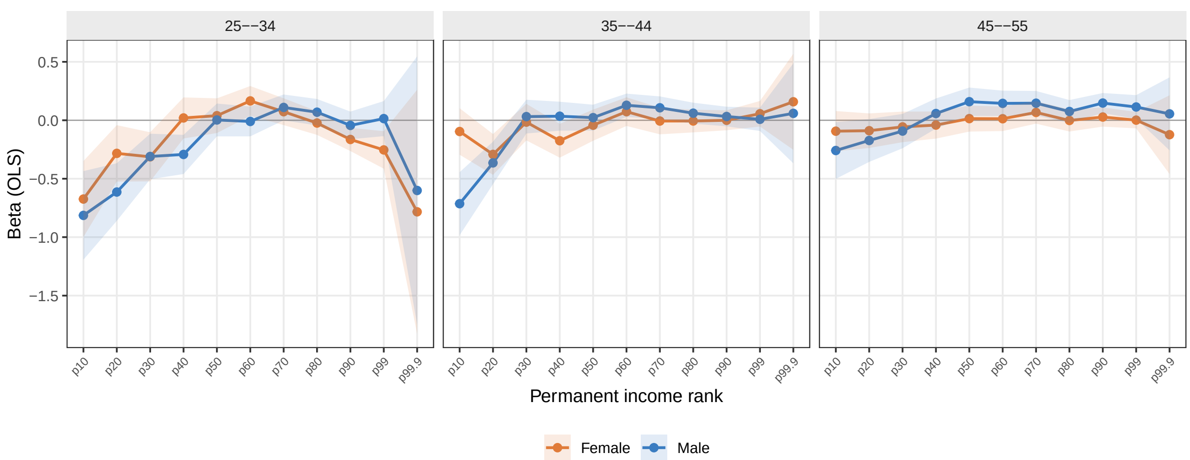

(b) Weeks Worked

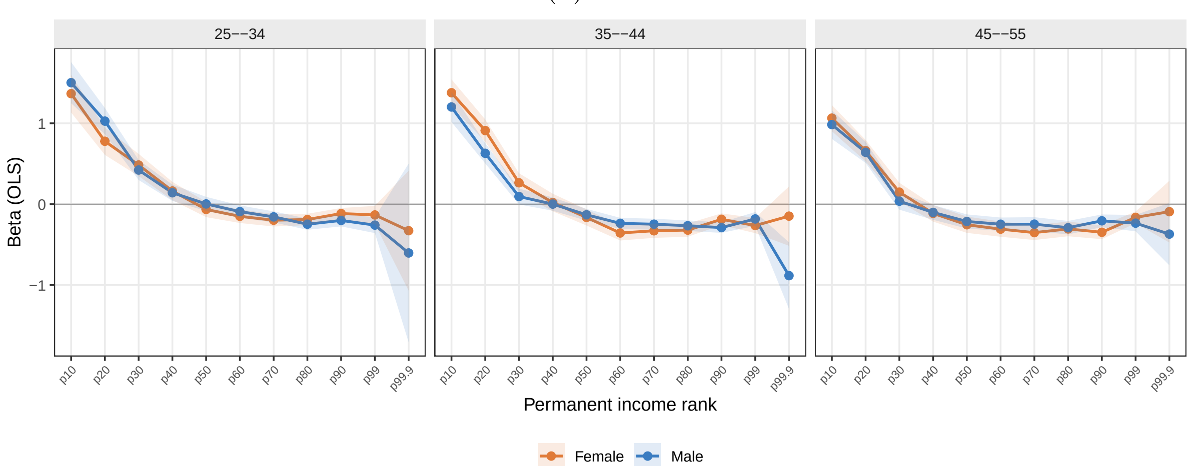

(c) Weekly Earnings

Figure C19: Average Responsiveness to Inflation (OLS Beta)

Note: Figure plots average (OLS) betas for annual earnings, weeks worked and weekly earnings.

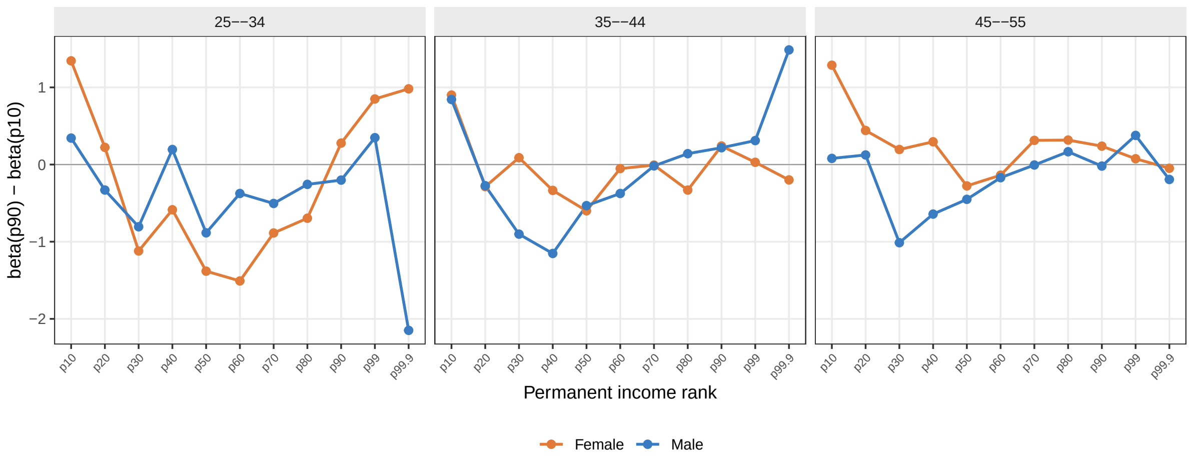

(a) Annual Earnings

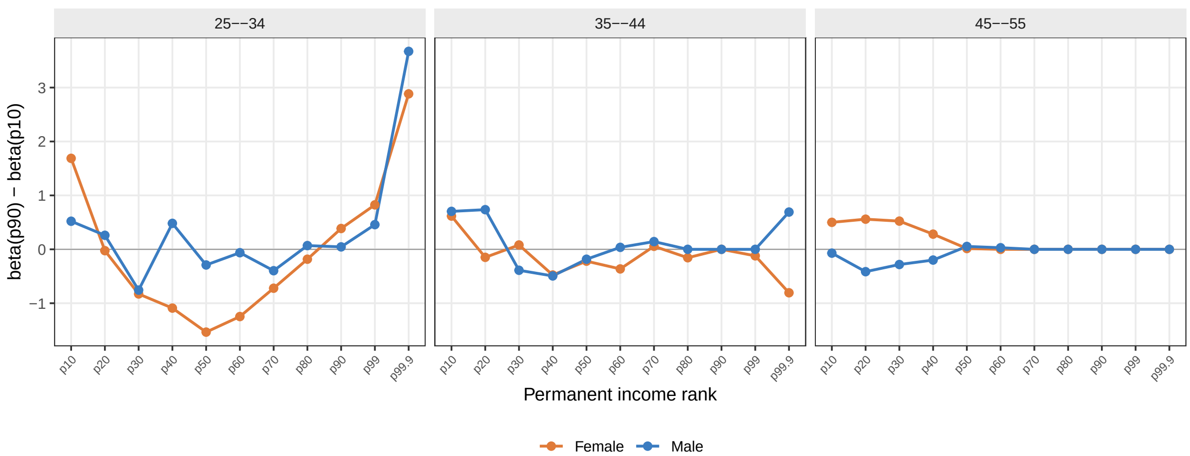

(b) Weeks Worked

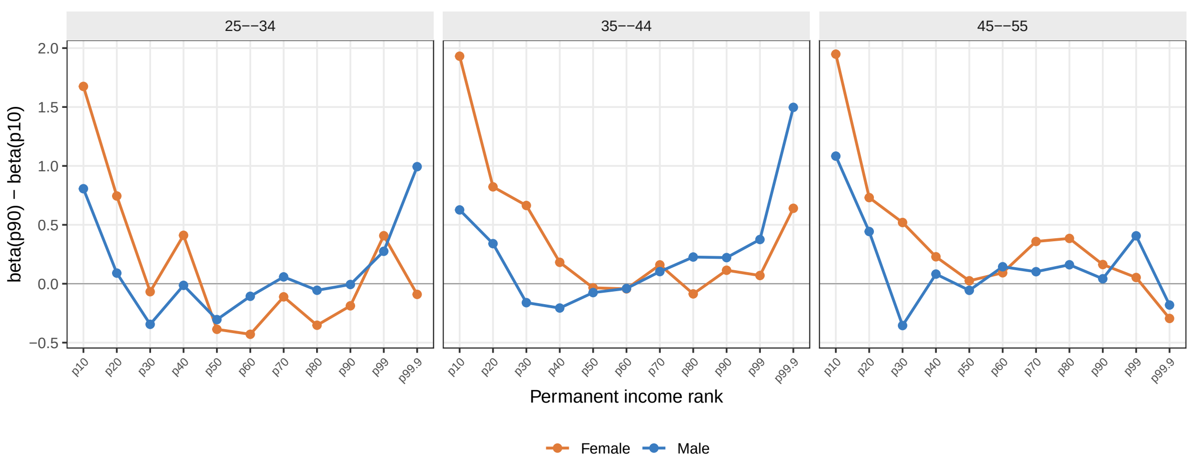

(c) Weekly Earnings

Figure C20: Responsiveness of Volatility (P90-P10) to Inflation

Note: Figure plots quantile betas for annual earnings, weeks worked and weekly earnings. We report $\beta_{\text{P90-P10}}$ as defined in Section 1.

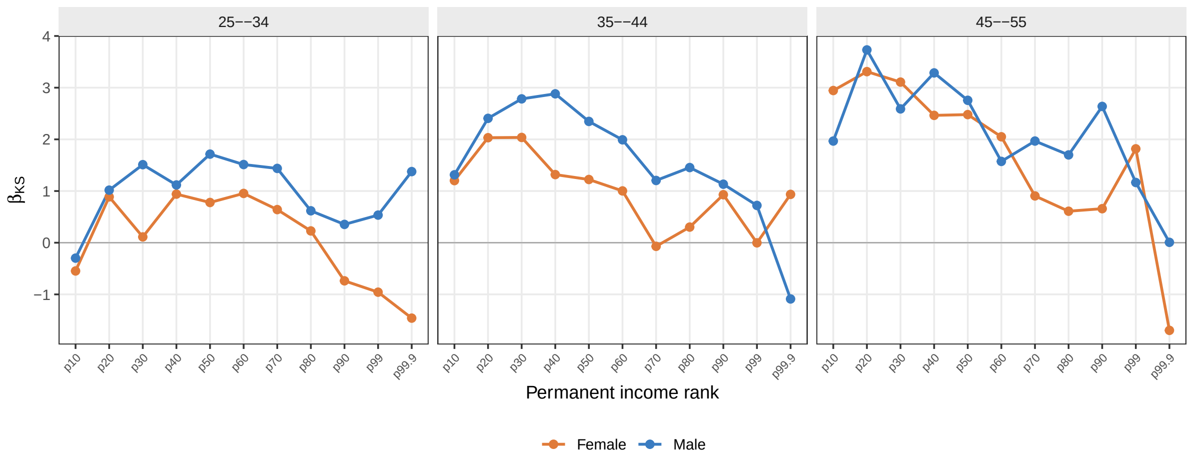

(a) Annual Earnings

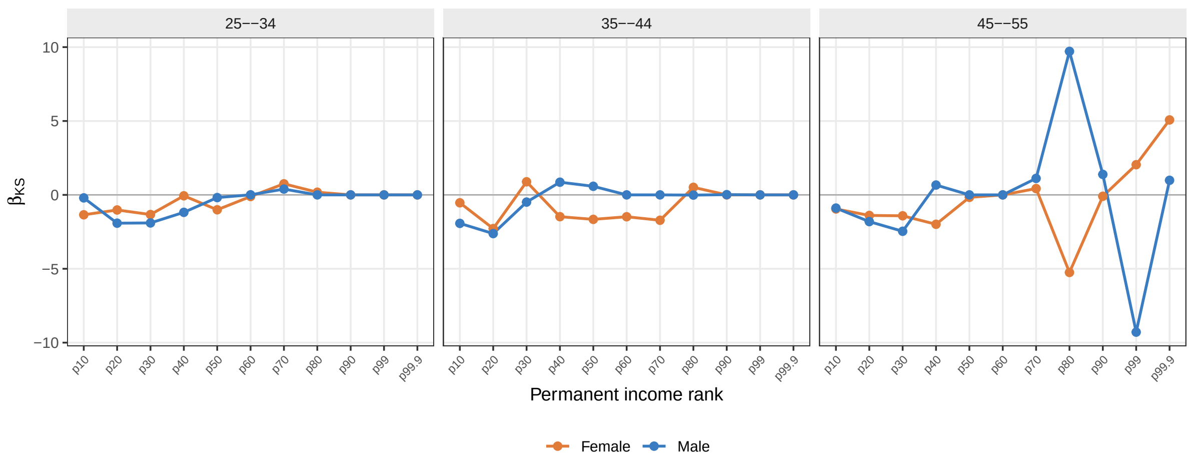

(b) Weeks Worked

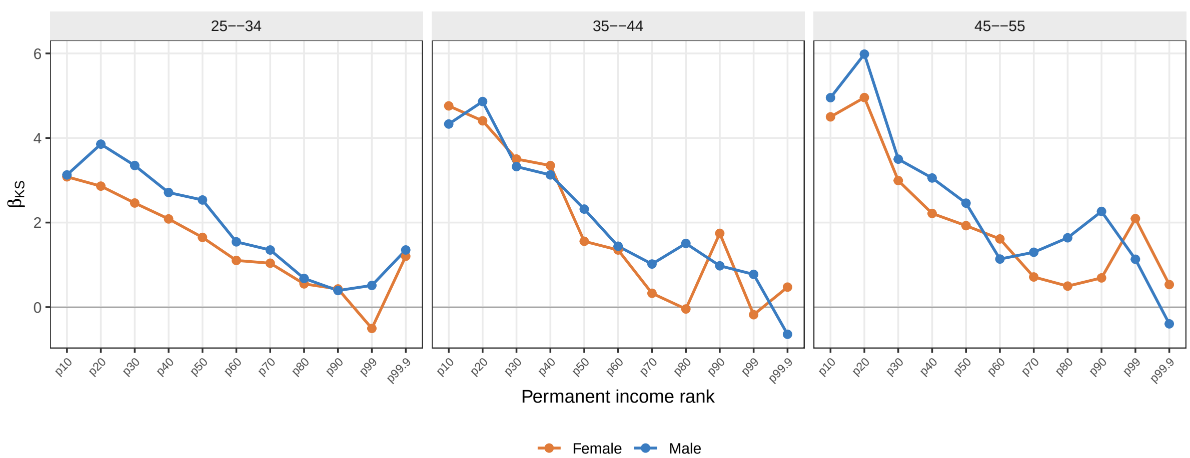

(c) Weekly Earnings

Figure C21: Responsiveness of Skewness to Inflation

Note: Figure plots quantile betas for annual earnings, weeks worked and weekly earnings. We report β_{KS} as defined in Section 1.

D Appendix to: Section 4

D.1 Migration and Selection

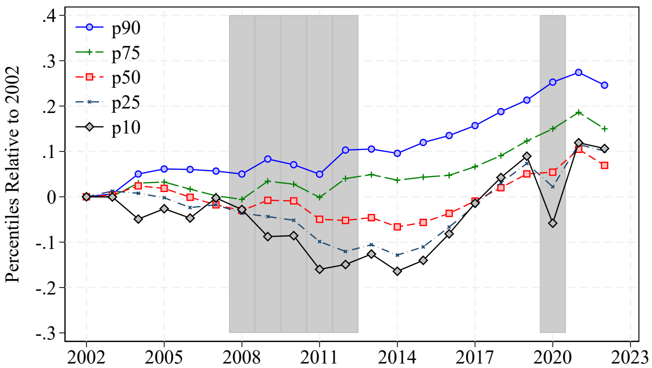

(a) Men - annual

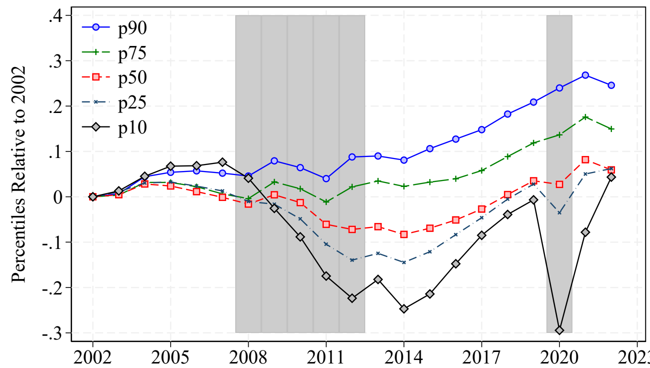

(b) Men - weekly

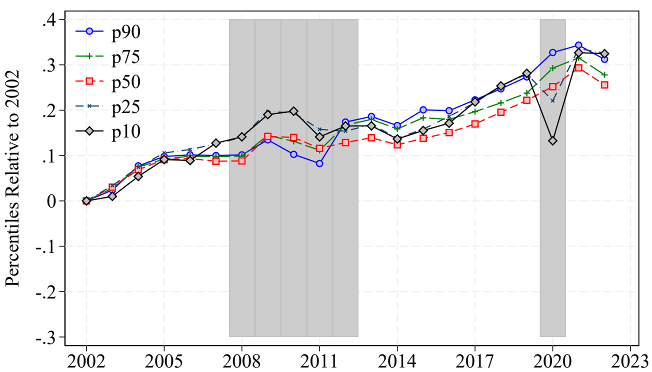

(c) Women - annual

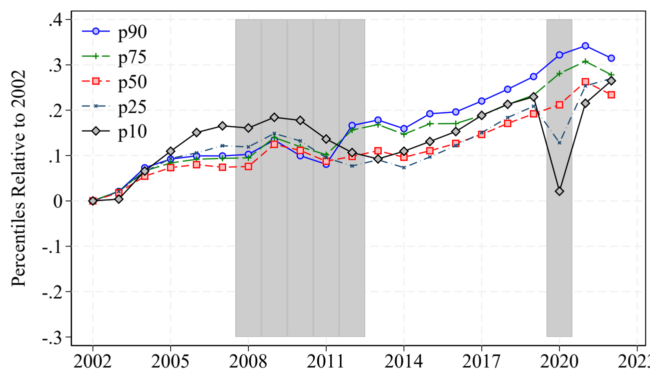

(d) Women - weekly

Figure D22: Annual vs Weekly Earnings

Note: Figure ranks workers using the relevant metric of earnings so the same individuals may not be at the same percentiles in each plot.

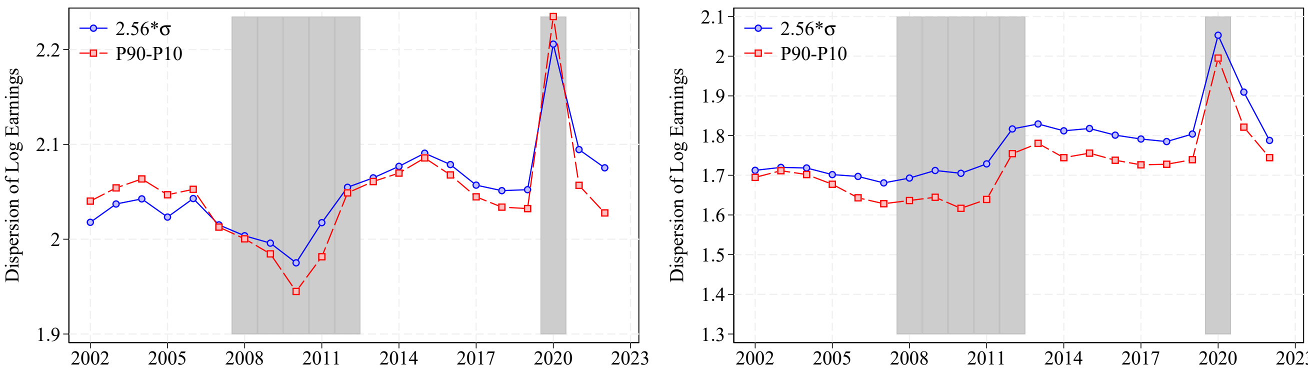

(a) Women - annual

(b) Women - weekly

Figure D23: Annual vs Weekly Earnings, Women, dispersion over time

Note: Figure ranks workers using the relevant metric of earnings so the same individuals may not be at the same percentiles in each plot.

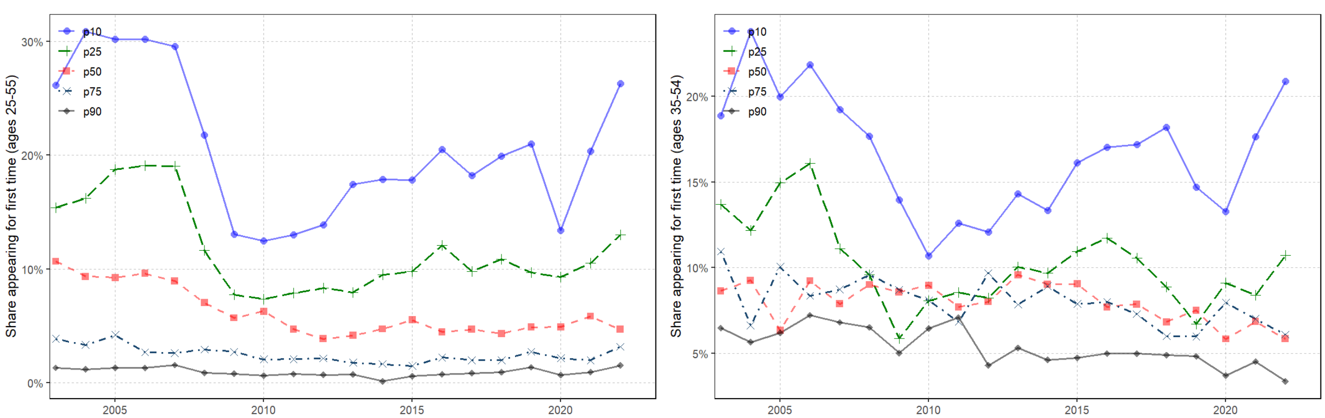

(a) New Entrants to WWLD Data - all

(b) New Entrants to WWLD Data- Age 35–54

Figure D24: Migrants and New Entrants to Data

Note: Left panel (A) shows the share of "new entrants" (first appearance in our administrative data) by percentile and year; right panel (B) shows the equivalent for new entrants aged 35 to 54.

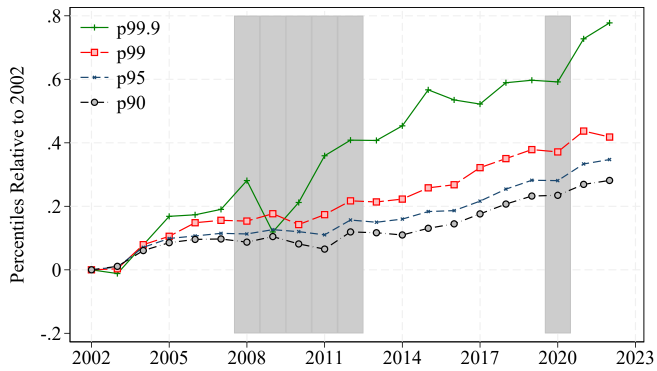

(a) Men, full year workers

(b) Men, survivors employed 2007-12

Figure D25: Selection (Men top percentiles)

Note: Log income is normalised to 0 in the base year 2002. Panels (A) and (B) show P90, P95, P99, P99.9. Shaded areas indicate recession years, based on declines in Modified Domestic Demand.

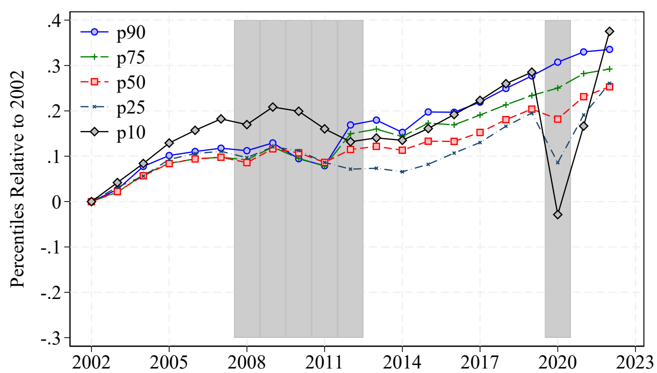

(a) Women, full year workers

(b) Women, survivors employed 2007-12

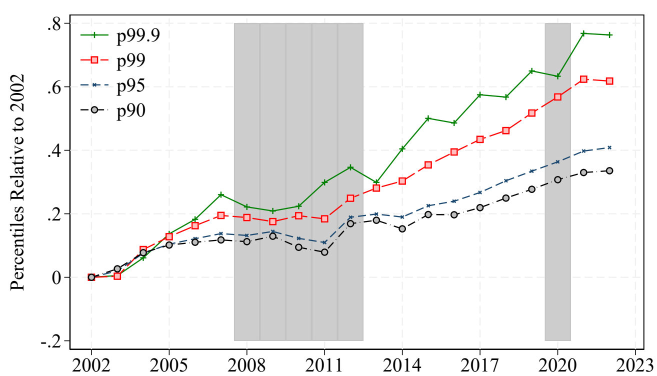

(c) Women, full year workers, top percentiles

(d) Women, survivors employed 2007-12, top

Figure D26: Selection (Female)

Note: Log income is normalised to 0 in the base year 2002. Top panels (A) and (B) show P10, P25, P50, P75, P90; Bottom panels (C) and (D) show P90, P95, P99, P99.9. Shaded areas indicate recession years, based on declines in Modified Domestic Demand.

D.2 Assessing Differences in Higher-Order Moments for Non-Migrants

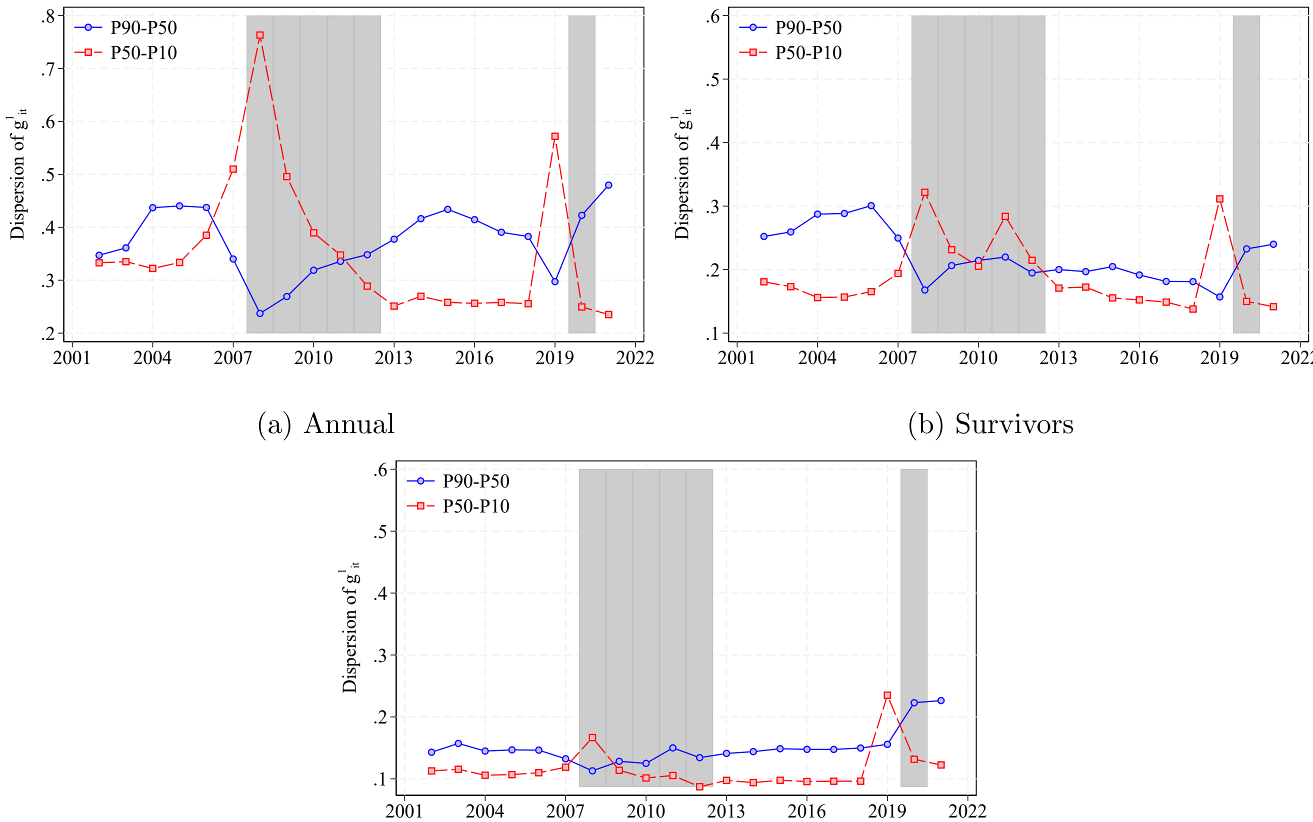

(a) Annual

(b) Survivors

(c) Full Year

Figure D27: Trends in Volatility for the main sample and sub-samples

Note: Volatility is the P90-P10 of earnings growth, skewness is Kelley skewness of earnings growth, and kurtosis is the 4th moment of the earnings growth distribution. Shaded areas indicate recession years, based on declines in Modified Domestic Demand.

E Role of Weeks Worked in Risk and Mobility

In this section, we introduce new evidence that changes in the number of weeks worked account for a large amount of earnings risk, as measured by the volatility, skewness and kurtosis of earnings growth.[41] Figure E29 compares the distribution of 1-year (residual) growth rates for annual earnings (left) and weekly earnings (right). It shows that changes in the number of weeks worked account for a large component of earnings risk in the Great Recession. The volatility (P90-P10) and kurtosis of earnings growth are almost half as large in terms of weekly earnings compared to annual earnings. The notable rise in earnings risk (Figure 8) was mostly driven by changes in the

[41]Figure E31 compares the densities.

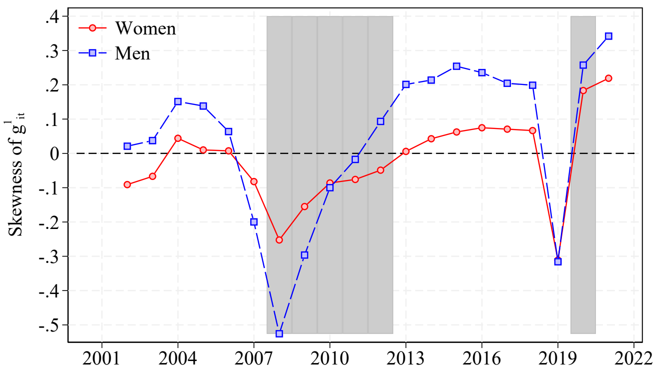

(a) Annual

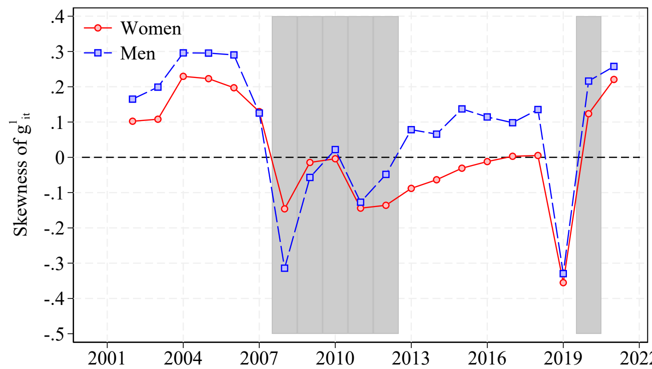

(b) Survivors

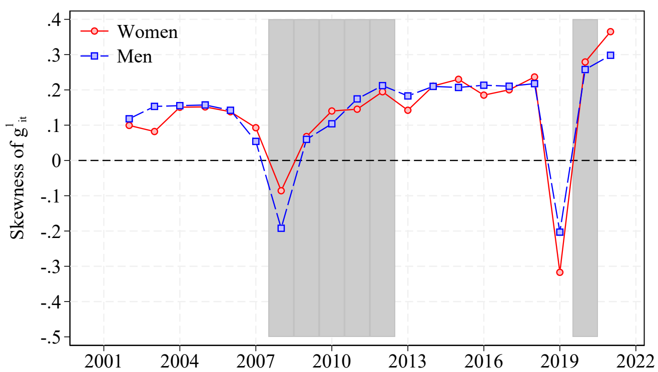

(c) Full Year

Figure D28: Trends in Skewness for the main sample and sub-samples

Note: Volatility is the P90-P10 of earnings growth, skewness is Kelley skewness of earnings growth, and kurtosis is the 4th moment of the earnings growth distribution. Shaded areas indicate recession years, based on declines in Modified Domestic Demand.

number of weeks worked, and thus is likely related to the patterns of inward migration documented in Section 4.

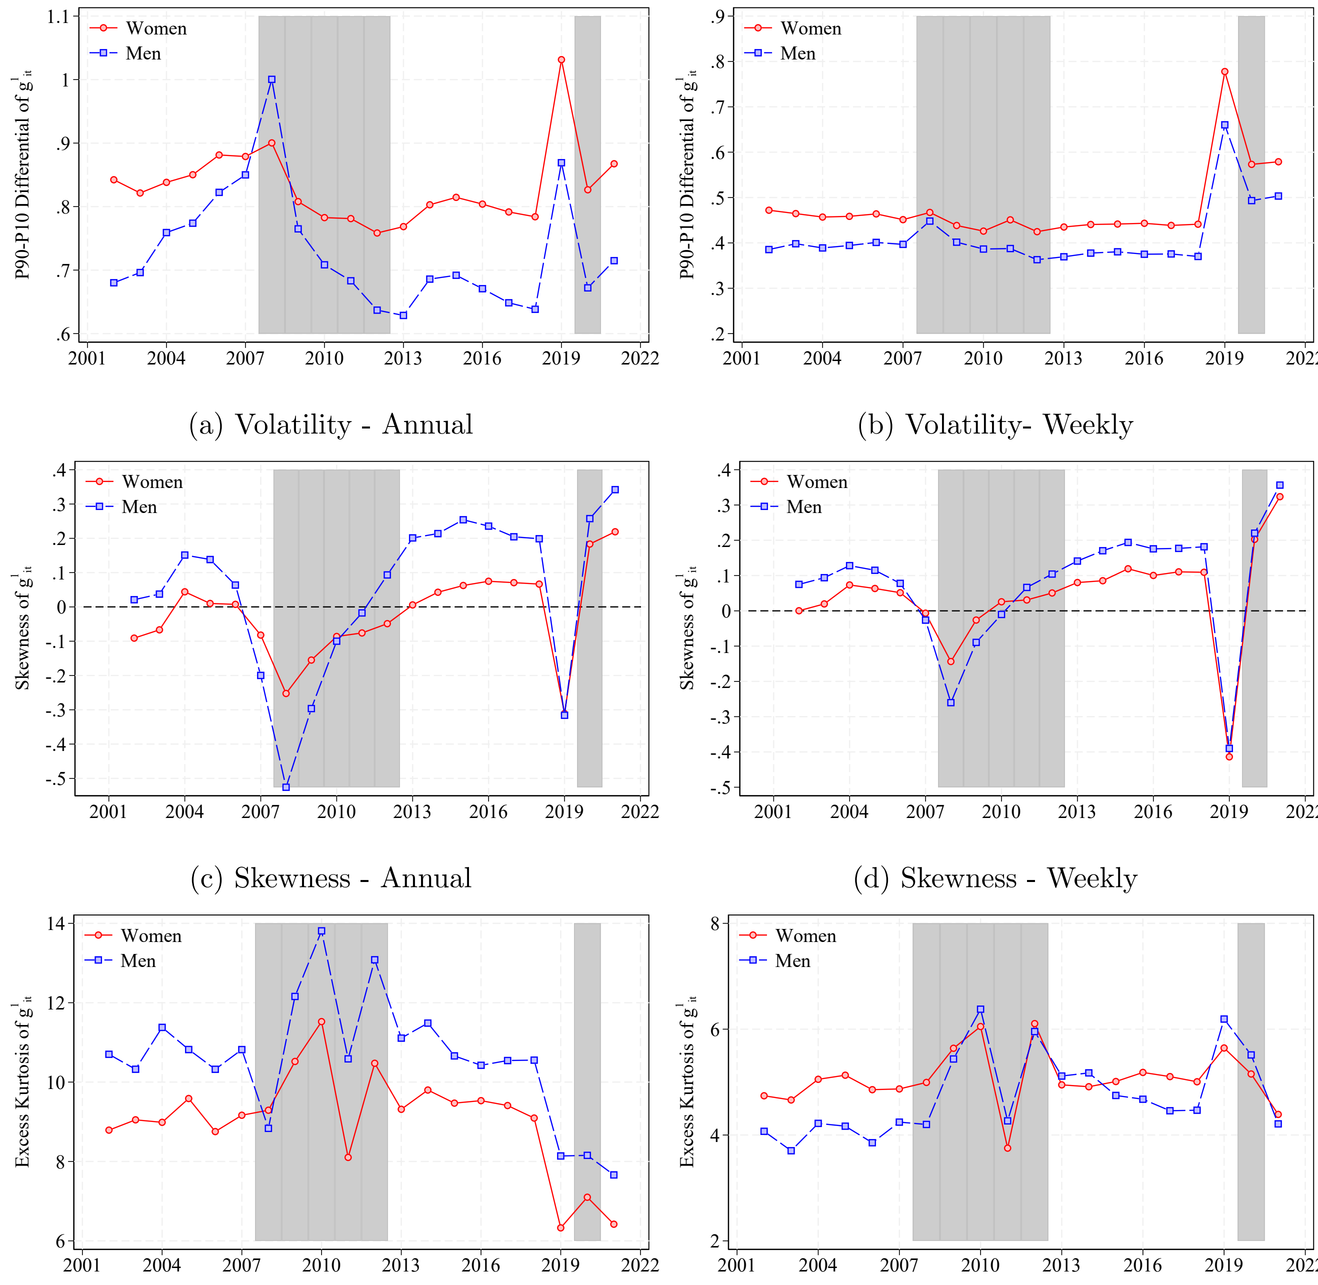

(a) Volatility - Annual

(b) Volatility- Weekly

(c) Skewness - Annual

(d) Skewness - Weekly

(e) Kurtosis - Annual

(f) Kurtosis - weekly

Figure E29: Trends in Volatility, Skewness and Kurtosis of Earnings Growth

Note: Volatility is the P90-P10 of earnings growth, skewness is Kelley skewness of earnings growth, and kurtosis is the 4th moment of the earnings growth distribution. Shaded areas indicate recession years, based on declines in Modified Domestic Demand.

There is also less skewness in weekly earnings compared to annual earnings. Busch et al.

(2022) argue that weeks worked do not change the fact that skewness is pro-cyclical. We also see that the skewness of weekly earnings continues to be pro-cyclical, but the absolute change is much smaller with the skewness of men's annual earnings falling to -.5 while weekly earnings fall to -0.25. Hoffmann and Malacrino (2019) also report that weekly earnings exhibit less cyclical skewness than annual earnings, though they conclude that the distribution of weekly earnings growth is close to symmetric and stable over the cycle. Figure E34 shows that skewness in weekly earnings is still associated with more negative changes and fewer positive changes.

Figure E29 also highlights a notable difference between the 2008 and Covid recessions. While much of the change in risk in the Great Recession was driven by changes in weeks worked, that does not appear to be the case for the Covid recession. There are two possibilities for this difference. First, the changes in demand during the Covid recession may have included more changes in hours by firms rather than contracted weeks. Second, some policy measures during the Covid recession subsidized firms to continue paying (part of) earnings; weeks worked may not have changed even while weekly earnings did fall.

Finally, Figure E30 shows novel evidence on how weekly earnings affect the relationship between permanent income and volatility of the 1-year earnings growth distribution. The hockey-stick relationship in the left-hand panel is a robust feature of the first wave of GRID. In every country, those at the bottom of the permanent income distribution have larger volatility in earnings growth going forward. Panel (b) shows that much of this volatility is coming from changes in the number of weeks worked at low permanent incomes. The pattern is flatter, with the higher volatility mostly occurring between the bottom and 10th percentile as opposed to the 30th percentile on the left. The magnitude is also smaller with about half the increase in volatility. This implies that higher volatility at the low end of the permanent income distribution is mostly driven by weeks worked, while at the very bottom there is still some volatility coming from changes in wages or hours within weeks.

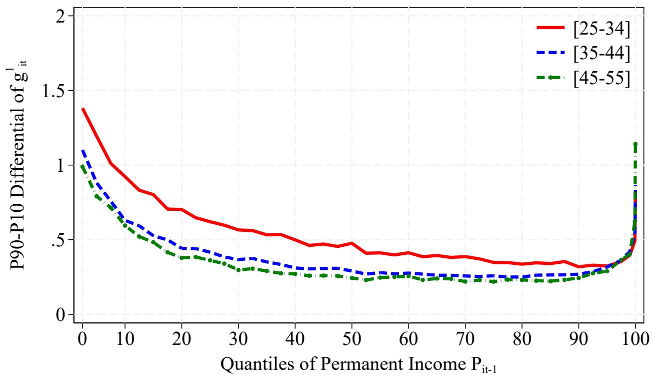

(a) Men - Annual

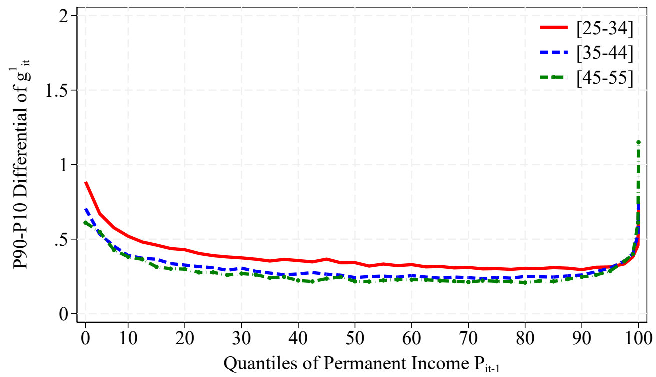

(b) Men - Weekly

Figure E30: Dispersion of earnings growth by permanent earnings and age

Note: Both panels show the relationship between permanent income and volatility of the 1-year earnings growth distribution for men. The left is annual earnings and the right is weekly earnings.

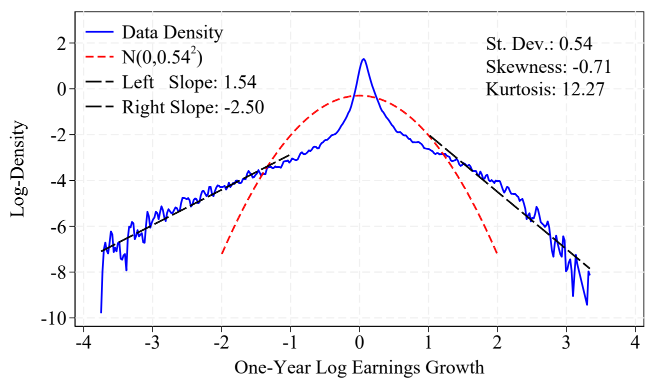

(a) Annual

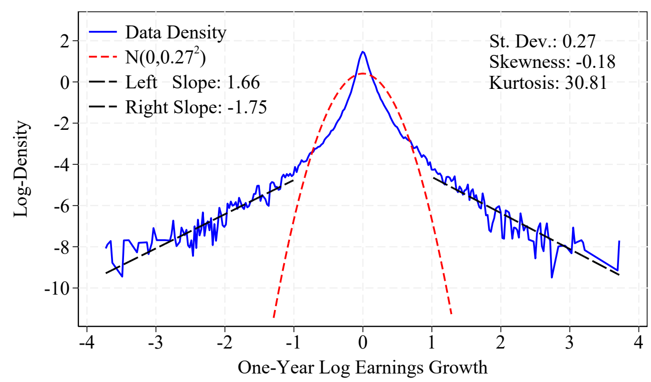

(b) Weekly

Figure E31: Comparing Weekly and Annual Earnings: Log-Densities of Earnings Growth

Note: Figure compares log density of one-year earnings growth using annual earnings (left) and weekly earnings (right).

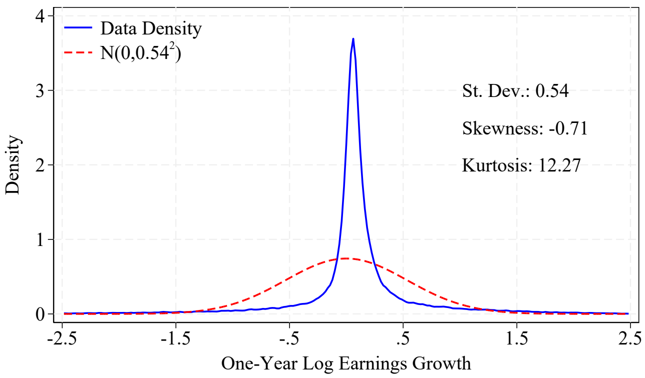

(a) Annual

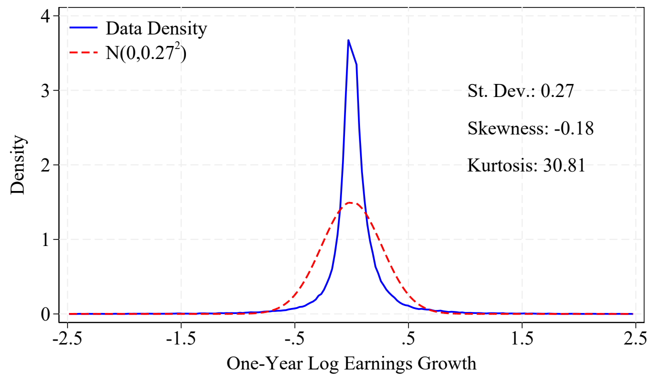

(b) Weekly

Figure E32: Comparing Weekly and Annual Earnings: Densities of Earnings Growth

Note: Figure compares density of one-year earnings growth using annual earnings (left) and weekly earnings (right).

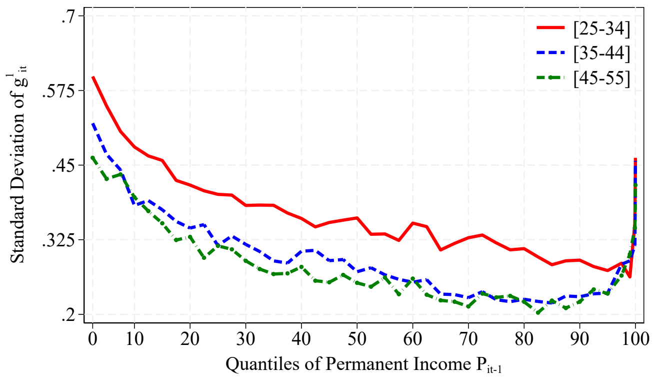

(a) Men

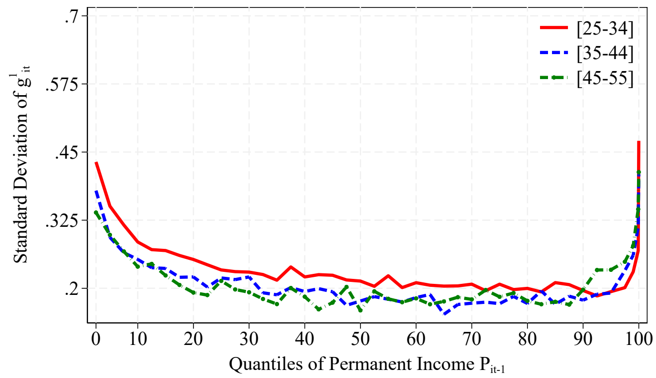

(b) Men - Weekly

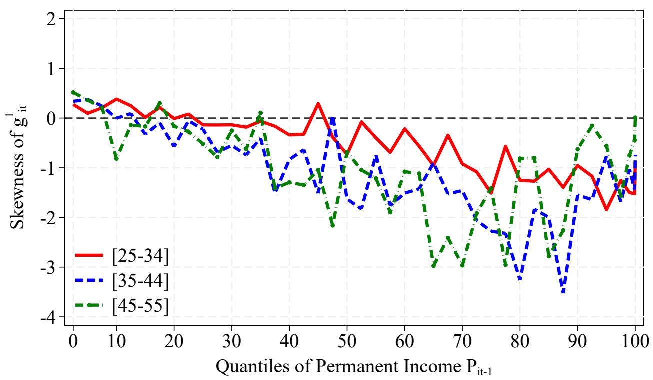

(c) Men

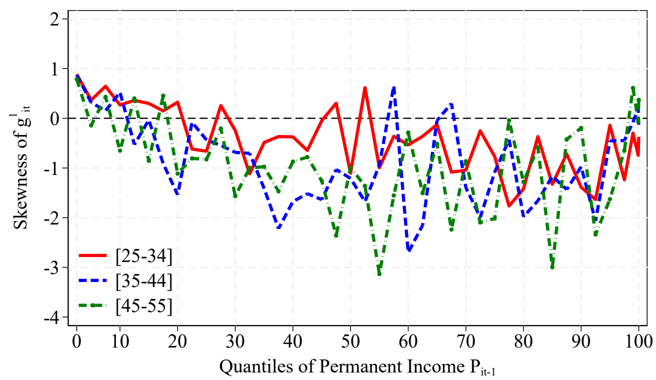

(d) Men - Weekly

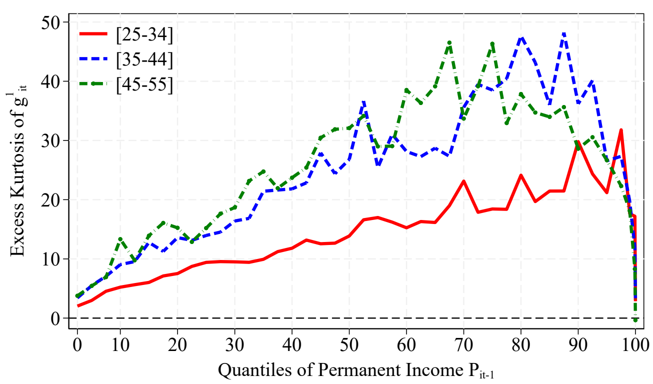

(e) Men

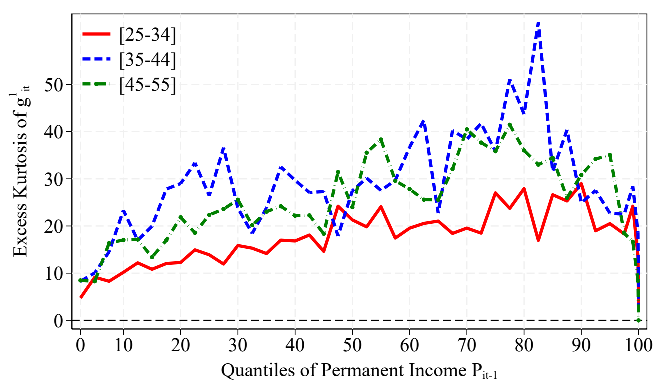

(f) Men - Weekly

Figure E33: Comparing Weekly and Annual Earnings: Earnings Growth by Permanent Earnings

Note: Figure compares metrics using annual earnings (left) and weekly earnings (right). Every panel shows the relationship between permanent income and moments of the 1-year earnings growth distribution. Top panels show the volatility of earnings growth; middle panels show the skewness of earnings growth; and bottom panels show the excess kurtosis of earnings growth.

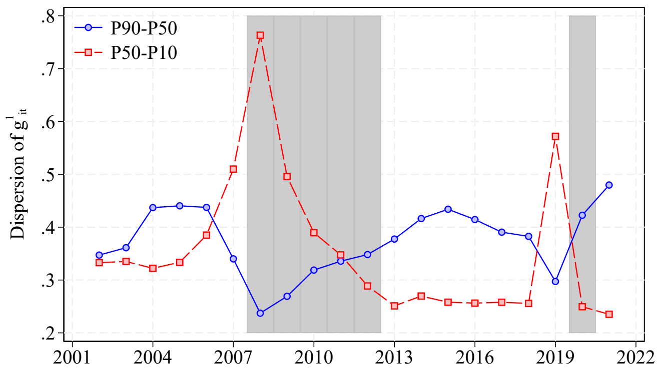

(a) Annual - Men

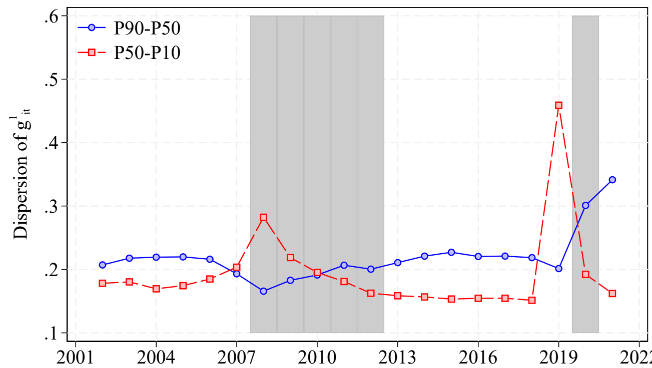

(b) Weekly - Men

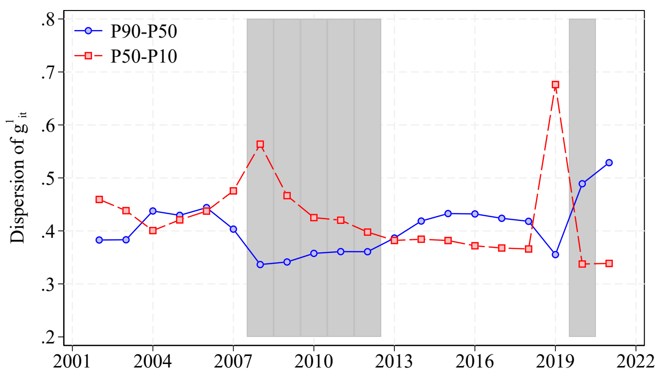

(c) Annual - Women

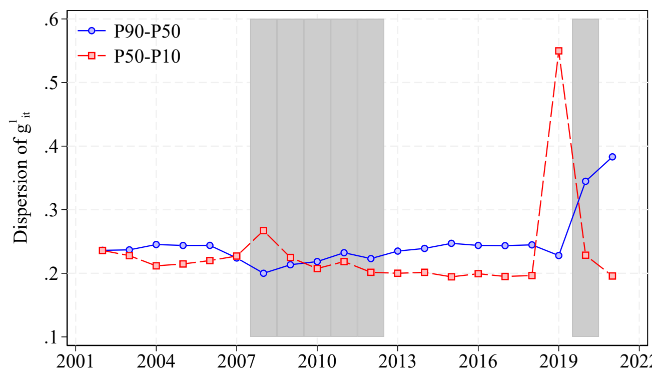

(d) Weekly - Women

Figure E34: Comparing Weekly and Annual: Good and Bad Shocks, and Skewness

Note: Figure decomposes skewness into positive shocks and negative shocks using both weekly earnings and annual earnings for men and women.

F EAADS Dataset

F.1 Data Description

Dataset. Our data comes from the Earnings Analysis using Administrative Data Sources (EAADS) dataset, which is an administrative panel dataset held by Ireland's Central Statistics Office (CSO). Earnings information in EAADS is based on tax records collected by the Revenue Commissioners — Ireland's tax authority. Records come from the universal tax return submitted by all registered employers. Until 2018, employers submitted the P35, an end-of-tax-year annual return form, while since 2019, records are based on real-time payroll reports for each payslip.

Timeframe. The data are available from 2011 to 2022.

Earnings Definition. Our measure of earnings is annual individual earnings before tax. This includes taxable bonuses and benefits-in-kind, and excludes pensions and severance payments.

Sampling Frame. The EAADS dataset is a 10 percent sample of employees, randomly sampled based on a digit from the unique personal public service (PPS) number. For each included employee we observe their entire employment history and observe earnings from all employments. The sample size ranges approximately from 170,000 to 230,000 individuals per year. After applying the sample restrictions required by GRID (described below), the sample is approximately between 120,000 and 160,000 per year.

The sample is top- and bottom-censored. At the bottom, employees earning less than €500 per annum, employments with duration less than two weeks, and secondary employments earning less than €4,000 are excluded. At the top, extremely high earnings values are dropped, as defined by earnings that are more than three times the inter-quartile range above the 75th percentile.

The sample includes both public and private sector employment. The sample excludes employments in NACE sectors A (Agriculture, forestry and fishing), T (Activities of households as employers) and U (Activities of extraterritorial organisations and bodies). In line with Eurostat requirements relating to Structure of Earnings Statistics the data used for this analysis has been restricted to employments that were active in the month of October.

GRID Covariates and Other Observables. We use 5-year age groups and gender as covariates to calculate residual earnings and for sub-group analysis. Unlike some of the other GRID countries, we do not observe education. We also observe employees' nationality and the county of residence of employers.

F.2 GRID Sample Restrictions

The previous description applies to all data available in the EAADS dataset. Next we describe our variable definitions and additional sample restrictions, where we do our best to be consistent with the other GRID countries, as described by Guvenen, Pistaferri and Violante (2022b).

We impose the same two sample restrictions imposed by other GRID countries. First, we focus on workers between 25 and 55 years old, a range within which most education choices are usually completed and after which workers tend to leave the labour force for retirement. Second, for most of the analysis we drop observations with earnings (defined next) below a threshold (call it $\underline{y}$) to avoid using records from workers without a meaningful attachment to the labour force or with very low earnings, which could skew log-based statistics. Specifically, we discard observations with earnings below what workers would earn if they were to work part-time for one quarter at the Irish national minimum wage.[42]

[42] Unlike countries in the first edition of GRID where labour income is top-coded (Italy and Mexico), we do not impute values of top-censored earnings. Our data is top-censored (i.e. set to missing) and not top-coded — Brazil featured similar top-censoring.

We construct the same three samples (CS, LX and H) as in the main text.

F.3 Variable Definitions

The variable definitions are the same as those in the main text.

F.4 Representativeness

In this section we show that the EAADS dataset is broadly representative of total employment and earnings. Figure F35 compares the time series for total earnings and total employment in national accounts with our estimate from aggregating EAADS. Relative to national accounts, the EAADS dataset captures 90 percent of total employment and 79 percent of total earnings in 2011. We can identify some differences in the sampling frame and definitions of earnings that may explain these discrepancies. In terms of employment, EAADS excludes (i) employment in agriculture, (ii) the top and bottom censoring noted above, and (iii) seasonal employment. The missing seasonal employment occurs because — in line with Eurostat standards — EAADS only includes employees working in October, and thus excludes any seasonal employments taking place during summer or over Christmas. In terms of earnings, the national accounts use a much broader definition of earnings. In addition to the income included in EAADS, the national accounts include employers' social insurance contributions (PRSI), defined contribution pension contributions actually made by employers, as well as imputed contributions for defined benefit schemes (including public sector schemes). Given the substantial differences in these definitions of earnings, we consider the 80 percent coverage rate of EAADS to be reasonably high.

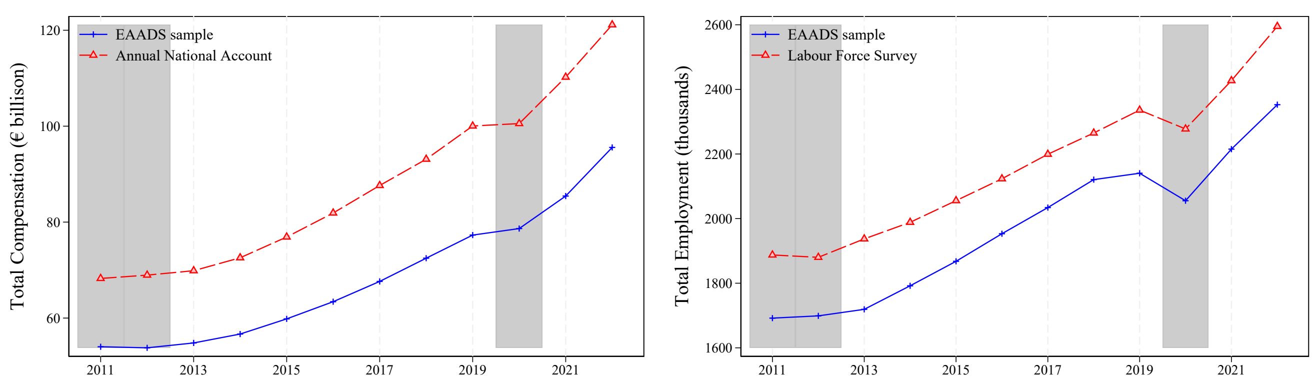

(a) Total Earnings

(b) Total Employment

Figure F35: Comparison of EAADS with Aggregate Statistics

Notes. Figure compares the coverage of the aggregated EAADS dataset with aggregate statistics. The left plot compares against total earnings in the national accounts. The right plot compares against total employment in the Labour Force Survey, which is the official measure of employment and unemployment in the state. See the text for a discussion of differences between the definitions in each data source.

While there are some differences in the *level* of coverage between EAADS and national accounts, the *trends* are almost identical. The shares of employment (90 percent) and earnings (79 percent) that are covered are identical at the beginning and end of the sample. Even during the fall in earnings that occurred during the 2020 Covid recession, EAADS captures the same share of employment and marginally less earnings (78 percent). This suggests that the trends identified in the GRID statistics are not affected by differential coverage over time.